

1 Intensity Frontier
2 Common Offline Documentation:
3 *art* Workbook and Users Guide

4 Alpha Release 0.90, working draft

5 June 2, 2016

6 This version of the documentation is written for version August2015 of the art-workbook
7 code.

8 Scientific Computing Division
9 Future Programs and Experiments Department
10 Scientific Software Infrastructure Group

11 Principal Author: Rob Kutschke

12 Other contributors: Marc Paterno, Mike Wang

13 Editor: Anne Heavey

14 *art* Developers: L. Garren, C. Green, K. Knoepfel,
15 J. Kowalkowski, M. Paterno and P. Russo

DRAFT

1 List of Chapters

2	Detailed Table of Contents	vii
3	List of Figures	xxvi
4	List of Tables	xxviii
5	List of Code and Output Listings	xxviii
6	I Introduction	2
7	1 How to Read this Documentation	3
8	2 Conventions Used in this Documentation	5
9	3 Introduction to the <i>art</i> Event Processing Framework	8
10	4 Unix Prerequisites	35
11	5 Site-Specific Setup Procedure	46
12	6 Get your C++ up to Speed	49

1	7 Using External Products in UPS	109
2	II Workbook	121
3	8 Preparation for Running the Workbook Exercises	122
4	9 Exercise 1: Running Pre-built <i>art</i> Modules	126
5	10 Exercise 2: Building and Running Your First Module	164
6	11 General Setup for Login Sessions	221
7	12 Keeping Up to Date with Workbook Code and Documentation	223
8	13 Exercise 3: Some other Member Functions of Modules	229
9	14 Exercise 4: A First Look at Parameter Sets	241
10	15 Exercise 5: Making Multiple Instances of a Module	267
11	16 Exercise 6: Accessing Data Products	274
12	17 Exercise 7: Making a Histogram	295
13	18 Exercise 8: Looping Over Collections	322
14	19 Exercise 9: Accessing and Using Particle Data	347
15	20 The <code>art::Ptr</code> Class Template	364
16	21 The Geometry Service	370

1	22 Instance Names of Data Products	371
2	23 InRun DataProducts	372
3	24 Provenance	373
4	25 Listing the Data Products in a File	374
5	26 Producer Module	375
6	27 Filter Module	376
7	28 Configuring Output Modules	377
8	29 Creating Your Own Data Products	378
9	30 Module Name Collisions	379
10	31 More FCL Concepts	380
11	32 art::Assns Connecting Hits to Intersections	381
12	33 Facade Pattern: Fitted Helices	383
13	34 art::View	384
14	35 Random Numbers	385
15	36 Repeating Random Numbers	386
16	37 Writing Your Own Service	387

1	38 Running the MC Chain	388
2	39 Reconstruction on Demand	389
3	40 Run the Existing 2D EventDisplay	390
4	41 Review of all Modules in toyExperiment	391
5	42 Running a Grid Job	392
6	43 Introducing SAM and dcache	393
7	44 3D Event Displays	394
8	45 Live Histogram Update	431
9	46 Checklist	432
10	47 Classes in the toyExperiment	436
11	48 Troubleshooting	440
12	III User's Guide	442
13	49 git	443
14	50 <i>art</i> Run-time and Development Environments	452
15	51 <i>art</i> Framework Parameters	460
16	52 Job Configuration in <i>art</i>: FHiCL	466

1	53 Data Products	486
2	54 Producer Modules	489
3	55 Analyzer Modules	490
4	56 Filter Modules	491
5	57 <i>art</i> Services	492
6	58 <i>art</i> Input and Output	501
7	59 <i>art</i> Misc Topics that Will Find Home	506
8	IV Appendices	521
9	A Obtaining Credentials to Access Fermilab Computing Resources	522
10	B Installing Locally	524
11	C <i>art</i> Completion Codes	529
12	D Viewing and Printing Figure Files	532
13	E CLHEP	534
14	F Include Guards	550

1	V Index	552
2	Index	553
3		

DRAFT

1 Detailed Table of Contents

2	Detailed Table of Contents	vii
3	List of Figures	xxvi
4	List of Tables	xxviii
5	List of Code and Output Listings	xxviii
6	I Introduction	2
7	1 How to Read this Documentation	3
8	1.1 If you are new to HEP Software...	3
9	1.2 If you are an HEP Software expert...	3
10	1.3 If you are somewhere in between...	4
11	2 Conventions Used in this Documentation	5
12	2.1 Terms in Glossary	5
13	2.2 Typing Commands	5
14	2.3 Listing Styles	6
15	2.4 Procedures to Follow	6
16	2.5 Important Items to Call Out	7
17	2.6 Site-specific Information	7
18	3 Introduction to the <i>art</i> Event Processing Framework	8
19	3.1 What is <i>art</i> and Who Uses it?	8
20	3.2 Why <i>art</i> ?	9

1	3.3	C++ and C++11	10
2	3.4	Getting Help	10
3	3.5	Overview of the Documentation Suite	10
4	3.5.1	The Introduction	12
5	3.5.2	The Workbook	12
6	3.5.3	Users Guide	13
7	3.5.4	Reference Manual	13
8	3.5.5	Technical Reference	13
9	3.5.6	Glossary	13
10	3.6	Some Background Material	13
11	3.6.1	Events and Event IDs	14
12	3.6.2	<i>art</i> Modules and the Event Loop	15
13	3.6.3	Module Types	19
14	3.6.4	<i>art</i> Data Products	20
15	3.6.5	<i>art</i> Services	21
16	3.6.6	Dynamic Libraries and <i>art</i>	23
17	3.6.7	Build Systems and <i>art</i>	23
18	3.6.8	External Products	24
19	3.6.9	The Event-Data Model and Persistency	26
20	3.6.10	Event-Data Files	27
21	3.6.11	Files on Tape	27
22	3.7	The Toy Experiment	28
23	3.7.1	Toy Detector Description	28
24	3.7.2	Workflow for Running the Toy Experiment Code	30
25	3.8	Rules, Best Practices, Conventions and Style	34
26	4	Unix Prerequisites	35
27	4.1	Introduction	35
28	4.2	Commands	35
29	4.3	Shells	37
30	4.4	Scripts: Part 1	37
31	4.5	Unix Environments	38
32	4.5.1	Building up the Environment	38
33	4.5.2	Examining and Using Environment Variables	40
34	4.6	Paths and \$PATH	40

1	4.7	Scripts: Part 2	42
2	4.8	bash Functions and Aliases	43
3	4.9	Login Scripts	44
4	4.10	Suggested Unix and bash References	44
5	5	Site-Specific Setup Procedure	46
6	6	Get your C++ up to Speed	49
7	6.1	Introduction	49
8	6.2	File Types Used and Generated in C++ Programming	50
9	6.3	Establishing the Environment	51
10	6.3.1	Initial Setup	51
11	6.3.2	Subsequent Logins	53
12	6.4	C++ Exercise 1: Basic C++ Syntax and Building an Executable	53
13	6.4.1	Concepts to Understand	53
14	6.4.2	How to Compile, Link and Run	54
15	6.4.3	Discussion	56
16	6.4.3.1	Primitive types, Initialization and Printing Output	56
17	6.4.3.2	Arrays	57
18	6.4.3.3	Equality testing	57
19	6.4.3.4	Conditionals	58
20	6.4.3.5	Some C++ Standard Library Types	58
21	6.4.3.6	Pointers	59
22	6.4.3.7	References	60
23	6.4.3.8	Loops	61
24	6.5	C++ Exercise 2: About Compiling and Linking	61
25	6.5.1	What You Will Learn	61
26	6.5.2	The Source Code for this Exercise	61
27	6.5.3	Compile, Link and Run the Exercise	63
28	6.5.4	Alternate Script <code>build2</code>	67
29	6.5.5	Suggested Homework	68
30	6.6	C++ Exercise 3: Libraries	70
31	6.6.1	What You Will Learn	70
32	6.6.2	Building and Running the Exercise	70
33	6.7	Classes	74

1	6.7.1	Introduction	74
2	6.7.2	C++ Exercise 4 v1: The Most Basic Version	76
3	6.7.3	C++ Exercise 4 v2: The Default Constructor	82
4	6.7.4	C++ Exercise 4 v3: Constructors with Arguments	84
5	6.7.5	C++ Exercise 4 v4: Colon Initializer Syntax	87
6	6.7.6	C++ Exercise 4 v5: Member functions	89
7	6.7.7	C++ Exercise 4 v6: Private Data and Accessor Methods	93
8	6.7.7.1	Setters and Getters	93
9	6.7.7.2	What's the deal with the underscore?	98
10	6.7.7.3	An example to motivate private data	99
11	6.7.8	C++ Exercise 4 v7: The <code>inline</code> Specifier	99
12	6.7.9	C++ Exercise 4 v8: Defining Member Functions within the Class Declaration	101
13	6.7.10	C++ Exercise 4 v9: The Stream Insertion Operator and Free Func- tions	102
14	6.7.11	Review	106
15	6.8	Overloading functions	107
16	6.9	C++ References	107
17			
18			
19	7	Using External Products in UPS	109
20	7.1	The UPS Database List: PRODUCTS	109
21	7.2	UPS Handling of Variants of a Product	111
22	7.3	The setup Command: Syntax and Function	111
23	7.4	Current Versions of Products	113
24	7.5	Environment Variables Defined by UPS	113
25	7.6	Finding Header Files	114
26	7.6.1	Introduction	114
27	7.6.2	Finding <i>art</i> Header Files	115
28	7.6.3	Finding Headers from Other UPS Products	116
29	7.6.4	Exceptions: The Workbook, ROOT and Geant4	117
30	II	Workbook	121
31	8	Preparation for Running the Workbook Exercises	122

1	8.1	Introduction	122
2	8.2	Getting Computer Accounts on Workbook-enabled Machines	122
3	8.3	Choosing a Machine and Logging In	123
4	8.4	Launching new Windows: Verify X Connectivity	124
5	8.5	Choose an Editor	125
6	9	Exercise 1: Running Pre-built <i>art</i> Modules	126
7	9.1	Introduction	126
8	9.2	Prerequisites	126
9	9.3	What You Will Learn	126
10	9.4	The <i>art</i> Run-time Environment	127
11	9.5	The Input and Configuration Files for the Workbook Exercises	128
12	9.6	Setting up to Run Exercise 1	129
13	9.6.1	Log In and Set Up	129
14	9.6.1.1	Initial Setup Procedure using Standard Directory	129
15	9.6.1.2	Initial Setup Procedure allowing Self-managed Working Directory	131
16	9.6.1.3	Setup for Subsequent Exercise 1 Login Sessions	132
17	9.7	Execute <i>art</i> and Examine Output	132
18	9.8	Understanding the Configuration	134
19	9.8.1	Some Bookkeeping Syntax	135
20	9.8.2	Some Physics Processing Syntax	136
21	9.8.3	<i>art</i> Command line Options	138
22	9.8.4	Maximum Number of Events to Process	138
23	9.8.5	Changing the Input Files	139
24	9.8.6	Skipping Events	141
25	9.8.7	Identifying the User Code to Execute	142
26	9.8.8	Paths and the <i>art</i> Workflow	144
27	9.8.8.1	Paths and the <i>art</i> Workflow: Details	146
28	9.8.8.2	Order of Module Execution	149
29	9.8.9	Writing an Output File	149
30	9.9	Understanding the Process for Exercise 1	151
31	9.9.1	Follow the Site-Specific Setup Procedure (Details)	152
32	9.9.2	Make a Working Directory (Details)	152
33	9.9.3	Setup the toyExperiment UPS Product (Details)	153
34			

1	9.9.4	Copy Files to your Current Working Directory (Details)	154
2	9.9.5	Source makeLinks.sh (Details)	154
3	9.9.6	Run <i>art</i> (Details)	155
4	9.10	How does <i>art</i> find Modules?	155
5	9.11	How does <i>art</i> find FHiCL Files?	157
6	9.11.1	The <code>-c</code> command line argument	157
7	9.11.2	<code>#include</code> Files	158
8	9.12	Review	158
9	9.13	Test your Understanding	159
10	9.13.1	Tests	159
11	9.13.2	Answers	162
12	10	Exercise 2: Building and Running Your First Module	164
13	10.1	Introduction	164
14	10.2	Prerequisites	165
15	10.3	What You Will Learn	166
16	10.4	Initial Setup to Run Exercises	167
17	10.4.1	“Source Window” Setup	167
18	10.4.2	Examine Source Window Setup	168
19	10.4.2.1	About git and What it Did	168
20	10.4.2.2	Contents of the Source Directory	170
21	10.4.3	“Build Window” Setup	171
22	10.4.3.1	Standard Procedure	171
23	10.4.3.2	Using Self-managed Working Directory	172
24	10.4.4	Examine Build Window Setup	173
25	10.5	The <i>art</i> Development Environment	177
26	10.6	Running the Exercise	180
27	10.6.1	Run <i>art</i> on <code>first.fcl</code>	180
28	10.6.2	The FHiCL File <code>first.fcl</code>	180
29	10.6.3	The Source Code File <code>First_module.cc</code>	182
30	10.6.3.1	The <code>#include</code> Statements	184
31	10.6.3.2	The Declaration of the Class <code>First</code> , an Analyzer Module	184
32	10.6.3.3	An Introduction to Analyzer Modules	186
33	10.6.3.4	The Constructor for the Class <code>First</code>	187

1	10.6.3.5	Aside: Omitting Argument Names in Function Declara-	
2		tions	188
3	10.6.3.6	The Member Function <code>analyze</code> and the Representa-	
4		tion of an Event	189
5	10.6.3.7	Representing an Event Identifier with <code>art::EventID</code>	190
6	10.6.3.8	<code>DEFINE_ART_MACRO</code> : The Module Maker Macros .	193
7	10.6.3.9	Some Alternate Styles	194
8	10.7	What does the Build System Do?	197
9	10.7.1	The Basic Operation	197
10	10.7.2	Incremental Builds and Complete Rebuilds	199
11	10.7.3	Finding Header Files at Compile Time	200
12	10.7.4	Finding Dynamic Library Files at Link Time	202
13	10.7.5	Build System Details	204
14	10.8	Suggested Activities	205
15	10.8.1	Create Your Second Module	205
16	10.8.2	Use <code>artmod</code> to Create Your Third Module	207
17	10.8.3	Running Many Modules at Once	209
18	10.8.4	Access Parts of the EventID	211
19	10.9	Final Remarks	212
20	10.9.1	Why is there no <code>First_module.h</code> File?	212
21	10.9.2	The Three-File Module Style	213
22	10.10	Flow of Execution from Source to FHiCL File	215
23	10.11	Review	216
24	10.12	Test Your Understanding	217
25	10.12.1	Tests	217
26	10.12.2	Answers	219
27	10.12.2.1	FirstBug01	219
28	10.12.2.2	FirstBug02	219
29	11	General Setup for Login Sessions	221
30	11.1	Source Window	221
31	11.2	Build Window	222
32	12	Keeping Up to Date with Workbook Code and Documentation	223
33	12.1	Introduction	223

1	12.2 Special Instructions for Summer 2014	223
2	12.3 How to Update	224
3	12.3.1 Get Updated Documentation	225
4	12.3.2 Get Updated Code and Build It	225
5	12.3.3 See which Files you have Modified or Added	227
6	13 Exercise 3: Some other Member Functions of Modules	229
7	13.1 Introduction	229
8	13.2 Prerequisites	230
9	13.3 What You Will Learn	230
10	13.4 Setting up to Run this Exercise	231
11	13.5 The Source File <code>Optional_module.cc</code>	231
12	13.5.1 About the <code>begin*</code> Member Functions	232
13	13.5.2 About the <code>art::*ID</code> Classes	232
14	13.5.3 Use of the <code>override</code> Identifier	233
15	13.5.4 Use of <code>const</code> References	233
16	13.5.5 The <code>analyze</code> Member Function	234
17	13.6 Running this Exercise	235
18	13.7 The Member Function <code>beginJob</code> versus the Constructor	236
19	13.8 Suggested Activities	237
20	13.8.1 Add the Matching <code>end</code> Member functions	237
21	13.8.2 Run on Multiple Input Files	237
22	13.8.3 The Option <code>--trace</code>	238
23	13.9 Review	238
24	13.10 Test Your Understanding	239
25	13.10.1 Tests	239
26	13.10.2 Answers	240
27	14 Exercise 4: A First Look at Parameter Sets	241
28	14.1 Introduction	241
29	14.2 Prerequisites	242
30	14.3 What You Will Learn	242
31	14.4 Setting up to Run this Exercise	243
32	14.5 The Configuration File <code>pset01.fcl</code>	244
33	14.6 The Source code file <code>PSet01_module.cc</code>	245

1	14.7 Running the Exercise	249
2	14.8 Member Function Templates and their Arguments	252
3	14.8.1 Types Known to <code>ParameterSet::get<T></code>	252
4	14.8.2 User-Defined Types	253
5	14.9 Exceptions (as in “Errors”)	253
6	14.9.1 Error Conditions	253
7	14.9.2 Error Handling	254
8	14.9.3 Suggested Exercises	255
9	14.10 Parameters and Data Members	256
10	14.11 Optional Parameters with Default Values	257
11	14.11.1 Policies About Optional Parameters	259
12	14.12 Numerical Types: Precision and Canonical Forms	259
13	14.12.1 Why Have Canonical Forms?	261
14	14.12.2 Suggested Exercises	261
15	14.12.2.1 Formats	261
16	14.12.2.2 Fractional versus Integral Types	262
17	14.13 Dealing with Invalid Parameter Values	262
18	14.14 Review	264
19	14.15 Test Your Understanding	265
20	14.15.1 Tests	265
21	14.15.2 Answers	266
22	15 Exercise 5: Making Multiple Instances of a Module	267
23	15.1 Introduction	267
24	15.2 Prerequisites	267
25	15.3 What You Will Learn	267
26	15.4 Setting up to Run this Exercise	268
27	15.5 The Source File <code>Magic_module.cc</code>	268
28	15.6 The FHiCL File <code>magic.fcl</code>	269
29	15.7 Running the Exercise	269
30	15.8 Discussion	270
31	15.8.1 Order of Analyzer Modules is not Important	270
32	15.8.2 Two Meanings of <i>Module Label</i>	271
33	15.9 Review	271
34	15.10 Test Your Understanding	272

1	15.10.1 Tests	272
2	15.10.2 Answers	273
3	16 Exercise 6: Accessing Data Products	274
4	16.1 Introduction	274
5	16.2 Prerequisites	274
6	16.3 What You Will Learn	275
7	16.4 Background Information for this Exercise	275
8	16.4.1 The Data Type <code>GenParticleCollection</code>	276
9	16.4.2 Data Product Names	277
10	16.4.3 Specifying a Data Product	279
11	16.4.4 The Data Product used in this Exercise	280
12	16.5 Setting up to Run this Exercise	280
13	16.6 Running the Exercise	281
14	16.7 Understanding the First Version, <code>ReadGens1</code>	281
15	16.7.1 The Source File <code>ReadGens1_module.cc</code>	281
16	16.7.2 Adding a Link Library to <code>CMakeLists.txt</code>	285
17	16.7.3 The FHiCL File <code>readGens1.fcl</code>	286
18	16.8 The Second Version, <code>ReadGens2</code>	286
19	16.9 The Third Version, <code>ReadGens3</code>	288
20	16.10 Suggested Activities	288
21	16.11 Review	291
22	16.12 Test Your Understanding	292
23	16.12.1 Tests	292
24	16.12.2 Answers	292
25	17 Exercise 7: Making a Histogram	295
26	17.1 Introduction	295
27	17.2 Prerequisites	296
28	17.3 What You Will Learn	296
29	17.4 Setting up to Run this Exercise	297
30	17.5 The Source File <code>FirstHist1_module.cc</code>	298
31	17.5.1 Introducing <code>art::ServiceHandle</code>	300
32	17.5.2 Creating a Histogram	302
33	17.5.3 Filling a Histogram	304

1	17.5.4 A Few Last Comments	304
2	17.6 The Configuration File <code>firstHist1.fcl</code>	305
3	17.7 The file <code>CMakeLists.txt</code>	305
4	17.8 Running the Exercise	307
5	17.9 Inspecting the Histogram File	308
6	17.9.1 A Short Cut: the browse command	310
7	17.9.2 Using CINT Scripts	312
8	17.10 Finding ROOT Documentation	316
9	17.10.1 Overwriting Histogram Files	316
10	17.10.2 Changing the Name of the Histogram File	317
11	17.10.3 Changing the Module Label	317
12	17.10.4 Printing From the TBrowse	317
13	17.11 Review	318
14	17.12 Test Your Understanding	318
15	17.12.1 Tests	318
16	17.12.2 Answers	320
17	18 Exercise 8: Looping Over Collections	322
18	18.1 Introduction	322
19	18.2 Prerequisites	322
20	18.3 What You Will Learn	322
21	18.4 Setting Up to Run Exercise	323
22	18.5 The Class <code>GenParticle</code>	324
23	18.5.1 The Included Header Files	324
24	18.5.2 Particle Parent-Child Relationships	326
25	18.5.3 The Public Interface for the Class <code>GenParticle</code>	326
26	18.5.4 Conditionally Excluded Sections of Header File	330
27	18.6 The Module <code>LoopGens1</code>	330
28	18.7 <code>CMakeLists.txt</code>	333
29	18.8 Running the Exercise	333
30	18.9 Variations on the Exercise	335
31	18.9.1 <code>LoopGens2_module.cc</code>	335
32	18.9.2 <code>LoopGens3_module.cc</code>	336
33	18.9.3 <code>LoopGens3a_module.cc</code>	339
34	18.10 Review	340

1	18.11 Test Your Understanding	341
2	18.11.1 Test 1	341
3	18.11.2 Test 2	343
4	18.11.3 Test 3	344
5	18.11.4 Answers	345
6	18.11.4.1 Test 1	345
7	18.11.4.2 Test 2	345
8	18.11.4.3 Test 3	346
9	19 Exercise 9: Accessing and Using Particle Data	347
10	19.1 Introduction	347
11	19.2 Prerequisites	349
12	19.3 What You Will Learn	349
13	19.4 Setting Up to Run Exercise	350
14	19.5 The Particle Data Table Service	351
15	19.5.1 The PDT FHiCL File	351
16	19.5.2 The PDT Service	352
17	19.5.3 Using the PDT Service	353
18	19.5.4 Configuring the PDT Service	354
19	19.5.5 The Class <code>ParticleInfo</code>	355
20	19.6 The File <code>CountGen1_module.cc</code>	356
21	19.7 Running the Exercise	359
22	19.8 The Class <code>IdCounter</code>	360
23	19.9 Discussion	362
24	19.10 Suggested Activities	363
25	20 The <code>art::Ptr</code> Class Template	364
26	20.1 Introduction	364
27	20.2 Prerequisites	367
28	20.3 What You Will Learn	367
29	20.4 Rough Notes	367
30	20.5 Running the Exercise	369
31	20.6 Discussion	369
32	20.7 Test Your Understanding	369
33	20.8 Review	369

1	21 The Geometry Service	370
2	21.1 Prerequisites	370
3	21.2 What You Will Learn	370
4	21.3 Running the Exercise	370
5	21.4 Discussion	370
6	21.5 Suggested Activities	370
7	22 Instance Names of Data Products	371
8	23 InRun DataProducts	372
9	24 Provenance	373
10	25 Listing the Data Products in a File	374
11	26 Producer Module	375
12	27 Filter Module	376
13	28 Configuring Output Modules	377
14	29 Creating Your Own Data Products	378
15	30 Module Name Collisions	379
16	31 More FCL Concepts	380
17	32 art::Assns Connecting Hits to Intersections	381
18	33 Facade Pattern: Fitted Helices	383
19	34 art::View	384
20	35 Random Numbers	385
21	36 Repeating Random Numbers	386
22	37 Writing Your Own Service	387

1	38 Running the MC Chain	388
2	39 Reconstruction on Demand	389
3	40 Run the Existing 2D EventDisplay	390
4	41 Review of all Modules in toyExperiment	391
5	42 Running a Grid Job	392
6	43 Introducing SAM and dcache	393
7	44 3D Event Displays	394
8	44.1 Introduction	394
9	44.2 Prerequisites	395
10	44.3 What You Will Learn	395
11	44.4 Setting up to Run this Exercise	395
12	44.5 Running the Exercise	396
13	44.5.1 Startup and General Layout	396
14	44.5.2 The Control Panel	396
15	44.5.2.1 The List-Tree Widget and Context-Sensitive Menus . .	396
16	44.5.2.2 The Event-Navigation Pane	403
17	44.5.3 Main EVE Display Area	404
18	44.6 Understanding How the 3D Event Display Module Works	406
19	44.6.1 Overview of the Source Code File <code>EventDisplay3D_module.cc</code>	407
20	44.6.2 Class Declaration and Constructor	408
21	44.6.3 Creating the GUI and Drawing the Static Detector Components	
22	in the <code>beginJob()</code> Member Function	412
23	44.6.3.1 The Default GUI	412
24	44.6.3.2 Adding the Global Elements	413
25	44.6.3.3 Customizing the GUI	413
26	44.6.3.4 Adding the Navigation Pane	419
27	44.6.4 Drawing the Generated Hits and Tracks in the <code>analyze()</code> Mem-	
28	ber Function	423
29	45 Live Histogram Update	431

1	46 Checklist	432
2	47 Classes in the toyExperiment	436
3	48 Troubleshooting	440
4	48.1 Updating Workbook Code	440
5	48.2 XWindows (xterm and Other XWindows Products)	440
6	48.2.1 Mac OSX 10.9	440
7	48.3 Trouble Building	440
8	48.4 <i>art</i> Won't Run	441
9	III User's Guide	442
10	49 git	443
11	49.1 Aside: More Details about git	444
12	49.1.1 Central Repository, Local Repository and Working Directory	444
13	49.1.1.1 Files that you have Added	445
14	49.1.1.2 Files that you have Modified	445
15	49.1.1.3 Files with Resolvable Conflicts	446
16	49.1.1.4 Files with Unresolvable Conflicts	446
17	49.1.2 git Branches	447
18	49.1.3 Seeing which Files you have Modified or Added	451
19	50 <i>art</i> Run-time and Development Environments	452
20	50.1 The <i>art</i> Run-time Environment	452
21	50.2 The <i>art</i> Development Environment	456
22	51 <i>art</i> Framework Parameters	460
23	51.1 Parameter Types	460
24	51.2 Structure of <i>art</i> Configuration Files	461
25	51.3 Services	464
26	51.3.1 System Services	464
27	51.3.2 FloatingPointControl	464
28	51.3.3 Message Parameters	464
29	51.3.4 Optional Services	464

1	51.3.5 Sources	464
2	51.3.6 Modules	464
3	52 Job Configuration in <i>art</i>: FHiCL	466
4	52.1 Basics of FHiCL Syntax	467
5	52.1.1 Specifying Names and Values	467
6	52.1.2 FHiCL-reserved Characters and Identifiers	469
7	52.2 FHiCL Identifiers Reserved to <i>art</i>	470
8	52.3 Structure of a FHiCL Run-time Configuration File for <i>art</i>	472
9	52.4 Order of Elements in a FHiCL Run-time Configuration File for <i>art</i>	475
10	52.5 The <i>physics</i> Portion of the FHiCL Configuration	477
11	52.6 Choosing and Using Module Labels and Path Names	479
12	52.7 Scheduling Strategy in <i>art</i>	480
13	52.8 Scheduled Reconstruction using Trigger Paths	482
14	52.9 Reconstruction On-Demand	484
15	52.10 Bits and Pieces FIXME:	485
16	53 Data Products	486
17	53.1 Overview	486
18	53.2 The Full Name of a Data Product	487
19	54 Producer Modules	489
20	55 Analyzer Modules	490
21	56 Filter Modules	491
22	57 <i>art</i> Services	492
23	57.1 About Services	492
24	57.2 Service Handles	494
25	57.3 Implementing Simple Services	495
26	57.4 Configuring a Service	495
27	57.5 Accessing a Service	496
28	57.6 Writing a Service	497
29	57.6.1 Declaring and Defining Services	498
30	57.7 Service Interfaces	499

1	58 <i>art</i> Input and Output	501
2	58.1 Input Modules	501
3	58.1.1 Configuring Input Modules to Read from Files	501
4	58.2 Output Filtering	504
5	58.3 Configuring Output Modules	505
6	59 <i>art</i> Misc Topics that Will Find Home	506
7	59.0.1 The Bookkeeping Structure and Event Sequencing Imposed by <i>art</i>	506
8	59.1 Rules for Module Names	508
9	59.2 Data Products and the Event Data Model	510
10	59.3 Basic <i>art</i> Rules	510
11	59.4 Compiling, Linking, Loading and Executing C++ Classes and <i>art</i> Modules	511
12	59.5 Dynamic Libraries and <i>art</i>	514
13	59.6 Namespaces, <i>art</i> and the Workbook	514
14	59.7 Orphans FIXME:	515
15	59.8 Inheritance	516
16	59.8.1 Introduction	516
17	59.8.2 Homework	517
18	59.8.3 Discussion	518
19	59.9 Inheritance Relic	518
20	59.10 Pointers	519
21	59.11 RootOutput and table of event IDs	519
22	59.12 Troubleshooting	520
23	IV Appendices	521
24	A Obtaining Credentials to Access Fermilab Computing Resources	522
25	A.1 Kerberos Authentication	522
26	A.2 Fermilab Services Account	523
27	B Installing Locally	524
28	B.1 Install the Binary Distributions: A Cheat Sheet	524
29	B.2 Preparing the Site Specific Setup Script	526
30	B.3 Links to the Full Instructions	528

1	C	art Completion Codes	529
2	D	Viewing and Printing Figure Files	532
3		D.1 Viewing Figure Files Interactively	532
4		D.2 Printing Figure Files	533
5	E	CLHEP	534
6		E.1 Introduction	534
7		E.2 Multiple Meanings of <i>Vector</i> in CLHEP	535
8		E.3 CLHEP Documentation	535
9		E.4 CLHEP Header Files	536
10		E.4.1 Naming Conventions and Syntax	536
11		E.4.2 <code>.icc</code> Files	536
12		E.5 The CLHEP Namespace	537
13		E.5.1 <code>using</code> Declarations and Directives	537
14		E.6 The Vector Package	538
15		E.6.1 <code>CLHEP::Hep3Vector</code>	539
16		E.6.1.1 Some Fragile Member Functions	543
17		E.6.2 <code>CLHEP::HepLorentzVector</code>	544
18		E.6.2.1 <code>HepBoost</code>	548
19		E.7 The Matrix Package	548
20		E.8 The Random Package	549
21	F	Include Guards	550
22	V	Index	552
23		Index	553

DRAFT

1 List of Figures

2	3.1	Principal components of the <i>art</i> documentation suite	11
3	3.2	Flowchart describing the <i>art</i> event loop	18
4	3.3	Geometry of the toy experiment's detector	28
5	3.4	Event display of a simulated event in the toy detector.	31
6	3.5	Event display of another simulated event in the toy detector	32
7	3.6	Invariant mass of reconstructed pairs of oppositely charged tracks	33
8	4.1	Computing environment hierarchies	39
9	6.1	Memory diagram at the end of a run of <code>Classes/v1/ptest.cc</code>	81
10	6.2	Memory diagram at the end of a run of <code>Classes/v6/ptest.cc</code>	97
11	9.1	Elements of the <i>art</i> run-time environment for the first exercise	128
12	10.1	Representation of reader's source directory structure	170
13	10.2	Representation of reader's build directory structure	174
14	10.3	Elements of the <i>art</i> development environment and information flow	178
15	10.4	Reader's directory structure once development environment is established	181
16	17.1	TBrowser window after opening <code>output/firstHist1.root</code>	311
17	17.2	TBrowser window after displaying the histogram <code>hNGens;1</code>	311
18	17.3	Figure made by running the CINT script <code>drawHist1.C</code>	314
19	18.1	Histograms made by <code>loopGens1.fcl</code>	334
20	19.1	Number of each type of Particle in the input stream (from <code>countGens1.C</code>)348	

1	44.1	The TEveBrowser is a specialization of the ROOT TBrowser for the	
2		ROOT EVE Event Visualization Environment. Shown above is the one	
3		used in this workbook exercise which is divided into three major regions:	
4		1) a control panel, 2) a main EVE display area, and 3) a ROOT command	
5		console.	397
6	44.2	Shown above are two different views of the list-tree widget, showing the	
7		top-level items in (a), and expanded to second-level items in (b).	398
8	44.3	Expanded views of the (a) <i>WindowManager</i> and (b) <i>Viewers</i> list-tree items.	399
9	44.4	Expanded views of the (a) <i>Scenes</i> and (b) <i>Event</i> list-tree items.	400
10	44.5	The context-sensitive menu below the list-tree widget changes in response	
11		to the selected list-tree item. Shown above is the menu for a viewer-type	
12		item. In this example, we have enabled an interactive clipping plane.	402
13	44.6	Shown above is the context-sensitive menu displayed below the list-tree	
14		widget when a track element is selected.	403
15	44.7	The <i>Event Nav</i> pane on the control panel.	404
16	44.8	The orthographic XY and RZ views in the <i>Ortho Views</i> tabbed pane of	
17		the main EVE display panel.	405
18	44.9	Hovering the mouse cursor over the lower edge of the title bar in a viewport	
19		reveals a pull-down menu bar with more options.	406
20	44.10	Tooltips with relevant information show up when the mouse cursor is	
21		hovered over (a) track and (b) hit elements	407
22	49.1	Illustration of git branches, simple	447
23	49.2	Illustration of git branches FIXME: needs descrip	450
24	50.1	<i>art</i> run-time environment (same as Figure 9.1) <small>fig:runtime-env-x</small>	453
25	50.2	<i>art</i> run-time environment (everything pre-built)	454
26	50.3	<i>art</i> run-time environment (with officially tracked inputs)	455
27	50.4	<i>art</i> development environment for Workbook (same as Figure 10.3) <small>fig:dev-wkbb</small>	457
28	50.5	<i>art</i> development environment (for building full code base)	458
29	50.6	<i>art</i> development environment (for building against prebuilt base)	459
30	59.1	Illustration of “regular” C++ classes used within <i>art</i> framework	512
31	59.2	Illustration of <i>art</i> modules, each built into a single dynamic library	513

1 List of Tables

2	3.1	Compiler flags for the optimization levels defined by cetbuildtools	24
3	3.2	Units used in the Workbook	29
4	5.1	Site-specific setup procedures for experiments that run <i>art</i>	48
5	7.1	Namespaces for selected UPS products	118
6	8.1	Experiment-specific Information for new users	123
7	8.2	Login machines for running the Workbook exercises	124
8	9.1	Input files provided for the Workbook exercises	129
9	10.1	Compiler and linker flags for a profile build	205
10	14.1	Canonical forms of numerical values in FHiCL files	260
11	51.1	<i>art</i> Floating Point Parameters	465
12	51.2	<i>art</i> Message Parameters	465
13	C.1	<i>art</i> completion status codes. The return code is the least significant byte	
14		of the status code.	531
15	E.1	Selected member functions of <code>CLHEP::Hep3Vector</code>	542

16

1 List of Code and Output Listings

2	5.1	NOvA setup procedure	47
3	6.1	Layout of a class.	75
4	6.2	File <code>Point.h</code> with the simplest version of the class <code>Point</code>	78
5	6.3	The contents of <code>v1/ptest.cc</code>	79
6	9.1	Sample output from running <code>hello.fcl</code>	132
7	9.2	Listing of <code>hello.fcl</code>	134
8	9.3	Example of the value of <code>LD_LIBRARY_PATH</code>	155
9	10.1	Example of output created by <code>setup_for_development</code>	175
10	10.2	The contents of <code>First_module.cc</code>	183
11	10.3	An alternate layout for <code>First_module.cc</code>	196
12	10.4	The file <code>art-workbook/FirstModule/CMakeLists.txt</code>	204
13	10.5	The physics parameter set for <code>all.fcl</code>	210
14	10.6	The contents of <code>First.h</code> in the three-file model	214
15	10.7	The contents of <code>First.cc</code> in the three-file model	214
16	10.8	The contents of <code>First_module.cc</code> in the three-file model	215
17	12.1	Example of the output produced by <code>git pull</code>	226
18	13.1	Output from <code>Optional_module.cc</code> with <code>optional.fcl</code>	235
19	14.1	Parameter set <code>psetTester</code> from <code>pset01.fcl</code>	244
20	14.2	First part of constructor in <code>PSet01_module.cc</code>	245
21	14.3	Part 1 of the remainder of the constructor in <code>PSet01_module.cc</code> . The values of <code>pset</code> are printed out in a standard way. Two other ways of printing the values <code>a</code> and <code>b</code> from the parameter set <code>f</code> (lines 13 and 14) will also be shown.	248

1	14.4	Part 2 of the remainder of the constructor in <code>PSet01_module.cc</code> .	
2		These lines use the <code>to_string()</code> and <code>to_indented_string()</code>	
3		member functions of the class <code>fhicl::ParameterSet</code> to print the	
4		values <code>a</code> and <code>b</code> from the parameter set <code>f</code>	248
5	14.5	Part 3 of the remainder of the constructor in <code>PSet01_module.cc</code> .	
6		This portion uses the <code>to_indented_string()</code> member function to	
7		print everything found in the parameter set <code>psetTester</code> , including <code>a</code>	
8		and <code>b</code> from the parameter set <code>f</code>	249
9	14.6	Output from <code>PSet01</code> with <code>pset01.fcl</code> (<i>art</i> -standard output not shown)	251
10	14.7	Output from <code>PSet01</code> with <code>pset01.fcl</code> (parameter <code>b</code> removed) . . .	255
11	14.8	Output from <code>PSet01</code> with <code>pset01.fcl</code> (parameter <code>c</code> misdefined) . .	256
12	14.9	Output from <code>PSet02</code> with <code>pset02.fcl</code>	257
13	14.10	Parameter-related portion of output from <code>PSet03</code> with <code>pset03.fcl</code> .	258
14	14.11	Output from <code>PSet04</code> with <code>pset04.fcl</code>	262
15	14.12	Output from <code>PSet04</code> with modified <code>pset04.fcl</code> (intentional error) .	263
16	15.1	Output using <code>magic.fcl</code>	270
17	16.1	Contents of <code>GenParticleCollection.h</code>	276
18	16.2	Output using <code>readGens1.fcl</code>	281
19	16.3	Include statements in <code>ReadGens1_module.cc</code>	282
20		<code>readgens1</code>	282
21		<code>readgens1</code>	282
22	16.4	Configuring the module label <code>read</code> in <code>readGens1.fcl</code>	286
23	16.5	Warning message for misspelled module label of data product	289
24	16.6	Exception message for <code>ProductNotFound</code> , default Exceptions disabled	291
25	17.1	Top portion of <code>FirstHist1_module.cc</code>	299
26	17.2	Implementation of the class <code>FirstHist1</code>	301
27	17.3	<code>firstHist1.fcl</code>	306
28	17.4	<code>CMakeLists.tex</code> in the directory <code>FirstHistogram</code>	308
29	17.5	Sample CINT file <code>DrawHist1.C</code>	313
30	18.1	The const accessors of class <code>GenParticle</code> . The order has been changed	
31		from that found in the header file to match the order of discussion in the	
32		text. The white space has also been altered to make the code fit better	
33		in the narrow page of this document.	328
34	18.2	The loop over the generated particles, from <code>LoopGens1_module.cc</code>	331
35		<code>f</code>	332

1	19.1 Particle data for this exercise; output from <code>countGens1.fcl</code>	348
2	19.2 Contents of <code>pdt.fcl</code>	351
3	19.3 FHiCL fragment to configure the PDT service	354
4	19.4 The contents of <code>output/countGens1.txt</code>	358
5	44.1 The declaration of the class <code>EventDisplay3D</code> from <code>EventDisplay3D_module.cc</code>	409
6	44.2 Implementation of the constructor for the <code>EventDisplay3D</code> class from	
7	<code>EventDisplay3D_module.cc</code>	411
8	44.3 Implementation of the <code>beginJob()</code> member function of the <code>EventDisplay3D</code>	
9	class from <code>EventDisplay3D_module.cc</code> (continued on listing 44.4).	414
10	44.4 Implementation of the <code>beginJob()</code> member function of the <code>EventDisplay3D</code>	
11	class from <code>EventDisplay3D_module.cc</code> (continued from listing 44.3	
12	and continued on listing 44.5).	415
13	44.5 Implementation of the <code>beginJob()</code> member function of the <code>EventDisplay3D</code>	
14	class from <code>EventDisplay3D_module.cc</code> (continued from listing 44.4).	418
15	44.6 Implementation of the <code>makeNavPanel()</code> member function of the <code>EventDisplay3D</code>	
16	class from <code>EventDisplay3D_module.cc</code> (continued on listing 44.7).	420
17	44.7 Implementation of the <code>makeNavPanel()</code> member function of the <code>EventDisplay3D</code>	
18	class from <code>EventDisplay3D_module.cc</code> (continued from listing 44.6).	422
19	44.8 Implementation of the <code>analyze()</code> member function of the <code>EventDisplay3D</code>	
20	class from <code>EventDisplay3D_module.cc</code> (continued on listing 44.8).	424
21	44.9 Implementation of the <code>analyze()</code> member function of the <code>EventDisplay3D</code>	
22	class from <code>EventDisplay3D_module.cc</code> (continued from listing 44.8	
23	and continued on listing 44.11).	426
24	44.10 The <code>drawHit()</code> helper function from <code>EventDisplay3D_module.cc</code> .	427
25	44.11 Implementation of the <code>analyze()</code> member function of the <code>EventDisplay3D</code>	
26	class from <code>EventDisplay3D_module.cc</code> (continued from listing 44.9	
27	and continued on listing 44.12).	429
28	44.12 Implementation of the <code>analyze()</code> member function of the <code>EventDisplay3D</code>	
29	class from <code>EventDisplay3D_module.cc</code> (continued from listing 44.11).	430
30	58.1 Reading in a ROOT data file	501
31	58.2 Reading in a ROOT data file	502
32	58.3 Reading in a ROOT data file	502
33	58.4 Reading in a ROOT data file	502
34	58.5 Reading in a ROOT data file	503
35	58.6 Reading in a ROOT data file	504

1	59.1 Module source sample	509
2	C.1 Error Codes when an <code>art::Exception</code> is caught in <code>main</code>	530
3	E.1 The constructors of <code>Hep3Vector</code>	539

DRAFT

1 **FIXME:** *ex 9; To do*

DRAFT

1

Part I

2

Introduction

1 How to Read this Documentation

ow-to-read

The *art* document suite, which is currently in an alpha release form, consists of an introductory section and the first few exercises of the Workbook¹, plus a glossary and an index. There are also some preliminary (incomplete and unreviewed) portions of the Users Guide included in the compilation.

The Workbook exercises require you to download some code to edit, execute and evaluate. Both the documentation and the code it references are expected to undergo continual development throughout 2013 and 2014. The latest is always available at the *art* Documentation website. Chapter 12 tells you how to keep up-to-date with improvements and additions to the Workbook code and documentation.



1.1 If you are new to HEP Software...

Read Parts I and II (the introductory material and the Workbook) from start to finish. The Workbook is aimed at an audience who is familiar with (although not necessarily expert in) Unix, C++ and Fermilab's UPS product management system, and who understands the basic *art* framework concepts. The introductory chapters prepare the "just starting out" reader in all these areas.

1.2 If you are an HEP Software expert...

Read chapters 1, 2 and 3: this is where key terms and concepts used throughout the *art* document suite get defined. Skip the rest of the introductory material and jump straight

¹The Workbook is expected to contain roughly 35 exercises when complete.

- 1 into running Exercise 1 in Chapter 9 of the Workbook. Take the approach of: Don't need
2 it? Don't read it.

3 **1.3 If you are somewhere in between...**

- 4 Read chapters 1, 2 and 3 and skim the remaining introductory material in Part I to glean
5 what you need. Along with the experts, you can take the approach of: Don't need it? Don't
6 read it.

2 Conventions Used in this Documentation

conventions

Most of the material in this introduction and in the Workbook is written so that it can be understood by those new to HEP computing; if it is not, please let us know (see [Section 3.4](#)! [sec:getting-help](#))

2.1 Terms in Glossary

The first instance of each term that is defined in the glossary is written in *italics* followed by a γ (Greek letter gamma), e.g., *framework*(γ).

2.2 Typing Commands

Unix commands that you must type are shown in the format `unix command`. Portions of the command for which you must substitute values are shown in slanted font within the command. e.g., you would type your actual username when you see `username`).

While *art* supports OS X as well as flavors of Linux, the instructions for using *art* are nearly identical for all supported systems. When operating-system specific instructions are needed they are noted in the exercises.

When an example Unix command line would overflow the page width, this documentation will use a trailing backslash to indicate that the command is continued on the next line. We indent the second line to make clear that it is not a separate command from the first line. For example:

```
mkdir -p $ART_WORKBOOK_WORKING_BASE/username/workbook-tutorial/  
directory1/directory2/directory3
```

1 You can type the entire command on a single line if it fits, without typing the backslash, *or*
2 on two lines *with* the backslash as the final character of the first line. Do not leave a space
3 before the backslash unless it is required in the command syntax, e.g., before an option, as
4 in

```
5 mkdir \  
6 -p mydir
```

7 2.3 Listing Styles

8 Code listings in C++ are shown as:

```
9 1 // This is a C++ file listing.  
10 2 float* pa = &a;
```

11 Code listings in FHiCL are shown as:

```
11 // This is a FHiCL file listing.  
12 source: {  
13     module_type : RootInput  
14     ...  
15 }
```

17 Other script or file content is denoted:

```
18 This represents script contents.
```

19 Computer output from a command is shown as:

```
20 This is output from a command.
```

21 2.4 Procedures to Follow

22 Step-by-step procedures that the reader is asked to follow are denoted in the following
23 way:

- 24 1. First step...
2. Commands inside procedures are denoted as:


```
mkdir -p mydir
```

1

2.5 Important Items to Call Out

Occasionally, text will be called out to make sure that you don't miss it. Important or tricky terms and concepts will be marked with an "pointing finger" symbol in the margin, as shown at right.



Items that are even trickier will be marked with a "bomb" symbol in the margin, as shown at right. You really want to avoid the problems they describe.



In some places it will be necessary for a paragraph or two to be written for experts. Such paragraphs will be marked with a "dangerous bends" symbol in the margin, as shown at right. Less experienced users can skip these sections on first reading and come back to them at a later time.



2.6 Site-specific Information

Text that refers in particular to Fermilab-specific information is marked with a Fermilab picture, as shown at right.



Text that refers in particular to information about using *art* at non-Fermilab sites is marked with a "generic site" picture, as shown at right. A *site* is defined as a unique combination of experiment and location, and is used to refer to a set of computing resources configured for use by a particular experiment at a particular location. Two examples of sites are the Fermilab supplied resources used by your experiment and the group computing resources an institution that collaborates on your experiment. If you have the necessary software installed on your own laptop, it is also a site. Similarly for your own desktop.



Experiment-specific information will be kept to an absolute minimum; wherever it appears, it will be marked with an experiment-specific icon, e.g., the Mu2e icon at right.



3 Introduction to the *art* Event Processing Framework

chap:intro

3.1 What is *art* and Who Uses it?

sec:what:is:art

art(γ) is an event-processing *framework*(γ) developed and supported by the Fermilab Scientific Computing Division (SCD). The *art* framework is used to build physics programs by loading physics algorithms, provided as plug-in modules. Each experiment or user group may write and manage its own modules. *art* also provides infrastructure for common tasks, such as reading input, writing output, provenance tracking, database access and run-time configuration.

The initial clients of *art* are the Fermilab Intensity Frontier experiments but nothing prevents other experiments from using it as well. The name *art* is always written in *italic lower case*; it is not an acronym.

art is written in C++ and is intended to be used with user code written in C++. (*User code* includes experiment-specific code and any other user-written, non-*art*, non-*external-product*(γ) code.)

art has been designed for use in most places that a typical HEP experiment might require a software framework, including:

- high-level software triggers
- online data monitoring
- calibration
- reconstruction

- 1 ○ analysis
 - 2 ○ simulation
- 3 *art* is not designed for use in real-time environments, such as the direct interface with
4 data-collection hardware.

5 The Fermilab SCD has also developed a related product named *artdaq*(γ), a layer that
6 lives on top of *art* and provides features to support the construction of data-acquisition
7 (*DAQ*(γ)) systems based on commodity servers. Further discussion of *artdaq* is outside
8 the scope of this documentation; for more information consult the *artdaq* home page:
9 <https://cdcv.sfnal.gov/redmine/projects/artdaq/wiki>.

10 A technical paper on *artdaq* is available at: <http://inspirehep.net/record/1229212?ln=en>;

11 The design of *art* has been informed by the lessons learned by the many High Energy
12 Physics (HEP) experiments that have developed C++ based frameworks over the past 20
13 years. In particular, it was originally forked from the framework for the CMS experiment,
14 *cmsrun*.

15 Experiments using *art* are listed at the *art* Documentation website under “Experiments
16 using *art*.”

17 3.2 Why *art*?

18 In all previous experiments at Fermilab, and in most previous experiments elsewhere, in-
19 frastructure software (i.e., the framework, broadly construed – mostly forms of bookkeep-
20 ing) has been written in-house by each experiment, and each implementation has been
21 tightly coupled to that experiment’s code. This tight coupling has made it difficult to share
22 the framework among experiments, resulting in both great duplication of effort and mixed
23 quality.

24 *art* was created as a way to share a single framework across many experiments. In partic-
25 ular, the design of *art* draws a clear boundary between the framework and the user code;
26 the *art* framework (and other aspects of the infrastructure) is developed and maintained
27 by software engineers who are specialists in the field of HEP infrastructure software; this
28 provides a robust, professionally maintained foundation upon which physicists can develop
29 the code for their experiments. Experiments use *art* as an *external package*. Despite some

1 constraints that this separation imposes, it has improved the overall quality of the frame-
2 work and reduced the duplicated effort.

3 **3.3 C++ and C++11**

sec:c++11

4 In 2011, the International Standards Committee voted to approve a new standard for C++,
5 called C++ 11.

6 Much of the existing user code was written prior to the adoption of the C++ 11 standard
7 and has not yet been updated. As you work on your experiment, you are likely to encounter
8 both code written the new way and code written the old way. Therefore, the Workbook will
9 often illustrate both practices.

10 A very useful compilation of what is new in C++ 11 can be found at

<https://cdcvs.fnal.gov/redmine/projects/gm2public/wiki/Cpp2011>



12 This reference material is written for advanced C++ users.

13 **3.4 Getting Help**

sec:getting-help

14 Please send your questions and comments to art-users@fnal.gov. More support informa-
15 tion is listed at <https://web.fnal.gov/project/ArtDoc/SitePages/Support.aspx>.

16 **3.5 Overview of the Documentation Suite**

erview:doc:suite

17 When complete, this documentation suite will contain several principal components, or
18 *volumes*: the introduction that you are reading now, a Workbook, a Users Guide, a Refer-
19 ence Manual, a Technical Reference and a Glossary. At the time of writing, drafts exist for
20 the Introduction, the Workbook, the Users Guide and the Glossary. The components in the
21 documentation suite are illustrated in Figure 3.1. fig:docsuite

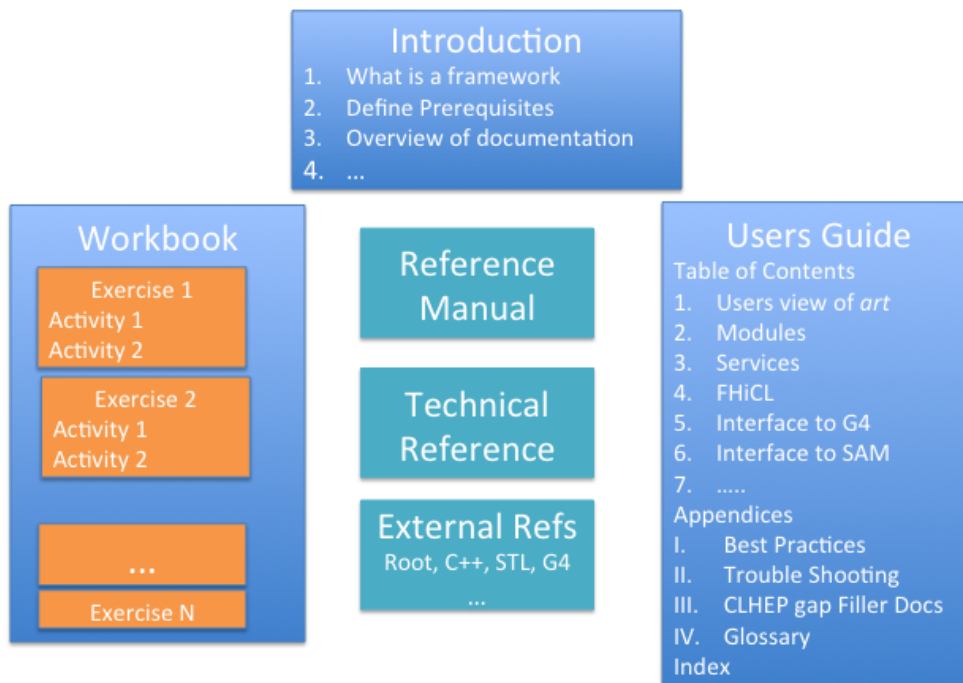


Figure 3.1: Principal components of the *art* documentation suite

g:docsuite

1 3.5.1 The Introduction

2 This introductory volume is intended to set the stage for using *art*. It introduces *art*, pro-
3 vides background material, describes some of the software tools on which *art* depends,
4 describes its interaction with related software and identifies prerequisites for successfully
5 completing the Workbook exercises.

6 3.5.2 The Workbook

7 The Workbook is a series of standalone, self-paced exercises that will introduce the build-
8 ing blocks of the *art* framework and the concepts around which it is built, show practical
9 applications of this framework, and provide references to other portions of the documen-
10 tation suite as needed. It is targeted towards physicists who are new users of *art*, with the
11 understanding that such users will frequently be new to the field of computing for HEP
12 and to C++.

13 One of the Workbook's primary functions is training readers how and where to find more
14 extensive documentation on both *art* and external software tools; they will need this in-
15 formation as they move on to develop and use the scientific software for their experi-
16 ment.

17 The Workbook assumes some basic computing skills and some basic familiarity with the
18 C++ computing language; Chapter 6 provides a tutorial/refresher for readers who need to
19 improve their C++ skills.

20 The Workbook is written using recommended best practices that have become current
21 since the adoption of C++ 11 (see Section 3.8).

22 Because *art* is being used by many experiments, the Workbook exercises are designed
23 around a *toy* experiment that is greatly simplified compared to any actual detector, but it
24 incorporates enough richness to illustrate most of the features of *art*. The goal is to enable
25 the physicists who work through the exercises to translate the lessons learned there into
26 the environment of their own experiments.

3.5.3 Users Guide

The Users Guide is targeted at physicists who have reached an intermediate level of competence with *art* and its underlying tools. It contains detailed descriptions of the features of *art*, as seen by the physicists. The Users Guide will provide references to the *external products*(γ) on which *art* depends, information on how *art* uses these products, and as needed, documentation that is missing from the external products' own documentation.

3.5.4 Reference Manual

The Reference Manual will be targeted at physicists who already understand the major ideas underlying *art* and who need a compact reference to the Application Programmer Interface (*API*(γ)). The Reference Manual will likely be generated from annotated source files, possibly using *Doxygen*(γ). **FIXME:** *reference*

3.5.5 Technical Reference

The Technical Reference will be targeted at the experts who develop and maintain *art*; few physicists will ever want or need to consult it. It will document the internals of *art* so that a broader group of people can participate in development and maintenance.

3.5.6 Glossary

The glossary will evolve as the documentation set grows. At the time of writing, it includes definitions of *art*-specific terms as well as some HEP, Fermilab, C++ and other relevant computing-related terms used in the Workbook and the Users Guide.

3.6 Some Background Material

This section defines some language and some background material about the *art* framework that you will need to understand before starting the Workbook.

3.6.1 Events and Event IDs

In almost all HEP experiments, the core idea underlying all bookkeeping is the *event*(γ). In a triggered experiment, an event is defined as all of the information associated with a single trigger; in an untriggered, spill-oriented experiment, an event is defined as all of the information associated with a single spill of the beam from the accelerator. Another way of saying this is that an event contains all of the information associated with some time interval, but the precise definition of the time interval changes from one experiment to another¹. Typically these time intervals are a few nanoseconds to a few tens of microseconds. The information within an event includes both the raw data read from the Data Acquisition System (DAQ) and all information that is derived from that raw data by the reconstruction and analysis algorithms. An event is the smallest unit of data that *art* can process at one time.

In a typical HEP experiment, the trigger or DAQ system assigns an event identifier (event ID) to each event; this ID uniquely identifies each event, satisfying a critical requirement imposed by *art* that each event be uniquely identifiable by its event ID. This requirement also applies to simulated events.

The simplest event ID is a monotonically increasing integer. A more common practice is to define a multi-part ID and *art* has chosen to use a three-part ID, including:

- *run*(γ) number
- *subRun*(γ) number
- *event*(γ) number

There are two common methods of using this event ID scheme and *art* allows experiments to chose either:

1. When an experiment takes data, the event number is incremented every event. When some predefined condition occurs, the event number is reset to 1 and the subRun number is incremented, keeping the run number unchanged. This cycle repeats until some other predefined condition occurs, at which time the event number is reset to

¹There is a second, distinct, sense in which the word *event* is sometimes used; it is used as a synonym for a *fundamental interaction*; see the glossary entry for *event (fundamental interaction)*(γ). Within this documentation suite, unless otherwise indicated, the word *event* refers to the definition given in the main body of the text.

1 1, the subRun number is reset to 0 (0 not 1 for historical reasons) and the run number
2 is incremented.

3 2. The second method is the same as the first except that the event number monotonically
4 increases throughout a run and does not reset to 1 on subRun boundaries. The
5 event number does reset to 1 at the start of each run.

6 *art* does not define what conditions cause these transitions; those decisions are left to each
7 experiment. Typically experiments will choose to start new runs or new subRuns when
8 one of the following happens: a preset number of events is acquired; a preset time interval
9 expires; a disk file holding the output reaches a preset size; or certain running conditions
10 change.

11 *art* requires only that a subRun contain zero or more events and that a run contain zero or
12 more subRuns.

13 When an experiment takes data, events read from the DAQ are typically written to disk
14 files, with copies made on tape. The events in a single subRun may be spread over sev-
15 eral files; conversely, a single file may contain many runs, each of which contains many
16 subRuns.

17 3.6.2 *art* Modules and the Event Loop

18 Users provide executable code to *art* in pieces called *art modules*(γ)² that are dynamically
19 loaded as plugins and that operate on event data. The concept of reading events and, in
20 response to each new event, calling the appropriate member functions of each module, is
21 referred to as the *event loop*(γ). The concepts of the *art module* and the *event loop* will be
22 illustrated via the following discussion of how *art* processes a job.

23 The simplest command to run *art* looks like:

```
24   art -c filename.fcl
```

25 The argument to *-c* is the *run-time configuration file*(γ), a text file that tells one run of
26 *art* what it should do. Run-time configuration files for *art* are written in the Fermilab

²Many programming languages have an idea named *module*; the use of the term *module* by *art* and in this documentation set is an *art*-specific idea that will be developed through the first few chapters of the Workbook.

1 Hierarchical Configuration Language *FHiCL*(γ) (pronounced “fickle”) and the filenames
2 end in `.fcl`. As you progress through the Workbook, this language and the conventions
3 used in the run-time configuration file will be explained; the full details are available in
4 Chapter 52 of the Users Guide. (The run-time configuration file is often referred to as
5 simply the *configuration file* or even more simply as just the *configuration*(γ .)

6 When *art* starts up, it reads the configuration file to learn what input files it should read,
7 what user code it should run and what output files it should write. As mentioned above, an
8 experiment’s code (including any code written by individual experimenters) is provided in
9 units called *art modules*. A module is simply a C++ class, provided by the experiment or
10 user, that obeys a set of rules defined by *art* and whose *source code*(γ) file gets compiled
11 into a *dynamic library*(γ) that can be loaded at run-time by *art*.

12 These rules will be explained as you work through the Workbook and they are summarized
13 in *a future chapter in the User’s Guide*.

14 The code base of a typical experiment will contain many C++ classes. Only a small fraction
15 of these will be modules; most of the rest will be ordinary C++ classes that are used within
16 modules³.

17 A user can tell *art* the order in which modules should be run by specifying that order in
18 the configuration file. A user can also tell *art* to determine, on its own, the correct order in
19 which to run modules; the latter option is referred to as *reconstruction on demand*.

20 Imagine the processing of each event as the assembly of a widget on an assembly line
21 and imagine each module as a worker that needs to perform a set task on each widget.
22 Each worker has a task that must be done on each widget that passes by; in addition some
23 workers may need to do some start-up or close-down jobs. Following this metaphor, *art*
24 requires that each module provide code that will be called once for every event and *art*
25 allows any module to provide code that will be called at the following times:

- 26 ○ at the start of the *art* job
- 27 ○ at the end of the *art* job
- 28 ○ at the start of each run

³*art* defines a few other specialized roles for C++ classes; you will encounter these in Sections 3.6.4 and 3.6.5.

- 1 ○ at the end of each run
- 2 ○ at the start of each SubRun
- 3 ○ at the end of each SubRun

4 For those of you who are familiar with *inheritance* in C++, a module class (i.e., a “mod-
5 ule”) must inherit from one of a few different module *base classes*. Each module class
6 must override one pure-virtual member function from the base class and it may override
7 other virtual member functions from the base class.



8 After *art* completes its initialization phase (intentionally not detailed here), it executes the
9 event loop, illustrated in Figure 3.2, and enumerated below.

fig:eventloop

10 The event loop

- 11 1. calls the *constructor*(γ) of every module in the configuration.
- 12 2. calls the *beginJob member function*(γ) of every module that provides one.
- 13 3. reads one event from the input source, and for that event
 - 14 (a) determines if it is from a run different from that of the previous event (true for
15 first event in loop);
 - 16 (b) if so, calls the *beginRun* member function of each module that provides one;
 - 17 (c) determines if the event is from a subRun different from that of the previous
18 event (true for first event in loop);
 - 19 (d) if so, calls the *beginSubRun* member function of each module that provides
20 one;
 - 21 (e) calls each module’s (required) per-event member function.
- 22 4. reads the next event and repeats the above per-event steps until it encounters a new
23 subRun.
- 24 5. closes out the current subRun by calling the *endSubRun* member function of each
25 module that provides one.
- 26 6. repeats steps 4 and 5 until it encounters a new run.

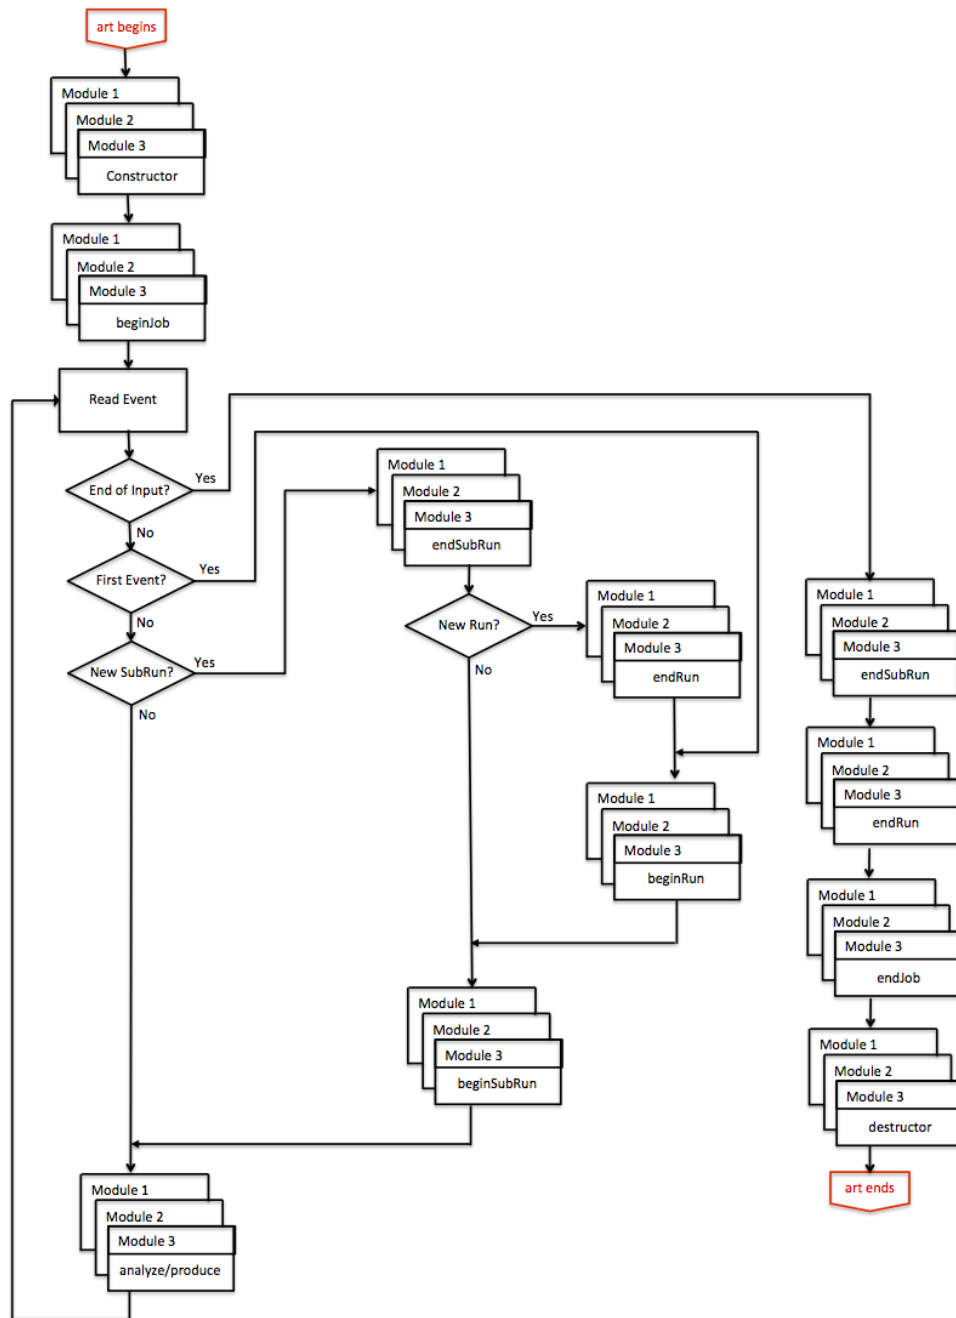


fig:eventloop

Figure 3.2: Flowchart describing the *art* event loop for an input file that contains at least one event. *art* begins at the box in the upper left and ends at the box in the lower right. On the first event, the tests for new subRun and new run are true. Not all features of the event loop are shown, just those that you will encounter in the early parts of the *art* workbook. The case of a file with no events is not shown because it has many subcases and is not of general interest.

- 1 7. closes out the current run by calling the `endRun` member function of each module
- 2 that provides one.
- 3 8. repeats steps 3 through 7 until it reaches the end of the input source.
- 4 9. calls the `endJob` member function of each module that provides one.
- 5 10. calls the `destructor(γ)` of each module.

6 This entire set of steps comprises the event loop. One of *art*'s most visible jobs is control-
7 ling the event loop.

8 3.6.3 Module Types

9 Every *art* module must be one of the following five types, which are defined by the ways
10 in which they interact with each event and with the event loop:

11 ***analyzer module***(γ) May inspect information found in the event but may not add new
12 information to the event. .

13 ***producer module***(γ) May inspect information found in the event and may add new infor-
14 mation to the event.

15 ***filter module***(γ) Same functions as a producer module but may also tell *art* to skip the
16 processing of some, or all, modules for the current event; may also control which
17 events are written to which output.

18 ***source module***(γ) Reads events, one at a time, from some source; *art* requires that every
19 *art* job contain exactly one source module. A source is often a disk file but other
20 options exist and will be described in the Workbook and Users Guide. **FIXME:** *ref*
21 *UG ch when made*

22 ***output module***(γ) Reads selected data products from memory and writes them to an out-
23 put destination; an *art* job may contain zero or more output modules. An output
24 destination is often a disk file but other options exist and will be described in the
25 Users' Guide. **FIXME:** *ref ch:ug-output*.

26 Note that no module may change information that is already present in an event.



27 What does an analyzer do if it may neither alter information in an event nor add to it?

1 Typically it creates printout and it creates ROOT files containing histograms, *trees*(γ) and
2 *nuples*(γ) that can be used for downstream analysis. (If you have not yet encountered these
3 terms, the Workbook will provide explanations as they are introduced.)

4 Most novice users will only write analyzer modules and filter modules; readers with a little
5 more experience may also write producer modules. The Workbook will provide examples
6 of all three. Few people other than *art* experts and each experiment's software experts will
7 write source or output modules, however, the Workbook will teach you what you need to
8 know about configuring source and output modules.

9 **3.6.4 *art* Data Products**

10 This section introduces more ideas and terms dealing with event information that you will
11 need as you progress through the Workbook.

12 The term *data product*(γ) is used in *art* to mean the unit of information that user code
13 may add to an event or retrieve from an event. Data products are created in a number of
14 ways.

- 15 1. The DAQ system will package the raw data into data products, perhaps one or two
16 data products for each major subsystem.
- 17 2. Each module in the reconstruction chain will create one or more data products.
- 18 3. Some modules in the analysis chain will produce data products; others may just
19 make histograms and write information in non-*art* formats for analysis outside of
20 *art*; they may, for example, write user-defined ROOT `Trees`.
- 21 4. The simulation chain will usually create many data products. Some will be simu-
22 lated event-data while others will describe the true properties of the simulated event.
23 These data products can be used to study the response of the detector to simulated
24 events; they can also be used to develop, debug and characterize the reconstruction
25 algorithms.

26 Because these data products are intrinsically experiment-dependent, each experiment de-
27 fines its own data products. In the Workbook, you will learn about a set of data products
28 designed for use with the toy experiment. There are a small number of data products that
29 are defined by *art* and that hold bookkeeping information; these will be described as you

1 encounter them in the Workbook.

2 A data product is just a C++ *type*(γ) (a class, *struct*(γ) or typedef) that obeys a set of rules
 3 defined by *art*; these rules are very different than the rules that must be followed for a class
 4 to be a module; when the sections that describe these rules in detail have been prepared,
 5 we will add references here. A data product can be a single integer, a large complex class
 6 hierarchy, or anything in between.



7 **FIXME:** *I removed the section reference that was in previous para because it was not ap-*
 8 *propriate. We need to reference two sections, either in the users guide or in future sections*
 9 *of the workbook: the one that describes the rules for modules and the one that describes*
 10 *the rules for data products. We do not yet have a section that discusses the later.*

11 *Add the missing references alluded to in the previous para.*

12 Very often, a data product is a *collection*(γ) of some experiment-defined type. The C++
 13 standard libraries define many sorts of collection types; *art* supports many of these and
 14 also provides a custom collection type named `cet::map_vector` **FIXME:** *Reference*
 15 *to map_vector when available; will be in art::ptr I think.* Workbook exercises will
 16 clarify the *data product* and *collection type* concepts.

17 3.6.5 *art* Services

rtServices

18 Previous sections of this Introduction have introduced the concept of C++ classes that have
 19 to obey a certain set of rules defined by *art*, in particular, modules in Section 3.6.2 and data
 20 products in Section 3.6.4. *art services*(γ) are yet other examples of this.

21 In a typical *art* job, two sorts of information need to be shared among the modules. The
 22 first sort is stored in the data products themselves and is passed from module to module
 23 via the event. The second sort is not associated with each event, but rather is valid for some
 24 aggregation of events, subRuns or runs, or over some other time interval. Three examples
 25 of this second sort include the geometry specification, the conditions information⁴ and, for
 26 simulations, the table of particle properties.

27 To provide managed access to the second sort of information, *art* supports an idea named

⁴The phrase “conditions information” is the currently fashionable name for what was once called “calibration constants”; the name change came about because most calibration information is intrinsically time-dependent, which makes “constants” a poor choice of name.

1 *art services* (again, shortened to *services*). Services may also be used to provide certain
2 types of utility functions. Again, a service in *art* is just a C++ class that obeys a set of
3 rules defined by *art*. The rules for services are different than those for modules or data
4 products.

5 *art* implements a number of services that it uses for internal functions, a few of which
6 you will encounter in the first couple of Workbook exercises. The *message service*(γ)
7 is used by both *art* and experiment-specific code to limit printout of messages with a
8 low severity level and to route messages to appropriate destinations. It can be configured
9 to provide summary information at the end of the *art* job. The *TFileService*(γ) and the
10 `RandomNumberGenerator` service are not used internally by *art*, but are used by most
11 experiments. Experiments may also create and implement their own services.

12 After *art* completes its initialization phase and before it constructs any modules (see Sec-
13 tion 3.6.2), it

- 14 1. reads the configuration to learn what services are requested, and
- 15 2. calls the constructor of each requested service.

16 Once a service has been constructed, any code in any module can ask *art* for a *smart*
17 *pointer*(γ) to that service and use the features provided by that service. Because services
18 are constructed before modules, they are available for use by modules over the full life
19 cycle of each module.

20 It is also legal for one service to request information from another service as long as the
21 dependency chain does not have any loops. That is, if Service A uses Service B, then
22 Service B may not use Service A, either directly or indirectly.

23 For those of you familiar with the C++ Singleton Design Pattern, an *art* service has some
differences and some similarities to a Singleton. The most important difference is that the
lifetime of a service is managed by *art*, which calls the constructors of all services at a
well-defined time in a well-defined order. Contrast this with the behavior of Singletons,
for which the order of initialization is undefined by the C++ standard and which is an
accident of the implementation details of the loader. *art* also includes services under the
umbrella of its powerful run-time configuration system; in the Singleton Design pattern
this issue is simply not addressed.



3.6.6 Dynamic Libraries and *art*

When code is executed within the *art* framework, *art*, not the experiment, provides the main executable. The experiment provides its code to the *art* executable in the form of dynamic libraries that *art* loads at run time; these libraries are also called *dynamic load libraries*, *shareable object libraries*, or *plugins*. On Linux, their filenames typically end in `.so`; on OS X, the suffixes `.dylib` and `.so` are both used. **FIXME:** For more information about dynamic libraries, see Section 59.5. (Still in misc topics chap)

3.6.7 Build Systems and *art*

To make an experiment's code available to *art*, the source code must be compiled and linked (i.e., *built*) to produce dynamic libraries (Section 3.6.6). The tool that creates the dynamic libraries from the C++ source files is called a *build system*(γ).

Experiments that use *art* are free to choose their own build systems, as long as the system follows the conventions that allow *art* to find the name of the `.so` file given the name of the module class, as discussed in Section ???. The Workbook will use a build system named *cetbuildtools*, which is a layer on top of *cmake*⁵.

The **cetbuildtools** system defines three standard compiler optimization levels, called “debug”, “profile” and “optimized”; the last two are often abbreviated “prof” and “opt”. When code is compiled with the “opt” option, it runs as quickly as possible but is difficult to debug. When code is compiled with the “debug” option, it is much easier to debug but it runs more slowly. When code is compiled with the “prof” option the speed is almost as fast as for an “opt” build and the most useful subset of the debugging information is retained. The “prof” build retains enough debugging information that one may use a profiling tool to identify in which functions the program spends most of its time; hence its name “profile”. The “prof” build provides enough information to get a useful traceback from a core dump. Most experiments using *art* use the “prof” build for production and the “debug” build for development.

The compiler options corresponding to the three levels are listed in Table 3.1.

⁵**cetbuildtools** is also used to build *art* itself.



Table 3.1: Compiler flags for the optimization levels defined by **cetbuildtools**; compiler options not related to optimization or debugging are not included in this table.

	Name	flags
tab:opt-levels	debug	-O0 -g
	prof	-O3 -g -fno-omit-frame-pointer -DNDEBUG
	opt	-O3 -DNDEBUG

1 3.6.8 External Products

external:products

2 As you progress through the Workbook, you will see that the exercises use some software
 3 packages that are part of neither *art* nor the toy experiment's code. The Workbook code, *art*
 4 and the software for your experiment all rely heavily on some external tools and, in order
 5 to be an effective user of *art*-based HEP software, you will need at least some familiarity
 6 with them; you may, in fact, need to become expert in some.

7 These packages and tools are referred to as *external products*(γ) (sometimes called simply
 8 *products*).

9 An initial list of the external products you will need to become familiar with includes:

10 *art* the event processing framework

11 FHiCL the run-time configuration language used by *art*

12 CETLIB a utility library used by *art*

13 *MF*(γ) a message facility that is used by *art* and by (some) experiments that use *art*

14 ROOT an analysis, data presentation and data storage tool widely used in HEP

15 *CLHEP*(γ) a set of utility classes; the name is an acronym for *Class Library for HEP*

16 *boost*(γ) a class library with new functionality that is being prototyped for inclusion in
 17 future C++ standards

18 gcc the GNU C++ compiler and run-time libraries; both the core language and the standard
 19 library are used by *art* and by your experiment's code.

20 *git*(γ) a source code management system that is used for the Workbook and by some
 21 experiments; similar in concept to the older CVS and SVN, but with enhanced func-
 22 tionality

- 1 *cetbuildtools*(γ) the build system that is used by the *art* Workbook (and by *art* itself).
- 2 *UPS*(γ) a Fermilab-developed system for accessing software products; it is an acronym
3 for *Unix Product Support*.
- 4 *UPD*(γ) a Fermilab-developed system for distributing software products; it is an acronym
5 for *Unix Product Distribution*.
- 6 *jobsub_tools*(γ) tools for submitting jobs to the Fermigrid batch system and monitoring
7 them.
- 8 *ifdh_sam*(γ) allows *art* to use *SAM*(γ) as an external run-time agent that can deliver re-
9 mote files to local disk space and can copy output files to tape. *SAM* is a Fermilab-
10 supplied resource that provides the functions of a file catalog, a replica manager and
11 some functions of a batch-oriented workflow manager

12 Any particular line of code in a Workbook exercise may use elements from, say, four or
13 five of these packages. Knowing how to parse a line and identify which feature comes from
14 which package is a critical skill. The Workbook will provide a tour of the above packages
15 so that you will recognize elements when they are used and you will learn where to find
16 the necessary documentation.

17 For the *art* Workbook, external products are made available to your code via a mechanism
18 called UPS, which will be described in Section 7. Many Fermilab experiments also use
19 UPS to manage their external products; this is not required by *art* and you may choose to
20 manage external products whichever way you prefer. UPS is, itself, just another external
21 product. From the point of view of your experiment, *art* is an external product. From the
22 point of view of the Workbook code, both *art* and the code for the toy experiment are
23 external products.

24 Finally, it is important to recognize an overloaded word, *products*. When a line of docu-
25 mentation simply says *products*, it may be referring either to data products or to external
26 products. If it is not clear from the context which is meant, please let us know (see Sec-
27 tion 3.4).



3.6.9 The Event-Data Model and Persistency

ssec:edm

ssec:artDataProducts

Section 3.6.4 introduced the idea of *art* data products. In a small experiment, a fully reconstructed event may contain on the order of ten data products; in a large experiment there may be hundreds.

While each experiment will define its own data product classes, there is a common set of questions that *art* users on any experiment need to consider:

1. How does my module access data products that are already in the event?
2. How does my module publish a data product so that other modules can see it?
3. How is a data product represented in the memory of a running program?
4. How does an object in one data product refer to an object in another data product?
5. What metadata is there to describe each data product? (Such metadata might include: the module that created it; the run-time configuration of that module; the data products read by that module; the code version of the module that created it.)
6. How does my module access the metadata associated with a particular data product?

The answers to these questions form what is called the *Event-Data Model*(γ) (EDM) that is supported by the framework.

A question that is closely related to the EDM is: what technologies are supported to write data products from memory to a disk file and to read them from the disk file back into memory in a separate *art* job? A framework may support several such technologies. *art* currently supports only one disk file format, a ROOT-based format, but the *art* EDM has been designed so that it will be straightforward to support other disk file formats as it becomes useful to do so.

A few other related terms that you will encounter include:

1. *transient representation*: the in-memory representation of a data product
2. *persistent representation*: the on-disk representation of a data product
3. *persistency*: the technology to convert data products back and forth between their persistent and transient representations

3.6.10 Event-Data Files

1

data-files

2 When you read data from an experiment and write the data to a disk file, that disk file is
3 usually called a *data file*.

4 When you simulate an experiment and write a disk file that holds the information pro-
5 duced by the simulation, what should you call the file? The Particle Data Group has rec-
6 ommended that this not be called a “data file” or a “simulated data file;” they prefer that
7 the word “data” be strictly reserved for information that comes from an actual experiment.
8 They recommend that we refer to these files as “files of simulated events” or “files of
9 Monte Carlo events”⁶. Note the use of “events,” not “data.”

10 This leaves us with a need for a collective noun to describe both data files and files of
11 simulated events. The name in current use is *event-data files*(γ); yes this does contain the
12 word “data” but the hyphenated word, “event-data”, is unambiguous and this has become
13 the standard name.

3.6.11 Files on Tape

14

-data-tape

15 Many experiments do not have access to enough disk space to hold all of their event-data
16 files, ROOT files and log files. The solution is to copy a subset of the disk files to tape and
17 to read them back from tape as necessary.

18 At any given time, a snapshot of an experiment’s files will show some on tape only, some
19 on tape with copies on disk, and some on disk only. For any given file, there may be
20 multiple copies on disk and those copies may be distributed across many *sites*(γ), some at
21 Fermilab and others at collaborating laboratories or universities.

22 Conceptually, two pieces of software are used to keep track of which files are where, a
23 *File Catalog* and a *Replica Manager*. One software package that fills both of these roles is
24 called SAM, which is an acronym for “Sequential data Access via Metadata.” SAM also
25 provides some tools for Workflow management. SAM is in wide use at Fermilab and you
26 can learn more about SAM at:

27 <https://cdcv.s.fnal.gov/redmine/projects/sam-main/wiki>.

⁶ In HEP almost all simulations codes use *Monte Carlo*(γ) methods; therefore simulated events are often referred to as *Monte Carlo events* and the simulation process is referred to as *running the Monte Carlo*.

3.7 The Toy Experiment

The Workbook exercises are based around a made-up (*toy*) experiment. The code for the toy experiment is deployed as a UPS product named *toyExperiment*. The rest of this section will describe the physics content of *toyExperiment*; the discussion of the code in the **toyExperiment** UPS product will unfold in the Workbook, in parallel to the exposition of *art*.

The software for the toy experiment is designed around a toy detector, which is shown in Figure 3.3. The *toyExperiment* code contains many C++ classes: some modules, some data products, some services and some plain old C++ classes. About half of the modules are producers that individually perform either one step of the simulation process or one step of the reconstruction/analysis process. The other modules are analyzers that make histograms and ntuples of the information produced by the producers. There are also event display modules.

3.7.1 Toy Detector Description




Figure 3.3: The geometry of the toy detector; the figures are described in the text. A uniform magnetic field of strength 1.5 T is oriented in the $+z$ direction.

The toy detector is a central detector made up of 15 concentric shells, with their axes

Table 3.2: Units used in the Workbook

Quantity	Unit
Length	mm
Energy	MeV
Time	ns
Plane Angle	radian
Solid Angle	steradian
Electric Charge	Charge of the proton = +1
Magnetic Field	Tesla

1 centered on the z axis; the left-hand part of Figure 3.3 shows an xy view of these shells
 2 and the right shows the radius vs z view. The inner five shells are closely spaced radially
 3 and are short in z ; the ten outer shells are more widely spaced radially and are longer in
 4 z . The detector sits in a uniform magnetic field of 1.5 T oriented in the $+z$ direction. The
 5 origin of the coordinate system is at the center of the detector. The detector is placed in a
 6 vacuum.

7 Each shell is a detector that measures (φ, z) , where φ is the azimuthal angle of a line from
 8 the origin to the measurement point. Each measurement has perfectly gaussian measure-
 9 ment errors and the detector always has perfect separation of hits that are near to each
 10 other. The geometry of each shell, its efficiency and resolution are all configurable at run-
 11 time.

12 All of the code in the toyExperiment product works in the set of units described in Ta-
 13 ble 3.2. Because the code in the Workbook is built on toyExperiment, it uses the same
 14 units. *art* itself is not unit-aware and places no constraints on which units your experiment
 15 may use.

16 The first six units listed in Table 3.2 are the base units defined by the CLHEP SystemOfUnits
 17 package. These are also the units used by Geant4.



3.7.2 Workflow for Running the Toy Experiment Code

The workflow of the toy experiment code includes five steps: three simulation steps, a reconstruction step and an analysis step:

1. event generation
2. detector simulation
3. hit-making
4. track reconstruction
5. analysis of the mass resolution

For each event, the event generator creates some signal particles and some background particles. The first signal particle is generated with the following properties:

- Its mass is the rest mass of the ϕ meson; the event generator does not simulate a natural width for this particle.
- It is produced at the origin.
- It has a momentum that is chosen randomly from a distribution that is uniform between 0 and 2000 MeV/ c .
- Its direction is chosen randomly on the unit sphere.

The event generator then decays this particle to K^+K^- ; the center-of-mass decay angles are chosen randomly on the unit sphere.

The background particles are generated by the following algorithm:

- Background particles are generated in pairs, one π^+ and one π^- .
- The number of pairs in each event is a random variate chosen from a Poisson distribution with a mean of 0.75.
- Each of the pions is generated as follows:
 - It is produced at the origin.
 - It has a momentum that is chosen randomly from a distribution that is uniform between 0 and 800 MeV/ c .




Figure 3.4: Event display of a simulated event in the toy detector.

- 1 – Its direction is chosen randomly on the unit sphere.
- 2 The above algorithm generates events with a total charge of zero but there is no concept of
- 3 momentum or energy balance. About 47% of these events will not have any background
- 4 tracks.
- 5 In the detector simulation step, particles neither scatter nor lose energy when they pass
- 6 through the detector cylinders; nor do they decay. Therefore, the charged particles follow
- 7 a perfectly helical trajectory. The simulation follows each charged particle until it either
- 8 exits the detector or until it completes the outward-going arc of the helix. When the sim-
- 9 ulated trajectory crosses one of the detector shells, the simulation records the true point
- 10 of intersection. All intersections are recorded; at this stage in the simulation, there is no
- 11 notion of inefficiency or resolution. The simulation does not follow the trajectory of the ϕ
- 12 meson because it was decayed in the generator.
- 13 Figure 3.4 shows an event display of a simulated event that has no background tracks.
- 14 In this event the ϕ meson was travelling close to 90° to the z axis and it decayed nearly
- 15 symmetrically; both tracks intersect all 15 detector cylinders. The left-hand figure shows
- 16 an xy view of the event; the solid lines show the trajectory of the kaons, red for K^+ and
- 17 blue for K^- ; the solid dots mark the intersections of the trajectories with the detector
- 18 shells. The right-hand figure shows the same event but in an rz view.




fig:toy-evt5

Figure 3.5: Event display of another simulated event in the toy detector; a K^- (blue) is produced with a very shallow trajectory and it does not intersect any detector shells while the K^+ (red) makes five hits in the inner detector and seven in the outer detector

- 1 Figure 3.5 shows an event display of another simulated event, one that has four background
 2 tracks, all drawn in green. In the xy view it is difficult to see the two π^- tracks, which
 3 have very low transverse momentum, but they are clear in the rz view. Look at the K^+
 4 track, draw in red; its trajectory just stops in the middle of the detector. Why does this
 5 happen? In order to keep the exercises focused on *art* details, not geometric corner cases,
 6 the simulation stops a particle when it completes the outward-going arc of the helix and
 7 starts to curl back towards the z axis; it does this even if the the particle is still inside the
 8 detector.
- 9 The third step in the simulation chain (hit-making) is to inspect the intersections produced
 10 by the detector simulation and turn them into data-like hits. In this step, a simple model of
 11 inefficiency is applied and some intersections will not produce hits. Each hit represents a
 12 2D measurement (φ, z) ; each component is smeared with a gaussian distribution.
- 13 The three simulation steps use tools provided by *art* to record the *truth information* (γ)
 14 about each hit. Therefore it is possible to navigate from any hit back to the intersec-
 15 tion from which it is derived, and from there back to the particle that made the inter-
 16 section.
- 17 The fourth step is the reconstruction step. The toyExperiment does not yet have properly




Figure 3.6: The invariant mass of all reconstructed pairs of oppositely charged tracks; for this all reconstructed tracks are assumed to be kaons.

1 working reconstruction code; instead it mocks up credible looking results. The output of
 2 this code is a data product that represents a fitted helix; it contains the fitted track parame-
 3 ters of the helix, their covariance matrix and collection of smart pointers that point to the
 4 hits that are on the reconstructed track. When we write proper tracking finding and track
 5 fitting code for the `toyExperiment`, the classes that describe the fitted helix will not change.
 6 Because the main point of the Workbook exercises is to illustrate the bookkeeping features
 7 in *art*, this is good enough for the task at hand. The mocked-up reconstruction code will
 8 only create a fitted helix object if the number of hits on a track is greater than some min-
 9 imum value. Therefore there may be some events in which the output data product is be
 10 empty.

11 The fifth step in the workflow does a simulated analysis using the fitted helices from the
 12 reconstruction step. It forms all distinct pairs of tracks and requires that they be oppositely
 13 charged. It then computes the invariant mass of the pair, under the assumption that both
 14 fitted helices are kaons.⁷ This module is an analyzer module and does not make any output

⁷The toy experiment does not have any particle identification system so analysis code cannot know if a

1 data product. But it does make some histograms, one of which is a histogram of the recon-
2 structed invariant mass of all pairs of oppositely charged tracks; this histogram is shown
3 in Figure 3.6. When you run the Workbook exercises, you will make this plot and can
4 compare it to Figure 3.6. In the figure you can see a clear peak that is created when the two
5 reconstructed tracks are the two true daughters of the generated φ meson. You can also see
6 an almost flat contribution that occurs when at least one of the reconstructed tracks comes
7 from one of the generated background particles.

8 3.8 Rules, Best Practices, Conventions and Style

9 In many places, the Workbook will recommend that you write fragments of code in a
10 particular way. The reason for any particular recommendation may be one of the follow-
11 ing:

- 12 ○ It is a hard rule enforced by the C++ language or by one of the external products.
- 13 ○ It is a recommended best practice that might not save you time or effort now but will
14 in the long run.
- 15 ○ It is a convention that is widely adopted; C++ is a rich enough language that it will
16 let you do some things in many different ways. Code is much easier to understand
17 and debug if an experiment chooses to always write code fragments with similar
18 intent using a common set of conventions.
- 19 ○ It is simply a question of style.

20 It is important to be able to distinguish between rules, best practices, conventions and
21 styles; you must follow the rules; it wise to use best practices and established conventions;
22 but style suggestions are just that, suggestions. This documentation will distinguish among
23 these options when discussing the recommendations that it makes.

24 If you follow the recommendations for best practices and common conventions, it will
25 be easier to verify that your code is correct and your code will be easier to understand,
26 develop and maintain.

reconstructed track is a pion or a kaon. A planned enhancement of the toy experiment is to add a particle identification device.

4 Unix Prerequisites

4.1 Introduction

You will work through the Workbook exercises on a computer that is running some version of the Unix operating system. This chapter describes where to find information about Unix and gives a list of Unix commands that you should understand before starting the Workbook exercises. This chapter also describes a few ideas that you will need immediately but which are usually not covered in the early chapters of standard Unix references.

If you are already familiar with Unix and the *bash*(γ) shell, you can safely skip this chapter.

4.2 Commands

In the Workbook exercises, most of the commands you will enter at the Unix prompt will be standard Unix commands, but some will be defined by the software tools that are used to support the Workbook. The non-standard commands will be explained as they are encountered. To understand the standard Unix commands, any standard Linux or Unix reference will do. Section 4.10 provides links to Unix references.

Most Unix commands are documented via the *man page* system (short for “manual”). To get help on a particular command, type the following at the command prompt, replacing *command-name* with the actual name of the command:

```
man command-name
```

In Unix, everything is case sensitive; so the command `man` must be typed in lower case. You can also try the following; it works on some commands and not others:

1 `command-name --help`

2 or

3 `command-name -?`

4 Before starting the Workbook, make sure that you understand the basic usage of the fol-
5 lowing Unix commands:

6 `cat, cd, cp, echo, export, gzip, head, less, ln -s, ls,`

7 `mkdir, more, mv, printenv, pwd, rm, rmdir, tail, tar`

8 **FIXME:** *Future: Consider adding a some material for less since many intro texts skip*
9 *it.*

10 You also need to be familiar with the following Unix concepts:


- 11 ○ filename vs pathname
- 12 ○ absolute path vs relative path
- 13 ○ directories and subdirectories (equivalent to folders in the Windows and Mac worlds)
- 14 ○ current working directory
- 15 ○ home directory (aka login directory)
- 16 ○ `.. /` notation for viewing the directory above your current working directory
- 17 ○ environment variables (discussed briefly in Section [4.5](#))
- 18 ○ *paths*(γ) (in multiple senses; see Section [4.6](#))
- 19 ○ file protections (read-write-execute, owner-group-other)
- 20 ○ symbolic links
- 21 ○ stdin, stdout and stderr
- 22 ○ redirecting stdin, stdout and stderr
- 23 ○ putting a command *in the background* via the `&` character
- 24 ○ pipes


4.3 Shells

sec:shells

2 When you type a command at the prompt, a command-line interpreter called a *Unix shell*,
3 or simply a *shell*, reads your command and figures out what to do. Most versions of Unix
4 support a variety of different shells, e.g., *bash* or *csh*. The *art* Workbook code expects to
5 be run in the *bash shell*. You can see which shell you're running by entering:

```
6 echo $SHELL
```

7 For those of you with accounts on a Fermilab machine, your login shell was initially set to
8 the *bash shell*¹. 


9 If you are working on a non-Fermilab machine and *bash* is not your default shell, consult
10 a local expert to learn how to change your login shell to *bash*. 

11 Some commands are executed internally by the shell but other commands are dispatched
12 to an appropriate program or script, and launch a child shell (of the same variety) called a
13 *subshell*.

4.4 Scripts: Part 1

s:part:one

15 In order to automate repeated operations, you may write multiple Unix commands into
16 a file and tell *bash* to run all of the commands in the file as if you had typed them se-
17 quentially. Such a file is an example of a *shell script* or a *bash script*. The *bash* scripting
18 language is a powerful language that supports looping, conditional execution, tests to learn
19 about properties of files and many other features.

20 Throughout the Workbook exercises you will run many scripts. You should understand
21 the big picture of what they do, but you don't need to understand the details of how they
22 work. 

23 If you would like to learn more about *bash*, some references are listed in Section 4.10. [|sec:unix:references](#)

¹ If you have had a Fermilab account for many years, your default shell might be something else. If your default shell is not *bash*, open a Service Desk ticket to request that your default shell be changed to *bash*.

4.5 Unix Environments

4.5.1 Building up the Environment

Very generally, a Unix *environment* is a set of information that is made available to programs so that they can find everything they need in order to run properly. The Unix operating system itself defines a generic environment, but often this is insufficient for everyday use. However, an environment sufficient to run a particular set of applications doesn't just pop out of the ether, it must be *established* or *set up*, either manually or via a script. Typically, on institutional machines at least, system administrators provide a set of login scripts that run automatically and enhance the generic Unix environment. This gives users access to a variety of system resources, including, for example:

- disk space to which you have read access
- disk space to which you have write access
- commands, scripts and programs that you are authorized to run
- proxies and tickets that authorize you to use resources available over the network
- the actual network resources that you are authorized to use, e.g., tape drives and DVD drives

FIXME: *This section is addressing many fewer issues that it did when it started out. To sharpen this up maybe we need to focus on the notion of the environment defined within a shell. (AH: please add the content you think should be here, Rob.*

This constitutes a basic *working environment* or *computing environment*. Environment information is largely conveyed by means of *environment variables* that point to various program executable locations, data files, and so on. A simple example of an environment variable is HOME, the variable whose value is the absolute path to your home directory. Environment variables are inherited by subshells, which is a child process launched by a shell or a shell script.

Particular programs (e.g., *art*) usually require extra information, e.g., paths to the program's executable(s) and to its dependent programs, paths indicating where it can find input files and where to direct its output, and so on. In addition to environment variables,

1 the *art*-enabled computing environment includes some aliases and *bash* functions that have
 2 been defined; these are discussed in Section [4.8](#).

3 In turn, the Workbook code, which must work for all experiments and at Fermilab as well as at
 4 collaborating institutions, requires yet more environment configuration – a *site-specific*
 5 configuration.

6 Given the different experiments using *art* and the variety of laboratories and universities
 7 at which the users work, a *site*(γ) in *art* is a unique combination of *experiment* and *insti-*
 8 *tution*. It is used to refer to a set of computing resources configured for use by a particular
 9 experiment at a particular institution. Setting up your site-specific environment will be
 10 discussed in Section [4.7](#).



11 When you finish the Workbook and start to run real code, you will set up your experiment-
 12 specific environment on top of the more generic *art*-enabled environment, in place of the
 13 Workbook's. To switch between these two environments, you will log out and log back in,
 14 then run the script appropriate for the environment you want. Because of potential naming
 15 “collisions,” it is not guaranteed that these two environments can be overlain and always
 16 work properly.

17 This concept of the environment hierarchy is illustrated in Figure [4.1](#).

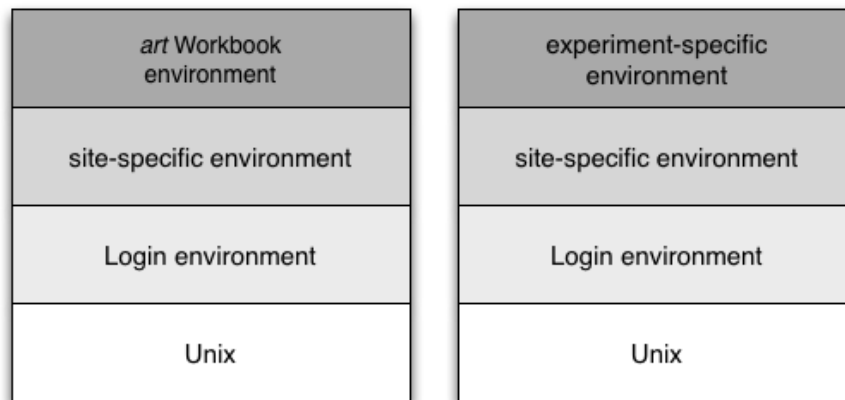


Figure 4.1: Components of the *art* Workbook (left) and experiment-specific (right) computing environments, shown in the order in which they are constructed, starting with the Unix environment

1 4.5.2 Examining and Using Environment Variables

2 One way to see the value of an environment variable is to use the `printenv` command:

```
3 printenv HOME
```

4 At any point in an interactive command or in a shell script, you can tell the shell that
5 you want the value of the environment variable by prefixing its name with the `$` charac-
6 ter:

```
7 echo $HOME
```

8 Here, `echo` is a standard Unix command that copies its arguments to its output, in this case
9 the screen.

10 By convention, environment variables are virtually always written in all capital letters².

11 There may be times when the Workbook instructions tell you to set an environment vari-
12 able to some value. To do so, type the following at the command prompt:

```
13 export ENVNAME=value
```

14 If you read *bash* scripts written by others, you may see the following variant, which ac-
15 complishes the same thing:

```
16 ENVNAME=value
```

```
17 export ENVNAME
```

18 4.6 Paths and \$PATH

sec:paths

19 *Path* (or *PATH*) is an overloaded word in computing. Here are the ways in which it is
20 used:

21 Lowercase `path` can refer to the location of a file or a directory; a path may be absolute or
22 relative, e.g.

```
23 /absolute/path/to/mydir/myfile or
```

²Another type of variable, *shell variables*, are local to the current shell. Whereas environment variables are inherited by subshells, shell variables are not; this is the only difference between them. By convention, shell variables are written in lower or mixed case. These conventions provide a clue to the programmer as to whether changing a variable's value might have consequences outside the current shell.

```
1     relative/path/to/mydir/myfile or
2     ../another/relative/path/to/mydir/myfile
```

3 PATH refers to the standard Unix environment variable set by your login scripts and up-
4 dated by other scripts that extend your environment; it is a colon-separated list of
5 directory names, e.g.,

```
6     /usr/bin:/bin:/usr/sbin:/sbin:/usr/local/bin.
```

7 It contains the list of directories that the shell searches to find programs/files required
8 by Unix shell commands (i.e., PATH is used by the shell to “resolve” commands).

9 “path” generically, refers to any environment variable whose value is a colon-separated
10 list of directory names e.g.,

```
11     /abs/path/a:/abs/path/b:rel/path/c
```

12 In addition, *art* defines a fourth idea, also called a path, that is unrelated to any of the
13 above; it will be described as you encounter it in the Workbook, e.g., Section 9.8.8. issec:prebuilt:paths

14 All of these path concepts are important to users of *art*. In addition to PATH itself, there
15 are three PATH-like environment variables (colon-separated list of directory names) that
16 are particularly important:

17 LD_LIBRARY_PATH (Linux only) used by *art* to resolve dynamic libraries

18 DYLD_LIBRARY_PATH (OS X only) used by *art* to resolve dynamic libraries

19 PRODUCTS used by UPS to resolve external products

20 FHiCL_FILE_PATH use by FHiCL to resolve #include directives.

21 When you source the scripts that setup your environment for *art*, these will be defined and
22 additional colon-separated elements will be added to your PATH. To look at the value of
23 PATH (or the others), enter:

```
24 printenv PATH
```

25 To make the output easier to read by replacing all of the colons with newline characters,
26 enter:

```
27 printenv PATH | tr : \\n
```

28 In the above line, the vertical bar is referred to as a *pipe* and tr is a standard Unix com-
29 mand. A pipe takes the output of the command to its left and makes that the input of the

1 command to its right. The `tr` command replaces patterns of characters with other patterns
2 of characters; in this case it replaces every occurrence of the colon character with the new-
3 line character. To learn why a double back slash is needed, read bash documentation to
4 learn about escaping special characters.

5 4.7 Scripts: Part 2

6 There are two ways to run a bash script (actually three, but two of them are the same).
7 Suppose that you are given a bash script named `file.sh`. You can run any of these
8 commands:

9 `file.sh`

10 `source file.sh`

11 `. file.sh`

12 The first version, `file.sh`, starts a new bash shell, called a subshell, and it executes
13 the commands from `file.sh` in that subshell; upon completion of the script, control re-
14 turns to the parent shell. At the startup of a subshell, the environment of that subshell is
15 initialized to be a copy of the environment of its parent shell. If `file.sh` modifies its
16 environment, then it will modify only the environment of the subshell, leaving the envi-
17 ronment of the parent shell unchanged. This version is called *executing* the script.

18 The second and third versions are equivalent. They do not start a subshell; they execute the
19 commands from `file.sh` in your current shell. If `file.sh` modifies any environment
20 variables, then those modifications remain in effect when the script completes and control
21 returns to the command prompt. This is called *sourcing* the script.

22 Some shell scripts are designed so that they must be sourced and others are designed so
23 that they must be executed. Many shell scripts will work either way.

24 If the purpose of a shell script is to modify your working environment then it must be
25 sourced, not executed. As you work through the Workbook exercises, pay careful attention
26 to which scripts it tells you to source and which to execute. In particular, the scripts
27 to setup your environment (the first scripts you will run) are bash scripts that must be
28 sourced because their purpose is to configure your environment so that it is ready to run
29 the Workbook exercises.

1 Some people adopt the convention that all bash scripts end in `.sh`; others adopt the con-
2 vention that only scripts designed to be sourced end in `.sh` while scripts that must be
3 executed have no file-type ending (no “.something” at the end). Neither convention is uni-
4 formly applied either in the Workbook or in HEP in general.

5 If you would like to learn more about bash, some references are listed in Section 4.10. [\[sec:unix:references\]](#)

6 **4.8 bash Functions and Aliases**

7 The bash shell also has the notion of a *bash function*. Typically bash functions are defined
8 by sourcing a bash script; once defined, they become part of your environment and they
9 can be invoked as if they were regular commands. The setup *product* “command” that you
10 will sometimes need to issue, described in Chapter 7, is an example. [\[ch:ups:setup\]](#) A bash function is
11 similar to a bash script in that it is just a collection of bash commands that are accessible
12 via a name; the difference is that bash holds the definition of a function as part of the
13 environment while it must open a file every time that a bash script is invoked.

14 You can see the names of all defined bash functions using:

```
15 declare -F
```

16 The bash shell also supports the idea of *aliases*; this allows you to define a new command
17 in terms of other commands. You can see the definition of all aliases using:

```
18 alias
```

19 You can read more about bash shell functions and aliases in any standard bash refer-
20 ence.

21 When you type a command at the command prompt, bash will resolve the command using
22 the following order:

- 23 1. Is the command a known alias?
- 24 2. Is the command a bash keyword, such as `if` or `declare`?
- 25 3. Is the command a shell function?
- 26 4. Is the command a shell built-in command?
- 27 5. Is the command found in `$PATH`?

1 To learn how bash will resolve a particular command, enter:

2 `type command-name`

3 4.9 Login Scripts

4 When you first login to a computer running the Unix operating system, the system will
 5 look for specially named files in your home directory that are scripts to set up your working
 6 environment; if it finds these files it will source them before you first get a shell prompt.
 7 As mentioned in Section 4.5, these scripts modify your PATH and define bash functions,
 8 aliases and environment variables. All of these become part of your environment.



9 When your account on a Fermilab computer was first created, you were given standard
 10 versions of the files `.profile` and `.bashrc`; these files are used by bash³. You can
 11 read about login scripts in any standard bash reference. You may add to these files but you
 12 should not remove anything that is present.



13 If you are working on a non-Fermilab computer, inspect the login scripts to understand
 14 what they do.

15 It can be useful to inspect the login scripts of your colleagues to find useful customiza-
 16 tions.

17 If you read generic Unix documentation, you will see that there are other login scripts with
 18 names like, `.login`, `.cshrc` and `.tcshrc`. These are used by the csh family of shells
 19 and are not relevant for the Workbook exercises, which require the bash shell.

20 4.10 Suggested Unix and bash References

21 The following cheat sheet provides some of the basics:

22 ○ <http://mu2e.fnal.gov/atwork/computing/UnixHints.shtml>

23 A more comprehensive summary is available from:

24 ○ <http://www.tldp.org/LDP/.../GNU-Linux-Tools-Summary.html>

³These files are used by the sh family of shells, which includes bash.

1 Information about writing bash scripts and using bash interactive features can be found
2 in:

- 3 ○ BASH Programming - Introduction HOW-TO
4 <http://tldp.org/HOWTO/Bash-Prog-Intro-HOWTO.html>
- 5 ○ Bash Guide for Beginners
6 <http://www.tldp.org/LDP/Bash-Beginners-Guide/html/Bash-Beginners-Guide.html>
- 7 ○ Advanced Bash Scripting Guide
8 <http://www.tldp.org/LDP/abs/html/abs-guide.html>

9 The first of these is a compact introduction and the second is more comprehensive.

10 The above guides were all found at the Linux Documentation Project: Workbook:

- 11 ○ <http://www.tldp.org/guides.html>

12 Books about Unix are numerous, of course. Examples include Mark Sobell's *A practi-*
13 *cal guide to the UNIX system* and Graham Glass' *UNIX for programmers and users:*
14 *a complete guide*, both of which are in the Fermilab library along with many others
15 (<http://ccd.fnal.gov/library/>).

5 Site-Specific Setup Procedure

ch:site-setup

Section 4.5 discussed the notion of a working environment on a computer. This chapter answers the question: How do I make sure that my environment is configured so that I can run the Workbook exercises or my experiment's code?

This chapter will explain how to do this in several different situations:

1. If you are logged in to one of your experiment's computers.
2. If you are logged in to one of the machines supported for the August 2015 *art*/LArSoft course; at this writing there are two machines named `alcourse.fnal.gov` and `alcourse2.fnal.gov`; more may be added.
3. If you install *art* and its tool chain to your own computer.

On every computer that hosts the Workbook, a procedure must be established that every user is expected to follow once per login session. In most cases (NO ν A being a notable exception), the procedure involves only sourcing a shell script (recall the discussion in Section 4.7). In this documentation, we refer to this procedure as the "site-specific setup procedure." It is the responsibility of the people who maintain the Workbook software for each *site*(γ) to ensure that this procedure does the right thing on all the site's machines.



As a user of the Workbook, you will need to know what the procedure is and you must remember to follow it each time that you log in.

For all of the Intensity Frontier experiments at Fermilab, the site-specific setup procedure defines all of the environment variables that are necessary to create the working environment for either the Workbook exercises or for the experiment's own code.



Table 5.1 lists the site-specific setup procedure for each experiment. You will follow the

- 1 procedure when you get to Section 9.6.
- 2 NOvA users should check that their login scripts do not setup any of the UPS products
3 related to *art*. Remove any lines that do; then log out and log in again. In particular, make
4 sure that nothing in your login scripts, either directly or indirectly, executes the following
5 line:

```
6 source /grid/fermiapp/nova/novaart/novasvn/srt/srt.sh
```

- 7 Once you have a clean login, follow the procedure given in Listing 5.1.

Listing 5.1: NOvA setup procedure

```
1 source /nusoft/app/externals/setups  
2 export PRODUCTS=$PRODUCTS:/grid/fermiapp/products/common/db  
3 export ART_WORKBOOK_WORKING_BASE=/nova/app/users  
4 export ART_WORKBOOK_QUAL=s12:e7:nu  
5 export ART_WORKBOOK_OUTPUT_BASE=/nova/app/users
```



Table 5.1: Site-specific setup procedures for Experiments that run *art*; for the equivalent procedure at a non-Fermi site, consult an expert from that site. The MicroBoone experiment maintains two copies of its setup scripts, one on `/grid/fermiapp` and the other in CVMFS space; if you can see both from your computer, they are equivalent. The table also includes the procedure for the machines that support the *art*/LArSoft course in August 2015 and a link to the procedure for private installations.

Experiment	Site-Specific Setup Procedure
ArgoNeut	See the instructions for MicroBoone
Darkside	<code>source /ds50/app/ds50/ds50.sh</code>
LArIAT	Will be available in a future release of the workbook
DUNE	<code>source /grid/fermiapp/lbne/software/setup_lbne.sh</code>
MicroBoone	<code>source /cvmfs/uboone.opensciencegrid.org/products/setup_uboone.sh</code> also, the following will work on the Fermilab site <code>source /grid/fermiapp/products/uboone/setup_uboone.sh</code>
Muon g-2	<code>source source /grid/fermiapp/gm2/setup</code>
Mu2e	<code>setup mu2e</code>
NO ν A	See Listing 5.1 <small><code>lst:site:specific:NOvA</code></small>
<i>art</i> LArSoft Course	<code>source /products/course_setup.sh</code>
Private machine	See Appendix B <small><code>app:ug:install:local</code></small>

6 Get your C++ up to Speed

quick-start

6.1 Introduction

start-intro

This change is for topic one.

There are two goals for this chapter. The first is to provide an overview of the features of C++ that will be important for users of *art*, especially those features that will be used in the Workbook exercises. It does not attempt to cover C++ comprehensively.

You will need to consult standard documentation to learn about any of the features that you are not already familiar with. The examples and exercises in this chapter will in many cases only skim the surface of C++ features that you will need to know how to manipulate as you work through the Workbook exercises and then use C++ code with *art* in your work.



The second goal is to explain the process of turning source code files into an *executable program*. The two steps in this process are *compiling* and *linking*. In informal writing, the word *build* is sometimes used to mean just compiling or just linking, but usually it refers to the two together.

This chapter is designed around a handful of exercises, each of which you will first build and run, then “pick apart” to understand how the results were obtained.

6.2 File Types Used and Generated in C++ Programming

A typical program consists of many source code files, each of which contains a human-readable description of one or more components of the program. In the Workbook, you will see source code files written in the C++ computer language; these files have names that end in `.cc`.¹ In C++, there is a second sort of source code file, called a *header file*. These typically have names that end in `.h`²; in most cases, but not all, the source file has an associated header file with the same base name but with the different suffix. A header file can be thought of as the “parts list” for its corresponding source file; you will see how these are used in Section 6.5.

In the compilation step each source file is translated into *machine code*, also called *binary code* or *object code*, which is a set of instructions, in the computer’s native language, to do the tasks described by the source code. The output of the compilation step is called an *object file*; in the examples you will see in the Workbook, object files always end in `.o`.³

But an object file, by itself, is not an *executable program*. It is not executable primarily because it lacks the instructions that tell the operating system how to start executing the instructions in the file.

It is often convenient to collect related groups of object files and put them into *libraries*.

There are two kinds of library files, static libraries and dynamic libraries. Static libraries are not used by *art*, and we do not discuss them further; when this document refers to a library, it means a dynamic library. Putting many object files into a single library allows you to use them as a single coherent entity. We will defer further discussion of libraries until more background information has been provided.

The job of the *linking* step is to read the information found in the various libraries and object files and form them into either a *dynamic library* or an *executable program*. When you run the linker, you tell it the name of the file it is to create. It is a common, but not universal, practice that the filename of an executable program has no extension (i.e., no `.something` at the end). Dynamic libraries on Linux typically have the extension `.so`, and on OS X they typically have the extension `.dylib`.

¹Other conventions include `.cpp` and `.cxx`.

²Other conventions are `.hh`, `.hpp` and `.hxx`.

³Other conventions include `.os`.

1 After the linker has finished, you can run your executable program typing the filename
2 of the program at the bash command prompt. If you do not have the current directory on
3 the PATH, you need to preface the filename of the program with `./` (a dot followed by
4 a forward slash). At this time, the loader does the final steps of linking the program, to
5 allow the use of instructions and data in the dynamic libraries to which your program is
6 linked.

7 A typical program links both to libraries that were built from the program's source code
8 and to libraries from other sources. Some of these other libraries might have been devel-
9 oped by the same programmer as general purpose tools to be used by his or her future
10 programs; other libraries are provided by third parties, such as *art* or your experiment.
11 Many C++ language features are made available to your program by telling the linker to
12 use libraries provided by the C++ compiler vendor. Other libraries are provided by the
13 operating system.

14 Now that you know about libraries, we can give a second reason why an object file, by
15 itself, is not an executable program: until it is linked, it does not have access to the func-
16 tions provided by any of the external libraries. Even the simplest program will need to be
17 linked against some of the libraries supplied by the compiler vendor and by the operating
18 system.

19 The names of all of the libraries and object files that you give to the linker is called the
20 *link list*.

21 6.3 Establishing the Environment

22 6.3.1 Initial Setup

23 To start these exercises for the first time, do the following:

1. Log into the node that you will use for Workbook exercises.
2. Follow the site-specific setup procedure from Table 5.1. [tab:expsetups](#)
3. Create an empty working directory and `cd` to it.
4. Run one of the following commands to download the gzipped tar file. You only need to run one of them. Try `wget` and, if your

24

system does not have `wget` installed, use `curl`. In both cases a single command is shown on two lines:

```
wget https://web.fnal.gov/project/\
ArtDoc/Shared%20Documents/C++UpToSpeed_v0_88.tar.gz
```

```
curl -O https://web.fnal.gov/project/\
ArtDoc/Shared%20Documents/C++UpToSpeed_v0_88.tar.gz
```

5. Run these commands to unpack the tar gzipped file, and get a directory listing.

```
tar xzf C++UpToSpeed_v0_88.tar.gz
```

```
rm C++UpToSpeed_v0_88.tar.gz
```

```
cd GettingYourC++UpToSpeed
```

```
ls
```

```
BasicSyntax Build Classes Libraries setup.sh
```

6. To select the correct compiler version and define a few environment variables that will be used later in these exercises, run:

```
source setup.sh
```

1

- 2 After these steps, you are ready to begin the exercise in Section [6.4](#). [|sec:quick-start-syntax](#)

6.3.2 Subsequent Logins

If you log out and log back in again, reestablish your environment by following these steps:

1. Log into the node that you will normally use.
2. Follow the site-specific setup procedure from Table 5.1.
3. `cd` to the working directory you created in Section 6.3.1.
4. Run the command: `source setup.sh`
5. `cd` to the directory that contains the exercise you want to work on.

6.4 C++ Exercise 1: Basic C++ Syntax and Building an Executable

6.4.1 Concepts to Understand

This section provides a program that illustrates the concepts in C++ that are assumed knowledge for the Workbook material. Brief explanations are provided, but in many cases you will need to consult other sources to gain the level of understanding that you will need. Several C++ references are listed in Section 6.9.

This sample program will introduce you to the following C++ concepts and features:

- how to indicate comments
- what is a *main program*
- how to compile, link and run the main program
- how to distinguish between source, object, library, and executable files
- how to print to standard output, `std::cout`
- what is a *type*
- how to declare and define *variables*(γ) of some of the frequently used built-in types: `int`, `float`, `double`, `bool`

- 1 ○ the `{ }` initializer syntax (in addition to other forms)
- 2 ○ assignment of values to variables
- 3 ○ what are arrays, and how to declare and define them
- 4 ○ several forms of looping
- 5 ○ comparisons: `==`, `!=`, `<`, `>`, `>=`, `<=`
- 6 ○ `if` and `if-else`
- 7 ○ what are pointers, and how to declare and define them
- 8 ○ what are references, and how to declare and define them
- 9 ○ `std::string` (a type from the C++ Standard Library (*std(γ)*)
- 10 ○ what is the *class template* from the standard library, `std::vector<T>`⁴

11 The above list explicitly does not include classes, objects and inheritance, which will be
 12 discussed in Sections 6.7 and a future section on inheritance. **FIXME:** *Reminder: it's in*
 13 *misc temp; at 59.8.*

14 6.4.2 How to Compile, Link and Run

15 In this section you will learn how to compile, link and run the small C++ program that
 16 illustrates the features of C++ that are considered prerequisites to the Workbook exercises.

17
 18 Run the following procedure. The idea here is for you to get used to the steps and see what
 19 results you get. Then in Section 6.4.3 you will examine the source file and output.

20 To compile, link and run the sample C++ program, called `t1`:

1. If not yet done, log in and establish the working environment (Section 6.3).
2. From the working directory you have made, `cd` to the directory for this exercise and look at its contents.

21

⁴You need to know how to use `std::vector<T>` but you do not need to understand how it works or how to write your own templates.


```
cd BasicSyntax/v1/
```

```
ls
```

```
build t1.cc t1_example.log
```

The file `t1.cc` contains the source code of the *main program*, which is a function called `main`. The file `build` is a script that will compile and link the code. The file `t1_example.log` is an example of the output expected when you run `t1`.

3. Compile and link the code (one step); then look at a directory listing:

```
./build
```

```
t1.cc: In function 'int main()':  
t1.cc:43:26: warning: 'k' may be  
used uninitialized in this function  
[-Wuninitialized]
```

```
ls
```

```
build t1 t1.cc t1_example.log
```

The script named `build` compiles and links the code, and produces the executable file `t1`. The warning message, issued by the compiler, also comes during this step.

4. Run the executable file sending output to a log file:

```
./t1 > t1.log
```

1

2 Just to see how the exercise was built, look at the script `BasicSyntax/v1/build` that
3 you ran to compile and link `t1.cc`; the following command was issued:

```
4 c++ -Wall -Wextra -pedantic -std=c++11 -o t1 t1.cc
```

5 This turned the source file `t1.cc` into an executable program, named `t1` (the argument to

1 the `-o` (for “output”) option). We will discuss compiling and linking in Section 6.5.

2 6.4.3 Discussion

3 Look at the file `t1.cc`, in particular the relationship between the lines in the program and
 4 the lines in the output, and see how much you understand. Remember, you will need to
 5 consult standard documentation to learn about any of the features that you are not already
 6 familiar with; some are listed in Section 6.9. Note that some questions may be answered
 7 in Section 6.4.3.

8 In the source file, it is important to first point out the function called the *main program*.
 9 Every program needs one, and execution of the program takes place within the braces of
 10 this function, which is written

```
11 int main() {
12     ...executable code...
13 }
```

14 Compare your output with the standard example:

```
15 diff t1.log t1_example.log
```

16 There will almost certainly be a handful of differences, which we will discuss in Sec-
 17 tion 6.4.3.1.

18 The following sections correspond to sections of the code in `BasicSyntax/v1/t1.cc`
 19 and provide supplementary information.

20 6.4.3.1 Primitive types, Initialization and Printing Output

21 All variables, parameters, arguments, and so on in C++ need to have a *type*, e.g., `int`,
 22 `float`, `bool`, or another so-called primitive (or built-in) type, or a more complicated
 23 type defined by a *class* or *structure*. The code in this exercise introduces the primitive
 24 types.

25 Now, about the handful of differences in the output of one run versus another. There are
 26 two main sources of the differences: (1) an uninitialized variable and (2) variation in object
 27 addresses from run to run.

1 In `t1.cc`, the line `int k;` declares that `k` is a variable whose type is `int` but it does not
2 initialize the variable. Therefore the value of the variable `k` is whatever value happened
3 to be sitting in the memory location that the program assigned to `k`. Each time that the
4 program runs, the operating system will put the program into whatever region of memory
5 makes sense to the operating system; therefore the address of any variable, and thus the
6 value returned, may change unpredictably from run to run.

7 This line is also the source of the warning message produced by the `build` script. This
8 line was included to make it clear what we mean by *initialized* variables and *uninitialized*
9 variables. Uninitialized variables are frequent sources of errors in code and therefore you
10 should *always* initialize your variables. In order to help you establish this good coding
11 habit, the remaining exercises in this series and in the Workbook include the compiler
12 option `-Werror`. This tells the compiler to promote warning messages to error level and
13 to stop compilation without producing an output file.

14 See Section 6.4.3.6 for other output that may vary between program runs.

15 6.4.3.2 Arrays

16 The next section of the example code introduces *arrays*, sometimes called *C-style* arrays
17 to distinguish them from `std::array`, a class template element of the C++ Standard
18 Library. Classes will be discussed in Section 6.7, and class templates will first be used in
19 Chapter 6.7.

20 While you might find use of arrays in existing code, we recommend avoiding them in new
21 code arrays, and using either `std::vector` or `std::array`. See Section 6.4.3.5 for
22 an introduction to these types.

23 6.4.3.3 Equality testing

24 Two variables which refer to different objects that contain the same value (either by design
25 or by coincidence) are *equal*. Equality is tested using the equality testing operator, `==`. It
26 is important to distinguish between the assignment operator (`=`) and the equality testing
27 operator. Using `=` where `==` is intended is a common mistake.

28 Another distinction to be made is that of two variables being *identical* versus *equal*. In
29 contrast to equality, two variables are identical if they refer to the same object, and thus

1 have the same memory address. How can two variables be identical? One common case
 2 can be seen in the call to a function like `maxSize`:

```
3 #include <algorithm>
4 #include <string>
5 std::size_t maxSize(std::string const& a,
6                     std::string const& b) {
7     return std::max(a.size(), b.size());
8 }
```

9 If we consider the call:

```
10 std::string s("cow");
11 auto sz = maxSize(s, s);
```

12 then, in the body of the function `maxSize`, and for this call, the variables `a` and `b` refer
 13 to the same object—so they are identical.

14 6.4.3.4 Conditionals

sec:conditionals

15 The primary conditional statements in C++ are `if` and `if-else`. There is also the ternary
 16 operator, `?:`. The ternary operator, which evaluates an expression as true or false then
 17 chooses a value based on the result, can replace more cumbersome code with a single line.
 18 It is especially useful in places where an expression, not a statement, is needed. It works
 19 this way (shown here on three lines to emphasize the syntax):

```
20 // Note: this is pseudocode, not C++
21 type variable-to-initialize (expression-to-evaluate) ?
22                             value-if-true :
23                             value-if-false;
```

24 An example is shown in the code.

25 6.4.3.5 Some C++ Standard Library Types

ne-library-types

26 The C++ Standard Library is quite large, and contains many classes, functions, class tem-
 27 plates, and function templates. Our sample code introduces only three: the class `std::string`,
 28 and the templates `std::vector<T>` and `std::array`.

1 A `std::vector<T>` behaves much like is an array of objects of some type `T` (e.g., `int`
2 or `std::string`). It has the extra capability that its size can change as needed, unlike a
3 C-style array, whose size is fixed.

4 The `std::array` is new in C++, and should be used in preference to the older C-style
5 array discussed in Section 6.4.3.2 due to its greater capabilities. Unlike the C-style ar-
6 ray, `std::array` knows its own size, can be copied, and can be returned from a func-
7 tion.

8 6.4.3.6 Pointers

9 A pointer is a variable whose value is the memory address of another object. The type of
10 pointer must correspond to the type of the object to which it points.

11 In addition to the sources of difference in the program output between runs discussed in
12 Section 6.4.3.1, another stems from the line:

```
13 float* pa = &a;
```

14 This line declares a variable `pa` and initializes it to be the memory address of the variable
15 `a`. The variable `a` is of type `float`; therefore `pa` must be declared as type *pointer*(γ) to
16 `float`.

17 Note that this line could have been written with the asterisk next to `pa`:

```
18 float *pa = &a;
```

19 This latter style is common in the C community. In the C++ community, the former style
20 is preferred, because it emphasizes the type of the variable `pa`, rather than the type of the
21 expression `*pa`.

22 Since the address may change from run to run, so may the printout that starts `pa =`.

23 The next line,

```
24 std::cout << "*pa = " << *pa << std::endl;
```

25 shows how to access the value to which a pointer points. This is called *deferencing*
26 the pointer, or *following* the pointer, and is done with the dereferencing operator, `*`. The

1 expression `*pa` dereferences the pointer `pa`, yielding the value of the object to which `pa`
 2 points, in this case, the value of `a`.

3 In Section [6.7](#) you will learn about *classes*. One attribute of classes is that they have sep-
 4 arately addressable parts, called *members*. Members of a class are selected using syntax
 5 like `classname.membername`. The combination of dereferencing a pointer and selecting a
 6 member of the pointed-to object (the *pointee*) can be done in two steps: first dereferencing
 7 then selecting, or in one step using the member selection operator `operator->()`. The
 8 following two expressions are equivalent:

```
9 (*panimal).size()
10 panimal->size()
```

11 In the example code, the lines

```
12 std::cout << "The size of animal is: "
13           << (*panimal).size() << std::endl;
14
15 std::cout << "The size of animal is: "
16           << panimal->size() << std::endl;
```

17 do exactly the same thing. Note that the parentheses in the first line are necessary because
 18 the precedence of `.` is higher than that of `*`.

19 Note that in many situations, the compiler is free to convert an array-of-`T` into a pointer-
 20 to-`T`. In such cases, the value of the pointer-to-`T` is the address of the initial element in the
 21 array.

22 6.4.3.7 References

[sssec:references](#)

23 A *reference* is a variable that acts as an alias for another object, and it can not be re-seated
 24 to reference a different object. It is not an object itself, and thus a reference does not have
 25 an address. The address-of operator `operator&`, when used on a reference, yields the
 26 address of the referent:

```
27 float a;
28 float& ra = a;
29 float* p = &a;
```

```
1 float* q = &ra;
```

2 The values of `p` and `q` will be the same. Because they print memory addresses determined
3 by the compiler and linker, the lines in the printout that start `&a =` and `&ra =` may also
4 change from run to run.

5 6.4.3.8 Loops

6 Loops, also called *iteration statements*, appear in several forms in C++. The most prevalent
7 is the `for` loop. New in C++11 is the *range-based for* loop; this is the looping construction
8 that should be preferred for cases to which it applies.

9 6.5 C++ Exercise 2: About Compiling and Linking


10 6.5.1 What You Will Learn

11 In the previous exercise, the user code was found in a single file and the build script
12 performed compiling and linking in a single step. For all but the smallest programs, this
13 is not practical. It would mean, for example, that you would need to recompile and relink
14 everything when you made even the smallest change anywhere in the code; generally this
15 would take much too long. To address this, some computer languages, including C++,
16 allow you to break up a large program into many smaller files and rebuild only a small
17 subset of files when you make changes in one.

18 There are two exercises in this section. In the first one the source code consists of three
19 files. This example has enough richness to discuss the details of what happens during
20 compiling and linking, without being overwhelming. The second exercise introduces the
21 idea of libraries.

22 6.5.2 The Source Code for this Exercise

23 The source code for this exercise is found in `Build/v1`, relative to your working direc-
24 tory. The relevant files are `t1.cc`, `times2.cc` and `times2.h`. Open these files and
25 read along with the discussion below.

1 The file `t1.cc` contains the source code for the function `main` for this exercise. Every
 2 C++ program must have one and only one function named `main`, which is where the pro-
 3 gram actually starts execution. Note that the term *main program* sometimes refers to this
 4 function, but other times refers to the source file that contains it. In either case, *main pro-*
 5  *gram* refers to this function, either directly or indirectly. For more information, consult any
 6 standard C++ reference. The file `times2.h` is a header file that declares a function named
 7 `times2`. The file `times2.cc` is another source code file; it provides the definition of
 8 that function.

9 Look at `t1.cc`: it both declares and defines the program's function `main`, with the sig-
 10 nature `int main()`: it takes no arguments, and returns an `int`. A function with this
 11 *signature*(γ) has special meaning to the compiler and the linker: they recognize it as a C++
 12 *main program*. There are other signatures that the compiler and linker will recognize as a
 13 C++ main program; consult the standard C++ documentation.



14 To be recognized as a main program, there is one more requirement: `main` must be de-
 15 clared in the global namespace.

16 The body of the main program (between the braces), declares and defines a variable `a`
 17 of type `double` and initializes it to the value of 3.0; it prints out the value of `a`. Then
 18 it calls a function that takes `a` as an argument and prints out the value returned by that
 19 function.

20 You, as the programmer using that function, need to know what the function does but the
 21 C++ compiler doesn't. It only needs to know the name, argument list and return type of the
 22 function — information that is provided in the header file, `times2.h`. This file contains
 23 the line

```
24 double times2(double);
```

25 This line is called the *declaration*(γ) of the function. It says (1) that the identifier `times2`
 26 is the name of a function that (2) takes an argument of type `double` (the “double” inside
 27 the parentheses) and (3) returns a value of type `double` (the “double” at the start of the
 28 line). The file `t1.cc` includes this header file, thereby giving the compiler these three
 29 pieces of information it needs to know about *function*.

30 The other three lines in `times2.h` make up an *include guard*, described in Appendix F.
 31 In brief, they deal with the following scenario: suppose that we have two header files, `A.h`

lapp:ug:include

1 and `B.h`, and that `A.h` includes `B.h`; there are many scenarios in which it makes good
2 sense for a third file, either `.h` or `.cc`, to include both `A.h` and `B.h`. The include guards
3 ensure that, when all of the includes have been expanded, the compiler sees exactly one
4 copy of `B.h`.

5 Finally, the file `times2.cc` contains the source code for the function named `times2`:

```
6  
7 double times2(double i) {  
8     return 2 * i;  
9 }
```

10 It names its argument `i`, multiplies this argument by two and returns that value. This code
11 fragment is called the *definition* of the function or the *implementation*(γ) of the function.
12 (The C++ standard uses the word *definition* but *implementation* is in common use.)

13 We now have a rich enough example to discuss a case in which the same word is frequently
14 used for two different things — instead of two words used for the same thing.

15 Sometimes people use the phrase “the source code of the function named `times2`” to
16 refer collectively to both `times2.h` and `times2.cc`; sometimes they use it to refer
17 exclusively to `times2.cc`. Unfortunately the only way to distinguish the two uses is
18 from context.

19 The phrase *header file* always refers unambiguously to the `.h` file. The term *implemen-*
20 *tation file* is used to refer unambiguously to the `.cc` file. This name follows from the its
21 contents: it describes how to implement the items declared in the header file.

22 Based on the above description, when this exercise is run, we expect it to print out:

```
23 a = 3  
24 times2(a) 6
```

25 6.5.3 Compile, Link and Run the Exercise

26 To perform this exercise, first log in and `cd` to your working directory if you haven't al-
27 ready, then

1. From the working directory you have made, `cd` to the directory for this exercise and look at its contents.

```
cd Build/v1
```

```
ls
```

```
build build2 t1.cc times2.cc times2.h
```

The two files, `build` and `build2` are scripts that show two different ways to build the code.

2. Compile and link this exercise, then get an updated directory listing:

```
./build
```

```
ls
```

```
build build2 t1 t1.cc t1.o times2.cc times2.h  
times2.o
```

Notice the new files `t1`, `t1.o` and `times2.o`.

3. Run the exercise:

```
./t1
```

```
a = 3  
times2(a) 6
```

1

2 This matches the expected printout.

3 Look at the file `build` that you just ran. It has three steps:4 1. It compiles the main program, `t1.cc`, into the object file (with the default name)
5 `t1.o` (which will now be the thing that the term *main program* refers to):6 `++ -Wall -Wextra -pedantic -Werror -std=c++11 -c t1.cc`7 2. It (separately) compiles `times2.cc` into the object file `times2.o`:

```
1      c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c times2.cc
```

2 3. It links `t1.o` and `times2.o` (and some system libraries) to form the executable
3 program `t1` (the name of the main program is the argument of the `-o` option):

```
4      c++ -std=c++11 -o t1 t1.o times2.o
```

5 You should have noticed that the same command, `c++`, is used both for compiling and
6 linking. The full story is that when you run the command `c++`, you are actually running a
7 program that parses its command line to determine which, if any, files need to be compiled
8 and which, if any, files need to be linked. It also determines which of its command line
9 arguments should be forwarded to the compiler and which to the linker. It then runs the
10 compiler and linker as many times as required.

11 If the `-c` option is present, it tells `c++` to compile only, and not to link. If `-c` is specified,
12 the source file(s) to compile must also be specified. Each of the files will be compiled to
13 create its corresponding object file and then processing stops. In our example, the first two
14 commands each compile a single source file. Note that if any object files are given on the
15 command line, `c++` will issue a warning and ignore them.

16 The third command (with no `-c` option) is the linking step. Even if the `-c` option is missing,
17 `c++` will first look for source files on the command line; if it finds any, it will compile
18 them and put the output into temporary object files. In our example, there are none, so it
19 goes straight to linking. The two just-created object files are specified (at the end, here, but
20 the order is not important); the `-o t1` portion of the command tells the linker to write its
21 output (the executable) to the file `t1`.

22 As it is compiling the main program, `t1.cc`, the compiler recognizes every function that
23 is defined within the file and every function that is called by the code in the file. It recog-
24 nizes that `t1.cc` defines a function `main` and that `main` calls a function named `times2`,
25 whose definition is not found inside `t1.cc`. At the point that `main` calls `times2`, the
26 compiler will write to `t1.o` all of the machine code needed to prepare for the call; it will
27 also write all of the machine code needed to use the result of `times2`. In between these
28 two pieces, the compiler will write machine code that says “call the function whose mem-
29 ory address is” but it must leave an empty placeholder for the address. The placeholder is
30 empty because the compiler does not know the memory address of that function.

31 The compiler also makes a table that lists all functions defined by the file and all functions
32 that are called by code within the file. The name of each entry in the table is called a

1 *linker symbol* and the table is called a *symbol table*. When the compiler was compiling
2 `t1.cc` and it found the definition of the main program, it created a linker symbol for
3 the main program and added a notation to say the this file contains the definition of that
4 symbol. When the compiler was compiling `t1.cc` and it encountered the call to `times2`,
5 it created a linker symbol for this function; it marked this symbol as an *undefined reference*
6 (because it could not find the definition of `times2` within `t1.cc`). The symbol table
7 also lists all of the places in the machine code of `t1.o` that are placeholders that must be
8 updated once the memory address of `times2` is known. In this example there is only one
9 such place.

10 When the compiler writes an object file, it writes out both the compiled code and the table
11 of linker symbols.

12 The symbol table in the file `times2.o` is simple; it says that this file defines a function
13 named `times2` that takes a single argument of type double and that returns a double. The
14 name in the symbol table encodes not only the function name, but also the number and
15 types of the function arguments. These are necessary for *overload resolution*(γ).

16 The job of the linker (also invoked by the command `c++`) is to play match-maker. First
17 it inspects the symbol tables inside all of the object files listed on the command line and
18 looks for a linker symbol that defines the location of the main program. If it cannot find
19 one, or if it finds more than one, it will issue an error message and stop. In this example
20

- 21 1. The linker will find the definition of a main program in `t1.o`.
- 22 2. It will start to build the executable (output) file by copying the machine code from
23 `t1.o` to the output file.
- 24 3. Then it will try to resolve the unresolved references listed in the symbol table of
25 `t1.o`; it does this by looking at the symbol tables of the other object files on the
26 command line. It also knows to look at the symbol tables from a standard set of
27 compiler-supplied and system-supplied libraries.
- 28 4. It will discover that `times2.o` resolves one of the external references from `t1.o`.
29 So it will copy the machine code from `times2.o` to the executable file.
- 30 5. It will discover that the the other unresolved references in `t1.o` are found in the
31 compiler-supplied dynamic libraries. It will put into the executable the necessary

- 1 information to resolve these references at the time when the program is run.
- 2 6. Once all of the machine code has been copied into the executable, the compiler
3 knows the memory address of every function, or where to find them at run-time.
4 The compiler can then go into the machine code, find all of the placeholders and
5 update them with the correct memory addresses.
- 6 Sometimes resolving one unresolved reference will generate new ones. The linker iterates
7 until (a) all references are resolved and no new unresolved references appear (success)
8 or (b) the same unresolved references continue to appear (error). In the former case, the
9 linker writes the output to the file specified by the `-o` option; if no `-o` option is specified the
10 linker will write its output to a file named `a.out`. In the latter case, the linker issues an
11 error message and does not write the output file.
- 12 After the link completes, the files `t1.o` and `times2.o` are no longer needed because
13 everything that was useful from them was copied into the executable `t1`. You may delete
14 the object files, and the executable will still run.

15 6.5.4 Alternate Script `build2`

16 The script `build2` shows an equivalent way of building `t1` that is commonly used for
17 small programs; it does it all on one line. To exercise this script:

1. Stay in the same directory as before, `Build/v1`.
2. Clean up from the previous build and look at the directory contents:

```
rm t1 t1.o times2.o
```

```
ls
```

```
build build2 t1.cc times2.cc times2.h
```

3. Run the `build2` script, and again look at directory contents:

```
./build2
```

18

```
ls
```

```
build build2 t1 t1.cc times2.cc times2.h
```

Note that `t1` was created but there are no object files.

4. Execute the program that you just built:

```
./t1
```

```
a = 3
```

```
times2(a) 6
```

1

2 Look at the script `build2`; it contains only one command:

```
3 c++ -Wall -Wextra -pedantic -Werror -std=c++11 -o t1 t1.cc times2.cc
```

4 This script automatically does the same operations as `build` but it knows that the object files are temporaries. Perhaps the command `c++` kept the contents of the two object files in memory and never actually wrote them out as disk files. Or, perhaps, the command `c++` did explicitly create disk files and deleted them when it was finished. In either case you don't see them when you use `build2`.

9 6.5.5 Suggested Homework

ilding-homework

10 It takes a bit of experience to decipher the error messages issued by a C++ compiler. The
11 three exercises in this section are intended to introduce you to them so that you (a) get
12 used to looking at them and (b) understand these particular errors if/when you encounter
13 them later.

14 Each of the following three exercises is independent of the others. Therefore, when you
15 finish with each exercise, you will need to undo the changes you made in the source file(s)
16 before beginning the next exercise.

17 1. In `Build/v1/t1.cc`, comment out the include directive for `times2.h`; rebuild
18 and observe the error message.

19 2. In `Build/v1/times2.cc`, change the return type to `float`; rebuild and observe

1 the error message.

2 3. In `Build/v1/t1.cc`, change `double a = 3.` to `{cppcodefloat a = 3.;` rebuild
3 and run. This will work without error and will produce the same output as before.

4 The first homework exercise will issue the diagnostic:

```
5 t1.cc: In function 'int_main()':  
6 t1.cc:9:40: error: 'times2' was not declared in this scope
```

7 When you see a message like this one, you can guess that either you have not included a
8 required header file or you have misspelled the name of the function.

9 The second homework exercise will issue the diagnostic (second and last lines split into
10 two here):

```
11 times2.cc: In function 'float_times2(double)':  
12 times2.cc:3:22: error: new declaration 'float_times2(double)'  
13 float times2(double i) {  
14 ^  
15 In file included from times2.cc:1:0:  
16 times2.h:4:8: error: ambiguates old declaration 'double_times2(double)'  
17 double times2(double);  
18 ^
```

19 This error message says that the compiler has found two functions that have the same sig-
20 nature but different return types. The compiler does not know which of the two functions
21 you want it to use.

22 The bottom line here is that you must ensure that the definition of a function is consistent
23 with its declaration; and you must ensure that the use of a function is consistent with its
24 declaration.

25 The third homework exercise illustrates the C++ idea of *implicit (type) conversion*; in this
26 case the compiler will make a temporary variable of type `double` and set its value to that
27 of `a`, as if the code included:

```
28 double tmp = a;  
29 ...  
30 std::cout << "times2(a) " << times2(tmp) << std::endl;
```

31 Consult the standard C++ documentation to understand when implicit type conversions
32 may occur; see Section 6.9. [\[sec.cpp:references\]](#)

6.6 C++ Exercise 3: Libraries

Multiple object files can be grouped into a single file known as a *library*, obviating the need to specify each and every object file when linking; you can reference the libraries instead. This simplifies the multiple use and sharing of software components. Libraries were introduced in Section 6.1); here we introduce the building of libraries.

6.6.1 What You Will Learn

In this section you will repeat the example of Section 6.5 with a variation. You will create a library from `times2.o` and use that library in the link step. This pattern generalizes easily to the case that you will encounter in your experiment software, where object libraries will typically contain many object files.

6.6.2 Building and Running the Exercise

To perform this exercise, do the following:

1. From the working directory you have made, `cd` to the directory for this exercise and look at its contents.

```
cd Libraries/v1
```

```
ls
```

```
build build2 build3 t1.cc times2.cc times2.h
```

The three files, `times2.cc`, `times2.h` and `t1.cc` are identical to those from the previous exercise. The three files, `build`, `build2` and `build3` are scripts that show three different ways to build the main program in this exercise.

2. Compile and link this exercise using `build`, then compare the directory listing to that taken pre-build:


```
./build
```

```
ls
```

```
build build3 libpackage1.so t1.cc
```

```
build2 t1 t1.o times2.cc times2.h times2.o
```

Note that on OS X, the library will be called `libpackage1.dylib` rather than `libpackage1.so`.

3. Execute the main program:

```
./t1
```

```
a = 3
```

```
times2(a) 6
```

1

2 This matches the expected printout. Now let's look at the script `build`. It has four parts
3 which do the following things:

4 1. Compiles `times2.cc`; the same as the previous exercise:

```
5 c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c times2.cc
```

6 2. Creates the library named `libpackage1.so` (on OS X, the standard suffix for a
7 dynamic library is different, and the the library is called `libpackage1.dylib`)
8 from `times2.o`:

```
9 c++ -o libpackage1.so -shared times2.o
```

```
10 or
```

```
11 c++ -o libpackage1.dylib -shared times2.o
```

12 Note that the name of the library must come before the name of the object file. The
13 flag `-shared` directs the linker to create a dynamic library rather than an executable
14 image; without this flag, this command would produce an error complaining about
15 the lack of a `main` function.

16 3. Compiles `t1.cc`; the same as the previous exercise:

```
17 c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c t1.cc
```

18 4. Links the main program against the dynamic library (either `libpackage1.so` or
19 `libpackage1.dylib`) and, implicitly, the necessary system libraries:

```

1      c++ -o t1 t1.o libpackage1.so
2      or
3      c++ -o t1 t1.o libpackage1.dylib

```



Note that from this point on, in order to reduce the verbosity of some library descriptions, we will use the Linux form of library names (e.g. `libpackage1.so`). If you are working on OS X, you will need to translate all these to the OS X form (e.g. `libpackage1.dylib`).

The two new features are in step 2, which creates the dynamic library, and step 4, in which `times2.o` is replaced in the link list with dynamic library. If you have many object files to add to the library, you may add them one at a time by repeating step 2 or you may add them all in one command. When you do the latter you may name each object file separately or may use a wildcard:

```
12 c++ -o libpackage1.so -shared *.o
```

In the filename `libpackage1.so` the string `package1` has no special meaning; it was an arbitrary name chosen for this exercise.



The other parts of the name, the prefix `lib` and the suffix `.so`, are part of a long-standing Unix convention. Some Unix tools presume that libraries are named following this convention, so you should always follow it. The use of this convention is illustrated by the scripts `build2` and `build3`.

To perform the exercise using `build2`, stay in the same directory and clean up and rebuild as follows:

1. Remove files built by `build1`:

```
rm times2.o t1.o libpackage1.* t1
```

2. Build the code with `build2` and look at the directory contents:

```
./build2
```

```
ls
```

```
build build3 times2.h libpackage1.so t1.cc
build2 times2.cc times2.o t1 t1.o
```

21

3. Run `./t1` as before.

1

2 The only difference between `build` and `build2` is the link line. The version from
3 `build` is:

```
4 c++ -o t1 t1.o libpackage1.so
```

5 while that from `build2` is:

```
6 c++ -o t1 t1.o -L. -lpackage1
```

7 In the script `build`, the path to the library, relative or absolute, is written explicitly on
8 the command line. In the script `build2`, two new elements are introduced. The command
9 line may contain any number of `-L` options; the argument of each option is the name of
10 a directory. The ensemble of all of the `-L` options forms a search path to look for named
11 libraries; the path is searched in the order in which the `-L` options appear on the line.
12 The names of libraries are specified with the `-l` options (this is a lower case letter L, not
13 the numeral one); if a `-l` option has an argument of `XXX` (or `package1`), then the linker
14 will search the path defined by the `-L` options for a file with the name `libXXX.so` (or
15 `libpackage1.so`).

16 In the above, the dot in `-L.` is the usual Unix pathname that denotes the current working
17 directory. And it is important that there be no whitespace after a `-L` or a `-l` option and its
18 value.

19 This syntax generalizes to multiple libraries in multiple directories as follows. Suppose
20 that the libraries `libaaa.so`, `libbbb.so` and `libccc.so` are in the directory `L1`
21 and that the libraries `libddd.so`, `libeee.so` and `libfff.so` are in the directory
22 `L2`. In this case, the link list would look like:

```
23 -Lpath-to-L1 -laaa -lbbb -lccc -Lpath-to-L2 -lddd -leee -lfff
```

24 The `-L -l` syntax is in common use throughout many Unix build systems: if your link list
25 contains many object libraries from a single directory then it is not necessary to repeatedly
26 specify the path to the directory; once is enough. If you are writing link lists by hand, this is
27 very convenient. In a script, if the path name of the directory is very long, this convention
28 makes a much more readable link list.

- 1 To perform the exercise using `build3`, stay in the same directory and clean up and rebuild
- 2 as follows:

1. Remove files built by `build2`:

```
rm times2.o t1.o libpackage1.* t1
```

2. Build the code with `build3` and look at the directory contents:

```
./build3
```

```
ls
```

```
build build3 times2.h libpackage1.so t1.cc  
build2 times2.cc times2.o t1
```

3. Run `./t1` as before

3

- 4 The difference between `build2` and `build3` is that `build3` compiles the main program
- 5 and links it all one one line instead of two.

6.7 Classes

6

ck-start-classes

6.7.1 Introduction

7

ct-classes-intro

- 8 The comments in the sample program used in Section 6.4 emphasized that every variable
- 9 has a type: `int`, `float`, `std::string`, `std::vector<std::string>`, and so
- 10 on. One of the basic building blocks of C++ is that users may define their own types; user-
- 11 defined types may be built-up from all types, including other user-defined types.

- 12 The language features that allow users to define types are the `class(γ)` and the `struct(γ)`.
- 13 As you work through the Workbook exercises, you will see classes that are defined by the
- 14 Workbook itself; you will also see classes defined by the toyExperiment UPS product; you
- 15 will see classes defined by `art` and you will see classes defined by the many UPS products
- 16 that support `art`. You will also write some classes of your own. When you work with the

impleclass

```
class MyClassName {  
    // required: declarations of all members of the class  
    // optional: definitions of some members of the class  
};
```

Listing 6.1: Layout of a class.

1 software for your experiment you will work with classes defined within your experiment's
2 software.

3 Classes and structures (types introduced by either `class` or `struct`) are called *user-*
4 *defined types*. `std::string`, etc., although defined by the Standard Library, are still
5 called user-defined types.

6 In general, a class is specified by both a *definition* (that describes what objects of that
7 class's type consist of) and an *implementation*(γ) (that describes how the class works).
8 The definition specifies the data that comprise each object of class; these data are call *data*
9 *members* or *member data*. The definition also specifies some functions (called *member*
10 *functions*) that will operate on that data. It is legal for a class declaration (and therefore, a
11 class) to contain only data or only functions. A class definition has the form shown in list-
12 lst:simpleclassing 6.1. The string *class* is a keyword that is reserved to C++ and may not be used for any
13 user-defined identifier. This construct tells the C++ compiler that `MyClassName` is the
14 name of a class; everything that is between the braces is part of the class definition.

15 A class *declaration* (which you will rarely use) presents the name of the newly defined
16 type, and states that the type is a class:

```
17 class MyClassName;
```

18 Class declarations are rarely used because a class definition also acts as a class declara-
19 tion.

20 sec:quick-start-classesThe remainder of Section 6.7 will give many examples of *members* of a class.

21 In a class definition, the semi-colon after the closing brace is important.



22 The upcoming sections will illustrate some features of classes, with an emphasis on fea-
23 tures that will be important in the first few Workbook exercises. This is not indended to be

1 a comprehensive description of classes. To illustrate, we will show nine versions of a type
2 named `Point` that represents a point in a plane. The first version will be simple and each
3 subsequent version will add features.

4 This documentation will use technically correct language so that you will find it easier
5 to read the standard reference materials. We will point out colloquial usage as neces-
6 sary.

7 Note that the C++ Standard uses the phrase *a type of class-type* to mean a type that is
8 either a class or a structure. In this document we will usually use *class* rather than the
9 more formal *a type of class-type*; we will indicate when you need to distinguish between
10 types that are classes and types that are structures.

11 6.7.2 C++ Exercise 4 v1: The Most Basic Version

12 Here you will see a very basic version of the class `Point` and an illustration of how
13 `Point` can be used. The ideas of *data members* (sometimes called *member data*), *objects*
14 and *instantiation* will be defined.

15 To build and run this example:

1. Log in and follow the steps in Section [6.3](#). sec:quick-start-env
2. `cd` to the directory for this exercise and examine it:

```
cd Classes/v1/
```

```
ls
```

```
Point.h ptest.cc
```

Within the subdirectory `v1` the main program for this exercise is the file `ptest.cc`. The file `Point.h` contains the first version of the class `Point`; shown in Listing [6.1](#). lst:simpleclass

3. Build the exercise. Note that the `../build` below is not a typograph-

16

ical error; the double-dot is necessary.

```
../build
```

```
ls
```

```
Point.h ptest ptest.cc
```

The file named `ptest` is the executable program.

4. Run the exercise:

```
./ptest
```

```
p0: (2.31827e-317, 0)
```

```
p0: (1, 2)
```

```
p1: (3, 4)
```

```
p2: (1, 2)
```

```
Address of p0: 0x7fff883fe680
```

```
Address of p1: 0x7fff883fe670
```

```
Address of p2: 0x7fff883fe660
```

1

2 The values printed out in the first line of the output may be different when you run the
3 program (remember initialization?). When you look at the code you will see that `p0` is not
4 initialized and therefore contains unpredictable data. The last three lines of output may
5 also differ when you run the program; they are memory addresses.

6 Look at the header file `Point.h`, reproduced in Listing 6.2, which shows the basic ver-
7 sion of the class `Point`.

8 The three lines starting with `#` make up an include guard, described in Appendix F. [\[app:ug:include:guards\]](#)

9 Line 4 introduces the name `Point`, and states that `Point` is a class.

10 The body of the class definition begins on line 4, with the opening brace; the body of the
11 class definition ends on line 8, with the closing brace. The definition of the class is followed
12 by a semicolon. Line 5 states that the following members of the class are `public`, which
13 means they are accessible from code outside the class (that is, they are accessible not only

lst:point1

```

1  #ifndef Point_h
2  #define Point_h

4  class Point {
5  public:
6      double x;
7      double y;
8  };

10 #endif /* Point_h */

```

Listing 6.2: File `Point.h` with the simplest version of the class `Point`.

1 by member functions of this class, but by member functions of other classes and by free
 2 functions). Members can also be `private` or `protected`. Section 6.7.7 addresses the
 3 meaning of `private`. The use of `protected` is beyond the scope of this introduction.
 4 Lines 6 and 7 declare the data members `x` and `y`, both of which are of type `double`.

5 In this exercise there is no file `Point.cc` because the class has no user-defined member
 6 functions to implement.

7 Look at the function `main` (the *main program*) in `ptest.cc`, reproduced in Listing 6.3.
 8 This function illustrates the use of the class `Point`.

9 `ptest.cc` includes `Point.h` so that the compiler will know about the class `Point`.
 10 It also includes the Standard Library header `<iostream>` which enables printing with
 11 `std::cout`.

12 When the first line of code in the `main` function,

```

13 Point p0;

```

14 is executed, the program will ensure that memory has been allocated⁵ to hold the data
 15 members of `p0`. If the class `Point` contained code to initialize data members then the
 16 program would also run that, but `Point` does not have any such code. Therefore the data

⁵ This is deliberately vague — there are many ways to allocate memory, and sometimes the memory allocation is actually done much earlier on, perhaps at link time or at load time.

t:pctest-v1

```
1 #include "Point.h"
2 #include <iostream>

4 int main() {
5     Point p0;
6     std::cout << "p0: (" << p0.x << ", " << p0.y << ")"
7         << std::endl;

9     p0.x = 1.0;
10    p0.y = 2.0;
11    std::cout << "p0: (" << p0.x << ", " << p0.y << ")"
12        << std::endl;

14    Point p1;
15    p1.x = 3.0;
16    p1.y = 4.0;
17    std::cout << "p1: (" << p1.x << ", " << p1.y << ")"
18        << std::endl;

20    Point p2 = p0;
21    std::cout << "p2: (" << p2.x << ", " << p2.y << ")"
22        << std::endl;

24    std::cout << "Address of p0: " << &p0 << std::endl;
25    std::cout << "Address of p1: " << &p1 << std::endl;
26    std::cout << "Address of p2: " << &p2 << std::endl;

28    return 0;
29 }
```

Listing 6.3: The contents of `v1/pctest.cc`.

1 members take on whatever values happened to preexist in the memory that was allocated
2 for them.

3 Some other standard pieces of C++ nomenclature can now be defined:

- 4 1. The identifier⁶ `p0` refers to an object in memory. Recall that the C++ meaning of
5 *object* is a region of memory.
- 6 2. The type of this identifier is `Point`. The compiler uses this type to interpret the
7 bytes stored in the object.
- 8 3. When the running program executes line 5 of the main program, it *constructs*(γ) the
9 *object*(γ) named by the identifier `p0`.
- 10 4. The object associated with the identifier `p0` is an *instance*(γ) of the class `Point`.

11 An important take-away from the above is that a *variable* is an identifier in a source code
12 file that refers to some object, while an *object* is something that exists in the computer
13 memory. Most of the time a one-to-one correspondence exists between variables in the
14 source code and objects in memory. There are exceptions, however, for example, some-
15 times a compiler needs to make anonymous temporary objects that do not correspond to
16 any variable in the source code, and sometimes two or more variables in the source code
17 can refer to the same object in memory.

18 We have now seen multiple meanings for the word *object*:

- 19 1. An object is a file containing machine code, the output of a compiler.
- 20 2. An object is a region of memory.
- 21 3. An object is an instance of a class.

22 Which is meant must be determined from context. In this Workbook, we will use “class
23 instance” rather than “object” to distinguish between the second and third meanings in any
24 place where such differentiation is necessary.

25 The last section of the main program (and of `p_test.cc` itself) prints the address of
26 each of the three objects, `p0`, `p1` and `p2`. The addresses are represented in hexadecimal
27 (base 16) format. On almost all computers, the size of a `double` is eight bytes. Therefore
28 an object of type `Point` will have a size of 16 bytes. If you look at the printout made

⁶An identifier that refers to an object is often called a *variable*.

p2.x	1.0
p2.y	2.0
p1.x	3.0
p1.y	4.0
p0.x	1.0
p0.y	2.0

Figure 6.1: Memory diagram at the end of a run of `Classes/v1/ptest.cc`

1 by `ptest` you will see that the addresses of `p0`, `p01` and `p2` are separated by 16 bytes;
 2 therefore the three objects are contiguous in memory.

3 `Figure 6.1` shows a diagram of the computer memory at the end of running `ptest`; the
 4 outer box (blue outline) represents the memory of the computer; each filled colored box
 5 represents one of the three class instances in this program. The diagram shows them in
 6 contiguous memory locations, which is not necessary; there could have been gaps between
 7 the memory locations.

8 Now, for a bit more terminology: each of the objects referred to by the variables `p0`, `p1`
 9 and `p2` has the three attributes required of an *object*:

- 10 1. a *state*(γ), given by the values of its data members;
- 11 2. the ability to have operations performed on it: e.g., setting/reading in value of a data
 12 member, assigning value of object of a given type to another of the same type;
- 13 3. a unique address in memory, and therefore a unique *identity*.

1 6.7.3 C++ Exercise 4 v2: The Default Constructor

es-default-ctor

- 2 This exercise expands the class `Point` by adding a user-written default *constructor*(γ).
- 3 To build and run this example:

1. Log in and follow the steps in Section [6.3](#). sec:quick-start-env

2. Go to the directory for this exercise:

```
cd Classes/v2
```

```
ls
```

```
Point.cc Point.h ptest.cc
```

In this example, `Point.cc` is a new file.

3. Build the exercise:

```
../build
```

```
ls
```

```
Point.cc Point.h ptest ptest.cc
```

4. Run the exercise:

```
./ptest
```

```
p0: (0, 0)
```

```
p0: (3.1, 2.7)
```


4

5 When you run the code, all of the printout should match the above printout exactly.

6 Look at `Point.h`. There is one new line in the body of the class definition:

```
7 Point();
```

1 The parentheses tell you that this new member is some sort of function. A C++ class may
2 have several different kinds of functions.

3 A function that has the same name as the class itself has a special role and is called a *con-*
4 *structor*; if a constructor can be called with no arguments it is called a *default constructor*⁷. 
5 In informal written material, the word constructor is sometimes written as *c'tor*.

6 `Point.h` declares that the class `Point` has a default constructor, but does not define
7 it (i.e., provide an implementation). The definition (implementation) of the constructor is
8 found in the file `Point.cc`.

9 Look at the file `Point.cc`. It `#includes` the header file `Point.h` because the com-
10 piler needs to know all about this class before it can compile the code that it finds in
11 `Point.cc`. The rest of the file contains a *definition* of the constructor. The syntax `Point::`
12 says that the function to the right of the `::` is part of (a member of) the class `Point`. The
13 body of the constructor gives initial values to the two data members, `x` and `y`:

```
14 Point::Point() {  
15     x = 0.;  
16     y = 0.;  
17 }
```

18 Look at the program `ptest.cc`. The first line of the `main` function is again

```
19 Point p0;
```

20 When the program executes this line, the first step is the same as before: it ensures that
21 memory has been allocated for the data members of `p0`. This time, however, it also calls
22 the default constructor of the class `Point` (declared in `Point.h`), which initializes the
23 two data members (per `Point.cc`) such that they have well defined initial values. This
24 is reflected in the printout made by the next line.

25 The next block of the program assigns new values to the data members of `p0` and prints
26 them out.

27 In the previous example, `Classes/v1/ptest.cc`, a few things happened behind the
28 scenes that will make more sense now that you know what a constructor is.

⁷Note that a constructor declaration that provides a default value for each argument can be called with no arguments, and thus qualifies as a default constructor.

- 1 1. Since the source code for class `Point` did not contain any user-defined constructor,
2 the compiler generated a default constructor for you; this is required by the C++
3 Standard and will be done for any class that has no user-written constructor.
- 4 2. The compiler puts the generated constructor code directly into the object file; it does
5 not affect the source file.
- 6 3. The generated default constructor will default construct each data member of the
7 class.
- 8 4. Default construction of an object of a primitive type leaves that object uninitialized;
9 this is why the data members `x` and `y` of version 1 of `Point` were uninitialized.

10 6.7.4 C++ Exercise 4 v3: Constructors with Arguments

11 This exercise introduces four new ideas:

- 12 1. constructors with arguments,
- 13 2. the copy constructor,
- 14 3. the implicitly generated constructor,
- 15 4. single-phase construction vs. two-phase construction.

16 To build and run this exercise, `cd` to the directory `Classes/v3` and follow the same
17 instructions as in Section 6.7.3. When you run the `ptest` program, you should see the
18 following output:

```
19 ./ptest
```

```
20 p0: (1, 2)  
21 p1: (1, 2)
```

22 Look at the file `Point.h`. This contains one new line:

```
23 Point( double ax, double ay);
```

24 This line declares a second constructor; we know it is a constructor because it is a function
25 whose name is the same as the name of the class. It is distinguishable from the default
26 constructor because its argument list is different than that of the default constructor. As

1 before, the file `Point.h` contains only the declaration of this constructor, not its *definition*
2 (*implementation*).

3 Look at the file `Point.cc`. The new content in this file is the implementation of the new
4 constructor; it assigns the values of its arguments to the data members. The names of the
5 arguments, `ax` and `ay`, have no meaning to the compiler; they are just identifiers. It is good
6 practice to choose names that bear an obvious relationship to those of the data members.
7 One convention that is sometimes used is to make the name of the argument be the same as
8 that of the data member, but with a prefix letter `a`, for `argument`. Whatever convention you
9 (or your experiment) choose(s), use it consistently. When you update code that was initially
10 written by someone else, we strongly recommend that you follow whatever convention
11 they adopted. Choices of style should be made to reinforce the information present in the
12 code, not to fight it.

13 Look at the file `ptest.cc`. The first line of the `main` function is now:

```
14 Point p0(1.,2.);
```

15 This line declares the variable `p0` and initializes it by calling the new constructor defined
16 in this section. The next line prints the value of the data members.

17 The next line of code

```
18 Point p1(p0);
```

19 uses the *copy constructor*. A copy constructor is used by code (like the above) that wants to
20 create a copy (e.g., `p1`) of an existing object (e.g., `p0`). The default meaning of copying is
21 data-member-by-data-member copying. Under the appropriate conditions (to be described
22 later), the compiler will implicitly generate a copy constructor, with public access, for a
23 class; this definition of `Point` meets the necessary conditions. As is done for a generated
24 default constructor, the compiler puts the generated code directly into the object file; it
25 does not affect the source file.

26 We recommend that for any class whose data members are either built-in types, of which
27 `Point` is an example, or simple aggregates of built-in types, you let the compiler write
28 the copy constructor for you.

29 **FIXME:** *Idea (for Anne): Create figure for the compile/link flow illustrating implicit*
30 *declaration of constructor from 7/3 meeting with Rob.*

1 If your class has data members that are pointers, or data members that manage some
2 external resource, such as a file that you are writing to, these pointers should be *smart*
3 *pointers*, such as `std::shared_ptr<T>` or `std::unique_ptr<T>`. This will al-
4 low the compiler-generated copy constructor to give the correct behavior. For a description
5 of smart pointers, consult the standard C++ references (listed in Section 6.9). There are rare
6 cases in which you will need to write your own copy constructor, but discussing them here
7 is beyond the scope of this document. When you need to write your own copy constructor,
8 you can learn how to do so from any standard C++ reference.


9 The next line in the file prints the values of the data members of `p1` and you can see that
10 the copy constructor worked as expected.

11 Notice that in the previous version of `pctest.cc`, the variable `p0` was initialized in three
12 lines:

```
13 Point p0;  
14 p0.x = 3.1;  
15 p0.y = 2.7;
```

16 This is called *two-phase construction*. In contrast, the present version uses *single-phase*
17 *construction* in which the variable `p0` is initialized in one line:

```
18 Point p0(1., 2.);
```

 19 We strongly recommend using single-phase construction whenever possible. Obviously it
20 takes less real estate, but more importantly:

- 21 1. Single-phase construction more clearly conveys the intent of the programmer: the
22 intent is to initialize the object `p0`. The second version says this directly. In the first
23 version you needed to do some extra work to recognize that the three lines quoted
24 above formed a logical unit distinct from the remainder of the program. This is
25 not difficult for this simple class, but it can become so with even a little additional
26 complexity.
- 27 2. Two-phase construction is less robust. It leaves open the possibility that a future
28 maintainer of the code might not recognize all of the follow-on steps that are part of
29 construction and will use the object before it is fully constructed. This can lead to
30 difficult-to-diagnose run-time errors.

- 1 3. Single-phase construction can be more efficient than two-phase construction.
- 2 4. Single-phase construction is the only way to initialize variables that are declared
- 3 `const`. It is good practice to declare `const` any variable that is not intended to be
- 4 changed.

5 **FIXME:** *Where do we mention the compiler written assignment operator?*

6 **FIXME: MFP:** *We haven't used one in any of the example programs. We should consider*
7 *adding an additional section on "special member functions", which includes constructors*
8 *(of all types), assignment (both normal and move), and destructors. This would allow us*
9 *to also introduce the information about when the compiler will generate special member*
10 *functions, and the "rule of 5" (see <http://www.stroustrup.com/C++11FAQ.html> (the sec-*
11 *tion titled "control of defaults: move and copy") which describes the issue but does not*
12 *use the phrase "rule of 5"). Anne and Marc would prefer to see the discussion of the rule*
13 *of 5 in an exercise, so it is not skipped by those who already know C++.*

14 6.7.5 C++ Exercise 4 v4: Colon Initializer Syntax

15 This version of the class `Point` introduces *colon-initializer syntax* for constructors.

16 To build and run this exercise, `cd` to the directory `Classes/v4` and follow the same
17 instructions as in the previous two sections. When you run the `ptest` program you should
18 see the following output:

```
19 ./ptest  
20 p0: (1, 2)  
21 p1: (1, 2)
```

22 The file `Point.h` is unchanged between this version and the previous one.

23 Now look at the file `Point.cc`, which contains the *definitions* of both constructors. The
24 first thing to look at is the default constructor, which has been rewritten using colon-
25 initializer syntax. The rules for the colon-initializer syntax are:

- 26 1. A colon must immediately follow the closing parenthesis of the argument list.
- 27 2. There must be a comma-separated list of data members, each one initialized by
- 28 calling one of its constructors.

- 1 3. Data members are guaranteed to be initialized in the order in which they appear
2 in the class declaration. Therefore it is good practice to use the same order for the
3 initialization list.
- 4 4. The body of the constructor, enclosed in braces, must follow the initializer list. The
5 body of the constructor will most often be empty.
- 6 5. If a data member is missing from the initializer list, that member will be default-
7 constructed. Thus data members that are of a primitive type and are missing from
8 the initialize list will not be initialized.
- 9 6. If no initializer list is present, the compiler will call the default constructor of every
10 data member, and it will do so in the order in which data members were specified in
11 the class declaration.

12 If you think about these rules carefully, you will see that in `Classes/v3/`
13 `Point.cc`, the compiler did the following when compiling the default constructor.

- 14 1. The compiler did not find an initializer list, so it generated machine code that created
15 uninitialized `x` and `y`.
- 16 2. It then wrote the machine code to make the assignments `x=0` and `y=0`.

17 On the other hand, when the compiler compiled the source code for the default constructor
18 in `Classes/v4/Point.cc`, it wrote the machine code to initialize `x` and `y` each to
19 zero.

20 Therefore, the machine code for the `v3` version might do more work than that for the `v4`
21 version. In practice `Point` is a sufficiently simple class that the compiler likely recognized
22 and elided all of the unnecessary steps in `v3`; it is likely that the compiler actually produced
23 identical code for the two versions of the class. For a class containing more complex data,
24 however, the compiler may not be able to recognize meaningless extra work and it will
25 write the machine code to do that extra work.

26 In some cases it does not matter which of these two ways you use to write a constructor;
27 but on those occasions that it does matter, the right answer is always the colon-initializer
28 syntax. So we strongly recommend that you always use the colon-initializer syntax. In the
29 Workbook, all classes are written with colon-initializer syntax.

30 Now look at the second constructor in `Point.cc`; it also uses colon-initializer syntax

- 1 but it is laid out differently. The difference in layout has no meaning to the compiler —
- 2 whitespace is whitespace. Choose which ever seems natural to you.
- 3 Look at `pctest.cc`. It is the same as the version `v3` and it makes the same printout.

4 **6.7.6 C++ Exercise 4 v5: Member functions**

5 This section will introduce *member functions*(γ), both *const member functions*(γ) and non-
6 const member functions. It will also introduce the header `<cmath>`. Suggested homework
7 for this material follows.

8 To build and run this exercise, `cd` to the directory `Classes/v5` and follow the same
9 instructions as in Section 6.7.3. When you run the `pctest` program you should see the
10 following output:

```
11 ./pctest
```

```
12 Before p0: (1, 2) Magnitude: 2.23607 Phi: 1.10715  
13 After p0: (3, 6) Magnitude: 6.7082 Phi: 1.10715
```

14 Look at the file `Point.h`. Compared to version `v4`, this version contains three additional
15 lines:

```
16 double mag() const;  
17 double phi() const;  
18 void scale(double factor);
```

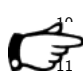
19 All three lines declare *member functions*. As the name suggests, a *member function* is
20 a function that can be called and it is a member of the class. Contrast this with a *data*
21 *member*, such as `x` or `y`, which are not functions. A member function may access any or
22 all of the member data of the class.

23 The first of these member functions is named `Point::mag`. The name indicates this
24 function is a member of class `Point`. `Point::mag` does not take any arguments and it
25 returns a `double`; you will see that the value of the `double` is the magnitude of the 2-
26 vector from the origin $(0, 0)$ to (x, y) . The qualifier `const` represents a contract between
27 the definition/implementation of `mag` and any code that uses `mag`; it “promises” that the
28 implementation of `Point::mag` will not modify the value of any data members. The

1 consequences of breaking the contract are illustrated in the homework at the end of this
2 subsection.

3 Similarly, the member function named `Point::phi` takes no arguments, returns a value
4 of type `double` and has the `const` qualifier. You will see that the value of the `double`
5 is the azimuthal angle of the vector from the origin $(0, 0)$ to the point (x, y) .

6 The third member function, `Point::scale`, takes one argument, `factor`. Its return
7 type is `void`, which means that it returns nothing. You will see that this member function
8 multiplies both `x` and `y` by `factor` (i.e., changing their values). This function declaration
9 does not have the `const` qualifier because it actually does modify member data.

 10 If a member function does not modify any data members, you should always declare it
11 `const` simply as a matter of course. Any negative consequences of not doing so might
12 only become apparent later, at which point a lot of tedious editing will be required to make
13 everything right.

14 Look at `Point.cc`. Near the top of the file an additional include directive has been added;
15 `<cmath>` is a header from the C++ standard library that declares a set of functions for
16 computing common mathematical operations and transformations. Functions from this
17 library are in the *namespace*(γ) `std`.

18 Later on in `Point.cc` you will find the definition of `Point::mag`, which computes the
19 magnitude of the 2-vector from the origin $(0, 0)$ to (x, y) . To do so, it uses `std::sqrt`,
20 a function declared in the `<cmath>` header. This function takes the square root of its
21 argument. The qualifier `const` that was present in the declaration of `Point::mag` must
22 also be present in its definition (it must be present in these two places, but not at calling
23 points).

24 The next part of `Point.cc` contains the definition of the member function `phi`. To do
25 its work, this member function uses the `atan2` function from the standard library.

26 The next part of `Point.cc` contains the definition of the member function `Point::scale`.
27 You can see that this member function simply multiplies the two data members by the value
28 of the argument.

29 The file `ptest.cc` contains a `main` function that illustrates these new features. The first
30 line of this function declares and initializes an object, `p0`, of type `Point`. It then prints out
31 the value of its data members, the value returned from calling the function `Point::mag`

1 and the value returned from calling `Point::phi`. This shows how to invoke a member
2 function: you write the name of the variable, followed by a dot (the *member selection*
3 *operator*), followed by the *unqualified name* of the member function, followed by the
4 argument list in the function call parentheses. The unqualified name of a member function
5 is the part of the name that follows the double-colon scope resolution operator (`::`). Thus
6 the unqualified name of `Point::phi` is just `phi`.

7 The next line calls the member function `Point::scale` with the argument `3`. The print-
8 out verifies that the call to `Point::scale` had the intended effect.

9 One final comment is in order. Many other modern computer languages have ideas very
10 similar to C++ classes and C++ member functions; in some of those languages, the name
11 *method* is the technical term corresponding to *member function* in C++. The name *method*
12 is not part of the formal definition of C++, but is commonly used nonetheless. In this
13 documentation, the two terms can be considered synonymous.

14 Here we suggest four activities as homework to help illustrate the meaning of `const` and
15 to familiarize you with the error messages produced by the C++ compiler. Before moving
16 to a subsequent activity, undo the changes that you made in the current activity.

- 17 1. In the definition of the member function `Point::mag()`, found in `Point.cc`,
18 before taking the square root, multiply the member datum `x` by 2.

```
19 double Point::mag() const {  
20     x *= 2.;  
21     return std::sqrt( x*x + y*y );  
22 }
```

23 Then build the code again; you should see the following diagnostic message:

```
24 Point.cc: In member function 'double Point::mag() const':  
25 Point.cc:13:8: error: assignment of member 'Point::x' in  
26 read-only object
```

- 27 2. In `ptest.cc`, change the first line to

```
28 Point const p0(1,2);
```

29 Then build the code again; you should see the following diagnostic message:

```
30 ptest.cc: In function 'int main()':
```

```

1      ptest.cc:13:14: error: no matching function for call to
2      'Point::scale(double)_const'
3      ptest.cc:13:14: note: candidate is:
4      In file included from ptest.cc:1:0:
5      Point.h:13:8: note: void Point::scale(double) <near match>
6      Point.h:13:8: note: no known conversion for implicit
7      'this' parameter from 'const_Point*' to 'Point*'

```

8 These first two homework exercises illustrate how the compiler enforces the contract de-
 9 fined by the qualifier `const` that is present at the end of the declaration of `Point::mag`
 10 and that is absent in the definition of the member function `Point::scale`. The con-
 11 tract says that the definition of `Point::mag` may not modify the values of any data
 12 members of the `Point` object on which it is called; users of the class `Point` may
 13 count on this behaviour. The contract also says that the definition of the member func-
 14 tion `Point::scale` may modify the values of data members of the class `Point`; users
 15 of the class `Point` must assume that `Point::scale` will indeed modify member data
 16 and act accordingly.⁸

17 In the first homework exercise, the value of a member datum is modified, thereby breaking
 18 the contract. The compiler detects it and issues a diagnostic message.

19 In the second homework exercise, the variable `p0` is declared `const`; therefore the code
 20 may not call non-`const` member functions of `p0`, only `const` member functions. When
 21 the compiler sees the call `p0.mag()` it recognizes that this is a call to `const` member
 22 function and compiles the call; when it sees the call `p0.scale(3.)` it recognizes that
 23 this is a call to a non-`const` member function and issues a diagnostic message.

24 3. In `Point.h`, remove the `const` qualifier from the declaration of the member func-
 25 tion `Point::mag`:

```

26      double mag();

```

27 Then build the code again; you should see the following diagnostic message:

```

28      Point.cc:12:8: error: prototype for 'double_Point::mag()
29      _const' does not match any in class 'Point'
30      In file included from Point.cc:1:0:
31      Point.h:11:10: error: candidate is: double Point::mag()

```

⁸ C++ has another keyword, the specifier `mutable`, that one can use to exempt individual data members from this contract. Its use is discouraged. The rare conditions under which its use is appropriate are beyond the scope of this introduction.

1 4. In `Point.cc`, remove the `const` qualifier in the definition of the member function
 2 `Point::mag`. Then build the code again; you should see the following diagnostic
 3 message:

```
4 Point.cc:12:8: error: prototype for 'double Point::mag()'
5     does not match any in class 'Point'
6 In file included from Point.cc:1:0:
7 Point.h:11:10: error: candidate is:
8     double Point::mag() const
```

9 The third and fourth homework exercises illustrate that the compiler considers two mem-
 10 ber functions that are identical except for the presence of the `const` identifier to be
 11 different functions⁹. In homework exercise 3, when the compiler tried to compile the
 12 `const`-qualified version of `Point::mag` in `Point.cc`, it looked at the class defini-
 13 tion in `Point.h` and could not find a matching member function declaration; this was a
 14 close, but not exact match. Therefore it issued a diagnostic message, telling us about the
 15 close match, and then stopped. Similarly, in homework exercise 4, it also could not find a
 16 match.

17 6.7.7 C++ Exercise 4 v6: Private Data and Accessor Methods

18 6.7.7.1 Setters and Getters

19 This version of the class `Point` is used to illustrate the following ideas:

- 20 1. The class `Point` has been redesigned to have private data members with access to
 21 them provided by *accessor functions* and *setter functions*.
- 22 2. The keyword `this`, which in the body of a (non-*static*) member function is an
 23 expression that has the value of the address of the object on which the function is
 24 called.
- 25 3. Even if there are many objects of type `Point` in memory, there is only one copy of
 26 the code.

27 A 2D point class, with member data in Cartesian coordinates, is not a good example of
 28 *why* it is often a good idea to have `private` data. But it does have enough richness to

⁹ Another way of saying the same thing is that the `const` identifier is part of the *signature*(γ) of a function.

1 illustrate the mechanics, which is the purpose of this section. Section 6.7.7.3 discusses an
 2 example in which having `private` data makes obvious sense.

3 To build and run this exercise, `cd` to the directory `Classes/v6` and follow the same
 4 instructions as in Section 6.7.3. When you run the `ptest` program you should see the
 5 following output:

```
6 ./ptest
7 Before p0: (1, 2) Magnitude: 2.23607 Phi: 1.10715
8 After p0: (3, 6) Magnitude: 6.7082 Phi: 1.10715
9 p1: (0, 1) Magnitude: 1 Phi: 1.5708
10 p1: (1, 0) Magnitude: 1 Phi: 0
11 p1: (3, 6) Magnitude: 6.7082 Phi: 1.10715
```

12 Look at `Point.h`. Compare it to the version in `v5`:

```
13 diff -wb Point.h ../v5/
```

14 Relative to version `v5` the following changes were made:

- 15 1. four new member functions have been declared,
 - 16 (a) `double x() const;`
 - 17 (b) `double y() const;`
 - 18 (c) `void set(double ax, double ay);`
 - 19 (d) `void set(Point const& p);`
- 20 2. the data members have been declared `private`
- 21 3. the data members have been renamed from `x` and `y` to `x_` and `y_`

22 Yes, there are two functions named `set`. At the site of any function call that uses the name
 23 `set`, the compiler makes use of the signature of the function to decide which function with
 24 that name to call. In C++ the signature of a member function encodes all of the following
 25 information:

- 26 1. the name of the class it is in;
- 27 2. the unqualified name of the member function;
- 28 3. the number, type and order of arguments in the argument list;

1 4. whether or not the function is qualified as `const`;

2 5. other qualifications (reference-qualification and `volatile`-qualification, both of
3 which are beyond the scope of this introduction).

4 The two member functions named `Point::set` are completely different member func-
5 tions with different signatures. A set of different functions with the same name but with
6 different signatures is called an *overload set*. As you work through the Workbook you
7 will encounter a many of these, and you should develop the habit of looking at the full
8 function signature (i.e., all the parts), not just the function name. In order to distinguish
9 between members of an overload set, C++ compilers typically rely on *name mangling*.
10 Name mangling is the process of decorating a function name with information that en-
11 codes the signature of the function. The mangled name associated with each function is
12 the symbol emitted by the compiler, and used by the linker to identify which member of
13 an overload set is associated with a specific function call. Each C++ compiler does this a
14 little differently.

15 If you want to see what mangled names are created for the class `Point`, you can do the
16 following



```
17 c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c Point.cc
```

```
18 nm Point.o
```

19 You can understand the output of `nm` by reading its man page.

20 In a class declaration, if any of the identifiers `public`, `private`, or `protected` ap-
21 pear, then all members following that identifier, and before the next such identifier, have
22 the named property. In `Point.h` the two data members are private and all other members
23 are public.

24 Look at `Point.cc`. Compare it to the version in `v5`:

```
25 diff -wb Point.cc ../v5/
```

26 Relative to version `v5` the following changes were made:

- 27 1. the data members have been renamed from `x` and `y` to `x_` and `y_`
- 28 2. an implementation is present for each of the four new member functions

1 Inspect the code in the implementation of each of the new member functions. The member
 2 function `Point::x` simply returns the value of the data member `x_`; similarly for the
 3 member function `Point::y`. These member functions are called *accessors*, *accessor*
 4 *functions*, or *getters*¹⁰. The notion of *accessor* is often extended to include any member
 5 function that returns the value of simple, non-modifying calculations on a subset of the
 6 member data; in this sense, the function `Point::mag` and `Point::phi` are considered
 7 *accessors*.

8 The member functions in the overload set for the name `Point::set` each set the member
 9 data of the `Point` object on which they are called. These are, not surprisingly, called
 10 *setters*, *setter functions* or *modifiers*.

11 There is no requirement that there be accessors and setters for every data member of a
 12 class; indeed, many classes provide no such member functions for many of their data mem-
 13 bers. If a data member is important for managing internal state but is of no direct interest
 14 to a user of the class, then you should certainly not provide an accessor or a setter.

15 Now that the data members of `Point` are private, only the code within `Point` is permitted
 16 to access these data members directly. All other code must access this information via the
 17 accessor and setter functions.

18 Look at `pctest.cc`. Compare it to the version in `v5`:

```
19 diff -wb pctest.cc ../v5/
```

20 Relative to version `v5` the following changes were made:

- 21 1. the printout has been changed to use the accessor functions, and
- 22 2. a new section has been added to illustrate the use of the two set methods.

23 fig:memory_v2 Figure 6.2 shows a diagram of the computer memory at the end of running this version of
 24 `pctest`. The two boxes with the blue outlines represent sections of the computer memory;
 25 the part on the left represents that part that is reserved for storing data (such as objects)
 26 and the part on the right represents the part of the computer memory that holds the exe-

¹⁰There is a coding style in which the function `Point::x` would have been called something like `Point::GetX`, `Point::getX` or `Point::get_x`; hence the name *getters*. Almost all of the code that you will see in the Workbook omits the `get` in the names of *accessors*; the authors of this code view the `get` as redundant. Within the Workbook, the exception is for accessors defined by ROOT. The Geant4 package also includes the `Get` in the names of its accessors.




Figure 6.2: Memory diagram at the end of a run of Classes/v6/ptest.cc

1 cutable code. This is a big oversimplification because, in a real running program, there are
 2 many parts of the memory reserved for different sorts of data and many parts reserved for
 3 executable code.

4 The key point in Figure 6.2 is that each object has its own member data but there is only
 5 one copy of the code. Even if there are thousands of objects of type `Point`, there will only
 6 be one copy of the code. When a line of code asks for `p0.mag()`, the computer will pass
 7 the address of `p0` as an argument to the function `Point::mag`, which will then do its
 8 work. When a line of code asks for `p1.mag()`, the computer will pass the address of `p1`
 9 as an argument to the function `Point::mag`, which will then do its work. This address is
 10 available in the body of the member function as the value of the expression `this`, which
 11 acts as a pointer to the object on which the function was called. In a member function
 12 declared as `const`, the expression acts as a pointer that is `const`-qualified.

13 Initially this sounds a little weird: the previous paragraph talks about passing an argument
 14 to the function `Point::mag` but, according to the source code, `Point::mag` does not
 15 take any arguments! The answer is that all member functions have an implied argument
 16 that always must be present — the address of the object that the member function will do

1 work on. Because it must always be there, and because the compiler knows that it must
 2 always be there, there is no point in actually writing it in the source code! It is by using
 3 this so called *hidden argument* that the code for `Point::mag` knew that `x_` means one
 4 thing for `p0` but that it means something else for `p1`.

5 For example, the accessor `Point::x` could have been written:

```
6 double x() const { return this->x_; }
```

7 This version of the syntax makes it much clearer how there can be one copy of the code
 8 even though there are many objects in memory; but it also makes the code harder to read
 9 once you have understood how the magic works. There are not many places in which you
 10 need to explicitly use the keyword `this`, but there will be some. For further information,
 11 consult standard C++ documentation (listed in Section 6.9). [sec:cpp:references](#)

12 6.7.7.2 What's the deal with the underscore?

-data-underscore

13 C++ will not permit you to use the same name for both a data member and its accessor.
 14 Since the accessor is part of the public interface, it should get the simple, obvious, easy-
 15 to-type name. Therefore the name of the data member needs to be decorated to make it
 16 distinct.

17 The convention used in the Workbook exercises and in the toyExperiment UPS product
 18 is that the names of member data end in an underscore character. There are some other
 19 conventions that you may encounter:

```
20 _name;  
21 m_name;  
22 mName;  
23 theName;
```



24 You may also see the choice of a leading underscore followed by a capital letter, or a double
 25 underscore. Never do this. Such names are reserved for use by C++ implementations; use
 26 of such names may produce accidental collisions with names used in an implementation,
 27 and cause errors that might be very difficult to diagnose. While this is a very small risk, it
 28 seems wise to adopt habits that guarantee that it can never happen.

1 It is common to extend the pattern for decorating the names of member data to all member
2 data, even those without accessors. One reason for doing so is just symmetry. A second
3 reason has to do with writing member functions; the body of a member function will, in
4 general, use both member data and variables that are local to the member function. If the
5 member data are decorated differently than the local variables, it can make the member
6 functions easier to understand.

7 6.7.7.3 An example to motivate private data

motivation

8 This section describes a class for which it makes sense to have private data: a 2D point class
9 that has data members `r` and `phi` instead of `x` and `y`. The author of such a class might
10 wish to define a standard representation in which it is guaranteed that `r` be non-negative
11 and that `phi` be on the domain $0 \leq \phi < 2\pi$. If the data are public, the class cannot make
12 these guarantees; any code can modify the data members and break the guarantee.

13 If this class is implemented with private data manipulated by member functions, then the
14 constructors and member functions can enforce the guarantees.

15 The language used in the software engineering texts is that a guaranteed relationship
16 among the data members is called an *invariant*. If a class has an invariant then the class
17 must have private data.

18 If a class has no invariant then one is free to choose public data. The Workbook and the
19 toyExperiment never make this choice. One reason is that classes that begin life without an
20 invariant sometimes acquire one as the design matures — we recommend that you plan for
21 this unless you are 100% sure that the class will never have an invariant. A second reason
22 is that many people who are just starting to learn C++ find it confusing to encounter some
23 classes with private data and others with public data.

24 6.7.8 C++ Exercise 4 v7: The `inline` Specifier

ses-inline

25 This section introduces the `inline` specifier.

26 To build and run this exercise, cd to the directory `Classes/v7` and follow the same
27 instructions as in Section 6.7.3. When you run the `ptest` program you should see the
28 following output:

1 `./ptest`

2 `p0: (1, 2) Magnitude: 2.23607 Phi: 1.10715`

3 Look at `Point.cc` and compare it to the version in `v6`. You will see that the implemen-
4 tations of the accessors `Point::x` and `Point::y` has been removed.

5 Comparing `Point.h` to the version in `v6`, you will see that it now contains the implemen-
6 tation of the accessor member functions — an almost exact copy of what was previously
7 found in the file `Point.cc`. Note that these accessors are defined outside of the class
8 declaration in `Point.h` and are now preceded by the specifier `inline`.

9 The `inline` specifier on a function sends a suggestion to the compiler to *inline* the func-
10 tion. If the compiler chooses to inline the function, the body of the function is substituted
11 directly into the machine code at the point of each call to it, instead of the program making
12 a run-time function call.

13 The specifier does not force inlining on the compiler. Why the option? In some cases
14 inlining is a net positive thing, in other cases it's a net negative; based on heuristics, the
15 compiler will determine which, and choose. For some functions, offering the option at
16 all (i.e., including the specifier `inline`) is a net negative no matter which option the
17 compiler would choose; this means that you as the programmer need to know when to use
18 it and when not to.

19 Specifying a function as `inline` is typically a good thing only for small and/or simple
20 functions, e.g., an accessor. The compiler will be likely to inline it because it this option is
21 likely to

- 22 ○ reduce the memory footprint¹¹
- 23 ○ execute more quickly than a function call
- 24 ○ allow additional compiler optimizations to be performed.

25 In the “decline-to-inline” case, the compiler will write a copy of the function once for
26 each source file in which a definition of the function appears¹². During linking, the copy

¹¹The instructions to set up a call stack, to make the call, and to tear down the call stack are not needed. For a small function (e.g. an accessor), this code might be large compared to the code that actually does the work of the function.

¹²More precisely, it is for each *compilation unit*, i.e., the unit of code that the compiler considers at one time. For most purposes, each source file is its own compilation unit.

1 of the compiled function in the same object file will be used to satisfy calls to the function.
2 Result: a larger memory footprint, but no reduction in execution time. Clearly, for a bigger
3 or more complex function, use of the `inline` specifier is disadvantageous.

4 C++ does not permit you to force inlining; an `inline` declaration is only a hint to the
5 compiler that a function is appropriate for inlining.

6 The bottom line is that you should always declare simple accessors and simple setters
7 `inline`. Here the adjective *simple* means that they do not do any significant computation
8 and that they do not contain any `if` statements or loops. The decision to inline anything
9 else should only follow careful analysis of information produced by a profiling tool.

10 Look at the definition of the member function `Point::y` in `Point.h`. Compared to the
11 definition of the member function `Point::x` there is only a small change in whitespace,
12 and of course the specifier `inline`. This whitespace difference is not meaningful to the
13 compiler.

14 **6.7.9 C++ Exercise 4 v8: Defining Member Functions within the Class Dec-** 15 **laration**

definition

16 The version of `Point` in this section introduces the feature that allows you to provide the
17 definition (implementation) of any member function inside the declaration of the class to
18 which it belongs, right at the point where it the function is declared. You will occasion-
19 ally see this syntax used in the Workbook. The definition of a non-member function (see
20 [ssec:quick-start-classes-stream-insertion](#) Section 6.7.10) must remain outside the class declaration.

21 To build and run this exercise, `cd` to the directory `Classes/v8` and follow the same
22 instructions as in Section [ssec:quick-start-classes-default-ctor](#) 6.7.3. When you run the `ptest` program you should see the
23 following output:

```
24 ./ptest
```

```
25 p0: ( 1, 2 ) Magnitude: 2.23607 Phi: 1.10715
```

26 This is the same output made by v7. The files `Point.cc` and `ptest.cc` are unchanged
27 with respect to v7, only `Point.h` has changed.

28 Relative to v7, the definition of the accessor methods `Point::x` and `Point::y` in
29 `Point.h` has been moved into the `Point` class declaration. Notice that the function

1 names are no longer prefixed with the class name and the `inline` specifiers have been
2 removed.



3 When you define a member function inside the class declaration, the function is implic-
4 itly declared `inline`. Section 6.7.8 discussed some cautions about inappropriate use of
5 inlining; those same cautions apply when a member function is defined inside the class
6 declaration.

7 When you define a member function within the class declaration, you must not prefix the
8 function name with the class name and the scope resolution operator; that is,

```
9 double Point::x() const { return x_; }
```



10 would produce a compiler diagnostic.

11 In summary, there are two ways to write inlined definitions of member functions. In most
12 cases, the two are entirely equivalent and the choice is simply a matter of style. The one
13 exception occurs when you are writing a class that will become part of an *art* data prod-
14 uct, due to limitations imposed by *art*. In this case it is recommended that you write the
15 definitions of member functions *outside* of the class declaration.



17 When writing an *art* data product, the code inside the associated header file is parsed by
18 software that determines how to write objects of that type to the output disk files and how to
19 read objects of that type from input disk files. The software that does the parsing has some
20 limitations and we need to work around them. The workarounds are easiest if any member
21 functions definitions in the header file are placed outside of the class declarations. For de-
22 tails see https://cdcvs.fnal.gov/redmine/projects/art/wiki/Data_Product_Design_Guide#Issues-mostly-related-to-ROOT.

23 6.7.10 C++ Exercise 4 v9: The Stream Insertion Operator and Free Func- 24 tions

stream-insertion

25 This section illustrates how to write a *stream insertion operator* for a type, in this case for
26 the class `Point`. This is the piece of code that lets you print an object of a given type
27 without having to print each data member by hand, for example:

```
28 Point p0(1,2);  
29 std::cout << p0 << std::endl;
```


1 instead of

```
2 Point p0(1, 2);  
3 std::cout << "p0: (" << p0.x() << ", " << p0.y() << ")"
```

4 To build and run this exercise, cd to the directory `Classes/v9` and follow the same
5 instructions as in Section 6.7.3. When you run the `ptest` program you should see the
6 following output:

```
7 ./ptest  
8 p0: ( 1, 2 ) Magnitude: 2.23607 Phi: 1.10715
```

9 This is the same output made by `v7` and `v8`.

10 Look at `Point.h`. The changes relative to `v8` are the following two additions:

- 11 1. an include directive for the header `<iosfwd>`
- 12 2. a declaration for the stream insertion operator, which appears in the file after the
13 declaration of the class `Point`.

14 Look at `Point.cc`. The changes relative to `v8` are the following two additions:


- 15 1. an include directive for the header `<ostream>`
- 16 2. the definition of the stream insertion operator, `operator<<`.

17 Look at `ptest.cc`. The only change relative to `v8` is that the printout now uses the
18 stream insertion operator for `p0` instead of inserting each data member of `p0` by hand.

```
19 std::cout << "p0:_ " << p0
```

20 In `Point.h`, the stream insertion operator is declared as (shown here on two lines)

```
21 std::ostream&  
22 operator<<(std::ostream& ost, Point const& p);
```

23 If the class whose type is used as second argument is declared in a namespace (which
24 it is not, in this case), then the stream insertion operator must be declared in the same
25 namespace. 

26 When the compiler sees the use of a `<<` operator that has an object of type
27 `std::ostream` on its left hand side and an object of type `Point` on its right hand side,

1 then the compiler will look for a function named `operator<<` whose first argument is
2 of type `std::ostream&` and whose second argument is of type `Point const&`. If
3 it finds such a function it will call that function to do the work; if it cannot find such a
4 function it will issue a compiler diagnostic.

5 We write `operator<<` with a return type of `std::ostream&` so that one may chain
6 together multiple uses of the `<<` operator:

```
7 Point p0(1,2), p1(3,4);  
8 std::cout << p0 << " " << p1 << std::endl;
```

9 The C++ compiler parses this left to right. First it recognizes the expression
10 `std::cout << p0`. Because `std::cout` is of type `std::ostream`, and because
11 `p0` is of type `Point`, the compiler calls our stream insertion operator to do this work.
12 The return type of the function call is `std::ostream&`, and so the next expression is
13 recognized as a call to the stream insertion operator for an array of characters (" "). The
14 next is another call to our stream insertion operator for class `Point`, this time using the
15 object `p1`. This also returns a `std::ostream&`, allowing the last part of the expression
16 to be recognized as a call to the stream insertion operator for `std::endl`, which writes
17 a newline and flushes the output stream.

18 Look at the implementation of the stream insertion operator in `Point.cc`:

```
19 std::ostream& operator<<(std::ostream& ost,  
20                          Point const& p) {  
21     ost << "( "  
22         << p.x() << ", "  
23         << p.y()  
24         << " )";  
25  
26     return ost;  
27 }
```

28 The first argument, `ost`, is a reference to an object of type `std::ostream`; the name
29 `ost` has no special meaning to C++. When writing the implementation for this operator we
30 don't know and don't care what the output stream will be connected to; perhaps a file; per-
31 haps standard output. In any case, you send output to `ost` just as you do to `std::cout`,

1 which is just another variable of type `std::ostream`. In this example we chose to en-
2 close the values of `x_` and `y_` in parentheses in the printout and to separate them with a
3 comma; this is simply our choice, not something required by C++ or by *art*.

4 In this example, the stream insertion operator does *not* end by inserting a newline into
5 `ost`. This is a very common choice as it allows the user of the operator to have full
6 control about line breaks. For a class whose printout is very long and covers many lines,
7 you might decide that this operator should end by inserting newline character; it's your
8 choice.

9 If you wish to write a stream insertion operator for another class, just follow the pattern
10 used here.

11 If you want to understand more about why the operator is written the way that it is, consult
12 the standard C++ references; see Section 6.9. [sec:cpp:references](#)

13 The stream insertion operator is a *free function*(γ), not a member function of the class
14 `Point`; the tie to the class `Point` is via its second argument. Because this function is
15 a free function, it could have been declared in its own header file and its implementation
16 could have been provided in its own source file. However that is not common practice.
17 Instead the common practice is as shown in this example: to include it in `Point.h` and
18 `Point.cc`.

19 The choice of whether to put the declaration of the stream insertion operator (or any other
20 free function) into (1) the header file containing a class declaration or (2) its own header
21 file is a tradeoff between the following two criteria:



- 22 1. It may be convenient to have it in the class header file; otherwise users would have
23 to remember to include an additional header file when they want to use this operator
24 (or function).
- 25 2. One can imagine many simple free functions that take an object of type `Point` as
26 an argument. If they are all inside `Point.h`, and if each is only infrequently used,
27 then the compiler will waste time processing the declarations every time `Point.h`
28 is included somewhere.

29 **FIXME:** *MFP: Rob, I'd like to remove the argument against declaring the function*
30 *associated with a class in the same header as the class, because the extra time we're talking*
31 *about here is miniscule.*

1 The definition of this operator is typically put into the implementation file, rather than
2 being inlined. Such functions are generally poor candidates for inlining.

3 Ultimately this is a judgment call and the code in this example follows the recommenda-
4 tions made by the *art* development team. Their recommendation is that the following sorts
5 of free functions, and only these sorts, should be included in header files containing a class
6 declaration:

- 7 1. the stream insertion operator for that class
- 8 2. out-of-class arithmetic and comparison operators

9 With one exception, if including a function declaration in `Point.h` requires the inclusion
10 of an additional header in `Point.h`, declare that function in a different header file. The
11 exception is that it is okay to include `<iosfwd>`.

12 6.7.11 Review

c-classes-review

13 The class `Point` is an example of a class that is primarily concerned with providing
14 convenient access to the data it contains. Not all classes are like this; when you work
15 through the Workbook, you will write some classes that are primarily concerned with
16 packaging convenient access to a set of related functions:

- 17 1. class
- 18 2. object
- 19 3. identifier
- 20 4. free function
- 21 5. member function

22 **FIXME:** *What about destructors and assignment and moves? Also `=delete` and `=default`.*
23 *(Rob to address)*

6.8 Overloading functions

A more complete description of overload sets, and an introduction to the rules for overload resolution, will go here. This should give an example illustrating the kinds of error messages given by the compiler when no suitable overload can be found and also an example of the kind of error message that results when the match from an overload set is ambiguous.

6.9 C++ References

This section lists some recommended C++ references, both text books and online materials.

The following references describe the C++ core language,

- Stroustrup, Bjarne: “The C++ Programming Language, 4th Edition”, Addison-Wesley, 2013. ISBN 0321563840.
- <http://www.cplusplus.com/doc/tutorial/>

The following references describe the C++ Standard Library,

- Josuttis, Nicolai M., “The C++ Standard Library: A Tutorial and Reference (2nd Edition)”, Addison-Wesley, 2012. ISBN 0321623215.
- <http://www.cplusplus.com/reference>

The following contains an introductory tutorial. Many copies of this book are available at the Fermilab library. It is a very good introduction to the big ideas of C++ and Object Oriented Programming but it is not a fast entry point to the C++ skills needed for HEP. It also has not been updated for the current C++ standard.

- Andrew Koenig and Barbara E. Moo, “Accelerated C++: Practical Programming by Example” Addison-Wesley, 2000. ISBN 0-201-70353-X.

The following contains a discussion of recommended best practices. It has not been updated for the current C++ standard.

- 1 ○ Herb Sutter and Andrei Alexandrescu, “C++ Coding Standards: 101 Rules, Guide-
- 2 lines, and Best Practices.”, Addison-Wesley, 2005. ISBN 0-321-11358-6.

7 Using External Products in UPS

`:ups:setup`
Section 3.6.8 introduced the idea of external products. For the Intensity Frontier experiments (and for Fermilab-based experiments in general), access to external products is provided by a Fermilab-developed product-management package called Unix Product Support (UPS). An important UPS feature – demanded by most experiments as their code evolves – is its support for multiple versions of a product and multiple builds (e.g., for different platforms) per version.

Another notable feature is its capacity to handle multiple databases of products. So, for example, on Fermilab computers, login scripts (see Section 4.9) set up the UPS system, providing access to a database of products commonly used at Fermilab.



The *art* Workbook and your experiment's code will require additional products (available in other databases). For example, each experiment will provide a copy of the `toyExperiment` product in its experiment-specific UPS database.

In this chapter you will learn how to see which products UPS makes available, how UPS handles variants of a given product, how you use UPS to initialize a product provided in one of its databases and about the environment variables that UPS defines.

7.1 The UPS Database List: PRODUCTS

`products:var`
The act of setting up UPS defines a number of environment variables (discussed in Section 7.5), one of which is `PRODUCTS`. This particularly important environment variable merits its own section.

The environment variable `PRODUCTS` is a colon-delimited list of directory names, i.e., it is a path (see Section 4.6). Each directory in `PRODUCTS` is the name of a *UPS database*,

1 meaning simply that each directory functions as a repository of information about one or
 2 more products. When UPS looks for a product, it checks each directory in PRODUCTS,
 3 in the order listed, and takes the first match.



4 If you are on a Fermilab machine, you can look at the value of PRODUCTS just after
 5 logging in, before sourcing your site-specific setup script. Run `printenv`:

```
6 printenv PRODUCTS
```

7 It should have a value of

```
8 /grid/fermiapp/products/common/db
```

9 This generic Fermilab UPS database contains a handful of software products commonly
 10 used at Fermilab; most of these products are used by all of the Intensity Frontier Experi-
 11 ments. This database does not contain any of the experiment-specific software nor does
 12 it contain products such as *ROOT*(γ), *Geant4*(γ), CLHEP or *art*. While these last few
 13 products are indeed used by multiple experiments, they are often custom-built for each
 14 experiment and as such are distributed via the experiment-specific (i.e., separate) UPS
 15 databases.

16 After you source your site-specific setup script, look at PRODUCTS again. It will probably
 17 contain multiple directories, thus making many more products available in your “site” en-
 18 vironment. For example, on the DS50+Fermilab site, after running the DS50 setup script,
 19 PRODUCTS contains:

```
20 /ds50/app/products/:grid/fermiapp/products/common/db
```

21 You can see which products PRODUCTS contains by running `ls` on its directories, one-by-
 22 one, e.g.,

```
23 ls /grid/fermiapp/products/common/db
```

```
24 afs    git      ifdhc    mu2e     python   ...
25 cpn    gitflow jobsub_tools oracle_tnsnames ...
26 encp   gits    login    perl     setpath  ...
```

```
27 ls /ds50/app/products
```

```
28 art          cetpkgsupport g4neutronxs libxml2   ...
29 artdaq       clhep          g4nucleonxs messagefacility ...
30 art_suite    cmake          g4photon    mpich     ...
```



```

1 art_workbook_base  cpp0x          g4pii          mvapich2       ...
2 boost             cppunit        g4radiative    python          ...
3 caencomm          ds50daq        g4surface      root            ...
4 ...

```

Each directory name in these listings corresponds to the name of a UPS product. If you are on a different experiment, the precise contents of your experiment's product directory may be slightly different. Among other things, both databases contain a subdirectory named `ups`¹; this is for the UPS system itself. In this sense, all these products, including *art*, *toyExperiment* and even the product(s) containing your experiment's code, regard UPS as just another external product.

7.2 UPS Handling of Variants of a Product

An important feature of UPS is its capacity to make multiple variants of a product available to users. This of course includes different versions, but beyond that, a given version of a product may be built more than one way, e.g., for use by different operating systems (what UPS distinguishes as *flavors*). For example, a product might be built once for use with SLF5 and again for use with SLF6. A product may be built with different versions of the C++ compiler, e.g., with the production version and with a version under test. A product may be built with full compiler optimization or with the maximum debugging features enabled. Many variants can exist. UPS provides a way to select a particular build via an idea named *qualifiers*.

The full identifier of a UPS product includes its product name, its version, its flavor and its full set of qualifiers. In Section 7.3, you will see how to fully identify a product when you set it up.

7.3 The `setup` Command: Syntax and Function

Any given UPS database contains several to many, many products. To select a product and make it available for use, you use the `setup` command.

In most cases the correct flavor can be automatically detected by `setup` and need not be

¹`ups` appears in both listings; as always, the first match wins!

1 specified. However, if needed, flavor, in addition to various qualifiers and options can
 2 be specified. These are listed in the UPS documentation referenced later in this section.
 3 The version, if specified, must directly follow the product name in the command line,
 4 e.g.,:

```
5 setup options product-name product-version -f flavor -q qualifiers
```

6 Putting in real-looking values, it would look something like:

```
7 setup -R myproduct v3_2 -f SLF5 -q BUILD_A
```

8 What does the setup command actually do? It may do any or all of the following:

- 9 ○ define some environment variables
- 10 ○ define some bash functions
- 11 ○ define some aliases
- 12 ○ add elements to your PATH
- 13 ○ setup additional products on which it depends

14 Setting up dependent products works recursively. In this way, a single setup command
 15 may trigger the setup of, say, 15 or 20 products.

16 When you follow a given site-specific setup procedure, the PRODUCTS environment vari-
 17 able will be extended to include your experiment-specific UPS repository.

18 setup is a bash function (defined by the UPS product when it was initialized) that shadows
 19 a Unix system-configuration command also named setup, usually found in /usr/bin/setup
 20 or /usr/sbin/setup. Running the right 'setup' should work automatically as long as
 21 UPS is properly initialized. If it's not, setup returns the error message:

```
22 You are attempting to run ``setup`` which requires administrative  

  23 privileges, but more information is needed in order to do so.
```

24 If this happens, the simplest solution is to log out and log in again. Make sure that you
 25 carefully follow the instructions for doing the site specific setup procedure.

26 Few people will need to know more than the above about the UPS system. Those who do
 27 can consult the full UPS documentation at:

28 <http://www.fnal.gov/docs/products/ups/ReferenceManual/index.html>

7.4 Current Versions of Products

1

up:current

2 For some UPS products, but not all, the site administrator may define a particular fully-
3 qualified version of the product as the default version. In the language of UPS this notion
4 of default is called the *current* version. If a current version has been defined for a product,
5 you can set up that product with the command:

```
6 setup product-name
```

7 When you run this, the UPS system will automatically insert the version and qualifiers of
8 the version that has been declared current.

9 Having a current version is a handy feature for products that add convenience features to
10 your interactive environment; as improvements are added, you automatically get them.

11 However the notion of a current version is very dangerous if you want to ensure that
12 software built at one site will build in exactly the same way on all other sites. For this
13 reason, the Workbook fully specifies the version number and qualifiers of all products that
14 it requires; and in turn, the products used by the Workbook make fully qualified requests
15 for the products on which they depend.



7.5 Environment Variables Defined by UPS

16

setup:vars

17 When your login script or site-specific setup script initializes UPS, it defines many envi-
18 ronment variables in addition to PRODUCTS (Section 7.1), one of which is UPS_DIR, the
19 root directory of the currently selected version of UPS. The script also adds \$UPS_DIR/bin
20 to your PATH, which makes some UPS-related commands visible to your shell. Finally,
21 it defines the bash function setup (see Sections 4.8 and 7.3). When you use the setup
22 command, as illustrated below, it is this bash function that does the work.

23 In discussing the other important variables, the toyExperiment product will be used as an
24 example product. For a different product, you would replace “toyExperiment” or “TOY-
25 EXPERIMENT” in the following text by the product’s name. Once you have followed
26 your appropriate setup procedure (Table 5.1) you can issue the following command this
27 is informational for the purposes of this section; you don’t need to do it until you start
28 running the first Workbook exercise):

1 `setup toyExperiment v0_00_29 -qe2:prof`

2 The version and qualifiers shown here are the ones to use for the Workbook exercises.
3 When the setup command returns, the following environment variables will be defined:

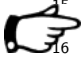
4 `TOYEXPERIMENT_DIR` defines the root DIRectory of the chosen UPS product

5 `TOYEXPERIMENT_INC` defines the path to the root directory of the C++ header files
6 that are provided by this product (so called because the header files are INcluded)

7 `TOYEXPERIMENT_LIB` defines the directory that contains all of the dynamic object
8 LIBraries (ending in `.so`) that are provided by this product

9 Almost all UPS products that you will use in the Workbook define these three environment
10 variables. Several, including `toyExperiment`, define many more. Once you're running the
11 exercises, you will be able to see all of the environment variables defined by the `toyExper-`
12 `iment` product by issuing the following command:

13 `printenv | grep TOYEXPERIMENT`

14 Many software products have version numbers that contain dot characters. UPS requires
15 that version numbers not contain any dot characters; by convention, version dots are re-
 16 placed with underscores. Therefore `v0.00.14` becomes `v0_00_14`. Also by convention,
17 the environment variables are all upper case, regardless of the case used in the product
18 names.

19 7.6 Finding Header Files

20 7.6.1 Introduction

21 *Header files* were introduced in Section 6.4.2. Recall that a header file typically contains
22 the “parts list” for its associated `.cc` source file and is “included” in the `.cc` file.

23 The software for the Workbook depends on a large number of external products; the same
24 is true, on an even larger scale, for the software in your experiment. The preceding sec-
25 tions in this chapter discussed how to establish a working environment in which all of
26 these software products are available for use.

1 When you are working with the code in the Workbook, and when you are working on your
2 experiment, you will frequently encounter C++ classes and functions that come from these
3 external products. An important skill is to be able to identify them when you see them and
4 to be able to follow the clues back to their source and documentation. This section will
5 describe how to do that.

6 An important aid to finding documentation is the use of *namespaces*; if you are not famil-
7 iar with namespaces, **FIXME:** *see Section 59.6 or* [consult the standard C++ documenta-](#)
8 [tion.](#)

9 7.6.2 Finding *art* Header Files

10 This subsection will use the example of the class `art::Event` to illustrate how to find
11 header files from the *art* UPS product; this will serve as a model for finding header files
12 from most other UPS products.

13 The class that holds the *art* abstraction of an HEP event is named, `art::Event`; that
14 is, the class `Event` is in the namespace `art`. In fact, all classes and functions defined by
15 *art* are in the namespace `art`. The primary reason for this is to minimize the chances of
16 accidental name collisions between *art* and other codes; but it also serves a very useful
17 documentation role and is one of the clues you can use to find header files.

18 If you look at code that uses `art::Event` you will almost always find that the file in-
19 cludes the following header file:

```
20 1 #include "art/Framework/Principal/Event.h"
```

21 The *art* UPS product has been designed so that the relative path used to include any *art*
22 header file starts with the directory `art`; this is another clue that the class or function of
23 interest is part of *art*.

24 When you setup the *art* UPS product, it defines the environment variable `ART_INC`, which
25 points to the root of the header file tree for *art*. You now have enough information to
26 discover where to find the header file for `art::Event`; it is at

```
27 $ART_INC/art/Framework/Principal/Event.h
```

28 You can follow this same pattern for any class or function that is part of *art*. This will only
29 work if you are in an environment in which `ART_INC` has been defined, which will be

1 described in Chapters [chap:orebuild:run:first:module](#) 9 and 10.

2 If you are new to C++, you will likely find this header file difficult to understand; you do
3 not need to understand it when you first encounter it but, for future reference, you do need
4 to know where to find it.

5 Earlier in this section, you read that if a C++ file uses `art::Event`, it would *almost*
6 *always* include the appropriate header file. Why *almost* always? Because the header file
7 `Event.h` might already be included within one of the other headers that are included in
8 your file. If `Event.h` is indirectly included in this way, it does not hurt also to include it
9 explicitly, but it is not required that you do so.²

10 We can summarize this discussion as follows: if a C++ source file uses `art::Event` it
11 must always include the appropriate header file, either directly or indirectly.

12 `art` does not rigorously follow the pattern that the name of file is the same as the name of
13 the class or function that it defines. The reason is that some files define multiple classes
14 or functions; in most such cases the file is named after the most important class that it
15 defines.

16 Finally, from time to time, you will need to dig through several layers of header files to
17 find the information you need.

18 There are two code browsing tools that you can use to help navigate the layering of header
19 files and to help find class declarations that are not in a file named for the class:

- 20 1. use the *art redmine*(γ) repository browser:
21 <https://cdcv.s.fnal.gov/redmine/projects/art/repository/revisions/master/show/art>
- 22 2. use the LXR code browser: <http://cdcv.s.fnal.gov/lxr/art/>

23 (In the above, both URLs are live links.)

24 7.6.3 Finding Headers from Other UPS Products

25 Section [sec:toy:experiment](#) 3.7 introduced the idea that the Workbook is built around a UPS product named
26 `toyExperiment`, which describes a made-up experiment. All classes and functions defined

² Actually there is small price to pay for redundant includes; it makes the compiler do unnecessary work, and therefore slows it down. But providing some redundant includes as a pedagogical tool is often a good trade-off; the Workbook will frequently do this.

1 in this UPS product are defined in the namespace `tex`, which is an acronym-like shorthand
 2 for `toyExperiment` (ToyEXperiment). (This shorthand makes it (a) easier to focus on the
 3 name of each class or function rather than the namespace and (b) quicker to type.)

4 One of the classes from the `toyExperiment` UPS product is `tex::GenParticle`, which
 5 describes particles created by the event generator, the first part of the simulation chain (see
 6 Section 3.7.2). The include directive for this class looks like

```
7 1 #include "toyExperiment/MCDataProducts/GenParticle.h"
```

8 As for headers included from `art`, the first element in the relative path to the included file
 9 is the name of the UPS product in which it is found. Similarly to `art`, the header file can be
 10 found using the environment variable `TOYEXPERIMENT_INC`:

```
11 $TOYEXPERIMENT_INC/toyExperiment/MCDataProducts/GenParticle.h
```

12 With a few exceptions, discussed in Section 7.6.4, if a class or function from a UPS product
 13 is used in the Workbook code, it will obey the following pattern:

- 14 1. The class will be in a namespace that is unique to the UPS product; the name of the
 15 namespace may be the full product name or a shortened version of it.
- 16 2. The lead element of the path specified in the include directive will be the name of
 17 the UPS product.
- 18 3. The UPS product setup command will define an environment variable named
 19 `PRODUCT-NAME_INC`, where `PRODUCT-NAME` is in all capital letters.

20 Using this information, the name of the header file will always be

```
21 $PRODUCT-NAME_INC/path-specified-in-the-include-directive
```

22 This pattern holds for all of the UPS products listed in Table 7.1.

23 A table listing git- and LXR-based code browsers for many of these UPS products can be
 24 found near the top of the web page:

```
25 https://cdcvs.fnal.gov/redmine/projects/art/wiki
```

26 7.6.4 Exceptions: The Workbook, ROOT and Geant4

27 There are three exceptions to the pattern described in Section 7.6.3:

Table 7.1: For selected UPS Products, this table gives the names of the associated namespaces. The UPS products that do not use namespaces are discussed in Section 7.6.4. †The namespace `tex` is also used by the *art* Workbook, which is not a UPS product.

UPS Product	Namespace
<code>art</code>	<code>art</code>
<code>boost</code>	<code>boost</code>
<code>cet</code>	<code>cetlib</code>
<code>clhep</code>	<code>CLHEP</code>
<code>fhiclcpp</code>	<code>fhicl</code>
<code>messagefacility</code>	<code>mf</code>
<code>toyExperiment</code>	<code>tex</code> †

- 1 ○ the Workbook itself
- 2 ○ ROOT
- 3 ○ Geant4

4 The Workbook is so tightly coupled to the `toyExperiment` UPS product that all classes
5 in the Workbook are also in its namespace, `tex`. Note, however, that classes from the
6 Workbook and the `toyExperiment` UPS product can still be distinguished by the leading
7 element of the relative path found in the include directives for their header files:

- 8 ○ `art-workbook` for the Workbook
- 9 ○ `toyExperiment` for the `toyExperiment`

10 The ROOT package is a CERN-supplied software package that is used by *art* to write data
11 to disk files and to read it from disk files. It also provides many data analysis and data
12 presentation tools that are widely used by the HEP community. Major design decisions for
13 ROOT were frozen before namespaces were a stable part of the C++ language, therefore
14 ROOT does not use namespaces. Instead ROOT adopts the following conventions:

- 15 1. All class names by defined by ROOT start with the capital letter T followed by
16 another upper case letter; for example, `TFile`, `TH1F`, and `TCanvas`.
- 17 2. With very few exceptions, all header files defined by ROOT also start with the same
18 pattern; for example, `TFile.h`, `TH1F.h`, and `TCanvas.h`.
- 19 3. The names of all global objects defined by ROOT start with a lower case letter `g`

1 followed by an upper case letter; for example `gDirectory`, `gPad` and `gFile`.

2 The rule for writing an include directive for a header file from ROOT is to write its name
3 without any leading path elements:

```
4 1 #include "TFile.h"
```

5 All of the ROOT header files are found in the directory that is pointed to by the envi-
6 ronment variable `$ROOT_INC`. For example, to see the contents of this file you could
7 enter:

```
8 less $ROOT_INC/TFile.h
```

9 Or you can learn about this class using the reference manual at the CERN web site:
10 <http://root.cern.ch/root/html534/ClassIndex.html>

11 You will not see the `Geant4` package in the Workbook but it will be used by the software
12 for your experiment, so it is described here for completeness. `Geant4` is a toolkit for mod-
13 eling the propagation particles in electromagnetic fields and for modeling the interactions
14 of particles with matter; it is the core of all detector simulation codes in HEP and is also
15 widely used in both the Medical Imaging community and the Particle Astrophysics com-
16 munity.

17 As with ROOT, `Geant4` was designed before namespaces were a stable part of the C++
18 language. Therefore `Geant4` adopted the following conventions.

- 19 1. The names of all identifiers begin with `G4`; for example, `G4Step` and `G4Track`.
- 20 2. All header file names defined by `Geant4` begin with `G4`; for example, `G4Step.h`
21 and `G4Track.h`.

22 Most of the header files defined by `Geant4` are found in a single directory, which is pointed
23 to by the environment variable `G4INCLUDE`.

24 The rule for writing an include directive for a header file from `Geant4` is to write its name
25 without any leading path elements:

```
26 #include "G4Step.h"
```

27 The workbook does not set up a version of `Geant4`; therefore `G4INCLUDE` is not defined.
28 If it were, you would look at this file by:

1 `less $G4INCLUDE/G4Step.h`



Both ROOT and Geant4 define many thousands of classes, functions and global variables.

3 In order to avoid collisions with these identifiers, do not define any identifiers that begin
4 with any of (case-sensitive):

- 5 ○ T, followed by an upper case letter
- 6 ○ g, followed by an upper case letter
- 7 ○ G4

1

Part II

2

Workbook

DRAFT

8 Preparation for Running the Workbook Exercises

8.1 Introduction

The Workbook exercises can be run in several environments:

1. on a computer that is maintained by your experiment, either at Fermilab or at another institution.
2. on one of the computers supplied for the *art/LArSoft* course.
3. on your own computer on which you install the necessary software. For details see Appendix B.

Many details of the working environment change from site to site¹ and these differences are parameterized so that (a) it is easy to establish the required environment, and (b) the Workbook exercises behave the same way at all sites. In this chapter you will learn how to find and log into the right machine remotely from your local machine (laptop or desktop), and make sure it can support your Workbook work.

8.2 Getting Computer Accounts on Workbook-enabled Machines


In order to run the exercises in the Workbook, you will need an account on a machine that can access your site's installation of the Workbook code. The experiments provide instruc-

¹Remember, a *site* refers to a unique combination of experiment and a location; your laptop can be its own site.

Table 8.1: Experiment-specific information for new users (pages are under <https://cdcv.s.fnal.gov/redmine/projects/> except for Mu2e and NOvA)

Experiment	Page for New Users
ArgoNeut	larsoftsvn/wiki/Using_LArSoft_on_the_GPVM_nodes
Darkside	darkside-public/wiki/Before_You_Arrive
LArSoft	larsoftsvn
DUNE	larsoftsvn/wiki/Using_LArSoft_on_the_GPVM_nodes
MicroBoone	larsoftsvn/wiki/Using_LArSoft_on_the_GPVM_nodes
Muon g-2	g-2/wiki/NewGm2Person
Mu2e	http://mu2e.fnal.gov/atwork/general/userinfo/index.shtml#comp
NOvA	http://www-nova.fnal.gov/NOvA_Collaboration_Information/index.html

1 tions for getting computer accounts on their machines (and various other information for
 2 new users) on web pages that they maintain, as listed in Table 8.1. The URLs in the table
 3 are live hyperlinks.

4 Currently, each of the experiments using *art* has installed the Workbook code on one of its
 5 experiment machines in the Fermilab General Purpose Computing Farm (GPCF). 

6 At time of writing, the new-user instructions for all LArSoft-based experiments are at the
 7 LArSoft site; there are no separate instructions for each experiment.

8 If you are planning to take the *art*/LArSoft course, see the course web site to learn how to
 9 get an account on the machines reserved for the course.

10 If you would like a computer account on a Fermilab computer in order to evaluate *art*,
 11 contact the *art* team (see Section 3.4).

12 8.3 Choosing a Machine and Logging In

13 The experiment-specific machines confirmed to host the Workbook code are listed in Ta-
 14 ble 8.2 In most cases the name given is not the name of an actual computer, but rather a
 15 round-robin alias for a cluster. For example, if you log into `mu2evm`, you will actually
 16 be connected to one of the five computers `mu2egpvm01` through `mu2egpvm05`. These
 17 Mu2e machines share all disks that are relevant to the Workbook exercises, so if you need

Table 8.2: Login machines for running the Workbook exercises

Experiment	Name of Login Node
ArgoNeut	argoneutvm.fnal.gov
Darkside	ds50.fnal.gov
DUNE	lbnevm.fnal.gov
MicroBoone	uboonevm.fnal.gov
Muon g-2	gm2gpvm.fnal.gov
Mu2e	mu2egpvm0x.fnal.gov, for x=1,2,3,4,5
NOvA	nova-offline.fnal.gov
art/LArSoft Course	alcourse.fnal.gov alcourse2.fnal.gov

tab:expMachines

- 1 to log in multiple times, it is perfectly OK if you are logged into two different machines;
- 2 you will still see all of the same files.
- 3 Each experiment's web page **FIXME:** *check* has instructions on how to log in to its com-
- 4 puters from your local machine.

5 8.4 Launching new Windows: Verify X Connectivity

- 6 Some of the Workbook exercises will launch an X window from the remote machine that
- 7 opens in your local machine. To test that this works, type `xterm &`:

8 `xterm &`



- 9 This should, without any messages, give you a new command prompt. After a few seconds,
- 10 a new shell window should appear on your laptop screen; if you are logging into a Fermilab
- 11 computer from a remote site, this may take up to 10 seconds. If the window does not
- 12 appear, or if the command issues an error message, contact a computing expert on your
- 13 experiment.

- 14 To close the new window, type `exit` at the command prompt in the new window:

15 `exit`



- 17 If you have a problem with `xterm`, it could be a problem with your Kerberos and/or `ssh` configurations. Try logging in again with `ssh -Y`.

8.5 Choose an Editor

1

ec:editors

2 As you work through the Workbook exercises you will need to edit files. Familiarize your-
3 self with one of the editors available on the computer that is hosting the Workbook. Most
4 Fermilab computers offer four reasonable choices: emacs, vi, vim and nedit. Of these,
5 nedit is probably the most intuitive and user-friendly. All are very powerful once you have
6 learned to use them. Most other sites offer at least the first three choices. You can always
7 contact your local system administrator to suggest that other editors be installed.

8 *A future version of this documentation suite will include recommended configurations for*
9 *each editor and will provide links to documentation for each editor.*

DRAFT

9 Exercise 1: Running Pre-built *art* Modules

chap:prebuilt

9.1 Introduction

sec:prebuilt:intro

In this first exercise of the Workbook, you will be introduced to the *FHiCL*(γ) configuration language and you will run *art* on several modules that are distributed as part of the toyExperiment UPS product. You will not compile or link any code.

9.2 Prerequisites

sec:prebuilt:prereq

Before running any of the exercises in this Workbook, you need to be familiar enough with the material discussed in Part I (Introduction) of this documentation set and with Chapter 8 to be able to find information as needed.

chap:robsGett.



If you are following the instructions below on an older Mac computer (OSX 10.6, Snow Leopard, or earlier), and if you are reading the instructions from a PDF file, be aware that if you use the mouse or trackpad to cut and paste text from the PDF file into your terminal window, the underscore characters will be turned into spaces. You will have to fix them before the commands will work.

9.3 What You Will Learn

sec:prebuilt:what

In this exercise you will learn:

- how to use the site-specific setup procedure, which you must do once at the start of each login session.
- a little bit about the *art* run-time environment (Section 9.4)

sec:prebuilt:runtime:end

- 1 ○ how to set up the toyExperiment UPS product (Section 9.6.1)
- 2 ○ how to run an *art* job (Section 9.6.1)
- 3 ○ how to control the number of events to process (Section 9.8.4)
- 4 ○ how to select different input files (Section 9.8.5)
- 5 ○ how to start at a run, subRun or event that is not the first one in the file (Section 9.8.6)
- 6 ○ how to concatenate input files (Section 9.8.5)
- 7 ○ how to write an output file (Section 9.8.9)
- 8 ○ some basics about the grammar and structure of a FHiCL file (Section 9.8)
- 9 ○ how *art* finds modules and configuration (FHiCL) files. (Sections 9.10 and 9.11)

10 9.4 The *art* Run-time Environment

11 This discussion is aimed to help you understand the process described in this chapter as a
 12 whole and how the pieces fit together in the *art* run-time environment. This environment
 13 is summarized in Figure 9.1. In this figure the boxes refer either to locations in memory or
 14 to files on a disk.

15 At the center of the figure is a box labelled “*art* executable;” this represents the *art* main
 16 program resident in memory after being loaded. When the *art* executable starts up, it reads
 17 its run-time configuration (FHiCL) file, represented by the box to its left. Following in-
 18 structions from the configuration file, *art* will load dynamic libraries from toyExperi-
 19 ment, from *art*, from ROOT, from CLHEP and from other UPS products. All of these
 20 dynamic libraries (.so or .dylib files) will be found in the appropriate UPS products
 21 in LD_LIBRARY_PATH (DYLD_LIBRARY_PATH for OS X), which points to directo-
 22 ries in the UPS products area (box at upper right). Also following instructions from the
 23 FHiCL file, *art* will look for input files (box labeled “Event-data input files” at right). The
 24 FHiCL file will tell *art* to write its event-data and histogram output files to a particular
 25 directory (box at lower right).

26 One remaining box in the figure (at right, second from bottom) is not encountered in the
 27 first Workbook exercise but has been provided for completeness. In most *art* jobs it is
 28 necessary to access experiment-related geometry and conditions information; in a mature

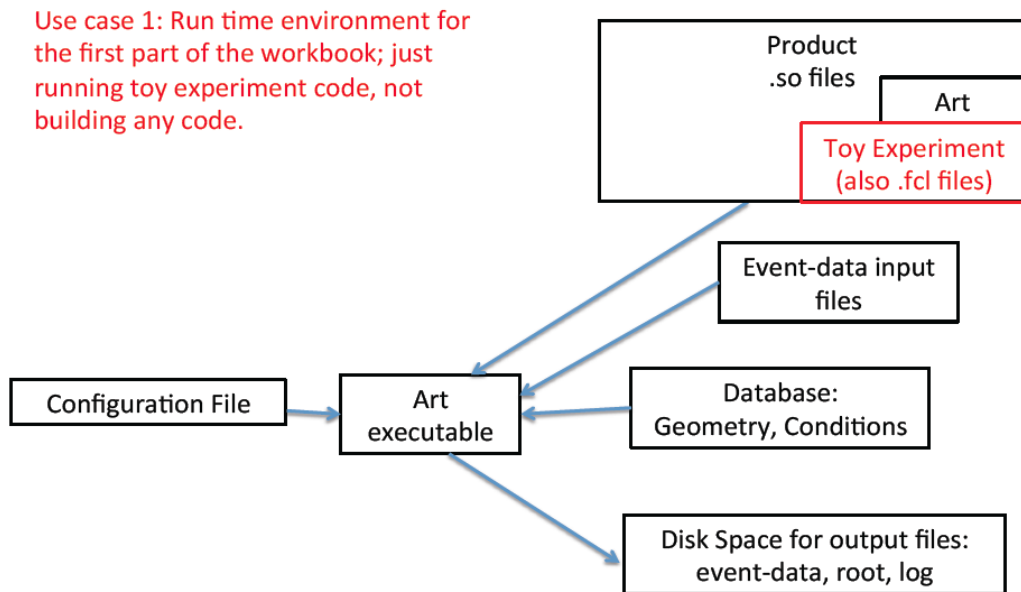


Figure 9.1: Elements of the *art* run-time environment for the first Workbook exercise

- 1 experiment, these are usually stored in a database that stands apart from the other elements
 2 in the picture.
- 3 The arrows in Figure 9.1 show the direction in which information flows. Everything but
 4 the output flows into the *art* executable.

5 9.5 The Input and Configuration Files for the Work- 6 book Exercises

- 7 Several event-data input files have been provided for use by the Workbook exercises. These
 8 input files are packaged as part of the toyExperiment UPS product. Table 9.1 lists the
 9 range of event IDs found in each file. You will need to refer back to this table as you
 10 proceed.
- 11 A run-time configuration (FHiCL) file has been provided for each exercise. For Exercise
 12 1 it is `hello.fcl`.

Table 9.1: Input files provided for the Workbook exercises

File Name	Run	SubRun	Range of Event Numbers
input01.art	1	0	1...10
input02.art	2	0	1...10
input03.art	3	0	1...5
	3	1	1...5
	3	2	1...5
input04.art	4	0	1...1000

9.6 Setting up to Run Exercise 1

9.6.1 Log In and Set Up

The intent of this section is for the reader to start from “zero” and execute an *art* job, without necessarily understanding each step, just to get familiar with the process. A detailed discussion of what these steps do will follow in Section 9.9.

Some steps are written as statements, others as commands to issue at the prompt. Notice that *art* takes the argument `-c hello.fcl`; this points *art* to the run-time configuration file that will tell it what to do and where to find the “pieces” on which to operate.

Most readers: Follow the steps in Section 9.6.1.1, then proceed directly to Section 9.7.

If you wish to manage your working directory yourself, skip Section 9.6.1.1, follow the steps in Section 9.6.1.2, then proceed to Section 9.7.

If you log out and wish to log back in to continue this exercise, follow the procedure outlined in Section 9.6.1.3.

9.6.1.1 Initial Setup Procedure using Standard Directory

1. Log in to the computer you chose in Section 8.3.
2. Follow the site-specific setup procedure; see Table 5.1.
3. Make the standard working directory then `cd` to it; substitute your

kerberos principal **FIXME:** *Does everyone use kerberos?* for the string *username*. These commands, shown on two lines, can each be typed on a single line.

```
mkdir -p $ART_WORKBOOK_WORKING_BASE/username/\  
workbook-tutorial/pre-built
```

```
cd $ART_WORKBOOK_WORKING_BASE/username/\  
workbook-tutorial/pre-built
```

4. Setup the toyExperiment UPS product:

```
setup toyExperiment v0_00_29 -q$ART_WORKBOOK_QUAL:prof
```

5. Copy the scripts into your working directory:

```
cp $TOYEXPERIMENT_DIR/HelloWorldScripts/* .
```

6. Use the provided script to create the symbolic links needed by the FHiCL files:

```
source makeLinks.sh
```

7. See what you have in the directory:

```
ls  
bug01.fcl bug04.fcl inputFiles output  
bug02.fcl helloExample.log inputs.txt  
skipEvents.fcl bug03.fcl hello.fcl  
makeLinks.sh writeFile.fcl
```

1

- 2 Proceed to Section [9.7](#). lssec:first-login-output

1 9.6.1.2 Initial Setup Procedure allowing Self-managed Working Directory



gin-expert

1. Log in to the computer you chose in Section [8.3](#).
2. Follow the site-specific setup procedure; see Table [5.1](#).
3. Make a working directory and cd to it.
4. Setup the toyExperiment UPS product:

```
setup toyExperiment v0_00_29 -q$ART_WORKBOOK_QUAL:prof
```

5. Copy the scripts into your working directory:

```
cp $TOYEXPERIMENT_DIR>HelloWorldScripts/* .
```

6. Make a subdirectory named `output`. If you prefer, you can make this on some other disk and put a symbolic link to it from the current working directory; name the link `output`.
7. Create a symbolic link to allow the FHiCL files to find the input files:

```
ln -s $TOYEXPERIMENT_DIR/inputFiles .
```

8. See what you have in the directory:

```
ls
bug01.fcl bug04.fcl inputFiles output
bug02.fcl helloExample.log inputs.txt
skipEvents.fcl bug03.fcl hello.fcl
makeLinks.sh writeFile.fcl
```

2

- 3 Proceed to Section [9.7](#).

9.6.1.3 Setup for Subsequent Exercise 1 Login Sessions

If you log out and later wish to log in again to work on this exercise, you need to do the following:

1. Log in to the computer you chose in Section 8.3.
2. Follow the site-specific setup procedure; see Section 5.
3. `cd` to your working directory, e.g., for the standard case (shown here on two lines):

```
cd $ART_WORKBOOK_WORKING_BASE/username/\
workbook-tutorial/pre-built
```

4. Setup the toyExperiment UPS product:

```
setup toyExperiment v0_00_29 -q$ART_WORKBOOK_QUAL:prof
```

Compare this with the list given in Section 9.6.1. You will see that three steps are missing because they only need to be done the first time.

You are now ready to run *art* as you were before.

9.7 Execute *art* and Examine Output

From your working directory, execute *art* on the FHiCL file `hello.fcl` and send the output to `output/hello.log`:

```
art -c hello.fcl >& output/hello.log
```

Compare the output you produced (in the file `output/hello.log`) against Listing 9.1; the only differences should be the timestamps and some line breaking. *art* will have processed the first file listed in Table 9.1.

Listing 9.1: Sample output from running `hello.fcl`

```

1 %MSG-i MF_INIT_OK: art 27-Apr-2013 21:22:13 CDT JobSetup
2 Messagellogger initialization complete.
3 %MSG
4 27-Apr-2013 21:22:14 CDT Initiating request to open file
5 inputFiles/input01.art
6 27-Apr-2013 21:22:14 CDT Successfully opened file
7 inputFiles/input01.art
8 Begin processing the 1st record. run: 1 subRun: 0 event: 1 at
9 27-Apr-2013 21:22:14 CDT
10 Hello World! This event has the id: run: 1 subRun: 0 event: 1
11 Begin processing the 2nd record. run: 1 subRun: 0 event: 2 at
12 27-Apr-2013 21:22:14 CDT
13 Hello World! This event has the id: run: 1 subRun: 0 event: 2
14 Hello World! This event has the id: run: 1 subRun: 0 event: 3
15 Hello World! This event has the id: run: 1 subRun: 0 event: 4
16 Hello World! This event has the id: run: 1 subRun: 0 event: 5
17 Hello World! This event has the id: run: 1 subRun: 0 event: 6
18 Hello World! This event has the id: run: 1 subRun: 0 event: 7
19 Hello World! This event has the id: run: 1 subRun: 0 event: 8
20 Hello World! This event has the id: run: 1 subRun: 0 event: 9
21 Hello World! This event has the id: run: 1 subRun: 0 event: 10
22 27-Apr-2013 21:22:14 CDT Closed file inputFiles/input01.art
23
24 TrigReport ----- Event Summary -----
25 TrigReport Events total = 10 passed = 10 failed = 0
26
27 TrigReport ----- Modules in End-Path: e1 -----
28 TrigReport Trig Bit# Visited Passed Failed Error Name
29 TrigReport 0 0 10 10 0 0 hi
30
31 TimeReport ----- Time Summary ---[sec]-----
32 TimeReport CPU = 0.004000 Real = 0.002411
33
34 Art has completed and will exit with status 0.

```

35 Every time you run *art*, the first thing to check is the last line in your output or log file. It
36 should be Art has completed and will exit with status 0. If the sta-
37 tus is not 0, or if this line is missing, it is an error; please contact the *art* team as described
38 in Section 3.4. [sec:getting-help](#)

39 *A future version of these instructions will specify how much disk space is needed, including*
40 *space for all output files.*

9.8 Understanding the Configuration

1

```
anding:hello:fcl
```

2

The file `hello.fcl` shown in Listing 9.2 gives *art* its run-time configuration.

```
lst:hellofcl
```

Listing 9.2: Listing of `hello.fcl`

```

1 #include "fcl/minimalMessageService.fcl"
2
3 process_name : hello
4
5 source : {
6     module_type : RootInput
7     fileNames   : [ "inputFiles/input01.art" ]
8 }
9
10 services : {
11     message : @local::default_message
12 }
13
14 physics :{
15     analyzers: {
16         hi : {
17             module_type : HelloWorld
18         }
19     }
20
21     e1      : [ hi ]
22     end_paths : [ e1 ]
23 }

```

26 This file is written in the Fermilab Hierarchical Configuration Language (FHiCL, pro-
 27 nounced “fickle”), a language that was developed at Fermilab to support run-time config-
 28 uration for several projects, including *art*. By convention, files written in FHiCL end in
 29 `.fcl`. As you work through the Workbook, the features of FHiCL that are relevant for
 30 each exercise will be explained.

31 *art* accepts some command line options that can be used in place of items in the FHiCL
 32 file. You will encounter some of these in this section.



34 The full details of the FHiCL language, plus the details of how it is used by *art*, are given in
 35 the Users Guide, Chapter 52. Most people will find it much easier to follow the discussion
 in the Workbook documentation than to digest the full documentation up front.

fcl-syntax


9.8.1 Some Bookkeeping Syntax

In a FHiCL file, the start of a comment is marked either by the hash sign character (#) or by a C++ style double slash (//); a comment may begin in any column.

The hash sign has one other use. If the first eight characters of a line are exactly `#include`, followed by whitespace and a quoted file path, then the line will be interpreted as an *include directive* and the line containing it will be replaced by the contents of the file named in the include directive.


The basic element of FHiCL is the *definition*, which has the form

```
name : value
```

A group of FHiCL definitions delimited by braces {} is called a *table*(γ). Within *art*, a FHiCL table gets turned into a C++ object called a *parameter set*(γ); this document set will often refer to a FHiCL table as a parameter set. 

The fragment of `hello.fcl` shown below contains the FHiCL table that configures the *source*(γ) of events that *art* will read in and operate on.

```
source : {
  module_type : RootInput
  fileNames   : [ "inputFiles/input01.art" ]
}
```

The name `source` (line 5, above) is an *identifier* in *art*; i.e., the name `source` has no special meaning to FHiCL but it does have a special meaning to *art*. To be precise, it only has a special meaning to *art* if it is at the outermost *scope*(γ) of a FHiCL file; i.e., not inside any braces {} within the file. **FIXME:** refer to discussion of scope once it's in Chapter 14 [ch:parameter-sets](#) 

When *art* sees a parameter set named `source` at the outermost scope, it will interpret that parameter set to be the description of the source of events for this run of *art*.

Within the `source` parameter set, `module_type` (line 6) is an identifier in *art* that tells *art* the name of a module that it should load and run, `RootInput` in this case. `RootInput` is one of the standard source modules provided by *art* and it reads disk files containing event-data written in an *art*-defined ROOT-based format. The default behavior of the `RootInput` module is to start at the first event in the first file and read to the end of the last event in the last file.¹

¹ In the Workbook, the only source `module_type` that you will see will be `RootInput`. Your exper-

1 The string `fileNames` (line 7) is again an identifier, but this time defined in the `RootInput`
 2 module. It gives the input module a list of filenames from which to read events. The list is
 3 delimited by square brackets and contains a comma-separated list of filenames. This exam-
 4 ple shows only one filename, but the square brackets are still required. The proper FHiCL
 5 name for a comma-separated list delimited by square brackets is a *sequence*(γ).

6 In most cases the filenames in the sequence must be enclosed in quotes. FHiCL, like many
 7 other languages has the following rule: if a string contains white space or any special
 8 characters, then quoting it is required, otherwise quotes are optional.

9 FHiCL has its own set of special characters; these include anything *except* all upper and
 10 lower case letters, the numbers 0 through 9 and the underscore character. *art* restricts the
 11 use of the underscore character in some circumstances; these will be discussed as they
 12 arise.

13 It is implied in the foregoing discussion that a FHiCL *value* need not be a simple thing,
 14 such as a number or a quoted string. For example, in `hello.fcl`, the value of `source`
 15 is a parameter set (of two parameters) and the value of `fileNames` is a (single-item)
 16 sequence.

17 9.8.2 Some Physics Processing Syntax

phys-fcl-syntax

18 The identifier *physics*(γ), when found at the outermost scope, is an identifier reserved to
 19 *art*. The `physics` parameter set is so named because it contains most of the information
 20 needed to describe the physics workflow of an *art* job.

21 The fragment of `hello.fcl` below shows a rather long-winded way of telling *art* to
 22 find a module named `HelloWorld` and execute it. Why so long-winded? *art* has very
 23 powerful features that enable execution of multiple complex chains of modules; the price
 24 is that specifying something simple takes a lot of keystrokes.

```
14 physics :{
15     analyzers: {
16         hi : {
17             module_type : HelloWorld
18         }
19     }
20 }
```

iment may have a source module that reads events from the live experiment and other source modules that read files written in experiment-defined formats.

```

20  e1      : [ hi ]
21  end_paths : [ e1 ]
22  }

```

```
cl:physics
```

4 At the outermost scope of the FHiCL file, *art* will interpret the physics parameter set as the
5 description of the physics workflow for this run of *art*. Within the `physics` parameter
6 set, notice the identifier `analyzers` on line 15. When found as a top-level identifier
7 within the `physics` scope, as shown here, it is recognized as a keyword reserved to *art*.
8 The `analyzers` parameter set defines the run-time configuration for all of the analyzer
9 modules that are part of the job — in this case, only `HelloWorld` (specified on line
10 17).

11 For our current purposes, the module `HelloWorld` does only one thing of interest,
12 namely for every event it prints one line (shown here as three):

```

13 Hello World! This event has the id: run: <RR>
14                                     subRun: <SS>
15                                     event: <EE>

```

16 where `RR`, `SS` and `EE` are substituted with the actual run, subRun and event number of
17 each event.

18 If you look back at Listing 9.1, you will see that this line appears ten times, once each
19 for events 1 through 10 of run 1, subRun 0 (as expected, according to Table 9.1). The
20 remainder of the listing is standard output generated by *art*.

21 On line 20, `e1` (an arbitrary identifier) is called a *path*; it is a FHiCL sequence of module
22 labels. On line 21, `end_paths` — an identifier reserved to *art* — is a FHiCL sequence of
23 path names. Together, these two identifiers specify the workflow; this will be discussed in
24 Section 9.8.8.

25 The remainder of the lines in `hello.fcl` appears below. Line 3 (different line number
26 than in Listing 9.2), starting with `process_name(γ)`, tells *art* that this job has a name and
27 that the name is “hello”; it has no real significance in these simple exercises. However
28 the name of the process must not contain any underscore characters; the reason for this
29 restriction will be explained in Section 16.4.2.

```

3b #include "fcl/minimalMessageService.fcl"
3c
3d   process_name : hello
3e ...

```

```

5     services : {
6         message : @local::default_message
7     }

```

```
ex1:fcl:services
```

4 The `services` parameter set (lines 5-7) provides the run-time configuration for all the
 5 required *art* services for the job, in this case only the message service. For our present pur-
 6 poses, it is sufficient to know that the configuration for the message service itself is found
 7 inside the file that is included in line 1. The message service controls the limiting and rout-
 8 ing of debug, informational, warning and error messages generated by *art* or by user code;
 9 it does not control information written directly to `std::cout` or `std::cerr`.

10 9.8.3 *art* Command line Options

```
lt:command:line
```

11 *art* supports some command line options. To see what they are, type the following com-
 12 mand at the bash prompt:

```
13 art --help
```

14 Note that some options have both a short form and a long form. This is a common con-
 15 vention for Unix programs; the short form is convenient for interactive use and the long
 16 form makes scripts more readable. It is also a common convention that the short form of
 17 an option begins single dash character, while the long form of an option begins with two
 18 dash characters, for example `--help` above.

19 9.8.4 Maximum Number of Events to Process

```
uilt:max:events
```

20 By default *art* will read all events from all of the specified input files. You can set a maxi-
 21 mum number of events in two ways, one way is from the command line:

```
22 art -c hello.fcl -n 5 >& output/hello-n5.log
```

```
23 art -c hello.fcl --nevts 4 >& output/hello-nevts4.log
```

24 Run each of these commands and observe their output.

25 The second way is within the FHiCL file. Start by making a copy of `hello.fcl`:

```
26 cp hello.fcl hi.fcl
```


1 Edit `hi.fcl` and add the following line anywhere in the source parameter set:

```
2 maxEvents : 3
```

3 By convention this is added after the `fileNames` definition but it can go anywhere inside
4 the source parameter set because the order of parameters within a FHiCL table is not
5 important. Run *art* again, using `hi.fcl`:

```
6 art -c hi.fcl >& output/hi.log
```

7 You should see output from the `HelloWorld` module for only the first three events.

8 To configure the file for *art* to process all the events, i.e., to run until *art* reaches the end
9 of the input files, either leave off the `maxEvents` parameter or give it a value of `-1`. 

10 If the maximum number of events is specified both on the command line and in the FHiCL
11 file, then the command line takes precedence. Compare the outputs of the following com-
12 mands:

```
13 art -c hi.fcl >& output/hi2.log
```


```
14 art -c hi.fcl -n 5 >& output/hi-n5.log
```

```
15 art -c hi.fcl -n -1 >& output/hi-neg1.log
```

16 9.8.5 Changing the Input Files

17 For historical reasons, there are multiple ways to specify the input event-data file (or the
18 list of input files) to an *art* job:

- 19 ○ within the FHiCL file's `source` parameter set
- 20 ○ on the *art* command line via the `-s` option (you may specify one input file only)
- 21 ○ on the *art* command line via the `-S` option (you may specify a text file that lists
22 multiple input files)
- 23 ○ on the *art* command line, after the last recognized option (you may specify one or
24 more input files)

25 If input file names are provided both in the FHiCL file and on the command line, the
26 command line takes precedence. 

input:files

1 Let's run a few examples.

2 We'll start with the `-s` command line option (second bullet). Run *art* without it (again), for
 3 comparison (or recall its output from Listing 9.1):

4 `art -c hello.fcl >& output/hello.log`

5 To see what you should expect given the following input file, check Table 9.1, then run:

6 `art -c hello.fcl -s inputFiles/input02.art >& output/hello-s.log`

7 Notice that the ten events in this output are from run 2 subRun 0, in contrast to the previous
 8 printout which showed events from run 1. Notice also that the command line specification
 9 overrode that in the FHiCL file. The `-s` (lower case) command line syntax will only permit
 10 you to specify a single filename.

11 This time, edit the source parameter set inside the `hi.fcl` file (first bullet); change it
 12 to:

```
13 source : {
14   module_type : RootInput
15   fileNames   : [ "inputFiles/input01.art",
16                   "inputFiles/input02.art" ]
17   maxEvents   : -1
18 }
```

19 (Notice that you also added `maxEvents : -1`.) The names of the two input files could
 20 have been written on a single line but this example shows that, in FHiCL, newlines are
 21 treated simply as white space.

22 Check Table 9.1 to see what you should expect, then rerun *art* as follows:

23 `art -c hi.fcl >& output/hi-2nd.log`

24 You will see 20 lines from the `HelloWorld` module; you will also see messages from
 25 *art* at the open and close operations on each input file.

26 Back to the `-s` command-line option, run:

27 `art -c hi.fcl -s inputFiles/input03.art >& output/run3.log`

28 This will read only `inputFiles/input03.art` and will ignore the two files specified
 29 in the `hi.fcl`. The output from the `HelloWorld` module will be the 15 events from the
 30 three subRuns of run 3.

1 There are several ways to specify multiple files at the command line. One choice is to use
2 the `-S` (upper case) `[--source-list]` command line option (third bullet) which takes as its
3 argument the name of a text file containing the filename(s) of the input event-data file(s).
4 An example of such as file has been provided, `inputs.txt`. Look at the contents of this
5 file:

```
6 cat inputs.txt  
7 inputFiles/input01.art  
8 inputFiles/input02.art  
9 inputFiles/input03.art
```

10 Now run *art* using `inputs.txt` to specify the input files:

```
11 art -c hi.fcl -S inputs.txt >& output/file010203.log
```

12 You should see the `HelloWorld` output from the 35 events in the three files; you should
13 also see the messages from *art* about the opening and closing of the three files.

14 Finally, you can list the input files at the end of the *art* command line (fourth bullet):

```
15 art -c hi.fcl inputFiles/input02.art inputFiles/input03.art >&\  
16 output/file0203.log
```

17 (Remember the Unix convention about a trailing backslash marking a command that con-
18 tinues on another line; see Chapter 2.) In this case you should see the `HelloWorld`
19 output from the 25 events in the two files.

20 In summary, there are three ways to specify input files from the command line; all of them
21 override any input files specified in the `FHiCL` file. Do not try to use two or more of these
22 methods on a single *art* command line; the *art* job will run without issuing any messages
23 but the output will likely be different than you expect.



24 9.8.6 Skipping Events

25 The `source` parameter set supports a syntax to start execution at a given event number
26 or to skip a given number of events at the start of the job. Look, for example, at the file
27 `skipEvents.fcl`, which differs from `hello.fcl` by the addition of two lines to the
28 source parameter set:

```

1  firstEvent : 5
2  maxEvents  : 3

```

3 *art* will process events 5, 6, and 7 of run 1, subRun 0. Try it:

```
4 art -c skipEvents.fcl >& output/skipevents1.log
```

5 An equivalent operation can be done from the command line in two different ways. Try
6 the following two commands and compare the output:

```
7 art -c hello.fcl -e 5 -n 3 >& output/skipevents2.log
```

```
8 art -c hello.fcl --nskip 4 -n 3 >& output/skipevents3.log
```

9 You can also specify the initial event to process relative to a given event ID (which, re-
10 call, contains the run, subRun and event number). Edit `hi.fcl` and edit the `source`
11 parameter set as follows:

```

12 source : {
13     module_type : RootInput
14     fileNames   : [ "inputFiles/input03.art" ]
15     firstRun    : 3
16     firstSubRun : 1
17     firstEvent  : 6
18 }

```

19 When you run this job, *art* will process events starting from run 3, subRun 2, event 1 —
20 because there are only five events in subRun 1.

```
21 art -c hi.fcl >& output/startatrun3.log
```

22 9.8.7 Identifying the User Code to Execute

23 Recall from Section 9.8.2 that the `physics` parameter set contains the physics content
24 for the *art* job. Within this parameter set, *art* must be able to determine which (user code)
25 modules to process. These must be referenced via `module labels(γ)`, which as you will see,
26 represent the pairing of a module name and a run-time configuration.

27 Look back at the listing on page 137, which contains the `physics` parameter set from
28 `hello.fcl`. The `analyzer` parameter set, nested inside the `physics` parameter set,
29 contains the definition:


```

1 hi : {
2   module_type : HelloWorld
3 }

```

The identifier `hi` is a module label (defined by the user, not by FHiCL or *art*) whose value must be a parameter set that *art* will use to configure a module. The parameter set for a module label must contain (at least) a FHiCL definition of the form:

```

7 module_type : best-module-name

```

Here `module_type` is an identifier reserved to *art* and `best-module-name` tells *art* the name of the module to load and execute. (Since it is within the analyzer parameter set, the module must be of type `EDAnalyzer`; i.e., the *base type* of `best-module-name` must be `EDAnalyzer`.)

Module labels are fully described in Section [52.5](#).

In this example *art* will look for a module named `HelloWorld`, which it will find as part of the `toyExperiment` product. Section [9.10](#) describes how *art* uses `best-module-name` to find the dynamic library that contains code for the `HelloWorld` module. A parameter set that is used to configure a module may contain additional lines; if present, the meaning of those lines is understood by the module itself; those lines have no meaning either to *art* or to FHiCL.

Now look at the FHiCL fragment below that starts with `analyzers`. We will use it to reinforce some of the ideas discussed in the previous paragraph.

art allows you to write a FHiCL file that uses a given module more than once. For example you may want to run an analysis twice, once with a loose mass cut on some intermediate state and once with a tight mass cut on the same intermediate state. In *art* you can do this by writing one module and making the cuts “run-time configurable.” This idea will be developed further in Chapter [15](#).

```

2b analyzers : {
2c
2d   loose: {
2e     module_type : MyAnalysis
2f     mass_cut : 20.
2g   }
2h
2i   tight: {
2j     module_type : MyAnalysis

```

```

10     mass_cut : 15.
11     }
12 }

```

4 When *art* processes this fragment it will look for a module named `MyAnalysis` (lines
5 4 and 9) and instantiate it twice, once using the parameter set labeled `loose` (line 3)
6 and once using the parameter set labeled `tight` (line 8). The two instances of the module
7 `MyAnalysis` are distinguished by their different module labels, `tight` and `loose`.



art requires that module labels be unique within a FHiCL file. Module labels may contain
9 only upper- and lower-case letters and the numerals 0 to 9.

10 In the FHiCL files in this exercise, all of the modules are analyzer modules. Since analyzers
11 do not make data products, these module labels are nothing more than identifiers inside
12 the FHiCL file. For producer modules, however, which *do* make data products, the module
13 label becomes part of the data product identifier and therefore has a real significance. All
14 module labels must conform to the same naming rules.

15 Within *art* there is no notion of reserved names or special names for module labels;
16 however your experiment will almost certainly have established some naming conven-
17 tions.

18 9.8.8 Paths and the *art* Workflow

19 In the `physics` parameter set in `hello.fcl` the two parameter definitions shown be-
20 low, taken together, specify the *workflow* of the *art* job. *Workflow* refers to the modules *art*
21 should run and the order in which to run them.²

```

21 physics {
22     ...
23     e1      : [ hi ]
24     end_paths : [ e1 ]
25 }

```

27 In this exercise there is only one module to run (the analyzer `HelloWorld` with the label
28 `hi` from Section ??), so the workflow is trivial: for each event, run the module with the

² The word *workflow* is used widely in the computing world. When a complete job comprises several discrete tasks, a workflow specifies the order in which the tasks should be performed. The full story of workflows is very rich and this section only covers the aspects needed to understand Workbook Exercise 1. **FIXME:** Add reference to the full chapter once it has been prepared

1 label `hi`. As you work through the Workbook you will encounter workflows that are more
2 complex and they will be described as you encounter them.

3 The FHiCL parameter `e1` is called a *path*. A path is simply a FHiCL sequence of module
4 labels. The name of a path can be any user-defined name that satisfies the following:

- 5 1. It must be defined as part of the `physics` parameter set, i.e., “at physics scope”.
- 6 2. It must be a valid FHiCL name.
- 7 3. It must be unique within the *art* job.
- 8 4. It must NOT be one of the following five names that are reserved to *art*: `analyzers`,
9 `filters`, `producers`, `end_paths` and `trigger_paths`.

10 An *art* job may contain many paths, each of which is a FHiCL sequence of module labels.
11 When many groups are working on a common project, this helps to maximize the independ-
12 ence of each work group. **FIXME:** *Needs a sentence about why different groups would*
13 *be using the same config file.*

14 Recall from Section [fcl-syntax](#) ?? the parameter `end_paths` is not itself a path. Rather it is a FHiCL
15 sequence of path names. It is the `end_paths` parameter that tells *art* the workflow it
16 should execute.

17 Note that any path listed in the `end_paths` parameter may only contain module labels
18 for analyzer and output modules. A similar mechanism is used to specify the workflow of
19 producer and filter modules; that mechanism will be discussed when you encounter it. If
20 you need a reminder about the types of modules, see Section [sec:intro-module-types](#) 3.6.3.



21 If the `end_paths` parameter is absent or defined as an empty FHiCL sequence,

```
22 end_paths : [ ]
```

23 both of which are allowable, *art* will understand that this job has no analyzer modules and
24 no output modules to execute.

25 As is standard in FHiCL, if the definition of `end_paths` appears more than once, the last
26 definition takes precedence.

27 The notion of *path* introduced in this section is the third thing in the *art* documentation
28 suite that is called a path. The other two, as you may recall from Section [sec:paths](#) 4.6, are the notion
29 of a path in a filesystem and the notion of an environment variable that is a colon-delimited

1 set of directory names. The use should be clear from the context; if it is not, please let the
 2 authors of the Workbook know; see Section [3.4](#).

3 The above description is intended to be sufficient for completing the Workbook exercises.
 4 If you want to learn more, now or later, Section [9.8.8.1](#) provides more detail.

5 **9.8.8.1 Paths and the *art* Workflow: Details**

6 This section is optional; it contains more details about the material just described in Sec-
 7 tion [9.8.8](#). It is not really a “dangerous bend” section for experts — just a side trip.

8 Exercise 1 is not rich enough to illustrate how to specify an *art* workflow, so let’s construct
 9 a richer example.

10 **FIXME:** *If the end paths is the same, doesn’t art execute the same stuff no matter who*
 11 *runs this configuration? If this is answered somewhere else, let’s reference it; I don’t find*
 12 *it.*

13 Suppose that there are two groups of people working on a large collaborative project,
 14 the project leaders are Anne and Rob. Each group has a workflow that requires running
 15 five or six module instances; some of the module instances may be in the workflow for
 16 both groups. Recall that an instance of a module refers to the name of a module plus
 17 its parameter set, and a module instance is specified by giving its module label. For this
 18 example let’s have eight module instances with the unimaginative names a through h. The
 19 workflow for this example might look something like:

```
20  anne      : [ a, b, c, d, e, h ]
21  rob       : [ a, b, f, c, g ]
22  end_paths : [ anne, rob ]
```

23 That is, Anne defines the modules that her group needs to run and Rob defines the modules
 24 that his group needs to run. Anne and Rob do not need to know anything about each other’s
 25 list. The parameter definitions `anne` and `rob` are called *paths*; each is a list of module
 26 labels. The rules for legal path names were given in Section [9.8.8](#).

27 The parameter named `end_paths` is not itself a path, rather it is a FHiCL sequence
 28 of paths. Moreover it has a special meaning to *art*. During *art*’s initialization phase, it
 29 needs to learn the workflow for the job. The first step is to find the parameter named

1 `end_paths`, defined within the `physics` parameter set. When *art* processes the defi-
 2 nition of `end_paths` it will form the set of all module labels found in the contributing
 3 paths, with any duplicates removed. For this example, the list might look something like:
 4 `[a, b, c, d, e, h, f, g]` . When *art* processes an event, this is the set of mod-
 5 ule instances that it will execute. The order in which the module instances are executed is
 6 discussed in Section 9.8.8.2.

7 The above machinery probably seems a little heavyweight for the example given. But con-
 8 sider a workflow like that needed to design the trigger for the CMS experiment, which
 9 requires about 200 paths and many hundreds of modules. Finding the set of unique mod-
 10 ules labels is not a task that is best done by hand! By introducing the idea of paths, the
 11 design allows each group to focus on its own work, unaffected by the other groups.

12 Actually, the above is only part of the story: the module labels given in the paths `anne` and
 13 `rob` may only be the labels of analyzer or output modules. There is a parallel mechanism
 14 to specify the workflow for producer and filter modules.

15 To illustrate this parallel mechanism let's continue the above example of two work groups
 16 led by Rob and Anne. In this case let there be filter modules with labels given by `f0`,
 17 `f1`, `f2` ... and producer modules with labels given by `p0`, `p1`, `p2` ... In this example, a
 18 workflow might look something like:

```

1b  t_anne      : [ p0, p1, p2, f0, p3, f1 ]
2b  t_rob      : [ p0, p1, f2, p2, f0, p4 ]
3b  trigger_paths : [ t_anne, t_rob ]
4b
5b  e_anne     : [ a, b, c, d, e ]
6b  e_rob      : [ a, b, f, c, g ]
7b  end_paths  : [ e_anne, e_rob ]

```

26 Here the parameters `t_anne`, `e_anne`, `t_rob`, and `e_rob` are all the names of paths.
 27 All must be legal FHiCL parameter names, be unique within an *art* job and not con-
 28 flict with identifiers reserved to *art* at physics scope. In this example the path names are
 29 prefixed with `t_` for paths that will be put into the `trigger_paths` parameter and with
 30 `e_` for paths that will be put into the `end_paths` parameter. This is just to make it easier
 31 for you to follow the example; the prefixes have no intrinsic meaning.

32 During *art*'s initialization phase it processes `trigger_paths` in the same way that
 33 it processes `end_paths`: it forms the set of all module labels found in the contribut-

1 ing paths, with duplicates removed. Again, the order of execution is discussed in Sec-
 2 tion 9.8.8.2.

3 Now, what happens if you define a path with a mix of modules from the two groups? It
 4 might look like this:

```

1  bad_path      : [ p0, p1, p2, f0, p3, f1, a, b ]
2  end_paths     : [ e_anne, e_rob, bad_path ]

```

7 In this case *art* (not FHiCL) will recognize that producer and filter modules are specified
 8 in a path that contributes to `end_paths`; *art* will then print a diagnostic message and
 9 stop. This will occur very early in *art*'s initialization phase so you will get reasonably
 10 prompt feedback. Similarly, if *art* discovers analyzer or output modules in any of the paths
 11 contributing to `trigger_paths`, it will print a diagnostic message and stop.

12 Furthermore, if you put a module label into either `end_paths` or `trigger_paths`,
 13 *art* will print a diagnostic message and stop. This is also true if you put a path name into
 14 the definition of another path *art*.

15 Now it's time to define two really badly chosen names:³ *trigger paths* and *end paths*, first
 16 *without underscores*. In the above fragment the paths prefixed with `t_` are called *trigger*
 17 *paths*, without an underscore; they are so named because they contain module labels for
 18 only producer and filter modules; therefore they are paths that satisfy the rules for inclusion
 19 in the definition of `trigger_paths` parameter. **FIXME:** *This doesn't say how the*
 20 *word trigger matches up with the words producer and filter. I think it's about what you*
 21 *need for triggering, then after everything's produced and filtered you 'end' by analyzing*
 22 *and producing output. Could you please put a sentence in about this?*

23 Similarly the paths prefixed with `e_` are called *end paths* because they satisfy the rules for
 24 inclusion in the definition of `end_paths` parameter.


25 This documentation will try to avoid avoid confusion between *trigger paths* and `trigger_paths`,
 26 and between *end paths* and `end_paths`.

³ It could have been worse. They could have been named on Opposite Day!

9.8.8.2 Order of Module Execution

1 `paths:order`
2 If the `trigger_paths` parameter contains a single *trigger path*, then *art* will execute
3 the modules in that *trigger path* in the order that they are specified.

4 When more than one *trigger path* is present in `trigger_paths`, *art* will choose one of
5 the *trigger paths* and execute its module instances in order. It will then choose a second
6 *trigger path*. If any module instances in this path were already executed in the first *trig-*
7 *ger path*, *art* will not execute them a second time; it will execute the remaining module
8 instances in the order specified by the second *trigger path*. And so on for any remaining
9 *trigger paths*.

10 The rules for order of execution of module instances named in an *end path* are different.
11 Since analyzer and output modules may neither add new information to the event nor
12 communicate with each other except via the event, the processing order is not important.
13 By definition, then, *art* may run analyzer and output modules in any order. In a simple *art*
14 job with a single path, *art* will, in fact, run the modules in the order of appearance in the 
15 path, but do not write code that depends on execution order because *art* is free to change
16 it.

9.8.9 Writing an Output File

17 `output:file`
18 The file `writeFile.fcl` gives an example of writing an output file. Open the file in an
19 editor and find the parts of the file that are discussed below.

20 Output files are written by output modules; one module can write one file. An *art* job may
21 run zero or more output modules.

22 If you wish to add an output module to an *art* job there three steps:

- 23 1. Create a parameter set named `outputs` at the outermost scope of the FHiCL file.
24 The name `outputs` is prescribed by *art*.
- 25 2. Inside the `outputs` parameter set, add a parameter set to configure an output mod-
26 ule. In `writeFile.fcl` this parameter set has the module label `output1`.
- 27 3. Add the module label of the output module to an *end path* (not to the `end_paths`
28 parameter but to one of the paths that is included in `end_paths`). In `writeFile.fcl`

1 the module label `output1` is added to the *end path* `e1`.

2 If you wish to add more output modules, repeat steps 2 and 3 for each additional output
3 file.

4 The parameter set `output1` tells *art* to make a module whose type is `RootOutput`. The
5 class `RootOutput` is a standard module that is part of *art* and that writes events from
6 memory to a disk file in an *art*-defined, `ROOT`-based format. The `fileName` parame-
7 ter specifies the name of the output file; this parameter is processed by the `RootOutput`
8 module. Files written by the module `RootOutput` can be read by the module `RootInput`.
9 The identifier `output1` is just another module label that obeys the rules discussed in Sec-
10 tion 9.8.7.

11 In the example of `writeFile.fcl` the output module takes its default behaviour: it will
12 write all of the information about each input event to the output file. `RootOutput` can
13 be configured to:

- 14 1. write only selected events
- 15 2. for each event write only a subset of the available data products.

16 *How to do this will be described in section that will be written later.*

17 **FIXME:** *Once that section is written, add a reference to it.*

18 Before running the exercise, look at the source parameter set of `writeFile.fcl`; note
19 that it is configured to read only events 4, 5, 6, and 7.

20 To run `writeFile.fcl` and check that it worked correctly:

```
21 art -c writeFile.fcl
```

```
22 ls -s output/writeFile.art
```

```
23 art -c hello.fcl -s output/writeFile.art
```

24 The first command will write the output file; the second will check that the output file was
25 created and will tell you its size; the last one will read back the output file and print the
26 event IDs for all of the events in the file. You should see the `HelloWorld` printout for
27 events 4, 5, 6 and 7.

9.9 Understanding the Process for Exercise 1

1

ding:steps

2 Section [9.6.1](#) contained a list of steps needed to run this exercise; this section will describe
3 each of those steps in detail. When you understand what is done in these steps, you will
4 understand *art*'s run-time environment. As a reminder, the steps are listed again here. The
5 commands that span two lines can be typed on a single line.

1. Log in to the computer you chose in Section [8.3](#).
2. Follow the site-specific setup procedure; see Chapter [5](#).
3.

```
mkdir -p $ART_WORKBOOK_WORKING_BASE/username/\
workbook-tutorial/pre-built
```

In the above and elsewhere as indicated, substitute your kerberos principal **FIXME:** *ref* for the string *username*.
4.

```
cd $ART_WORKBOOK_WORKING_BASE/username/\
workbook-tutorial/pre-built
```
5.

```
setup toyExperiment v0_00_29 -q$ART_WORKBOOK_QUAL:prof
```
6.

```
cp $TOYEXPERIMENT_DIR>HelloWorldScripts/* .
```
7.

```
source makeLinks.sh
```
8. Run *art*:

```
art -c hello.fcl >& output/hello.log
```

6

7 Steps 1 and 4 should be self explanatory and will not be discussed further.

8 When reading this section, you do not need to run any of the commands given here; this is
9 commentary on commands that you have already run.



9.9.1 Follow the Site-Specific Setup Procedure (Details)

The site-specific setup procedure, described in Chapter 5, ensures that the UPS system is properly initialized and that the UPS database (containing all of the UPS products needed to run the Workbook exercises) is present in the PRODUCTS environment variable.

This procedure also defines two environment variables that are defined by your experiment to allow you to run the Workbook exercises on their computer(s):

ART_WORKBOOK_WORKING_BASE the top-level directory in which users create their working directory for the Workbook exercises

ART_WORKBOOK_OUTPUT_BASE the top-level directory in which users create their output directory for the Workbook exercises; this is used by the script `makeLinks.sh`

If these environment variables are not defined, ask a system admin on your experiment.

9.9.2 Make a Working Directory (Details)

On the Fermilab computers the home disk areas are quite small so most experiments ask that their collaborators work in some other disk space. This is common to sites in general, so we recommend working in a separate space as a best practice. The Workbook is designed to require it.

This step, shown on two lines as:

```
mkdir -p $ART_WORKBOOK_WORKING_BASE/username/workbook-tutorial/\
pre-built
```

creates a new directory to use as your working directory. It is defined relative to an environment variable described in Section 9.9.1. It only needs to be done the first time that you log in to work on Workbook exercises.

If you follow the rest of the naming scheme, you will guarantee that you have no conflicts with other parts of the Workbook.

As discussed in Section 9.6.1.2, you may of course choose your own working directory on any disk that has adequate disk space.

9.9.3 Setup the toyExperiment UPS Product (Details)

This step is the main event in the eight-step process.

```
setup toyExperiment v0_00_14 -q$ART_WORKBOOK_QUAL:prof
```

This command tells UPS to find a product named `toyExperiment`, with the specified version and qualifiers, and to *setup* that product, as described in Section 7.3.

The required qualifiers may change from one experiment to another and even from one site to another within the same experiment. To deal with this, the site specific setup procedure defines the environment variable `ART_WORKBOOK_QUAL`, whose value is the qualifier string that is correct for that site.

The complete UPS qualifier for `toyExperiment` has two components, separated by a colon: the string defined by `ART_WORKBOOK_QUAL` plus a qualifier describing the compiler optimization level with which the product was built, in this case “prof”; see Section 3.6.7 for information about the optimization levels.

Each version of the `toyExperiment` product knows that it requires a particular version and qualifier of the *art* product. In turn, *art* knows that it depends on particular versions of `ROOT`, `CLHEP`, `boost` and so on. When this recursive setup has completed, over 20 products will have been setup. All of these products define environment variables and about two-thirds of them add new elements to the environment variables `PATH` and `LD_LIBRARY_PATH`.

If you are interested, you can inspect your environment before and after doing this setup. To do this, log out and log in again. Before doing the setup, run the following commands:

```
printenv > env.before
```

```
printenv PATH | tr : \\n > path.before
```

```
printenv LD_LIBRARY_PATH | tr : \\n > ldpath.before
```

Then setup `toyExperiment` and capture the environment afterwards (`env.after`). Compare the before and after files: the after files will have many, many additions to the environment. (The fragment `| tr : \\n` tells the bash shell to take the output of `printenv` and

1 replace every occurrence of the colon character with the newline character; this makes the
2 output much easier to read.)

3 **9.9.4 Copy Files to your Current Working Directory (Details)**

4 The step:

```
5 cp $TOYEXPERIMENT_DIR/HelloWorldScripts/* .
```

6 only needs to be done only the first time that you log in to work on the Workbook.

7 In this step you copied the files that you will use for the exercises into your current working
8 directory. You should see these files:

```
9 hello.fcl makeLinks.sh skipEvents.fcl writeFile.fcl
```

10 **9.9.5 Source makeLinks.sh (Details)**

11 This step:

```
12 source makeLinks.sh
```

13 only needs to be done only the first time that you log in to work on the Workbook. It
14 created some symbolic links that *art* will use.

15 The FHiCL files used in the Workbook exercises look for their input files in the subdi-
16 rectory `inputFiles`. This script made a symbolic link, named `inputFiles`, that points
17 to:

```
18 $TOYEXPERIMENT_DIR/inputFiles
```

19 in which the necessary input files are found.

20 This script also ensures that there is an output directory that you can write into when you
21 run the exercises and adds a symbolic link from the current working directory to this output
22 directory. The output directory is made under the directory `$ART_WORKBOOK_OUTPUT_BASE`;
23 this environment variable was set by the site-specific setup procedure and it points to disk
24 space that will have enough room to hold the output of the exercises.

1 9.9.6 Run *art* (Details)

2 Issuing the command:

```
3 art -c hello.fcl
```

4 runs the *art* main program, which is found in \$ART_FQ_DIR/bin. This directory was
5 added to your PATH when you setup toyExperiment. You can inspect your PATH to see
6 that this directory is indeed there.

7 9.10 How does *art* find Modules?

8 When you ran `hello.fcl`, how did *art* find the module `HelloWorld`?

9 It looked at the environment variable `LD_LIBRARY_PATH`, which is a colon-delimited
10 set of directory names defined when you setup the `toyExperiments` product. We saw the
11 value of `LD_LIBRARY_PATH` in Section 9.9.3; to see it again, type the following:

```
12 printenv LD_LIBRARY_PATH | tr : \\n
```

13 The output should look similar to that shown in Listing 9.3.

Listing 9.3: Example of the value of `LD_LIBRARY_PATH`

```
14 /ds50/app/products/tbb/v4_1_2/Linux64bit+2.6-2.12-e2-prof/lib
15 /ds50/app/products/sqlite/v3_07_16_00/Linux64bit+2.6-2.12-prof/lib
16 /ds50/app/products/libsigcpp/v2_2_10/Linux64bit+2.6-2.12-e2-prof/lib
17 /ds50/app/products/cppunit/v1_12_1/Linux64bit+2.6-2.12-e2-prof/lib
18 /ds50/app/products/clhep/v2_1_3_1/Linux64bit+2.6-2.12-e2-prof/lib
19 /ds50/app/products/python/v2_7_3/Linux64bit+2.6-2.12-gcc47/lib
20 /ds50/app/products/libxml2/v2_8_0/Linux64bit+2.6-2.12-gcc47-prof/lib
21 /ds50/app/products/fftw/v3_3_2/Linux64bit+2.6-2.12-gcc47-prof/lib
22 /ds50/app/products/root/v5_34_05/Linux64bit+2.6-2.12-e2-prof/lib
23 /ds50/app/products/boost/v1_53_0/Linux64bit+2.6-2.12-e2-prof/lib
24 /ds50/app/products/cpp0x/v1_03_15/slf6.x86_64.e2.prof/lib
25 /ds50/app/products/cetlib/v1_03_15/slf6.x86_64.e2.prof/lib2
26 /ds50/app/products/fhiclcpp/v2_17_02/slf6.x86_64.e2.prof/lib
27 /ds50/app/products/messagefacility/v1_10_16/slf6.x86_64.e2.prof/lib
28 /ds50/app/products/art/v1_06_00/slf6.x86_64.e2.prof/lib
29 /ds50/app/products/toyExperiment/v0_00_14/slf6.x86_64.e2.prof/lib
30 /grid/fermiapp/products/common/prd/git/v1_8_0_1/Linux64bit-2/lib
```

1 Of course the leading element of each directory name, `/ds50/app` will be replaced by
2 whatever is correct for your experiment. The last element in `LD_LIBRARY_PATH` is not
3 relevant for running *art* and it may or may not be present on your machine, depending on
4 details of what is done inside your site-specific setup procedure.

5 If you compare the names of the directories listed in `LD_LIBRARY_PATH` to the names
6 of the directories listed in the `PRODUCTS` environment variable, you will see that all of
7 these directories are part of the UPS products system. Moreover, for each product, the
8 version, flavor and qualifiers are embedded in the directory name. In particular, both *art*
9 and *toyExperiment* are found in the list.

10 If you `ls` the directories in `LD_LIBRARY_PATH` you will find that each directory contains
11 many dynamic object library (`.so` files).

12 When *art* looks for a module named `HelloWorld`, it looks through the directories de-
13 fined in

14 `LD_LIBRARY_PATH` and looks for a file whose name matches the pattern,

```
15 lib*HelloWorld_module.so
```

16 where the asterisk matches (zero or) any combination of characters. *art* finds that, in all
17 of the directories, there is exactly one file that matches the pattern, and it is found in the
18 directory (shown here on two lines):

```
19 /ds50/app/products/toyExperiment/v0_00_14/  
20     slf6.x86_64.e2.prof/lib/
```

21 The name of the file is:

```
22 libtoyExperiment_Analyzers_HelloWorld_module.so
```

23 If *art* had found no files that matched the pattern, it would have printed a diagnostic mes-
24 sage and stopped execution. If *art* had found more than one file that matched the pattern,
25 it would have printed a different diagnostic message and stopped execution. If this second
26 error occurs it is possible to tell *art* which of the matches to choose. *How to do this will*
27 *be covered in a future chapter.*

28 **FIXME:** *Add reference to that chapter when it is written.*

29 **FIXME:** *Add a para to address Anne's question about why not accept the first match.*

1 **FIXME:** *Note to Anne: I agree that the commented out text does not belong here. Note to*
2 *Rob: Could we put it at the end of sec ?? maybe?*

3 9.11 How does *art* find FHiCL Files?

4 This section will describe where *art* looks for FHiCL files. There are two cases: looking
5 for the file specified by the command line argument `-c` and looking for files that have been
6 included by a `#include` directive within a FHiCL file.

7 9.11.1 The `-c` command line argument

8 When you issued the command

```
9 art -c hello.fcl
```

10 *art* looked for a file named `hello.fcl` in the current working directory and found it.
11 You may specify any absolute or relative path as the argument of the `-c` option. If *art* had
12 not found `hello.fcl` in this directory it would have looked for it relative to the path
13 defined by the environment variable `FHICL_FILE_PATH`. This is just another path-type
14 environment variable, like `PATH` or `LD_LIBRARY_PATH`. You can inspect the value of
15 `FHICL_FILE_PATH` by:

```
16 printenv FHICL_FILE_PATH
```

```
18 .:<some-directory-structure>products//toyExperiment/v0_00_29
```

19 In this case, the output will show the translated value of the environment variable `TOY-`
20 `EXPERIMENT_DIR`. The presence of the current working directory (`.`) in the path is
21 redundant when processing the command line argument but it is significant in the case
22 discussed in the next section.

23 Some experiments have chosen to configure their version of the *art* main program so that it
24 will not look for the command line argument FHiCL file in `FHICL_FILE_PATH`. It is also
25 possible to configure *art* so that only relative paths, not absolute paths, are legal values of
26 the `-c` argument. This last option can be used to help ensure that only version-controlled



1 files are used when running production jobs. Experiments may enable or disable either of
 2 these options when their main program is built.

3 **9.11.2 #include Files**

4 Section 9.8 discussed the listing on page 138, which contains the fragments of `hello.fcl`
 5 that are related to configuring the message service. The first line in that listing is an
 6 include directive. *art* will look for the file named by the include directive relative to
 7 `FHICL_FILE_PATH` and it will find it in:

8 `$TOYEXPERIMENT_DIR/fcl/minimalMessageService.fcl`

9 This is part of the toyExperiment UPS product.

10 The version of *art* used in the Workbook does not consider the argument of the include
 11 directive as an absolute path or as a path relative to the current working directory; it only
 12 looks for files relative to `FHICL_FILE_PATH`. This is in contrast to the choice made when
 processing the `-c` command line option.



14 When building *art*, one may configure *art* to first consider the argument of the include
 15 directive as a path and to consider `FHICL_FILE_PATH` only if that fails.

16 **9.12 Review**

17 **FIXME:** Add a section called *Review* that looks at trigger paths, end paths, etc and works
 18 backwards. list from 'what you will learn' repeated here for reference

- 19 ○ how to use the site-specific setup procedure, which you must do once at the start of
 20 each login session.
- 21 ○ a little bit about the *art* run-time environment (Section 9.4) [|ssec:prebuilt:runtime:end](#)
- 22 ○ how to set up the toyExperiment UPS product (Section 9.6.1) [|ssec:first-login-run](#)
- 23 ○ how to run an *art* job (Section 9.6.1) [|ssec:first-login-run](#)
- 24 ○ how to control the number of events to process (Section 9.8.4) [|ssec:prebuilt:max:events](#)
- 25 ○ how to select different input files (Section 9.8.5) [|ssec:prebuilt:changing:input:files](#)

- 1 ○ how to start at a run, subRun or event that is not the first one in the file (Section 9.8.6)
- 2 ○ how to concatenate input files (Section 9.8.5)
- 3 ○ how to write an output file (Section 9.8.9)
- 4 ○ some basics about the grammar and structure of a FHiCL file (Section 9.8)
- 5 ○ how *art* finds modules and configuration (FHiCL) files. (Sections 9.10 and 9.11)

6 9.13 Test your Understanding

7 When you got set up to run Exercise 1 in Section 9.6.1, you may have noticed the four
8 FHiCL files with *bug* in their names: `bug01.fcl`, `bug02.fcl`, `bug03.fcl` and `bug04.fcl`.
9 These are files into which errors have expressly been inserted. You will now find and fix
10 these errors. Answers are provided at the end.

11 9.13.1 Tests

12 For each of these exercises, your job is to figure out what's wrong, fix it and rerun the
13 command. With practice you will learn to decipher the error messages that *art* throws your
14 way and solve the problems. Completion codes are listed in Appendix ??.

15 Note that in the output of each command the `MSG-i` messages are just informational – they
16 may or may not be relevant to the actual error. On the other hand, the `MSG-s` messages
17 – that appear only when the FHiCL files have errors – are severe and cause *art* to shut
18 down.

19 The `MSG-s` message is what actually caused *art* to shutdown.

- 20 1. `art -c bug01.fcl`

21
22
23 This will make an error message that includes the following:
24
25

```
1   Input file not found: input01.art.  
2   ...  
3   Art has completed and will exit with status 24.  
4   ...  
5   %MSG-s  
6   ...  
7
```

```
8  2. art -c bug02.fcl  
9  
10
```

11 This will make an error message that includes the following:

```
12  
13  
14 The following module label is not assigned to any path:  
15 'hi'  
16 ...  
17 Path configuration: The following were encountered while  
18 processing path configurations:  
19 ERROR: Entry hello in path e1 refers to a module label hello  
20 which is not configured.  
21 ...  
22 Art has completed and will exit with status 9.  
23 ...  
24 %MSG-s  
25 ...  
26
```

```
27 3. art -c bug03.fcl  
28  
29
```

30 This will make an error message that includes the following:

```
31  
32
```

```
1      ERROR: Configuration of module with label hi encountered
2      the following error:
3      --- Configuration BEGIN
4      UnknownModule
5      --- Configuration BEGIN
6      Library specificaton "Helloworld": does not correspond to
7      any library in LD_LIBRARY_PATH of type "module"
8      --- Configuration END
9      Module Helloworld with version v1_09_03 was not registered.
10     Perhaps your module type is misspelled or is not a framework
11     plugin.
12     ...
13     Art has completed and will exit with status 9.
```

4. `art -c bug04.fcl`

This will make an error message that includes the following:

```
21     ERROR: Entry e1 in path end_path refers to a module label
22     e1 which is not configured.
23     ...
24     Art has completed and will exit with status 9.
```

The answers are intentionally placed on a new page (remember to try before you read further!).

9.13.2 Answers

1

prebuilt:bug:ans

2

1. The critical part of the error message is:

3

4

5

```
Input file not found: input01.art
```

6

7

8

This says to look for a file with that name in the current working directory. It's not there because all of the input files are in the subdirectory `inputFiles`. The fix is to give a correct file name in the configuration of the source:

9

10

11

12

13

```
fileNames : [ "inputFiles/input01.art" ]
```

14

2. The information message tells you that the module label `hi` is never used anywhere; this level of message is just informational, which is less than a warning.

15

16

17

18

The second message tells you that something wrong in the path named `e1`. A path is supposed to be a sequence of module labels and there is no module labelled `hello`; it is labelled `hi`. The fix is to change the entry in `e1` to read `hi` or to change the name of the module label to `hello`.

19

20

21

22

3. The key parts of this error message are:

23

24

25

```
Library specificaton "Helloworld": does not correspond to  
any library in LD_LIBRARY_PATH of type "module"
```

26

27

28

and

29

30

```
Perhaps your module type is misspelled or is not a framework
```

1 plugin.
2
3

4 The answer is that in the definition of the module label `hi` the value of the `module_type`
5 parameter is misspelled. It has a lower case `w` when it should have upper case.

6 4. This fourth one is tricky. The `MSG-s` message is the same in `bug02.fc1`! The clue
7 is that it thinks `end_path` should contain module labels not a sequence of paths.
8 The answer is that `end_path` is misspelled. It is missing the final `s` and should be
9 `end_paths`.

DRAFT

10 Exercise 2: Building and Running Your First Module


10.1 Introduction

In this exercise you will build and run a simple *art* module. Section 3.6.7 introduced the idea of a build system, a software package that compiles and links your source code to turn it into machine code that the computer can execute. In this chapter you will be introduced to the *art* development environment, which adds the following to the run-time environment (discussed in Section 9.4):

1. a build system
2. a source code repository
3. a working copy of the Workbook source code
4. a directory containing dynamic libraries created by the build system

In this and all subsequent Workbook exercises, you will use the build system used by the *art* development team, **cetbuildtools**. Note that individual experiments may choose their build system, **cetbuildtools** or a different one. For the purposes of completing the *art* Workbook, we will provide all the information about it that you need. **FIXME: Is this true? Later in a fixme you say we may need to write some...** The **cetbuildtools** system will require you to open two shell windows your local machine and, in each one, to log into the remote machine ¹. The windows will be referred to as the *source window* and the *build window*:

¹**cetbuildtools** requires what are called *out-of-source builds*; this means that the source code and the working space for the build system must be in separate directories.

- 1 ○ In the *source window* you will check out and edit source code.
- 2 ○ In the *build window* you will build and run code.
- 3 Exercise 2 and all subsequent Workbook exercises will use the setup instructions found in
- 4 Sections 10.4 and 11. 

5 10.2 Prerequisites

6 Before running this exercise, you need to be familiar with the material in Part I (Introduc-
7 tion) of this documentation set and Chapter 9 from Part II (Workbook). Concepts that this
8 chapter refers to include:

- 9 ○ namespace
- 10 ○ `#include` directives
- 11 ○ header file
- 12 ○ class
- 13 ○ constructor
- 14 ○ destructor
- 15 ○ the C preprocessor
- 16 ○ member function (aka method)
- 17 ○ const vs non-const member function
- 18 ○ argument list of a function
- 19 ○ signature of a function
- 20 ○ *declaration vs definition* of a class
- 21 ○ arguments passed by reference
- 22 ○ arguments passed by const reference
- 23 ○ notion of *type*: e.g., a class, a struct, a free function or a typedef

1 In this chapter you will also encounter the C++ idea of *inheritance*. Understanding in-
2 heritance is not a prerequisite; it will be described as you encounter it in the Workbook
3 exercises. Inheritance includes such ideas as,

- 4 ○ base class
- 5 ○ derived class
- 6 ○ virtual function
- 7 ○ pure virtual function
- 8 ○ concrete class

9 **10.3 What You Will Learn**

10 In this exercise you will learn:

- 11 ○ how to establish the *art* development environment
- 12 ○ how to checkout the Workbook exercises from the `git` source code management
13 system
- 14 ○ how to use the **cetbuildtools** build system to build the code for the Workbook exer-
15 cises
- 16 ○ how include files are found
- 17 ○ what a *link list* is
- 18 ○ where the build system finds the link list
- 19 ○ what the `art::Event` is and how to access it
- 20 ○ what the `art::EventID` is and how to access it
- 21 ○ what makes a class an *art module*
- 22 ○ where the build system puts the `.so` files that it makes

10.4 Initial Setup to Run Exercises

10.4.1 “Source Window” Setup

Up through step 4 of the procedure in this section, the results should look similar to those of Exercise 1. Note that the directory name chosen here in the `mkdir` step is different than that chosen in the first exercise; this is to avoid file name collisions.

If you want to use a self-managed working directory, in steps 3 and 4, make a directory of your choosing and `cd` to it rather than to the directory shown.



In your source window do the following:

1. Log in to the computer you chose in Section 8.3.
2. Follow the site-specific setup procedure; see Table 5.1.
3. Make a new working directory. Remember that you can type this command, and all subsequent commands in the Workbook that are shown on two lines for formatting reasons, on a single line.

```
mkdir -p $ART_WORKBOOK_WORKING_BASE/username/\
workbook
```

4. `cd` to the new directory.

```
cd $ART_WORKBOOK_WORKING_BASE/username/workbook
```

5. Set up the source code management system `git` and use it to pull down the workbook code to the directory `art-workbook`, which will be referred to as your *source* directory. The output for each step is explained in Section 10.4.2.1:

- (a) `setup git`

(b) `git clone http://cdcvs.fnal.gov/projects/art-workbook`

(c) `cd art-workbook`

(d) `git checkout -b work origin/August2015`

6. Source the script that sets up the environment properly:

```
source ups/setup_deps -p $ART_WORKBOOK_QUAL
```

Ignore the last two lines of the output from this command; they tell you to use a different script than `setup_deps` and are not relevant to the source window.

1

2 The git commands are discussed in Section 10.4.2.1. The final step sources a script that
3 defines a lot of environment variables — the same set that will be defined in the build
4 window.

5 10.4.2 Examine Source Window Setup

6 10.4.2.1 About git and What it Did

7 Git is a source code management system² that is used to hold the source code for the
8 Workbook exercises. A source code management system is a tool that looks after the book-
9 keeping of the development of a code base; among many other things it keeps a complete
10 history of all changes and allows one to get a copy of the source code as it existed at any
11 time in the past. Because of git's many advanced features, many HEP experiments are
12 moving to git. git is fully described in the git manual³.

13 Some experiments set up git in their site-specific setup procedure; others do not. In running

²Other source code management systems with which you may be familiar are CVS and SVN.

³Several references for git can be found online; the “official” documentation is found at <http://git-scm.com/documentation>.

- 1 setup git, you have ensured that a working copy of git is in your PATH⁴.
- 2 The git clone and git checkout commands produce a working copy of the Workbook source
3 files in your source directory. Figure 10.1 shows a map of the source directory structure
4 created by the git commands. It does not show all the contents in each subdirectory. Note
5 that the `.git` (hidden) directory under the source directory is colored differently; this is
6 done to distinguish it from the rest of the contents of the source directory structure:
- 7 ○ When you ran git clone in Section 10.4.1, it copied the entire contents of the remote
8 repository into the `.git` directory. It then created the rest of the directories and files
9 under the source directory (what we call your “working area”). These files are the
10 most recent versions in the repository.
 - 11 ○ When you ran git checkout, it updated the files and directories in your working area
12 to be the versions specified by the head of the branch August2015 . This is important
13 because the instructions in this document are known to be correct for the head of the
14 branch August2015 ; there is no guarantee that they are correct for the most recent
15 versions.
 - 16 ○ In the git checkout command the name work is an arbitrary name that you get to
17 choose. It is the name of a new git branch and any changes that you make to this
18 code will become part of this branch, not part of the branch that you checked out.
19 You can read about git branches in any git documentation.

20 git clone should produce the following output:

```
21 Cloning into 'art-workbook'...
```

22 Executing the git checkout command should produce the following output:

```
23 Switched to a new branch 'work'
```

24 If you wish to learn about git branches, for the time being, you will need to consult a
25 git manual. **FIXME: FIND OUR TEXT IN U.G. WHEN IT'S THERE**

26 If you do not see the expected output, contact the *art* team as described in Section 3.4.

⁴No version needs to be supplied because the git UPS product has a current version declared; see Section 7.4.

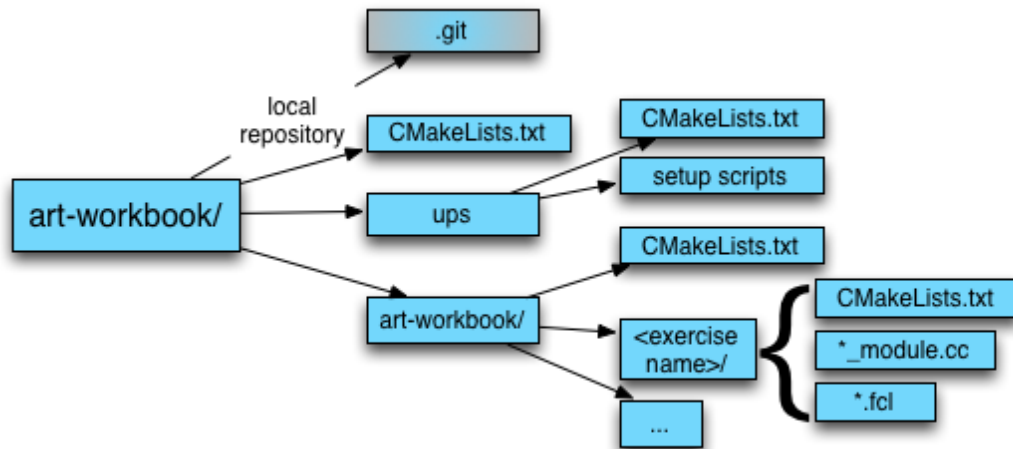


Figure 10.1: Representation of the reader's source directory structure (an `admin` directory is not shown)

10.4.2.2 Contents of the Source Directory

Figure 10.1 shows roughly what your source directory contains at the end of the setup procedure. You can see the correspondence between it and the output of the `ls -a` command:

```
cd $ART_WORKBOOK_WORKING_BASE/username/workbook/art-workbook
```

```
ls -a
```

```
. .. admin art-workbook CMakeLists.txt .git ups}
```

Notice that it contains a subdirectory of the same name as its parent, `art-workbook`.

- The `admin` directory (not shown in Figure 10.1) contains some scripts used by **cetbuildtools** to customize the configuration of the development environment.
- The `art-workbook` directory contains the main body of the source code for the Workbook exercises.
- The file `CMakeLists.txt` is the file that the build system reads to learn what steps it should do.

- 1 o The `ups` directory contains information about what UPS products this product de-
2 pends on; it contains additional information used to configure the development en-
3 vironment.

4 Look inside the `art-workbook` (sub)directory (via `ls`) and see that it contains several
5 files and subdirectories. The file `CMakeLists.txt` contains more instructions for the
6 build system. Actually, every directory contains a `CMakeLists.txt`; each contains ad-
7 ditional instructions for the build system. The subdirectory `FirstModule` contains the
8 files that will be used in this exercise; the remaining subdirectories contain files that will
9 be used in subsequent Workbook exercises.

10 If you look inside the `FirstModule` directory, you will see

11 `ls FirstModule`

```
12 CMakeLists.txt      FirstAnswer01_module.cc  First_module.cc
13 firstAnswer01.fcl  first.fcl
```

14 The file `CMakeLists.txt` in here contains yet more instructions for the build system
15 and will be discussed later. The file `First_module.cc` is the first module that you will
16 look at and `first.fcl` is the FHiCL file that runs it. This exercise will suggest that you
17 try to write some code on your own; the answer is provided in
18 `FirstAnswer01_module.cc` and the file `firstAnswer01.fcl` runs it. These
19 files will be discussed at length throughout the exercises.

20 10.4.3 “Build Window” Setup

21 Again, advanced users wanting to manage their own working directory may skip to Sec-
22 tion 10.4.3.2.



23 10.4.3.1 Standard Procedure

24 Now go to your build window and do the following:

1. Log in to the computer you chose in Section 8.3. [|sec:selectMachine](#)
2. Follow the site-specific setup procedure; see Chapter 5. [|ch:site-setup](#)

3. `cd $ART_WORKBOOK_WORKING_BASE/username/\`
`workbook`

4. `mkdir build-prof`

The `build-prof` directory will be your *build* directory.

5. `cd build-prof`

6. `source ../art-workbook/ups/setup_for_development \`
`-p $ART_WORKBOOK_QUAL`

The space before the backslash is required here; there must be a space before the `-p`. The output from this command will tell you to take some additional steps; *do not do those steps*.

7. `buildtool -j 4`

This step may take a few minutes.

1

2 Skip Section 10.4.3.2 and move on to Section 10.4.4.



10.4.3.2 Using Self-managed Working Directory

4 The steps in this procedure that are the same as for the “standard” procedure are explained
 5 in Section 10.4.4.

6 Now go to your build window and do the following:

1. Log in to the computer you chose in Section 8.3.

2. Follow the site-specific setup procedure; see Chapter 5.

7

3. Make a directory to hold the code that you will build and `cd` to it; this will be your *build* directory in your *build* window.
4. Make another directory, *outside of the hierarchy rooted at your build directory*, to hold output files created by the workbook exercises. (Don't `cd` to it.)
5. `In -s directory-for-output-files output`
6. `source your-SOURCE-directory/ups/setup_for_development \
-p $ART_WORKBOOK_QUAL`

The space before the backslash is required here; there must be a space before the `-p`. The output from this command (Listing 10.1) will tell you to take some additional steps; *do not do those steps*.

7. `buildtool -j 4`

1

2 10.4.4 Examine Build Window Setup

steps:build

3 Logging in and sourcing the site-specific setup script should be clear by now. Notice that
 4 next you are told to `cd` to the same *workbook* directory as in Step 4 of the instructions for
 5 the source window. From there, you make a directory in which you will run builds (your
 6 *build* directory), and `cd` to it. (The name `build-prof` can be any legal directory name
 7 but it is suggested here because this example performs a profile build; this is explained in
 8 Section 3.6.7). Figure 10.2 shows roughly what the build directory contains.

9 Step 6 sources a script called `setup_for_development` found in the `ups` subdirec-
 10 tory of the source directory. This script, run exactly as indicated, defines `build-prof` to
 11 be your build directory. This command selects a profile build (via the option `-p`); it also re-
 12 quests that the UPS qualifiers defined in the environment variable `ART_WORKBOOK_QUAL`
 13 be used when requesting the UPS products on which it depends; this environment variable
 14 was discussed in Section 9.9.3. The expected output is shown in Listing 10.1.

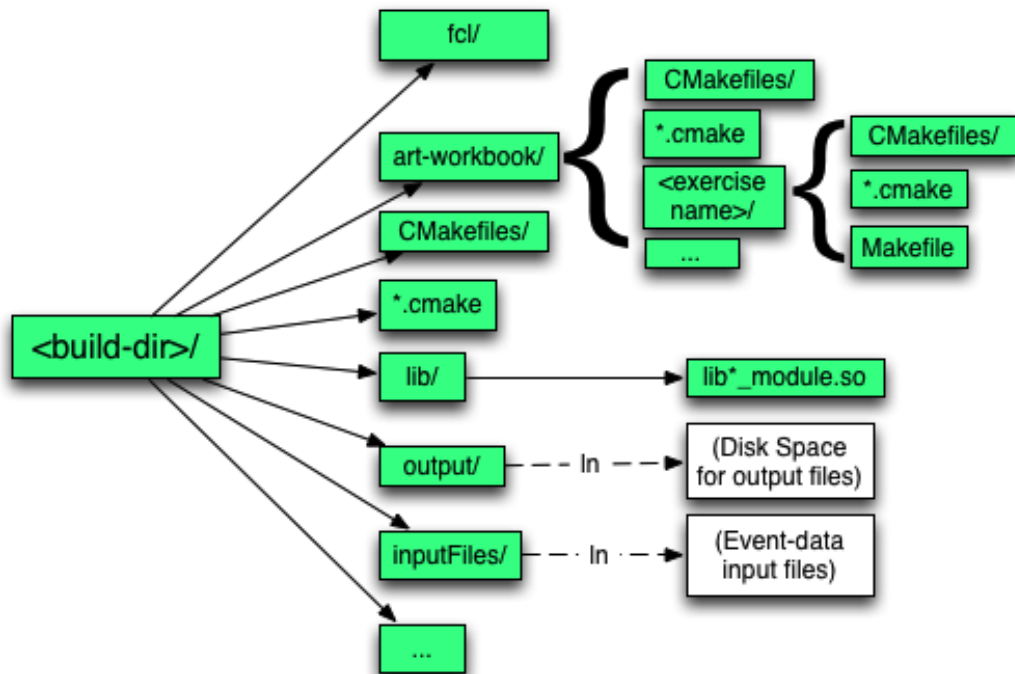


Figure 10.2: Representation of the reader's build directory structure (the `fcl/` directory is a symlink to `art-workbook/art-workbook/` in the source area)

- 1 Check that there are no error messages in the indicated block. The listing concludes with
 2 a request for you to run a `cmake` command; *do not run* `cmake` (this line is an artifact of
 3 layering **cetbuildtools** on top of `cmake`).



Listing 10.1: Example of output created by `setup_for_development`

```

4 The working build directory is
5     /ds50/app/user/kutschke/workbook/build-prof
6 The source code directory is
7     /ds50/app/user/kutschke/workbook/art-workbook
8 ----- check this block for errors -----
9 -----
10 /ds50/app/user/kutschke/workbook/build-prof/lib
11     has been added to LD_LIBRARY_PATH
12 /ds50/app/user/kutschke/workbook/build-prof/bin
13     has been added to PATH
14
15 CETPKG_SOURCE=/ds50/app/user/kutschke/workbook/art-workbook
16 CETPKG_BUILD=/ds50/app/user/kutschke/workbook/build-prof
17 CETPKG_NAME=art_workbook
18 CETPKG_VERSION=v0_00_15
19 CETPKG_QUAL=e2:prof
20 CETPKG_TYPE=Prof
21
22 Please use this cmake command:
23 cmake -DCMAKE_INSTALL_PREFIX=/install/path
24     -DCMAKE_BUILD_TYPE=$CETPKG_TYPE $CETPKG_SOURCE

```

- 25 This script sets up all of the UPS products on which the Workbook depends; this is anal-
 26 ogous to the actions taken by Step 6 in the first exercise (Section 9.6.1.1) when you were
 27 working in the *art* run-time environment. This script also creates several files and directo-
 28 ries in your `build-prof` directory; these comprise the working space used by **cetbuild-**
 29 **tools**.

- 30 After sourcing this script, the contents of `build-prof` will be

31 `ls buildprof`

```

32 art_workbook-August2015  bin    lib
33 cetpkg_variable_report  diag_report

```

- 34 At this time the two subdirectories `bin` and `lib` will be empty. The other files are used
 35 by the build system to keep track of its configuration.

1 Step 7 (buildtool) tells **cetbuildtools** to build everything found in the source directory; this
 2 includes all of the Workbook exercises, not just the first one. The build process will take
 3 two or three minutes on an unloaded (not undergoing heavy usage) machine. Its output
 4 should end with the lines:

```
5 -----  
6 INFO: Stage build successful.  
7 -----
```

8 After the build has completed do an ls on the directory lib; you will see that it contains
 9 a large number of dynamic library (.so) files; for August2015 there will be about 30
 10 .so files (subject to variation as versions change); these are the files that *art* will load as
 11 you work through the exercises.

12 Also do an ls on the directory bin; these are scripts that are used by **cetbuildtools** to
 13 maintain its environment; if the Workbook contained instructions to build any executable
 14 programs, they would have been written to this directory.

15 After running buildtool, the build directory will contain:

16 **ls buildprof**

```
17 admin                CMakeFiles                fcl  
18 art-workbook          cmake_install.cmake       inputFile  
19 art_workbook-August2015  CPackConfig.cmake        lib  
20 bin                   CPackSourceConfig.cmake  Makefile  
21 cetpkg_variable_report CTestTestfile.cmake      output  
22 CMakeCache.txt        diag_report                ups
```

23 Most of these files are standard files that are explained in the **cetbuildtools** documentation,
 24 <https://cdcv.s.fnl.gov/redmine/projects/cetbuildtools/wiki>. However, three of these items
 25 need special attention here because they are customized for the Workbook.

26 **FIXME:** *Note to Anne from Rob: the above reference for cetbuildtools is the best refer-*
 27 *ence that I know of. It is clearly not adequate. I expect that we will end up writing our own.*
 28 *One thing to be aware of: we don't want to fully explain cetbuildtools or even suggest that*
 29 *people should learn it well. We need a barebones "how to use it" and "how to fix com-*
 30 *mon issues". We will never write a deep how and why - I guess that is another 50 pages.*
 31 *I don't think that many experiments are using cetbuildtools as their build system. I think*
 32 *that it is a particularly foolish choice for a build system for the workbook but I was given*
 33 *a direct instruction by upper management to use it. Someday I will pick the fight but not*

1 *soon. In the mean time anyone who cares will always be dissatisfied with our description*
2 *of cetbuild tools.*

3 An `ls -l` on the files `fcl`, `inputFiles` and `output` will reveal that they are symbolic
4 links to

```
5 inputFiles -> ${TOYEXPERIMENT_DIR}/inputFiles  
6 output -> ${ART_WORKBOOK_OUTPUT_BASE}/  
7             username/art_workbook_output  
8 fcl -> your source directory/art-workbook
```

9 These links are present so that the FHiCL files for the Workbook exercises do not need to
10 be customized on a per-user or per-site basis.

- 11 ○ The link `inputFiles` points to the directory `inputFiles` present in the toyEx-
12 periment UPS product; this directory contains the input files that *art* will read when
13 you run the first exercise. These are the same files used in the first exercise; if you
14 need a reminder of the contents of these files, see Table 9.1. These input files will
15 also be used in many of the subsequent exercises.
- 16 ○ The link `outputFiles` points to a directory that was created to hold your output
17 files; the environment variable `ART_WORKBOOK_OUTPUT_BASE` was defined
18 by your site-specific setup procedure.
- 19 ○ The symlink `fcl` points into your source directory hierarchy; it allows you to ac-
20 cess the FHiCL files that are found in that hierarchy with the convenience of tab
21 completions.

22 10.5 The *art* Development Environment

23 In the preceding sections of this chapter you established what is known as the *art develop-*
24 *ment environment*; this is a superset of the *art* run-time environment, which was described
25 in Section 9.4. This section summarizes the new elements that are part of the development
26 environment but not part of the run-time environment.

27 In Section 10.4.1, step 4b (`git clone ...`) contacted the central source code repository for
28 the *art* Workbook code and made a clone of the repository in your source area under
29 `art-workbook`; the clone contains the complete history of the repository, including all
30 versions of `art-workbook`. Step 4d (`git checkout ...`) examined your clone of the repository,




fig:dev-wkbk

Figure 10.3: Elements of the *art* development environment as used in most of the Workbook exercises; the arrows denote information flow, as described in the text.

1 found the requested version of the code and put a working copy of that version into your
2 source directory. The central repository is hosted on a Fermilab server and is accessed via
3 the network. The upper left box in Figure 10.3 denotes the central repository and the box
4 below it denotes the clone of the repository in your disk space; the box below that denotes
5 the checked out working copy of the Workbook code. The flow of information during the
6 clone and checkout processes is indicated by the green arrows (at left) in the figure.

7 In step 7 of Section 10.4.3, you ran `buildtool` in your build area, which read the source
8 code files from your working copy of the Workbook code and turned them into dynamic
9 libraries. The script `buildtool` is part of the build system, which is denoted as the box
10 in the center left section of Figure 10.3. When you ran `buildtool`, it wrote dynamic library
11 files to the `lib` subdirectory of your build directory; this directory is denoted in the figure
12 as the box in the top center labeled `<build-directory>/lib`. The orange arrows in
13 the figure denote the information flow at build-time. In order to perform this task, `buildtool`
14 also needed to read header files and dynamic libraries found in the UPS products area,
15 hence the orange arrow leading from the UPS Products box to the build system box.

16 In the figure, information flow at run-time is denoted by the blue lines. When you ran the
17 `art` executable, it looked for dynamic libraries in the directories defined by `LD_LIBRARY_PATH`.
18 In the `art` development environment, `LD_LIBRARY_PATH` contains

- 19 1. the `lib` subdirectory of your build directory
- 20 2. all of the directories previously described in Section 9.10

21 In all environments, the `art` executable looks for FHiCL files in

- 22 1. in the file specified in the `-c` command line argument
- 23 2. in the directories specified in `FHICL_FILE_PATH`

24 The first of these is denoted in the figure by the box labeled “Configuration File.” In the
25 `art` development environment, `FHICL_FILE_PATH` contains

- 26 1. some directories found in your checked out copy of the source
- 27 2. all of the directories previously described in Section 9.11

28 The remaining elements in Figure 10.3 are the same as described for Figure 9.1.

29 Figure 10.4, a combination of Figures 10.1 and 10.2), illustrates the distinct source and

- 1 build areas, and the relationship between them. It does not show all the contents in each
- 2 subdirectory.

3 10.6 Running the Exercise

4 10.6.1 Run *art* on `first.fcl`

- 5 In your build window, make sure that your current working directory is your build direc-
- 6 tory. From here, run the first part of this exercise by typing the following:

```
7 art -c fcl/FirstModule/first.fcl >& output/first.log
```

- 8 (As a reminder, we suggest you get in the habit of routing your output to the output
- 9 directory.) The output of this step will look much like that in Listing 9.1, but with two
- 10 significant differences. The first difference is that the output from `first.fcl` contains
- 11 an additional line

```
12 Hello from First::constructor.
```

- 13 The second difference is that the words printed out for each event are a little different; the
- 14 printout from `first.fcl` looks like

```
15 Hello from First::analyze. Event id: run: 1 subRun: 0 event: 1
```

- 16 while that from `hello.fcl` looked like

```
17 Hello World! This event has the id: run: 1 subRun: 0 event: 1
```

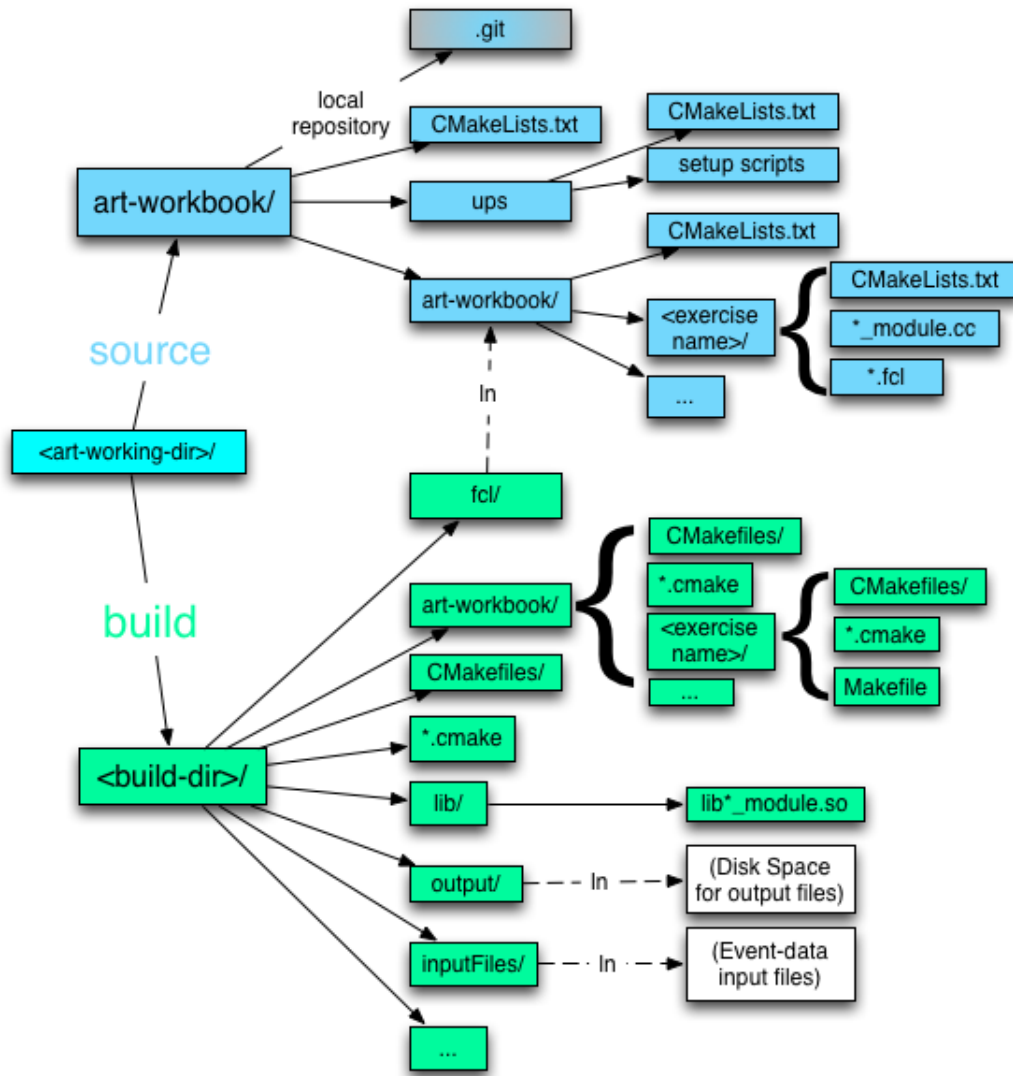
- 18 The reason for changing this printout is so that you can identify, from the printout, which
- 19 module was run.

20 10.6.2 The FHiCL File `first.fcl`

- 21 Compare the FHiCL file used in this exercise, `fcl/FirstModule/first.fcl`, with
- 22 `hello.fcl` from the first exercise (i.e., run `cat` or `diff` on them). Other than comments,
- 23 the only difference is that the `module_type` has changed from `HelloWorld` to `First`:

```
24
```

```
25 diff $TOYEXPERIMENT_DIR/HelloWorldScripts/hello.fcl fcl/FirstModule/first.fcl
```



dir-struct

Figure 10.4: Representation of the reader’s directory structure once the development environment is established.

```

1 ...
2 <     module_type : HelloWorld
3 ---
4 >     module_type : First

```

5 The file `first.fcl` tells *art* to run a module named `First`. As described in Section 9.10, *art* looks through the directories defined in `LD_LIBRARY_PATH` and looks for a file whose name matches the pattern `lib*First_module.so`. This module happens to be found at this location, relative to your build directory:

```
9 lib/libart-workbook_FirstModule_First_module.so
```

10 This dynamic library file was created when you ran `buildtool`.

11 10.6.3 The Source Code File `First_module.cc`

12 This section will describe the source code for the module `First` and will use it as a model to describe modules in general. The source code for this module is found in the following file, relative to your source directory (go to your source window!):

```
15 art-workbook/FirstModule/First_module.cc
```

16 When you ran `buildtool`, it compiled and linked this source file into the following dynamic library, relative to your build directory (go to your build window!):

```
18 lib/libart-workbook_FirstModule_First_module.so
```

19 This is the dynamic library that was loaded by *art* when you ran code for this exercise, in Section 10.6.2.

21 Look at the file `First_module.cc`, shown in Listing 10.2. In broad strokes, it does the following:

- 23 ○ declares a class named `First`
- 24 ○ provides the implementation for the class
- 25 ○ contains a call to the C-Preprocessor macro named `DEFINE_ART_MODULE`, discussed in Section 10.6.3.8

27 All module files that you will see in the Workbook share these “broad strokes.” Some experiments that use *art* have chosen to split the source code for one module into three

- 1 separate files; the *art* team does not recommend this practice, but it is in use and it will be
2 discussed in Section [10.9.2](#).

:module:cc

```
3 1 #include "art/Framework/Core/EDAnalyzer.h"
4 2 #include "art/Framework/Core/ModuleMacros.h"
5 3 #include "art/Framework/Principal/Event.h"
6
7 5 #include <iostream>
8
9 7 namespace tex {
10
11 9     class First : public art::EDAnalyzer {
12
13 11     public:
14
15 13         explicit First(fhicl::ParameterSet const& );
16
17 15         void analyze(art::Event const& event) override;
18
19 17     };
20
21 19 }
22
23 21 tex::First::First(fhicl::ParameterSet const& pset) :
24 22     art::EDAnalyzer(pset) {
25 23     std::cout << "Hello from First::constructor."
26 24         << std::endl;
27 25 }
28
29 27 void tex::First::analyze(art::Event const& event){
30 28     std::cout << "Hello from First::analyze. Event id: "
31 29         << event.id()
32 30         << std::endl;
33 31 }
34
```

```
1 33 DEFINE_ART_MODULE (tex::First)
```

Listing 10.2: The contents of `First_module.cc`

2 10.6.3.1 The `#include` Statements

3 The first three lines of code in the file `First_module.cc` are *include directives* that pull
4 in header files. All three of these files are included from the *art* UPS product (determining
5 the location of included header files is discussed in Section 7.6).

```
6 1 #include "art/Framework/Core/EDAnalyzer.h"
7 2 #include "art/Framework/Core/ModuleMacros.h"
8 3 #include "art/Framework/Principal/Event.h"
```

9 The next line includes the C++ header that enables this code to write output to the screen;
10 for details, see any standard C++ documentation.

```
11 5 #include <iostream>
```

12 Those of you with some C++ experience may have noticed that there is no file named
13 `First_module.h` in the directory `art-workbook/FirstModule`. The explana-
14 tion for this will be given in Section 10.9.1.



16 If you are a C++ beginner you will likely find these header files difficult to understand; you
17 do not need to understand them at this time but you do need to know where to find them
for future reference.

18 10.6.3.2 The Declaration of the Class `First`, an Analyzer Module

19 Let's start with short explanations of each line and follow up with more information.

```
20 7 namespace tex {
21
22 9     class First : public art::EDAnalyzer {
23
24 11     public:
25
26 13         explicit First(fhicl::ParameterSet const& );
```

```
1
2 15     void analyze(art::Event const& event) override;
3
4 17     };
```

5 The first line opens a namespace named `tex`. All of the code in the toyExperiment UPS
6 product was written in the namespace `tex`; the name `tex` is an acronym-like shorthand for
7 the toyExperiment (ToyEXperiment) UPS product. In order to keep things simple, all of
8 the classes in the Workbook are also declared in the namespace `tex`. For more information
9 about this choice, see Section 7.6.4. If you are not familiar with namespaces, consult the
10 standard C++ documentation.

11 The next line begins the declaration of the class `First`. In this line, the fragment
12 `": public art::EDAnalyzer"` tells the compiler that the class `First` is a *derived*
13 *class* that *inherits publicly* from the *base class* `art::EDAnalyzer`. At this time it is
14 not necessary to understand inheritance, which is fortunate, because it takes a long, long
15 time to explain. You just need to recognize and follow the pattern. You can read about C++
16 inheritance in the standard C++ documentation.

17 The line `public:` states that any member functions appearing directly below it are public;
18 i.e., accessible to any code, not only to members of the same class. You can see that two
19 functions are declared as public member functions of the class `First`.

20 The first one is a constructor. Since `First` is intended as a module, the constructor's
21 argument must follow `art`'s prescription. `art` will call the constructor once at the start of
22 each job.

```
23 13     explicit First(fhicl::ParameterSet const& );
```

24 The second declares the `analyze` member function, which also has an argument list
25 prescribed by `art`. This function gets called once per event. The `override` contextual
26 identifier is an important safety feature; please use it!

```
27 15     void analyze(art::Event const& event) override;
```

10.6.3.3 An Introduction to Analyzer Modules

Section 3.6.3 discussed the idea of *module types*: analyzer, producer, filter and so on. For a class to be a valid *art* analyzer module, it must follow a set of rules defined by *art*:

1. It must inherit publicly from `art::EDAnalyzer`.
2. It must provide a constructor with the argument list:

```
fhicl::ParameterSet const& pset
```

(Only the type of the argument is prescribed, not its name. You can use any name you want but the same name must be used in item 3.)
3. The initializer list of the constructor must call the constructor of the base class; and it must pass the parameter set to the constructor of the base class:

```
art::EDAnalyzer (pset)
```

4. It must provide a member function named `analyze`, with the signature⁵:

```
analyze( art::Event const&)
```

5. If the name of a module class is `ClassName` then the source code for the module must be in a file named `ClassName_module.cc` and this file must contain the lines:

```
#include "art/Framework/Core/ModuleMacros.h"
and
DEFINE_ART_MODULE (namespace::ClassName)
```

6. It may optionally provide other member functions with signatures prescribed by *art*; if these member functions are present in a module class, then *art* will call them at the appropriate times. Some examples are provided in Chapter 13.

A module may also contain any other member data and any other member functions that are needed to do its job. You can see from Listing 10.2 that the class `First` follows all of the above rules and that it does not contain any of the optional member functions.

The requirement that the class name match the filename (minus the `_module.cc` portion) is enforced by *art*'s system for run-time loading of dynamic libraries. The re-

⁵ In C++ the *signature* of a member function is the name of the class of which the function is a member, the name of the function, the number, types and order of the arguments, and whether the member function is marked as `const` or `volatile`. The signature does not include the return type; nor does it include the names of any of the arguments.

1 requirement that the class provide the prescribed constructor is enforced by the macro
 2 `DEFINE_ART_MODULE`, which will be described in Section [10.6.3.8](#). sssec:first:build:maker:macro

3 The declaration of the constructor begins with the keyword `explicit`. This is a safety
 4 feature this relevant only for constructors that have exactly one argument. A proper expla-
 5 nation would take too long so just follow a simple guideline: all constructors that have ex-
 6 actly one argument should be declared `explicit`. There will be rare circumstances in which
 7 you need to go against this guideline but you will not encounter any in the Workbook.

9 The *override* contextual identifier in the `analyzer` member function definition is a fea-
 10 ture that is new in C++ 11 so older references will not discuss it. It is a new safety feature
 11 that we recommend you use; we cannot give a proper explanation until we have had a
 12 chance to discuss inheritance further. For now, just consider it a rule that, in all `analyzer`
 13 modules, you should provide this identifier as part of the declaration of `analyze`.

14 For those who are knowledgeable about C++, the base class `art::EDAnalyzer` de-
 15 clares the member function `analyze` to be pure virtual; so it must be provided by the
 16 derived class. The optional member functions of the base class are declared virtual but
 17 not pure virtual; do-nothing versions of these member functions are provided by the base
 18 class.



19 *In a future version of this documentation suite, more information will be available in the*
 20 *Users Guide. **FIXME:** in Chapter ?? Reference to UG*

21 10.6.3.4 The Constructor for the Class `First`

22 The next code in the source file (Listing 10.2) is the *definition* of the constructor for the
 23 class `First`. This constructor simply prints some information (via `std::cout`) to let
 24 the user know that it has been called. lst:First:module:cc

```
25 21 tex::First::First(fhicl::ParameterSet const& pset) :
26 22     art::EDAnalyzer(pset) {
27 23     std::cout << "Hello from First::constructor."
28 24         << std::endl;
29 25 }
```

1 The fragment `tex::First::First` says that this definition is for a constructor of the
2 class `First` from the namespace `tex`.

3 The argument to the constructor is of type `fhiCL::ParameterSet const&` as re-
4 quired by *art*. The class `ParameterSet`, found in the namespace `fhiCL`, is a C++
5 representation of a FHiCL parameter set (aka FHiCL *table*). You will learn how to use this
6 parameter set object in Chapter 14.

7 The argument to the constructor is passed by `const` reference, `const&`. This is a re-
8 quirement specified by *art*; if you write a constructor that does not have exactly the correct
9 argument type, then the compiler will issue a diagnostic and will stop compilation.

10 The first line of the constructor contains the fragment “`: art::EDAnalyzer(pset)`”.
11 This is the constructor’s initializer list and it tells the compiler to call the constructor of
12 the base class `art::EDAnalyzer`, passing it the parameter set as an argument. This is
13 required by rule 3 in the list in Section 10.6.3.3.

14 The requirement that the constructor of an analyzer module pass the parameter set to the
15 constructor of `art::EDAnalyzer` started in *art* version 1.08.09. If you are using an
16 earlier version of *art*, constructors of analyzer modules must NOT call the constructor of
17 `art::EDAnalyzer`.

18 10.6.3.5 Aside: Omitting Argument Names in Function Declarations

19 In the declaration of the class `First`, you may have noticed that the declaration of the
20 member function `analyze` supplied a name for its argument (`event`) but the declaration
21 of the constructor did not supply a name for its argument.

22 In the declaration of a function, a name supplied for an argument is ignored by the com-
23 piler. So code will compile correctly with or without a name. Remember that a constructor
24 is just a special kind of function so the rule applies to constructors too. It is very common
25 for authors of code to provide an argument name as a form of documentation. You will
26 code written both with and without named arguments in declarations.

27 The above discussion only applied to the *declarations* of functions, not to their definition
28 (aka implementation).

10.6.3.6 The Member Function `analyze` and the Representation of an Event

1


`art::Event`

2 The definition of the member function `analyze` comes next in the source file and is
3 reproduced here.

```
4 27 void tex::First::analyze(art::Event const& event) {  
5 28     std::cout << "Hello from First::analyze. Event id: "  
6 29         << event.id()  
7 30         << std::endl;  
8 31 }
```

9 If the type of the argument is not exactly correct, including the `const&`, the compiler
10 will issue a diagnostic and stop compilation. The compiler is able to do this because of
11 one of the features of *inheritance*; the details of how this works is beyond the scope of this
12 discussion.

13 Note that the `override` contextual identifier that was present in the declaration of this
14 member function is not present in its definition; this is standard C++ usage.

15 Section 3.6.1 discussed the HEP idea of an event and the *art* idea of a three-part event
16 identifier. The class `art::Event` is the representation within *art* of the HEP notion of 
17 an event. For the present discussion it is safe to consider the following over-simplified
18 view of an event: it contains an event identifier plus a collection of data products (see
19 Section 3.6.4). The name of the argument `event` has no meaning either to *art* or to the
20 compiler — it is just an identifier — but your code will be easier to read if you choose a
21 meaningful name.

22 At any given time in a running *art* program there is only ever one `art::Event` object;
23 in the rest of this paragraph we will call this object *the event*. It is owned and managed by
24 *art*, but *art* lets analyzer modules see the contents of the event; it does so by passing the
25 event by `const` reference when it calls the `analyze` member function of analyzer mod-
26 ules. Because the event is passed by reference (indicated by the `&`), the member function
27 `analyze` does not get a copy of the event; instead it is told where to find the event. This
28 makes it efficient to pass an event object even if the event contains a lot of information.
29 Because the argument is a `const` reference, if your code tries to change the contents of
30 the event, the compiler will issue a diagnostic and stop compilation.

31 As described in Section 3.6.3, analyzer modules may only inspect data in `event`, not



modify it. This section has shown how *art* institutes this policy as a hard rule that will be enforced rigorously by the compiler:

1. The compiler will issue an error if an analyzer module does not contain a member function named `analyze` with exactly the correct signature.
2. In the correct signature, the argument `event` is a `const` reference.
3. Because `event` is `const`, the compiler will issue an error if the module tries to call any member function of `art::Event` that will modify the event.

You can find the header file for `art::Event` by following the guidelines described in Section 7.6.2. A future version of this documentation will contain a chapter in the *Users Guide* that provides a complete explanation of `art::Event`. Here, and in the rest of the Workbook, the features of `art::Event` will be explained as needed.

The body of the function is almost trivial: it prints some information to let the user know that it has been called. In Section 10.6.1, when you ran *art* using `first.fcl`, the printout from the first event was

```
Hello from First::analyze. Event id: run: 1 subRun: 0 event: 1
```

If you compare this to the source code you can see that the fragment `<< event.id()` creates the following printout

```
run: 1 subRun: 0 event: 1
```

This fragment tells the compiler to do the following:

1. In the class `art::Event`, find the member function named `id()` and call this member function on the object `event`. This returns an object of type `art::EventID`, which is the class that represents an *art* event identifier, which was described in Section 3.6.1. You will learn more about `art::EventID` in Section 10.6.3.7.
2. Print the event identifier.

10.6.3.7 Representing an Event Identifier with `art::EventID`

Section 3.6.1 discussed the idea of an event identifier, which has three components, a run number, a subRun number and event number. In this section you will learn where to find the class that *art* uses to represent an event identifier. Rather than simply telling you

1 the answer, this section will guide you through the process of discovering the answer for
2 yourself.

3 Before you work through this section, you may wish to review Section 7.6 which discusses
4 how to find header files.

5 **FIXME:** *(Anne addressed by saying it's like 'more', below. In the next para we tell people
6 to use less. Most beginners and a surprising nubmer of intermediates don't know what less
7 is. Do we have a section in the intro that tells people how to move around in less and how
8 to exit less? If not maybe we should tell them to open the file in their editor? Or we need to
9 add the relevant section and refer to it the first few times we mention less. Or we can just
10 tell people to look at the file and leave it up to them.*

11 In Section 10.6.3.6 you learned that the member function `art::Event::id()` returns
12 an object that represents the event identifier. To see this for yourself, look at the header
13 file for `art::Event`. Instead of `cat`, try the command `less` which is like `more` but more
14 functional:

```
15 less $ART_INC/art/Framework/Principal/Event.h
```

16 or use one of the code browsers discussed in 7.6.2. In this file you will find the definition
17 of the member function `id()`:⁶

```
18     EventID  
19     id() const {return aux_.id();}
```

20 The important thing to look at here is the return type, `EventID`; you do not need to (or
21 want to) know anything about the data member `aux_`. If you look near the beginning of
22 `Event.h` you will see that it has the line:

```
23 #include "art/Persistency/Provenance/EventID.h"
```

24 which is the header file that declares `EventID`. Look at this file, e.g.,

```
25 less $ART_INC/art/Persistency/Provenance/EventID.h
```

⁶In C++, newlines are treated the same as any other white space; so this could have been written on a single line but the authors of `Event.h` have adopted a style in which return types are always written on their own line.

1 and find the declaration for `EventID`; you will see that the class `EventID` is within the
2 namespace `art`, making its full name `art::EventID`. Near the top of the file you will
3 also see the comments:

```
4 // An EventID labels an unique readout of the data  
5 // acquisition system, which we call an 'event'.
```

6 Look again at `EventID.h`; you will see that it has accessor methods that permit you see
7 the three components of the event identifier:

```
8 RunNumber_t run() const;  
9 SubRunNumber_t subRun() const;  
10 EventNumber_t event() const;
```

11 Earlier in `EventID.h` the C++ *type*⁷ `EventNumber_t` was defined as:

```
12 namespace art {  
13     typedef std::uint32_t EventNumber_t;  
14 }
```

15 meaning that the event number is represented as a 32-bit unsigned integer. A *typedef*(γ) is
16 a different name, or an alias, by which a type can be identified. If you are not familiar with
17 the C++ concept of *typedef*, or if you are not familiar with the definite-length integral types
18 defined by the `<cstdint>` header, consult any standard C++ documentation. If you dig
19 deeper into the layers included in the `art::EventID` header, you will see that the `run`
20 number and `subRun` number are also implemented as 32-bit unsigned integers.

21 The authors of `art` might have chosen an alternate definition of `EventNumber_t`

```
22 namespace art {  
23     typedef unsigned EventNumber_t;  
24 }
```

25 The difference is the use of `unsigned` rather than `std::uint32_t`. This alternate
26 version was not chosen because it runs the risk that some computers might consider this
27 type to have a length of 32 bits while other computers might consider it to have a length


⁷In C++ the collective noun *type*, refers to both the built-in types, such as `int` and `float`, plus user defined types, which include classes, structs and typedefs.

1 of 16 or 64 bits. In the definition that is used by *art*, an event number is guaranteed to be
 2 exactly 32 bits on all computers.

3 Why did the authors of *art* insert the extra level of indirection and not simply define the
 4 following member function inside `art::EventID`?

```
5     std::unit32_t event() const;
```

6 The answer is that it makes it easy to change the definition of the type should that be
 7 necessary. If, for example, an experiment requires that event numbers be of length 64 bits,
 8 only one change is needed, followed by a recompilation.

9 It is good practice to use typedefs for every concept for which the underlying data type is
 10 not absolutely certain. 

11 It is a very common, but not universal, practice within the HEP C++ community that type-
 12 defs that are used to give context-specific names to the C++ built-in types (`int`, `float`,
 13 `char`, etc.) end in `_t`.

14 One last observation about `EventID.h`. Near the top of this file you can find the follow-
 15 ing fragment, with a few lines omitted for clarity:

```
16 namespace art {
17     std::ostream &
18     operator<<(std::ostream & os, EventID const & iID);
19 }
```

20 This tells the compiler that the class `art::EventID` has a stream insertion operator (see
 21 Section 6.7.10). Because this operator exists, the compiler knows how to use `std::cout`
 22 to print an object of type `art::EventID`. You have already used this capability — near
 23 the end of Section 10.6.3.6 see the discussion of the line

```
24     << event.id()
```

25 10.6.3.8 DEFINE_ART_MACRO: The Module Maker Macros

26 The final line in `First_module.cc`,

```
27     DEFINE_ART_MODULE(tex::First)
```

1 invokes a C preprocessor macro. This macro is defined in the header file that was pulled
2 in by

```
3 #include "art/Framework/Core/ModuleMacros.h"
```

4 If you are not familiar with the C preprocessor, don't worry; you do not need to look under
5 the hood. But if you would like to learn about it, consult any standard C++ reference.

6 The `DEFINE_ART_MODULE` macro instructs the compiler to put some additional code
7 into the dynamic library made by `buildtool`. This additional code provides the glue that
8 allows `art` to create instances of the class `First` without ever seeing the header or the
9 source for the class; it only gets to see the `.so` or `.dylib` file and nothing else.



10 The `DEFINE_ART_MODULE` macro adds two pieces of code to the `.so` file. It adds a
11 factory function that, when called, will create an instance of `First` and return a pointer
12 to the base classes `art::EDAnalyzer`. In this way, `art` never sees the derived type of
13 any analyzer module; it sees all analyzer modules via pointer to base. When `art` calls the
14 factory function, it passes as an argument the parameter set specified in the FHiCL file
15 for this module instance. The factory function passes this parameter set through to the
16 constructor of `First`. The second piece of code put into the `.so` file is a static object
17 that will be instantiated at load time; when this object is constructed, it will contact the
18 `art` module registry and register the factory function under the name `First`. When the
19 FHiCL file says to create a module of type `First`, `art` will simply call the registered
20 factory function, passing it the parameter set defined in the FHiCL file. This is the last step
21 in making the connection between the source code of a module and the `art` instantiation of
22 a module.

23 10.6.3.9 Some Alternate Styles

24 C++ allows some flexibility in syntax, which can be seen as either powerful or confusing,
25 depending on your level of expertise. Here we introduce you to a few alternate styles that
26 you will need to recognize and may want to use.

27 Look at the `std::cout` line in the `analyze` method of Listing 10.2: [lst:First:module:cc](#)

```
28     std::cout << "Hello from First::analyze. Event id: "  
29             << event.id()  
30             << std::endl;
```

```
1 }
```

2 This could have been written:

```
3     art::EventID id = event.id();
4     std::cout << "Hello from First::analyze. Event id: "
5               << id
6               << std::endl;
```

7 This alternate version explicitly creates a temporary object of type `art::EventID`,
8 whereas the original version created an *implicit* temporary object. When you are first
9 learning C++ it is often useful to break down compound ideas by introducing *explicit tem-*
10 *poraries*. However, the recommended best practice is to not introduce explicit temporaries
11 unless there is a good reason to do so.

12 You will certainly encounter the first line of the above written in a different style, too,
13 i.e.,

```
14     art::EventID id(event.id());
```


15 Here `id` is initialized using *constructor syntax* rather than using *assignment syntax*. For
16 almost all classes these two syntaxes will produce exactly the same result.

17 You may also see the argument list of the `analyze` function written a little differently,
18

```
19     void analyze( const art::Event& );
```

20 instead of

```
21     void analyze( art::Event const& );
```

22 The position of the `const` has changed. These mean exactly the same thing and the com-
23 piler will permit you to use them interchangeably. In most cases, small differences in the
24 placement of the `const` identifier have very different meanings but, in a few cases, both 
25 variants mean the same thing. When C++ allows two different syntaxes that mean the same
26 thing, this documentation suite will point it out.

27 Finally, Listing 10.3 shows the same information as Listing 10.2 but using a style in which
28 the namespace remains open after the class declaration. In this style, the leading `tex::`

1 is no longer needed in the definitions of the constructor and of `analyze`. Both layouts of
2 the code have the same meaning to the compiler. Many experiments use this style in their
3 source code.

module:alternate

```
4
5 #include "art/Framework/Core/EDAnalyzer.h"
6 #include "art/Framework/Core/ModuleMacros.h"
7 #include "art/Framework/Principal/Event.h"
8
9 #include <iostream>
10
11 namespace tex {
12
13     class First : public art::EDAnalyzer {
14
15     public:
16
17         explicit First(fhicl::ParameterSet const& );
18
19         void analyze(art::Event const& event) override;
20
21     };
22
23     First::First(fhicl::ParameterSet const& pset ) :
24         art::EDAnalyzer(pset) {
25         std::cout << "Hello from First::constructor."
26                 << std::endl;
27     }
28
29     void First::analyze(art::Event const& event){
30         std::cout << "Hello from First::analyze. Event id: "
31                 << event.id()
32                 << std::endl;
33     }
34
```

```
1 }  
2  
3 DEFINE_ART_MODULE(tex::First)
```

Listing 10.3: An alternate layout for `First_module.cc`

10.7 What does the Build System Do?

10.7.1 The Basic Operation

In Section 10.4.3 you issued the command `buildtool`, which *built* `First_module.so`. The purpose of this section is to provide some more details about building modules.

When you ran `buildtool` it performed the following steps:

1. It *compiled* `First_module.cc` to create an object file (ending in `.o`).
2. It *linked* the object file against the libraries on which it depends and inserted the result into a dynamic library (ending in `.so`).

The object file contains the machine code for the class `tex::First` and the machine code for the additional items created by the `DEFINE_ART_MODULE C` preprocessor macro. The dynamic library contains the information from the object file plus some additional information that is beyond the scope of this discussion. This process is called *building* the module.

The verb *building* can mean different things, depending on context. Sometimes it just means compiling; sometimes it just means linking; more often, as in this case, it means both.

To be complete, when you ran `buildtool` it built all of code in the Workbook, both modules and non-modules, but this section will only discuss how it built `First_module.so` starting from `First_module.cc`.

How did `buildtool` know what to do? The answer is that it looked in your source directory, where it found a file named `CMakeLists.txt`; this file contains instructions for **cmake**. Yes, when you ran `buildtool` in your build directory, it did look in your source directory; it knew to do this because, when you sourced `setup_for_development`,

1 it saved the name of the source directory. The instructions in `CMakeLists.txt` tell
2 **cetbuildtools** to look for more instructions in the subdirectory `ups` and in the file `art-workbook/
3 CMakeLists.txt`, which, in turn, tells it to look for more instructions in the `CMakeLists.txt`
4 files in each subdirectory of `art-workbook`.

5 When **cetbuildtools** has digested these instructions it knows the rules to build everything
6 that it needs to build.

7 The object file created by the compilation step is a temporary file and, once it has been
8 inserted into the dynamic library, it is not used any more. Therefore the name of the object
9 file is not important.

10 On the other hand, the name of the dynamic library file is very important. *art* requires
11 that for every module source file (ending in `_module.cc`) the build system must cre-
12 ate exactly one dynamic library file (ending in `_module.so`). It also requires that the
13 name of each `_module.so` file conform to a pattern. Consider the example of the file
14 `First_module.cc`; *art* requires that the dynamic library for this file match the pat-
15 tern

16 `lib*First_module.so`

17 where the `*` wildcard matches 0 or more characters.

18 When naming dynamic libraries, *buildtool* uses the following algorithm, which satisfies
19 the *art* requirements and adds some additional features; the algorithm is illustrated using
20 the example of `First_module.cc`:

- 21 1. find the relative path to the source file, starting from the source directory
22 `art-workbook/FirstModule/First_module.cc`
- 23 2. replace all slashes with underscores
24 `art-workbook_FirstModule_First_module.cc`
- 25 3. change the trailing `.cc` to `.so`
26 `art-workbook_FirstModule_First_module.so`
- 27 4. add the prefix `lib`
28 `libart-workbook_FirstModule_First_module.so`
- 29 5. put the file into the directory `lib`, relative to the build directory
30 `lib/libart-workbook_FirstModule_First_module.so`

1 You can check that this file is there by issuing the following command from your build
2 directory:

```
3 ls -l lib/libart-workbook_FirstModule_First_module.so
```

4 This algorithm guarantees that every module within `art-workbook` will have a unique
5 name for its dynamic library.

6 The experiments using *art* have a variety of build systems. Some of these follow the min-
7 imal *art*-conforming pattern, in which the wildcard is replaced with zero characters. If the
8 Workbook had used such a build system, the name of the dynamic library file would have
9 been

```
10 lib/libFirst_module.so
```

11 Both names are legal. **FIXME:** *Some additional features that are enabled by including the*
12 *full path in the name will be discussed in Section need reference.*

13 10.7.2 Incremental Builds and Complete Rebuilds

14 When you edit a file in your source area you will need to rebuild that file in order for those
15 changes to take effect. If any other files in your source area depend on the file that you
16 edited, they too will need to be rebuilt. To do this, reissue the command:

```
17 buildtool
```

18 Remember that this command must be executed from your build directory and that, before
19 executing it, you must have setup the environment in your build window. When you run
20 this command, **cetbuildtools** will automatically determine which files need to be rebuilt
21 and will rebuild them; it will not waste time rebuilding files that do not need to be rebuilt.
22 This is called an *incremental build* and it will usually complete much faster than the initial
23 build.

24 If you want to clean up everything in your build area and rebuild everything from scratch,
25 use the following command:

```
26 buildtool -c
```

27 This command will give you five seconds to abort it before it starts removing files; to abort,
28 type `ctrl-C` in your build window. It will take about the same time to execute as did your

1 initial build of the Workbook. The name of the option `-c` is a mnemonic for “clean”.

2 When you do a clean build it will remove all files in your build directory that are not man-
3 aged by **cetbuildtools**. For example, if you redirected the output of *art* as follows,



```
art -c fcl/FirstModule/first.fcl >& first.log
```

5 then, when you do a clean build, the file `first.log` will be deleted. This is why the
6 instructions earlier in this chapter told you to redirect output to a log file by

```
art -c fcl/FirstModule/first.fcl >& output/first.log
```

8 When you ran `buildtool`, it created a directory to hold your output files and you created
9 a symbolic link, named `output`, from your build directory to this new directory. Both the
10 other directory and the symbolic link survive clean builds and your output files will be
11 preserved. The Workbook exercises write all of their root and event-data output files to
12 this directory.

13 If you edit certain files in the `ups` subdirectory of your source directory, rebuilding re-
14 quires an extra step. If you edit one of these files, the next time that you run `buildtool`, it will
15 issue an error message saying that you need to re-source `setup_for_development`.

16 If you get this message, make sure that you are in your build directory, and

```
17 source ../art-workbook/ups/setup_for_development \  
18 -p $ART_WORKBOOK_QUAL  
19 buildtool
```

20 10.7.3 Finding Header Files at Compile Time

21 When `setup_for_development` establishes the working environment for the build
22 directory, it does a UPS setup on the UPS products that it requires; this triggers a chain of
23 additional UPS setups. As each UPS product is set up, that product defines many envi-
24 ronment variables, two of which are `PRODUCT-NAME_INC` and `PRODUCT-NAME_LIB`.
25 The first of these points to a directory that is the root of the header file hierarchy for that
26 version of that UPS product. The second of these points to a single directory that holds all
27 of the dynamic library files for that UPS product.

28 You can spot-check this by doing, for example,

1 `ls $TOYEXPERIMENT_INC/*`

2 `ls $TOYEXPERIMENT_LIB`

3 `ls $ART_INC/*`

4 `ls $ART_LIB`

5 You will see that the `_INC` directories have a subdirectory tree underneath them while the
6 `_LIB` directories do not.

7 There are a few small perturbations on this pattern. The most visible is that the `ROOT`
8 product puts most of its header files into a single directory, `$ROOT_INC`. The `Geant4`
9 product does a similar thing.

10 When the compiler compiles a `.cc` file, it needs to know where to find the files specified
11 by the `#include` directives. The compiler looks for included files by first looking for
12 arguments on the command line, of the form

13 `-Ipath-to-a-directory`

14 There may be many such arguments on one command line. If you compiled the code by
15 hand, you would add the `-I` options yourself. In the Workbook, one of the jobs of `buildtool`
16 is to figure out which `-I` options are needed when. The compiler assembles the set of all
17 `-I` arguments and uses it as an include path; that is, it looks for the header files by trying
18 the first directory in the path and if it does not find it there, it tries the second directory
19 in the path, and so on. The choice of `-I` for the name of the argument is a mnemonic for
20 Include.

21 **FIXME:** *FYI; the `tt` for `-I` is on purpose; in the ‘unix’ font it looked like lowercase `L`*
22 *(AH)*

23 When `buildtool` compiles a `.cc` file it adds many `-I` options to the command
24 line; it adds one for each UPS product that was set up when you sourced
25 `setup_for_development`. When building `First_module.cc`, `buildtool` added
26 `-$ART_INC`, `-$TOYEXPERIMENT_INC` and many more.

27 A corollary of this discussion is that when you wish to include a header file from a UPS
28 product, the `#include` directive must contain the relative path to the desired file, starting
29 from the `_INC` environment variable for that UPS product.



1 This system illustrates how the Workbook can work the same way on many different com-
2 puters at many different sites. As the author of some code, you only need to know paths of
3 include files relative to the relevant `_INC` environment variable. This environment variable
4 may have different values from one computer to another but the setup and build systems
5 will ensure that the site-specific information is communicated to the compiler using envi-
6 ronment variables and the `-I` option.

7 This system has the potential weakness that if two products each have a header file with
8 exactly the same relative path name, the compiler will get confused. Should this happen,
9 the compiler will always choose the file from the earlier of the two `-I` arguments on the
10 command line, even when the author of the code intended the second choice to be used. To
11 mitigate this problem, the *art* and UPS teams have adopted the convention that, whenever
12 possible, the first element of the relative path in an `#include` directive will be the UPS
13 package name. It is the implementation of this convention that led to the repeated directory
14 name `art-workbook/art-workbook` that you saw in your source directory. There
15 are a handful of UPS products for which this pattern is not followed and they will be
16 pointed out as they are encountered.

17 The convention of having the UPS product name in the relative path of `#include` direc-
18 tives also tells readers of the code where to look for the included file.

19 10.7.4 Finding Dynamic Library Files at Link Time


20 The module `First_module.cc` needs to call methods of the class `art::Event`.
21 Therefore the compiler left a notation in the object file saying “to use this object file you
22 need to tell it where to find `art::Event`.” **FIXME:** *footnote causes error* The techni-
23 cal way to say this is that the object file contains a list of *undefined symbols* or *undefined*
24 *external references*. When the linker makes the dynamic library

25 `libart-workbook_FirstModule_First_module.so`

26 it must resolve all of the undefined symbols from all of the object files that go into the li-
27 brary. To resolve a symbol, the linker must learn what dynamic library defines that symbol.
28 When it discovers the answer, it will write the name of that dynamic library into something
29 called the *dependency list* that is kept inside the dynamic library. **cetbuildtools** tells the
30 linker that the dependency list should contain only the filename of each dynamic library,

1 not the full path to it. If, after the linker has finished, there remain unresolved symbols,
2 then the linker will issue an error message and the build will fail.

3 If library A depends on library B and library B depends on library C, but library A does
4 not directly depend on library C, then the dependency list of library A should contain only
5 library B. In other words, the dependency should contain only *direct dependencies* (also
6 called *first order dependencies*).

7 To learn where to look for symbol definitions, the linker looks at its command line to
8 find something called the *link list*. The link list can be specified in several different ways
9 and the way that **cetbuildtools** uses is simply to write the link list as the absolute path to
10 every `.so` file that the linker needs to know about. The link list can be different for every
11 dynamic library that the build system builds. However it is very frequently true that if a
12 directory contains several modules, then all of the modules will require the same link list.
13 The bottom line is that the author of a module needs to know the link list that is needed to
14 build the dynamic library for that module. 

15 For these Workbook exercises, the author of each exercise has determined the link list for
16 each dynamic library that will be built for that exercise. In the **cetbuildtools** system, the
17 link list for `First_module.cc` is located in the `CMakeLists.txt` file from same
18 directory as `First_module.cc`; the contents of this file are shown in Listing 10.4. st:First:CMakeLists
19 This `CMakeLists.txt` file says that all modules found in this directory should be built
20 with the same link list and it gives the link list; the link list is the seven lines that begin
21 with a dollar sign; these lines each contain one `cmake` variable. Recall that **cetbuildtools** is
22 a build system that lives on top of `cmake`, which is another build system. A `cmake` variable
23 is much like an environment variable except that is only defined within the environment of
24 the running build system; you cannot look at it with `printenv`.

25 **FIXME:** *Find an appropriate place to discuss using the `cmake` message function to display*
26 *the value of a `cmake` variable.*

27 The five `cmake` variables beginning with `ART_` were defined when `buildtool` set up the
28 `UPS art` product. Each of these variables defines an absolute path to a dynamic library in
29 `$ART_LIB`. For example `${ART_FRAMEWORK_CORE}` resolves to

30 `$ART_LIB/libart_Framework_Core.so`

31 Almost all *art* modules will depend on these five libraries. Similarly the other two variables

First:CMakeLists

Listing 10.4: The file `art-workbook/FirstModule/CMakeLists.txt`

```

1 art_make (MODULE_LIBRARIES
2   ${ART_FRAMEWORK_CORE}
3   ${ART_FRAMEWORK_PRINCIPAL}
4   ${ART_PERSISTENCY_COMMON}
5   ${ART_FRAMEWORK_SERVICES_REGISTRY}
6   ${ART_FRAMEWORK_SERVICES_OPTIONAL}
7   ${FHICL_CPP}
8   ${CETLIB}
9 )

```

1 resolve to dynamic libraries in the **fhiclcpp** and **cetlib** UPS products.

2 When **cetbuildtools** constructs the command line to run the linker, it copies the link list
3 from the `CMakeLists.txt` file to the command linker line.

4 The experiments that use *art* use a variety of build systems. Some of these build systems
5 do not require that all external symbols be resolved at link time; they allow some external
6 symbols to be resolved at run-time. This is legal but it can lead to certain difficulties. A
7 *future version of this documentation suite will contain a chapter in the Users Guide that*
8 *discusses linkage loops and how use of closed links can prevent them. This section will*
9 *then just reference it.*

10 Consult the `cmake` and **cetbuildtools** documentation to understand the remaining details
11 of this file.

12 10.7.5 Build System Details

d-system-details



15 This section provides the next layer of details about the build system; *in a future version of this documentation set, the Users Guide will have a chapter with all of the details.* This entire section contains expert material.

16 If you want to see what buildtool is actually doing, you can enable verbose mode by issuing
17 the command:

```
18 buildtool VERBOSE=TRUE
```

19 For example, if you really want to know the name of the object file, you can find it in
20 the verbose output. For this exercise, the object file is (the path is shown here on two

Table 10.1: Compiler and linker flags for a profile build

Step	Flags
Compiler	-Dart_workbook_FirstModule_First_module_EXPORTS -DNDEBUG
Linker	-Wl,--no-undefined -shared
Both	-O3 -g -fno-omit-frame-pointer -Werror -pedantic -Wall -Werror=return-type -Wextra -Wno-long-long -Winit-self -Woverloaded-virtual -std=c++11 -D_GLIBCXX_USE_NANOSLEEP -fPIC

1 lines):

```
2 ./art-workbook/FirstModule/CMakeFiles/
3     art-workbook_FirstModule_First_module.dir/First_module.cc.o
```

4 Also, you can read the verbose listing to discover the flags given to the compiler and
5 linker. The more instructive compiler and linker flags valid at time of writing are given in
6 Table 10.1. The C++ 11 features are selected by the presence of the `-std=c++11` flag and
7 a high level of error checking is specified. The linker flag,

```
8 -Wl,--no-undefined
```

9 tells the linker that it must resolve all external references at link time. This is sometime
10 referred to as a *closed link*.

11 10.8 Suggested Activities

12 This section contains some suggested exercises in which you will make your own modules
13 and learn more about how to use the class `art::EventID`.

14 10.8.1 Create Your Second Module

15 In this exercise you will create a new module by copying `First_module.cc` and mak-
16 ing the necessary changes; you will build it using `buildtool`; you will copy `first.fcl`
17 and make the necessary changes; and you will run the new module using the new `FHiCL`
18 file.

1 Go to your source window and `cd` to your source directory. If you have logged out, out
2 remember to re-establish your working environment; see Section [II](#) Type the following
3 commands:

```
4 cd art-workbook/FirstModule
```

```
5 cp First_module.cc Second_module.cc
```

```
6 cp first.fcl second.fcl
```

7 Edit the files `Second_module.cc` and `second.fcl`. In both files, change every oc-
8 currence of the string “First” to “Second”; there are eight places in the source file and two
9 in the FHiCL file, one of which is in a comment.

10 The new module needs the same link list as did `First_module.cc` so there is no need
11 to edit `CMakeLists.txt`; the instructions in `CMakeLists.txt` tell `buildtool` to build
12 all modules that it finds in this directory and to use the same link list for all modules.

13 Go to your build window and `cd` to your build directory. Again, remember to re-establish
14 your working environment as necessary. Rebuild the Workbook code:

```
15 buildtool
```

16 This should complete with the message:

```
17 -----  
18 INFO: Stage build successful.  
19 -----
```

20 If you get an error message, consult a local expert or the *art* team as described in Sec-
21 tion [3.4](#).

22 When you run `buildtool` it will perform an incremental build (see Section [10.7.2](#)) during
23 which it will detect `Second_module.cc` and build it.

24 You can verify that `buildtool` created the expected dynamic library:

```
25 ls lib/*Second*.so
```

```
26 lib/libart-workbook_FirstModule_Second_module.so
```

27 Stay in your build directory and run the new module:

```
28 art -c fcl/FirstModule/second.fcl >& output/second.log
```


- 1 Compare `output/second.log` with `output/first.log`. You should see that “First”
- 2 has been replaced by “Second” everywhere and the date/time lines are different.

3 10.8.2 Use *artmod* to Create Your Third Module

4 This exercise is much like the previous one; the difference is that you will use a tool named
5 *artmod* to create the source file for the module.

6 Go to your source window and `cd` to your source directory. If you have logged out, re-
7 member to re-establish your working environment; see Section II

8 The command *artmod* creates a file containing the skeleton of a module. It is supplied
9 by the UPS product **cetpkgssupport**, which was set up when you performed the last step
10 of establishing the environment in the source window, sourcing `setup_deps`. You can
11 verify that the command is in your path by using the bash built-in command `type` (output
12 shown on two lines):

```
13 type artmod
```

```
14 artmod is hashed (/ds50/app/products/cetpkgssupport/  
15 v1_02_00/bin/artmod)
```

16 The leading elements of the directory name will reflect your UPS products area, and may
17 be different from what is shown here. The version number, `v1_02_00`, may also change
18 with time.

19 From your source directory, type the following commands:

```
20 cd art-workbook/FirstModule
```

```
21 artmod analyzer tex::Third
```

```
22 cp first.fcl third.fcl
```

23 The second command tells *artmod* to create a source file named `Third_module.cc`
24 that contains the skeleton for an ‘analyzer’ module, to be named `Third` in the namespace
25 `tex`.

26 If you compare `Third_module.cc` to `First_module.cc` you will see a few differ-
27 ences:

- 1 1. `Third_module.cc` is longer and has more comments
- 2 2. the layout of the class is a little different but the two layouts are equivalent
- 3 3. there are some extra `#include` directives
- 4 4. the include for `<iostream>` is missing
- 5 5. in the `analyze` member function, the name of the argument is different (`event`
- 6 `vs e`)
- 7 6. `artmod` supplies the skeleton of a destructor (`~Third`)

8 The `#include` directives provided by `artmod` are a best guess, made by the author of
9 `artmod`, about which ones will be needed in a “typical” module. Other than slowing down
10 the compiler by an amount you won’t notice, the extra `#include` directives do no harm;
11 keep them or leave them as you see fit.

12 Edit `Third_module.cc`

- 13 1. add the `#include` directive for `<iostream>`
- 14 2. copy the bodies of the constructor and the `analyze` member function from `First_module.cc`;
15 change the string “First” to “Third”
- 16 3. in the definition of the member function `analyze`, change the name of the argu-
17 ment to `event`.

18 When you built `First_module.cc`, the compiler wrote a destructor for you that is
19 identical to the destructor written by `artmod`; so you can leave the destructor as `artmod`
20 wrote it, i.e., with an empty body. Or you can delete it; if you decide to do so, you must
21 delete both the declaration and the implementation.

22 Edit `third.fcl` Change every occurrence of the string “First” to “Third”; there are two
23 places, one of which is in a comment.

24 Go to your build window and `cd` to your build directory. If you have logged, out re-
25 member to re-establish your working environment; see Section 11. Rebuild the Workbook
26 code:

27 `buildtool`

1 Refer to the previous section to learn how to identify a successful build and how to verify
2 that the expected library was created.

3 Stay in your build directory and run the third module:

```
4 art -c fcl/FirstModule/third.fcl >& output/third.log
```

5 Compare `output/third.log` with `output/first.log`. You should see that the
6 printout from `First_module.cc` has been replaced by that from `Third_module.cc`.

7 `artmod` has many options that you can explore by typing:

```
8 artmod --help
```

9 10.8.3 Running Many Modules at Once

10 In this exercise you will run four modules at once, the three made in this exercise plus the
11 `HelloWorld` module from Chapter 9.

12 Go to your source window and `cd` to your source directory. Type the following com-
13 mands:

```
14 cd art-workbook/FirstModule
```

```
15 cp first.fcl all.fcl
```

16 Edit the file `all.fcl` and replace the `physics` parameter set with the contents of List-
17 ing 10.5. This parameter set:

- 18 1. defines four module labels and
- 19 2. puts all four module labels into the `end_paths` sequence.

20 When you run `art` on this FHiCL file, `art` will first look at the definition of `end_paths`
21 and learn that you want it to run four module labels. Then it will look in the `analyzers`
22 parameter set to find the definition of each module label; in each definition `art` will find
23 the class name of the module that it should run. Given the class name and the environment
24 variable `LD_LIBRARY_PATH`, `art` can find the right dynamic library to load. If you need
25 a refresher on module labels and `end_paths`, refer to Sections 9.8.7 and 9.8.8.

`st:First:all:fcl`**Listing 10.5:** The `physics` parameter set for `all.fcl`

```

1 physics :{
2   analyzers: {
3     hello : {
4       module_type : HelloWorld
5     }
6     first : {
7       module_type : First
8     }
9     second : {
10    module_type : Second
11  }
12  third : {
13    module_type : Third
14  }
15 }
16
17 e1      : [ hello, first, second, third ]
18 end_paths : [ e1 ]
19
20 }

```

1 Go to your build window and `cd` to your build directory. If you have logged out, remember
 2 to re-establish your working environment; see Section [11](#). You do not need to build any
 3 code for this exercise.

4 Run the exercise:

```
5 art -c fcl/FirstModule/all.fcl >& output/all.log
```

6 Compare `output/all.log` with the log files from the previous exercises. The new log
 7 file should contain printout from each of the four modules. Once, near the start of the file,
 8 you should see the printout from the three constructors; remember that the `HelloWorld`
 9 module does not make any printout in its constructor. For each event you should see the
 10 printout from the four `analyze` member functions.

11 Remember that `art` is free to run analyzer modules in any order; this was discussed in
 12 Section [9.8.8](#).

10.8.4 Access Parts of the EventID

1 `ed:eventid`
2 In this exercise, you will access the individual parts of the event identifier.

3 Before proceeding with this section, review the material in Section 10.6.3.7 which dis-
4 cusses the class `art::EventID`. The header file for this class is:

5 `$ART_INC/art/Persistency/Provenance/EventID.h`

6 In this exercise, you are asked to rewrite the file `Second_module.cc` so that the print-
7 out made by the `analyze` method looks like the following (lines split here due to space
8 restrictions):

```
9 Hello from FirstAnswer01::analyze.  run number: 1  
10     sub run number: 0 event number: 1  
11 Hello from FirstAnswer01::analyze.  run number: 1  
12     sub run number: 0 event number: 2
```

13 and so on for each event.

14 To do this, you will need to reformat the text in the `std::cout` statement and you will
15 need to separately extract the run, subRun and event numbers from the `art::EventID`
16 object.

17 You will do the editing in your source window, in the subdirectory `art-workbook/`
18 `FirstModule`.

19 When you think that you have successfully rewritten the module, you can test it by going
20 to your build window and `cd`'ing to your build directory. Then:

21 `buildtool`

22 `art -c fcl/FirstModule/second.fcl >& output/eventid.log`

23 If you have not figured out how to do this exercise after about 15 minutes, you can find
24 one possible answer in the file `FirstAnswer01_module.cc`, in the same directory as
25 `First_module.cc`.

26 To run the answer module and verify that it makes the requested output, run:

27 `art -c fcl/FirstModule/firstAnswer01.fcl >& output/firstAnswer01.log`

1 (The command can be typed on a single line.) You did not need to build this module
 2 because it was already built the first time that you ran buildtool; that run of buildtool built
 3 all of the modules in the Workbook.

4 There is a second correct answer to this exercise. If you look at the header file for `art::Event`,
 5 you will see that this class also has member functions

```
6     EventNumber_t    event()    const {return aux_.event();}
7     SubRunNumber_t  subRun()   const {return aux_.subRun();}
8     RunNumber_t     run()      const {return id().run();}
```

9 So you could have called these directly,

```
10     std::cout << "Hello from FirstAnswer01::analyze. "
11             << " run number: "      << event.run()
12             << " sub run number: " << event.subRun()
13             << " event number: "   << event.event()
14             << std::endl;
```

15 instead of

```
16     std::cout << "Hello from FirstAnswer01::analyze. "
17             << " run number: "      << event.id().run()
18             << " sub run number: " << event.id().subRun()
19             << " event number: "   << event.id().event()
20             << std::endl;
```

21 But the point of this exercise was to learn a little about how to dig down into nested header
 22 files to find the information you need.

23 10.9 Final Remarks

24 10.9.1 Why is there no `First_module.h` File?

`c:First:why:no:h`

25 When you performed the exercises in this chapter, you saw, for example, the file `First_module.cc`
 26 but there was no corresponding `First_module.h` file. This section will explain why.

1 In a typical C++ programming environment there is a header file (.h) for each source
2 file (.cc). As an example, consider the files `Point.h` and `Point.cc` that you saw in
3 Section 6.7.10.

4 The reason for having `Point.h` is that the implementation of the class, `Point.cc`, and
5 the users of the class need to agree on what the class `Point` is. In the Section 6.7.10
6 example, the only user of the class is the main program, `pctest.cc`. The file `Point.h`
7 serves as the unique, authoritative declaration of what the class is; both `Point.cc` and
8 `pctest.cc` rely on on this declaration.

9 If you think carefully, you are already aware of a very common exception to the pattern of
10 one .h file for each .cc file: there is never a header file for a main program. For example,
11 in the examples that exercised the class `Point`, `pctest.cc` had no header file. Why not?
12 No other piece of user-written code needs to know about any classes or functions declared
13 or defined inside `pctest.cc`.

14 The `First_module.h` file is omitted simply because every entity that needs to see the
15 declaration of the class `First` is already inside the file `First_module.cc`. There is
16 no reason to have a separate header file. Recall the “dangerous bend” paragraph at the end
17 of Section 10.6.3.8 that described how *art* is able to use modules without needing to know
18 about the declaration of the module class.

19 *art* is designed such that only *art* may construct instances of module classes and only *art*
20 may call member functions of module classes. In particular, modules may not construct
21 other modules and may not call member functions of other modules. The absence of a
22 `First_module.h`, provides a physical barrier that enforces this design.

23 10.9.2 The Three-File Module Style

24 In this chapter, the source for the module `First` was written in a single file. You may also
25 write it using three files, `First.h`, `First.cc` and `First_module.cc`.

26 Some experiments use this three-file style. The authors of *art* do not recommend it, how-
27 ever, because it exposes the declaration of `First` in a way that permits it to be misused
28 (as was discussed in Section 10.9.1). The build system distributed with the Workbook has
29 not been configured to build modules written in this style.



1 In this style, `First.h` contains the class declaration plus any necessary `#include` di-
 2 rectives; it now also requires include guards; this is shown in Listing 10.6.

```

3 #ifndef art-workbook_FirstModule_First_h
4 #define art-workbook_FirstModule_First_h
5
6 #include "art/Framework/Core/EDAnalyzer.h"
7 #include "art/Framework/Principal/Event.h"
8
9 namespace tex {
10
11     class First : public art::EDAnalyzer {
12
13     public:
14
15         explicit First(fhicl::ParameterSet const& );
16
17         void analyze(art::Event const& event) override;
18
19     };
20
21 }
22 #endif

```

Listing 10.6: The contents of `First.h` in the three-file model

23 The file `First.cc` contains the definitions of the constructor and the `analyze` member
 24 function, plus the necessary `#include` directives; this is shown in Listing 10.7.

```

25 #include "art-workbook/FirstModule/First.h"
26
27 #include <iostream>
28
29 tex::First::First(fhicl::ParameterSet const& pset ) :
30     art::EDAnalyzer(pset) {
31     std::cout << "Hello from First::constructor."
32     << std::endl;

```



```

1  }
2
3  void tex::First::analyze(art::Event const& event) {
4      std::cout << "Hello from First::analyze. Event id: "
5                  << event.id()
6                  << std::endl;
7  }

```

Listing 10.7: The contents of `First.cc` in the three-file model

8 And `First_module.cc` is now stripped down to the invocation of the `DEFINE_ART_MODULE`
9 macro plus the necessary `#include` directives; this is shown in Listing 10.8.

```

10 #include "art-workbook/FirstModule/First.h"
11 #include "art/Framework/Core/ModuleMacros.h"
12
13 DEFINE_ART_MODULE(tex::First)

```

Listing 10.8: The contents of `First_module.cc` in the three-file model

14 10.10 Flow of Execution from Source to FHiCL File

15 The properties that a class must have in order to be an analyzer module are summarized
16 in Section 10.6.3.2 for reference. This section reviews how the source code found in an
17 analyzer module, e.g., `First_module.cc`, is executed by `art`:

- 18 1. The script `setup_for_development` defines many environment variables that
19 are used by `buildtool`, `art` and `toyExperiment`.
- 20 2. `LD_LIBRARY_PATH`, an important environment variable, contains the directory
21 `lib` in your build area plus the `lib` directories from many UPS products, including
22 `art`.
- 23 3. `buildtool` compiles `First_module.cc` to a temporary object file.
- 24 4. `buildtool` links the temporary object file to create a dynamic library in the `lib` sub-
25 directory of your build area:
26 `lib/libart-workbook_FirstModule_First_module.so`

- 1 5. When you run *art* using file `first.fcl`, this file tells *art* to find and load a module
- 2 with the “module_type” `First`.
- 3 6. In response to this request, *art* will search the directories in `LD_LIBRARY_PATH`
- 4 to find a dynamic library file whose name matches the pattern:
- 5 `lib*First_module.so`
- 6 7. If *art* finds either zero or more than one match to this pattern, it will issue an error
- 7 message and stop.
- 8 8. If *art* finds exactly one match to this pattern, it will load the dynamic library.
- 9 9. After *art* has loaded the dynamic library, it has access to a function that can, on
- 10 demand, create instances of the class `First`.

11 The last bullet really means that the dynamic library contains a factory function that can
 12 construct instances of `First` and return a pointer to the base class, `art::EDANalyzer`.
 13 The dynamic library also contains a static object that, at load-time, will contact the *art*
 14 module registry and register the factory function under the `module_type First`.



15 10.11 Review

16 **FIXME:** *To be added; list from ‘what you will learn’ repeated here for reference*

- 17 ○ how to establish the *art* development environment
- 18 ○ how to checkout the Workbook exercises from the `git` source code management
- 19 system
- 20 ○ how to use the **cetbuildtools** build system to build the code for the Workbook exer-
- 21 cises
- 22 ○ how include files are found
- 23 ○ what a *link list* is
- 24 ○ where the build system finds the link list
- 25 ○ what the `art::Event` is and how to access it
- 26 ○ what the `art::EventID` is and how to access it

- 1 ○ what makes a class an *art module*
- 2 ○ where the build system puts the `.so` files that it makes

3 10.12 Test Your Understanding

4 10.12.1 Tests

5 Two modules are provided that intentionally contain bugs and fail to build:

6 `FirstModule/FirstBug01_module.cc.nobuild`

7 `FirstModule/FirstBug02_module.cc.nobuild`

8 Your job in each case is to figure out what's wrong with the module and fix it. The build
9 system will ignore these files until the `.nobuild` is removed. It's a good idea to work on
10 these two files one at a time. In the procedure below, start with the one labeled 01. Repeat
11 the procedure for the second file, changing 01 to 02 as needed. Answers are provided at
12 the end of the chapter, in Section 10.12.2.

13 Pay attention to which window you do each operation in; there's a fair bit of back-and-forth
14 between the source and build windows!



15 In your source window, run:

```
cp FirstBug01_module.cc.nobuild FirstBug01_module.cc
```

Go to your build directory and give the command:

```
buildtool
```

This will include the first module in the build. For both files the compilation will fail and it will generate a long set of error messages. The best way to attack the problem is to find the first error message, understand it, fix that problem and repeat. Often, but not always, when you fix the first error, all of the other error messages will go away too! You may want

16

to redirect the buildtool output to a file and search on `error`, `warning`, `FirstBug01` or some combination to find the first error associated with this module.

Once you find the error and determine its cause, go to your source window to correct it.

Go to your build window and rebuild:

```
buildtool
```

Go back to your source window and prepare a FHiCL file to run this module. (See Section 10.6.2.)

Go to your build window and run *art*:

```
art -c fcl/FirstModule/your-file-name.fcl >& output/your-file-name.log
```

1

2 Repeat the procedure for the second file. Because this is the first time you are doing this,
3 we will give you a hint: you only need to fix one line in the first file but you need to fix
4 two in the second.

5 The answers are intentionally placed on a new page (remember to try before you read
6 further!).

10.12.2 Answers

10.12.2.1 FirstBug01

In `FirstModule/FirstBug01_module.cc` the error is in:

```
void analyze(art::Event& event) override;
```

There is a missing `const` in the type of the argument. Recall from Section 10.6.3.6 that the argument list of an analyzer member function is prescribed by `art` to be `art::Event const&`. The relevant error text from the printout, shown here on two lines, is:

```
error: 'void_tex::FirstBug01::analyze(art::Event&)' marked override,  
but does not override
```

Because of the `override` keyword, the signature of the function must exactly match one of the virtual functions in the base class. Because `const` was removed, there is no exact match.

10.12.2.2 FirstBug02

In `FirstModule/FirstBug02_module.cc` there are errors in two places. The first line is in the class declaration,

```
void analyze(art::Event& event);
```

and the second is in the implementation:

```
void tex::FirstBug02::analyze(art::Event& event) {
```

In the first line, both the `const` and the `override` were missing. In the second line the `const` was missing. Section 10.6.3.6 discusses the member function `analyze`. The relevant error text from the printout (shown on two lines here) is:

```
error: 'virtual_void_art::EDAnalyzer::analyze(const_art::Event&)'  
was hidden [-Werror=overloaded-virtual]
```

Because the `override` is gone, there is no longer a requirement that the function match any of the virtual functions in the base class. Therefore the error message from the previous

1 file is no longer present. However, the compiler does detect that the base class has a virtual
2 function with the same name but a different argument list; this is what the error text means
3 by “hidden”.

4 While there are legitimate reasons to hide a virtual function, most cases are, in fact, errors.
5 Therefore the build system for the Workbook is configured as follows:

- 6 1. When the compiler detects a hidden virtual function it will issue a warning (not an
7 error).
- 8 2. The compiler will promote all warnings to errors, which will stop the build. (hence
9 the `[-Werror=overloaded-virtual]` message at the end).

10 This forces you to think about each case and decide if it is really what you intended.
11 To learn how to allow hiding of virtual functions consult the `cetbuildtools` documenta-
12 tion.

13 **FIXME:** *Reference to the `cetbuildtools` documentation when it is ready*

11 General Setup for Login Sessions

g:in:again

2 After you've done the initial setup described in Section 10.4, there are some steps that
3 don't need to be repeated for subsequent login sessions. To begin a general login session
4 for Exercise 2 or any subsequent exercise, you need to follow the instructions in this chap-
5 ter.

6 If during your initial setup you chose to manage your own directory names, then the names
7 of your source and build directories will be different than those shown here.



11.1 Source Window

9 In your source window:

1. Log in to the computer you chose in Section 8.3.
2. Follow the site-specific setup procedure; see Table 5.1
3. cd to your source directory:

```
cd $ART_WORKBOOK_WORKING_BASE/username/\  
workbook/art-workbook
```

4. Set up the environment:

```
source ups/setup_deps -p $ART_WORKBOOK_QUAL
```

10

- 1 The contents of the source directory is discussed in Section [10.4.2.2](#). [|ssec:source-dir-contents](#)

2 **11.2 Build Window**

- 3 In your build window:

1. Log in to the computer you chose in Section [8.3](#). [|sec:selectMachine](#)
2. Follow the site-specific setup procedure; see Chapter [5](#). [|ch:site-setup](#)
3. cd to your build directory:

```
cd $ART_WORKBOOK_WORKING_BASE/username/\  
workbook/build-prof
```

4. Source the setup file:

```
source ../art-workbook/ups/setup_for_development \  
-p $ART_WORKBOOK_QUAL
```

- 4
- 5 The build window setup is discussed in Section [10.4.4](#) and the *art* development environ- [|sec:first:build:understanding:steps:build](#)
- 6 ment is described in Section [10.5](#). [|sec:art:dev:env](#)

12 Keeping Up to Date with Workbook Code and Documentation

up:to:date

12.1 Introduction

date:intro

As you well know by now, the Workbook exercises require you to download some code to edit, build, execute and evaluate. Both the documentation and the code it references are expected to undergo continual development throughout 2014. The latest is always available at the *art* Documentation website.

Announcements of new releases are made on the `art-users@fnal.gov` mailing list. Please subscribe!



Until the full set of exercises is written, you will have to update occasionally just to get the latest exercises. Come back to this chapter whenever you reach the end of the available exercises. Or come back and update whenever a new release is announced; it may include improvements to existing exercises.

12.2 Special Instructions for Summer 2014

Summer 2014: Until further notice, if you need to obtain updated Workbook code, you will need to reinstall the Workbook code from scratch. The procedures below will usually work but there are some circumstances in which they won't. Until the workbook team can document how you should deal with the exceptional cases, please reinstall from scratch. To do so, use the following procedure:



1. Save your existing work so that you can refer to it later.
 - (a) Go to the directory that is two above your source and build directories and get a directory listing:

```
cd $ART_WORKBOOK_WORKING_BASE/username
ls
```
 - (b) You should see a directory named `workbook` that contains your source and build directories `art-workbook` and `build-prof`. You may also see other files and directories.
 - (c) Choose a new name for the `workbook` directory, perhaps `workbook_sav1`. The suffix `_sav1` is just a suggestion — the only requirement is that the new name not conflict with existing ones.
 - (d) Rename the `workbook` directory

```
mv workbook workbook_sav1
```
2. Follow the instructions to install the Workbook code from scratch in Section 10.4.

12.3 How to Update

This chapter will show you how to update. The steps include:

For the moment, please restart from scratch. See the previous section

1. Determine whether an updated release is available, and what release it is.
2. Switch to the updated documentation.
3. In your source window, use `git` to update your working version of the code in the (higher-level) `art-workbook` directory
4. In your build window, build the new version of the code.

1 12.3.1 Get Updated Documentation

2 First, check which documentation release you’re currently using: it’s noted on the title
3 page of this document¹. Then go to the *art* Documentation website and compare your
4 documentation release number to the latest available.

5 Download a new copy of the documentation, as needed.

6 12.3.2 Get Updated Code and Build It

7 Also noted on the title page of the documentation is the release² of the *art-workbook* code
8 that the documentation is intended for. Recall from Figure 10.1 that `git` commands are used
9 to clone the code in the remote repository into your local copy, then copy the requested
10 release from that local copy into your working area. The `git` system is described in more
11 detail in Chapter 49.

12 Chances are that you’re using the code release that goes with the documentation you have
13 been using. You can check by looking in the file `art-workbook/ups/product_deps`.
14 From your source directory run:

```
15 grep art_workbook ups/product_deps  
16 parent art_workbook v0_00_13
```

17 This shows version `v0_00_13` as an example. If your version is earlier than the one listed
18 on the cover of the latest documentation, you will need to get new code and build it.

19 These instructions illustrate updating the working version of the *art-workbook* code from
20 version `v0_00_13` to version `v0_00_15`. There is nothing special about these two versions;
21 the instructions serve as a model for a change between any pair of versions.

22 1. Start from (or `cd` to) your source directory (see Section 10.4.1):
23 `cd $ART_WORKBOOK_WORKING_BASE/username/workbook/art-workbook`

¹Versions of the *art* documentation prior to 0.51 do not have this information on the front page; for these versions, the required version of *art-workbook* can be found in the section “Setting up to Run Exercises” in Exercise 2.

²The terms “release” and “version” are used interchangeably here.

- 1 2. Use `git status` and make a note of the files that you have modified and/or added (see
2 Section 12.3.3 for instructions).

3 `git status [-s]`

- 4 3. Switch from your tagged version branch back to the `develop` branch (“branches” are
5 discussed in Chapter 49, you don’t need to understand them at this stage³).

6 `git checkout develop`

7 Switched to branch ‘develop’

- 8 4. Update your local copy of the repository (the `.git` directory)

9 `git pull`

10 The output from this command is shown in Listing 12.1.

- 11 5. Switch your working code to the new branch:

12 `git checkout -b v0_00_15 v0_00_15`

13 Switched to a new branch ‘v0_00_15’

14 Use the new version number twice in this command. In the messages produced in
15 this step, watch for the names of files that you have modified. Check for conflicts
16 that git did not merge correctly.

Listing 12.1: Example of the output produced by `git pull`

```

17 From http://cdcvs.fnal.gov/projects/art-workbook
18    e79d9ef..81d2a76 develop -> origin/develop
19    6435ecc..c0c1af5 master  -> origin/master
20 From http://cdcvs.fnal.gov/projects/art-workbook
21 * [new tag]          v0_00_14 -> v0_00_14
22 * [new tag]          v0_00_15 -> v0_00_15
23 Updating e79d9ef..81d2a76
24 Fast-forward
25 art-workbook/ModuleInstances/magic.fcl      | 26 ++++++-----
26 art-workbook/ParameterSets/PSet01_module.cc | 36 ++++++-----
27 art-workbook/ParameterSets/PSet02_module.cc | 53 ++++++-----

```

³If you are familiar with git concepts, you may want to know this: The authors of the art-workbook follow the convention that they make a new git branch for every release of art-workbook and the name of the branch matches the version number of art-workbook. In the current example, the local working environment knows about two branches, the `develop` branch and the branch for version `v0_00_13`. The `develop` branch is the name of the branch that always contains the most recent art-workbook code.

```

1 art-workbook/ParameterSets/PSet03_module.cc | 28 ++++++-----
2 art-workbook/ParameterSets/PSet04_module.cc | 44 ++++++-----
3 art-workbook/ParameterSets/pset01.fcl      | 6 ++---
4 art-workbook/ParameterSets/pset02.fcl      | 14 ++++++----
5 art-workbook/ParameterSets/pset03.fcl      | 6 ++---
6 art-workbook/ParameterSets/pset04.fcl      | 7 +++---
7 ups/product_deps                          | 2 +-
8 10 files changed, 109 insertions(+), 113 deletions(-)

```

9 To rebuild your updated working code: **FIXME:** *We need to say what to do with any*
10 *updated files that you have; the setupfordev script fails if you just leave them there; there*
11 *are conflicts. Actually, I'm not sure that's the source of my present problem... 6/20/14*
12 *AH*

13 1. In your build window, cd to your build directory

```
14 cd $ART_WORKBOOK_WORKING_BASE/username/workbook/build-prof
```

15 2. Tell **cetbuildtools** to look for, and act on, any changes in your checked out version
16 of the code (command shown on two lines):

```
17 source ../art-workbook/ups/setup_for_development \-p $ART_WORKBOOK_QUAL
```

18 **FIXME:** *add example output*

19 3. Rebuild:

```
20 buildtool
```

21 If this step does not complete successfully, the first thing to try is a clean rebuild:

```
22 buildtool -c
```

23 **FIXME:** *add example output*

24 12.3.3 See which Files you have Modified or Added

```
git:status
```

25 At any time you can check to see which files you have modified and which you have added.
26 The code is structured in such a way that when you checkout a new version, these files will
27 remain in your working directory and will not be modified or deleted. The git checkout
28 command will generate some informational messages about them, but you do not need to
29 take any action.

30 To see the new/modified files, cd to your source directory and issue the `git status`
31 command. Suppose that you have checked out version `v0_00_13`, modified `first.fcl`
32 and added `second.fcl`. The `git status` command will produce the following output:

```
1 git status
2 # On branch v0_00_13
3 # Changes not staged for commit:
4 #   (use "git_add_<file>..." to update what will be committed)
5 #   (use "git_checkout_--<file>..." to discard changes in
6 working directory)
7 #
8 #       modified:   first.fcl
9 #
10 # Untracked files:
11 #   (use "git_add_<file>..." to include in what will be committed)
12 #
13 #       second.fcl
14 no changes added to commit (use "git_add" and/or "git_commit_-a")
```



16 Do not issue the git add or git commit commands that are suggested in the command output above.

17 In the rare case that you have neither modified nor added any files, the output of git status
18 will look like:

```
19 git status
20 # On branch v0_00_13
```

13 Exercise 3: Some other Member Functions of Modules

13.1 Introduction

Recall the discussion in Section 3.6.2 about widget-making workers on an assembly line. All workers have a task to perform on each widget as it passes by and some workers may also need to perform start-up or shut-down tasks. If a module has something that it must do at the start of the job, then the author of the module can write a member function named `beginJob()` that performs these tasks. Similarly the author of a module can write a member function named `endJob` to do tasks that need to be performed at the end of the job. *art* will call both of these member functions at the appropriate time.

The author of a module may also provide member functions to perform actions at the start of a subRun, at start of a run, at the end of a subRun or at the end of a Run.

These member functions are *optional*; i.e., they are always allowed in a module but never required. They have prescribed names and argument lists.



In this exercise you will build and execute an analyzer module that illustrates three of these member functions: `beginJob`, `beginRun` and `beginSubRun`. These member functions are called, respectively, once at the start of the *art* job, once for each new run and once for each new subRun.

You may also perform a suggested exercise to add the three corresponding member functions `endJob`, `endRun` and `endSubRun`.

1 13.2 Prerequisites

2 The prerequisites for this chapter include all of the material in Part I (Introduction) and all
3 of the material up to this point in Part II (Workbook).

4 In particular, make sure that you understand the *event loop* (see Section 3.6.2). [ssec:eventloop](#)

5 13.3 What You Will Learn

6 This chapter will show you *how* to provide the optional member functions in your *art*
7 modules to execute special functionality at the beginning and end of jobs, runs and/or
8 subRuns. These include

9 1. `beginJob()`

10 2. `beginRun(art::Run const&)`

11 3. `beginSubRun(art::SubRun const&)`

12 4. `endJob()`

13 5. `endRun(art::Run const&)`

14 6. `endSubRun(art::SubRun const&)`

15 As you gain experience, you will gain proficiency at knowing *when* to provide them.

16 You will also be introduced to the classes

17 1. `art::RunID`

18 2. `art::Run`

19 3. `art::SubRunID`

20 4. `art::SubRun`

21 that are analogous to the `art::EventID` and `art::Event` classes that you have al-
22 ready encountered.

13.4 Setting up to Run this Exercise

Follow the instructions in Chapter 11 if you are logging in after having closed an earlier session. If you are continuing on directly from the previous exercise, keep both your source and build windows open.

In your source window, look at the contents of the directory for this exercise, called `OptionalMethods`:

```
ls art-workbook/OptionalMethods
```

```
CMakeLists.txt      OptionalAnswer01_module.cc
Optional_module.cc  optionalAnswer01.fcl    optional.fcl
```

In your build window, just make sure that you are in your build directory. All the code for this exercise is already built; this happened the first time that you ran `buildtool`.

The source code for the module you will run is `Optional_module.cc` and the FHiCL file to run it is `optional.fcl`. The file `CMakeLists.txt` is identical to that used by the previous exercise since the new features introduced by this module do not require any modifications to the link list. The other two files relate to the exercise you will be asked to do in Section 13.8.

13.5 The Source File `Optional_module.cc`

In your source window, look at the source file `Optional_module.cc` and compare it to `First_module.cc`. The differences are

1. it has two new include directives, for `Run.h` and `SubRun.h`
2. the name of the class has changed from `First` to `Optional`
3. the `Optional` class declaration declares three new member functions

```

1     void beginJob () override;
2     void beginRun ( art::Run const& run ) override;
3     void beginSubRun( art::SubRun const& subRun ) override;

```

- 4 4. the text printed by the constructor and analyze member functions has changed
- 5 5. the file contains the definitions of the three new member functions, each of which
- 6 simply makes some identifying printout

7 13.5.1 About the `begin*` Member Functions

8 The optional member functions `beginJob`, `beginRun` and `beginSubRun`, described
9 in the Introduction to this chapter (Section 13.1), must have exactly the argument list
10 prescribed by `art` as shown in list item 3 above.

11 `art` knows to call the `beginJob` member function of each module, if present, once at
12 the start of the job; it knows to call `beginRun`, if present, at the start of each run and,
13 likewise, `beginSubRun` at the start of each subRun.

14 13.5.2 About the `art::*ID` Classes

15 In Section 10.6.3.7 you learned about the class `art::EventID`, which describes the
16 three-part event identifier. `art` also provides two related classes:

- 17 ○ `art::RunID`, a one-part identifier for a run number
- 18 ○ `art::SubRunID`, a two-part identifier for a subRun

19 The header files for these classes are found at:

```

20 $ART_INC/art/Persistency/Provenance/RunID.h
21 $ART_INC/art/Persistency/Provenance/SubRunID.h

```

22 Similar to the `art::Event` class discussed in Section 10.6.3.6, `art` provides `art::Run`
23 and `art::subRun`. These contain the IDs, e.g., `art::RunID`, plus the data products
24 for the entire run or subRun, respectively. You can find their header files at:

```

25 $ART_INC/art/Framework/Principal/Run.h
26 $ART_INC/art/Framework/Principal/SubRun.h


```


1 In the call to `beginSubRun` the argument is of type `art::SubRun const&`. A sim-
2 plified description of this object is that it contains an `art::SubRunID` plus a collection
3 of data products that describe the subRun. All of the comments about the class `art::Run`
4 in the preceding few paragraphs apply to `art::SubRun`. You can find the header file for
5 `art::SubRun` at:

```
6 less $ART_INC/art/Framework/Principal/SubRun.h
```

7 13.5.3 Use of the `override` Identifier

8 The `override` identifier on each of these member functions instructs the compiler to
9 check that both the name (and spelling) of the member function and its argument list are
10 correct; if not, the compiler will issue an error message and stop. This is a very handy
11 feature. Without it, a misspelled function name or incorrect argument list would cause the
12 compiler to assume that you intended to define a *new* member function unrelated to one of
13 these optional *art*-defined member functions. This would result in a difficult-to-diagnose
14 run-time error: *art* would simply not recognize your member function and would never
15 call it.

16 Always provide the `override` identifier when using any of the optional *art*-defined
17 member functions. 

18 For those with some C++ background, the three member functions `beginJob`,
19 `beginRun` and `beginSubRun` are declared as virtual in the base class,
20 `art::EDAnalyzer`. The `override` identifier is new in C++-11 and will not be de-
21 scribed in older text books. It instructs the compiler that this member function is intended
22 to override a virtual function from the base class; if the compiler cannot find such a func-
23 tion in the base class, it will issue an error. 

24 13.5.4 Use of `const` References

25 In `Optional_module.cc` the argument to the `beginRun` member function is a `const`
26 reference to an object of type `art::Run` that holds the current run ID and the collection
27 of data products that together describe the run. If you take a snapshot of a running *art* job
28 you will see that, at any time, there is exactly one object of type `art::Run`. This object is

1 owned by *art*. *art* gives modules access to it when it (*art*) calls the modules' `beginRun`
 2 and `endRun` member functions.

3 Because the object is passed by reference, the `beginRun` member function does not
 4 get a copy of the object; instead it is given access to it. Because it is passed by `const`
 5 reference in this example, your analyzer module may look at information in the object
 6 but it may not add or change information to the `art::Run` object. **FIXME:** *But the*
 7 *analyze function takes `art::Event` as an argument, not `art::Run`. I need some connecting*
 8 *information. FIXME: Does it have to be passed by `const ref` for every analyzer but not for*
 9 *a producer module, for example?*

10 There is a very important habit that you need to develop as a user of *art*. Many member
 11 functions in *art*, in the Workbook code and very likely in your experiment's code, will
 12 return information by `&` or by `const&`. If you receive these by value, not by reference,
 13 then you will make copies that waste both CPU and memory; in some cases these can be
 14 significant wastes. Unfortunately there is no way to tell the compiler to catch this mistake.
 15 The only solution is your own vigilance.

16 To access the `art::Run` and `art::SubRun` objects through, for example, an `art::Event`
 17 named `event`, you can use

```
18 art::SubRun const& subRun = event.getSubRun();
```

19 for the `subRun` and

```
20 art::Run const& run = subRun.getRun();
```

21 for the `run`.


22 13.5.5 The analyze Member Function

23 In your `analyze` member function, if you have an `art::Event`, named `event`, you
 24 can access the associated run information by:

```
25 art::Run const& run = event.getRun();
```

26 You may sometimes see this written as:

```
27 auto const& run = event.getRun();
```

1 Both versions mean exactly the same thing. When a type is long and awkward to write,
2 the `auto` identifier is very useful; however it is likely to be very confusing to beginners. 
3 When you encounter it, check the header files for the classes on the right hand side of the
4 assignment; from there you can learn the return type of the member function that returned
5 the information.

6 13.6 Running this Exercise

7 Look at the file `optional.fcl`. This FHiCL file runs the module `Optional` on the the
8 input file `inputFiles/input03.art`. Consult Table 9.1 and you will see that this file
9 contains 15 events, all from run 3. It contains events 1 through 5 from each of subRuns 0, 1
10 and 2. With this knowledge, and the knowledge of the source file `Optional_module.cc`,
11 you should have a clear idea of what this module will print out.

12 In your build directory, run the following command

```
13 art -c fcl/OptionalMethods/optional.fcl >& output/optional.log
```

14 The part of the printed output that comes from the module `Optional` is given in List-
15 ing 13.1. Is this what you expected to see? If not, understand why this module made the
16 printout that it did. If you did not get this printout, double check that you followed the
17 instructions carefully; if that still does not fix it, ask for help (see Section 3.4).

Listing 13.1: Output from `Optional_module.cc` with `optional.fcl`

```
18 Hello from Optional::constructor.  
19 Hello from Optional::beginJob.  
20 Hello from Optional::beginRun: run: 3  
21 Hello from Optional::beginSubRun: run: 3 subRun: 0  
22 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 1  
23 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 2  
24 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 3  
25 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 4  
26 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 5  
27 Hello from Optional::beginSubRun: run: 3 subRun: 1  
28 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 1  
29 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 2  
30 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 3  
31 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 4  
32 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 5  
33 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 5
```

```

1 Hello from Optional::beginSubRun: run: 3 subRun: 2
2 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 1
3 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 2
4 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 3
5 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 4
6 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 5

```

7 13.7 The Member Function `beginJob` versus the Constructor

8 `methods:beginRun`

9 The member function `beginJob` gets called once at the start of the job. The constructor
10 of the each module is also called once at the start of the job. This brings up the question:
11 What code belongs in the constructor and what code belongs in the `beginJob` member
12 function?

13 A small number of things must be done in the constructor — see below. Other tasks can be
14 done in either place but most experiments have found it useful to follow the rough guide-
15 line that you should put initializers and code related to *art* bookkeeping in the constructor
16 and that you should put physics-related code in `beginJob`. Hopefully the meaning of
17 this advice will become clear as you work through the Workbook. Your experiment may
18 have additional, more specific, guidelines.

19 The correct place to initialize data members is in the constructor and, whenever possible,
20 you should use the initializer list syntax. Never defer initialization of a data member to the
21 `beginJob` member function or later. When you encounter producer modules, you will
22 learn about some more tasks that must be performed in the constructor. *This chapter has*
23 *not yet been written.*



26 For those of you familiar with ROOT, we can provide an example of something physics-
27 related. You should create histograms, ntuples and trees in one of the `begin` member func-
28 tions, not in the constructor. In many cases you can create them in `beginJob` but there
29 are cases in which you will need to defer creation until `beginRun` or `beginSubRun`.
30 For example, conditions data is intrinsically time dependent and may not be available
at `beginJob`-time. If creating a histogram requires access to conditions information you
will need to create that histogram in `beginRun`, or `beginSubRun`, not in `beginJob`.

13.8 Suggested Activities

13.8.1 Add the Matching `end` Member functions

`art` defines the following three member functions:

```
void endJob      () override;
void endRun      ( art::Run const& run      ) override;
void endSubRun   ( art::SubRun const& subRun ) override;
```

Go to your source window. In the file `Optional_module.cc`, add these member functions to the declaration of the class `Optional` and provide an implementation for each. In your implementation, just copy the printout created in the corresponding `begin` function and, in that printout, change the string “begin” to “end”.

Then go to your build window and make sure that your current directory is your build directory. Then rebuild this module and run it:

```
buildtool
```

```
art -c fcl/OptionalMethods/optional.fcl >& output/optional2.log
```

Consult Chapter 10 if you need to remember how to identify that the build completed successfully. Compare the output from this run of `art` with that of the previous run: do you see the additional printout from the member functions that you added?

The solution to this activity is provided as the file `OptionalAnswer01_module.cc`. It is already built. You can run it with:

```
art -c fcl/OptionalMethods/optionalAnswer01.fcl >& output/optionalAnswer01.log
```

Does the output of your code match the output from this code?

13.8.2 Run on Multiple Input Files

In a single run of `art`, run your modified version of the module `Optional` on all of the three of the following input files:

```
inputFiles/input01.art
inputFiles/input02.art
```

1 `inputFiles/input03.art`

2 If you need a reminder about how to tell *art* to run on three input files in one job, consult
3 [Section 9.8.5](#).

4 Make sure that the printout from this job matches the description of the event loop found
5 in [Section 3.6.2](#).

6 **13.8.3 The Option `--trace`**

7 The *art* command supports a command line option named `--trace`. This creates addi-
8 tional printout that identifies every step in the event loop. Use this option to trace what *art*
9 is doing when you run this exercise. For example

```
10 art -c fcl/OptionalMethods/optional.fcl --trace >& output/trace.log
```

11 You should be able to identify your printout among the printout from *art* and see that your
12 printout appears in the expected place.

13 When you are getting an error from *art* and you don't understand which module is causing
14 the problem, you can use `--trace` to narrow your search.

15 **13.9 Review**

16 **FIXME:** *To be added; list from 'what you will learn' repeated here for reference*

17 1. `beginJob()`

18 2. `beginRun(art::Run const&)`

19 3. `beginSubRun(art::SubRun const&)`

20 4. `endJob()`

21 5. `endRun(art::Run const&)`

22 6. `endSubRun(art::SubRun const&)`

23 As you gain experience, you will gain proficiency at knowing *when* to provide them.

24 You will also be introduced to the classes

- 1 1. `art::RunID`
- 2 2. `art::Run`
- 3 3. `art::SubRunID`
- 4 4. `art::SubRun`

5 13.10 Test Your Understanding

6 13.10.1 Tests

7 Two files are provided, one of which intentionally contain bugs. The module should com-
8 pile and run, but some printout will be missing. The files are:

9 `OptionalMethods/OptionalBug01_module.cc`

10 `OptionalMethods/bug01.fcl`

11 Your job is to figure out what's missing and why, then fix it. The answer is provided at the
12 end.

13

Go to your build window and run *art*:

```
art -c fcl/OptionalMethods/bug01.fcl >& output/your-file-name.log
```

Once you figure out the problem, go to your source window and fix it. If it turns out to be in the module file, you will need to go to your build window and run:

```
buildtool
```

before you rerun the *art* command.

14

15 The answer is intentionally placed on a new page (remember to try before you read fur-
16 ther!).

1 **13.10.2 Answers**

2 The code compiles and runs but the printout from the `beginRun` member function is
3 missing. Why? Because `beginRun` was misspelled as `beginrun` (lower case R) in
4 both the declaration and the definition. They matched, but they weren't recognized by
5 the `std::cout` statement that goes with the definition.

6 **FIXME:** *Any reason the `beginRun` doesn't have a `const&` like the others?*

7 The warning that we tripped over in the Exercise 2 test only comes up if the name of the
8 member function is exactly the same and the argument list is different.

14 Exercise 4: A First Look at Parameter Sets

14.1 Introduction

In the previous few chapters you have used FHiCL files to configure *art* jobs. From Section 9.8 recall the definition of a FHiCL *table*: it is a group of FHiCL definitions delimited by braces { }. When *art* reads its run-time configuration FHiCL file, it transforms the FHiCL file into a C++ representation; in that representation, each FHiCL table becomes an object of type `fhiicl::ParameterSet`, which we refer to as a *parameter set*(γ).

Among other things, you have learned how to define a module label and its corresponding parameter set, the simplest case looking like:

```
moduleLabel : {  
    module_type : ClassName  
}
```

where the `moduleLabel` is an identifier that you define and `ClassName` is the name of a module class. *art* requires that the `module_type` parameter be present.

When you define a module label, you may enter additional FHiCL definitions (i.e., *parameters*) between the braces to form a larger parameter set. For example:

```
moduleLabel : {  
    module_type      : ClassName  
    thisParameter   : 1  
    thatParameter   : 3.14159  
    anotherParameter : "a_string"  
    arrayParameter  : [ 1, 3, 5, 7, 11 ]  
    nestedPSet      : {  
        // ...  
    }  
}
```

```

8         a : 1
9         b : 2
10        }
11 }

```

5 This functionality allows you to write modules whose behaviour is run-time configurable.
6 For example, if you have a reconstruction algorithm that depends on some cuts, the values
7 of those cuts can be provided in this way.

8 14.2 Prerequisites

9 The prerequisite for this chapter is all of the material in Part I (Introduction) and the ma-
10 terial in Part II (Workbook) up to and including Chapter 10. You can read this chapter
11 without necessarily having read Chapter 12 or 13.

12 14.3 What You Will Learn

13 In Section 10.6.3.4 you saw that the constructor of a module is required to take an argument
14 of type `fhiicl::ParameterSet const&`.

15 In this chapter you will learn how to use this argument to read additional parameters in
16 a parameter set. In particular, you will learn about the class `fhiicl::ParameterSet`
17 and after working through the exercises in this section, you should know how to:

- 18 1. read parameter values from a FHiCL file into a module
- 19 2. require that a particular parameter be present in a parameter set
- 20 3. use data members to communicate information from the constructor to other mem-
21 ber functions of a module
- 22 4. print a parameter set
- 23 5. use the colon initializer syntax
- 24 6. provide a default value for a parameter (if the parameter is absent from a parameter
25 set)
- 26 7. modify the precision of the printout of floating point types

- 1 8. recognize the error messages for a missing parameter or for a value that cannot be
2 converted to the requested type

3 You will also learn:

- 4 1. that you should find out your experiment's policy about what sorts of parameters are
5 allowed to have default values
- 6 2. an extra parameter automatically added by *art*, but only in parameter sets that are
7 used to configure modules
- 8 3. the canonical forms of parameters

9 Finally, you will learn a small amount about C++ templates and C++ exceptions, just
10 enough to understand the exercise.

11 14.4 Setting up to Run this Exercise

setting-up

12

Follow the instructions in Chapter 11 if you are logging in after having
closed an earlier session. If you are continuing on directly from the previ-
ous exercise, keep both your source and build windows open.

In your source window, look at the contents of the directory for this exer-
cise, called `ParameterSets`:

```
ls art-workbook/ParameterSets
```

```
bug01.fcl          pset03.fcl          pset07.fcl
bug02.fcl          PSet03_module.cc
PSet07_module.cc.nobuild
bug03.fcl          pset04.fcl          pset08.fcl
CMakeLists.txt    PSet04_module.cc
PSet08_module.cc
pset01.fcl         pset05.fcl          pset09.fcl
PSet01_module.cc  PSet05_module.cc.nobuild
PSet09_module.cc
pset02.fcl         pset06.fcl
PSet02_module.cc  PSet06_module.cc.nobuild
```

13

`PSet:pset01:fcl`**Listing 14.1:** Parameter set `psetTester` from `pset01.fcl`

```

1  analyzers: {
2    psetTester : {
3      module_type : PSet01
4      a : "this_is_quoted_string"
5      b : 42
6      c : 3.14159
7      d : true
8      e : [ 1, 2, 3 ]
9      f : {
10     a : 4
11     b : 5
12   }
13 }
14 }

```

In your build window, just make sure that you are in your build directory. All the code for this exercise is already built; this happened the first time that you ran `buildtool`.

1

2 The source code for the first module you will run is `PSet01_module.cc` and the FHiCL
 3 file to run it is `pset01.fcl`. The file `CMakeLists.txt` is identical to that used by the
 4 previous two exercises. The remaining files are the source and FHiCL files for additional
 5 steps in this exercise.

6 14.5 The Configuration File `pset01.fcl`

`-sets-pset01-fcl`

7 The FHiCL file that you will run in this exercise is `pset01.fcl`. Look at this file in your
 8 source window. You will see that `pset01.fcl` defines a parameter set `psetTester`,
 9 shown in Listing 14.1, that configures an analyzer module named `PSet01`.

10 The parameter `module_type` is processed by *art*. All of the other parameters are pro-
 11 cessed by code in the module class `PSet01`. Additional definitions like these in a FHiCL
 12 file have the following properties:

- 13 1. The module specified by the `module_type` parameter defines which parameters
 14 must be present in this list, and which parameters are optional.

- 1 2. Each definition must be a legal FHiCL definition.
- 2 3. These definitions have no meaning, per se, to *art* or to FHiCL; they only have mean-
3 ing to the C++ code in `PSet01_module.cc`.
- 4 4. Each definition may use the full power of FHiCL and may contain nested parameter
5 sets to arbitrary depth.

6 Looking at the parameter set, it appears that the parameter `a` has a value that is a string
7 of text, parameter `b`'s value is an integer number, parameter `c` is a floating point number,
8 parameter `d` is one of the two possible boolean values, parameter `e` is an array of integers
9 and that parameter `f` is a nested parameter set. You will learn in Section 14.6 that, from
10 the point of view of the code in `PSet01_module.cc`, this intuition is correct. But there
11 is one subtlety: FHiCL itself has no notion of *type* and, inside FHiCL, all parameter values
12 are just strings. The interpretation of a parameter value as a particular type is done by
13 code inside `PSet01_module.cc`. The computer-science-speak for this is that FHiCL is
14 a *type-free* language; this is in contrast to C++ which is a *strongly-typed* language.

15 14.6 The Source code file `PSet01_module.cc`

16 The source code for this exercise is found in the file `PSet01_module.cc`. The new
17 features seen in this exercise are all in the definition of the constructor.

18 When *art* starts up, it reads the file `pset01.fcl` and, among many other things, copies
19 the FHiCL table `psetTester` into an object of type `fhiCL::ParameterSet`. When
20 *art* calls the constructor of `PSet01`, it passes this `fhiCL::ParameterSet` as the
21 argument of the constructor, named `pset`. That is, the table named `psetTester` in
22 the FHiCL file appears in the module as a parameter set named `pset`.

23 Let's examine the first part of the constructor; see Listing 14.2.

```

24 1  tex::PSet01::PSet01(fhiCL::ParameterSet const& pset ):
25 2    art::EDAnalyzer(pset) {
26
27 4    std::string      a=pset.get<std::string>("a");
28 5    int              b=pset.get<int>    ("b");
29 6    double           c=pset.get<double>("c");
30 7    bool             d=pset.get<bool>  ("d");

```

```

1 8     std::vector<int>      e=pset.get<std::vector<int>>("e");
2 9     fhicl::ParameterSet f=pset.get<fhicl::ParameterSet>("f");
3
4 11     int                fa=f.get<int>("a");
5 12     int                fb=f.get<int>("b");
6
7 14     std::string module_type =
8 15                pset.get<std::string>("module_type");
9
10 17     std::string module_label =
11 18                pset.get<std::string>("module_label");

```

Listing 14.2: First part of constructor in PSet01_module.cc

12 Recall from Section [14.5](#) that the object `pset` internally represents the value of each
13 parameter as a string. If you ask that the value of a parameter be returned as a string, `pset`
14 will simply return a copy of its internal representation of that parameter. On the other hand,
15 if you ask that the value of a parameter be returned as any other type, then `pset` needs to
16 do some additional work. For example, if you ask that a parameter be returned as an `int`,
17 then `pset` must first find its internal string representation of that parameter; it must then
18 convert that string into a temporary variable of the requested type and return the temporary
19 variable. Therefore, when your code asks `pset` to return the value of a parameter, it must
20 tell `pset` two things:

- 21 1. the name of the parameter
- 22 2. the type to which the string representation should be converted

23 The angle bracket syntax `<>` is the signature of a feature of C++ called *templates*(γ). *art*
24 and *FHiCL* use templates in several prominent places. You do not need to fully understand
25 templates — just how to use them when you encounter them. The following pages describe
26 how to use templates when getting the value of a parameter from `pset`.

27 When you ask for the value of a parameter, the name of the parameter is specified as a
28 familiar function argument while the return type is specified between the angle brackets.
29 The name between the angle brackets is called a *template argument*. If you do not supply
30 a template argument, then your code will not compile.

1 For example, look at the line that reads the parameter `a`:

```
2 4 std::string a = pset.get<std::string>("a");
```

3 It first declares a local variable named `a` that is of type `std::string` and then asks
4 `pset` to do the following:

- 5 1. Check if it has a parameter named `a`.
- 6 2. If it has this parameter, return it as a string.


7 The returned value is used to initialize the local variable, `a`. Section 14.13 will describe
8 what happens if `pset` does not have a parameter named `a`.

9 It is not required that the local variable, `a`, have the same name as the FHiCL parameter
10 `a`. But, with rare exceptions, it is a good practice to make them either exactly the same or
11 close to the same.

12 The following line, that sets the parameter `b`,

```
13 5 int b = pset.get<int>("b");
```

14 is similar to the previous line; the main difference is that `pset` will convert the string to an
15 `int` before returning it. `pset` knows that it must perform the conversion to `int` because
16 the template argument tells it to. Section 14.13 will describe what happens if the string
17 cannot be converted to an `int`.

18 It is beyond the scope of this chapter to discuss how the template mechanism is used
19 to trigger automatic type conversions. It is sufficient to remember the following: when
20 you use the `get` member function of the class `fhiCL::ParameterSet`, the template
21 argument must always match the type of the variable on the left-hand side. Templates will
22 be discussed in Section 14.8. 

23 The authors of FHiCL could have designed a different interface, such as:

```
24 std::string a = pset.get_as_string("a");  
25 std::string b = pset.get_as_int ("b");
```

26 Instead they chose to write it using templates. The reason for this choice is that it allows
27 one to add new types to FHiCL without needing to recompile FHiCL. How you do this is

1 beyond the scope of this chapter. You now know everything that you need to know about
 2 templates in order to use `fhicl::ParameterSet` effectively.

3 The rest of the lines in the section of code shown in Listing 14.2 extract the remaining
 4 parameters from `pset` and make copies of them in local variables. The remainder of the
 5 constructor, shown in Listings 14.3, 14.4 and 14.5, prints the values of these parameters.
 6 The output, split into parts 1, 2 and 3, is shown in Section 14.7, Listing 14.6.

```

7 1   std::cout << "\n-----\nPart 1:\n";
8 2   std::cout << "a : " << a << std::endl;
9 3   std::cout << "b : " << b << std::endl;
10 4   std::cout << "c : " << c << std::endl;
11 5   std::cout << "d : " << d << std::endl;
12
13 7   std::cout << "e :";
14 8   for ( int i: e ){
15 9       std::cout << " " << i;
16 10  }
17 11  std::cout << std::endl;
18
19 13  std::cout << "f.a : " << fa << std::endl;
20 14  std::cout << "f.b : " << fb << std::endl;
21
22 16  std::cout << "module_type: " << module_type << std::endl;
23 17  std::cout << "module_label: " << module_label << std::endl;

```

Listing 14.3: Part 1 of the remainder of the constructor in `PSet01_module.cc`. The values of `pset` are printed out in a standard way. Two other ways of printing the values `a` and `b` from the parameter set `f` (lines 13 and 14) will also be shown.

```

24 1   std::cout << "\n-----\nPart 2:\n";
25 2   std::cout << "f as string:          "
26 3       << f.to_string()
27 4       << std::endl;
28 5   std::cout << "f as indented-string:\n"
29 6       << f.to_indented_string()

```

```
1 7         << std::endl;
```

Listing 14.4: Part 2 of the remainder of the constructor in `PSet01_module.cc`. These lines use the `to_string()` and `to_indented_string()` member functions of the class `fhiicl::ParameterSet` to print the values `a` and `b` from the parameter set `f`.

-listing2c

```
2 1     std::cout << "\n-----\nPart 3:\n";
3 2     std::cout << "pset:\n"
4 3         << pset.to_indented_string()
5 4         << std::endl;
```

Listing 14.5: Part 3 of the remainder of the constructor in `PSet01_module.cc`. This portion uses the `to_indented_string()` member function to print everything found in the parameter set `psetTester`, including `a` and `b` from the parameter set `f`.

6 Your code may ask for the values of parameters from a `ParameterSet` in any order, and
7 any number of times, including zero.

8 We offer two final comments on `PSet01_module.cc`. First, the `analyze` member
9 function is empty. Nevertheless, it must be present because *art* requires all analyzer mod-
10 ules to provide a member function named `analyze`. If we removed this member function
11 from the class `PSet01`, then the module would not compile. Second, the argument of the
12 `analyze` member function is not used; therefore it is not given a name. Were it given
13 a name, the compiler would complain that the argument was never used. When no name
14 is given the compiler understands that it is your intention not to use the argument. Even
15 though the code does not use the argument, its type must be present because the number,
16 type and order of the arguments are all parts of the signature of a function.

17 14.7 Running the Exercise

running-it

18 Now let's see what happens when you run the job. In your build directory, run the following
19 command

```
20 art -c fcl/ParameterSets/pset01.fcl >& output/pset01.log
```

21 The expected output from this command is shown in Listing 14.6.

22 The module reads in the parameter set and then prints out each of the values in several
23 different ways as explained in Section 14.6. Check that the printout matches the definitions

- 1 of the parameters from `pset01.fc1`. Understand the relationship between the printout
- 2 and the lines in the source file `PSet01_module.cc`.

t01:output

Listing 14.6: Output from PSet01 with `pset01.fcl` (*art*-standard output not shown)

```
-----  
Part 1:  
a : this is quoted string  
b : 42  
c : 3.14159  
d : 1  
e : 1 2 3  
f.a : 4  
f.b : 5  
module_type: PSet01  
module_label: psetTester  
  
-----  
Part 2:  
f as string:          a:4 b:5  
f as indented-string:  
a: 4  
b: 5  
  
-----  
Part 3:  
pset:  
a: "this_is_quoted_string"  
b: 42  
c: 3.14159  
d: true  
e: [ 1  
    , 2  
    , 3  
  ]  
f: { a: 4  
     b: 5  
   }  
module_label: "psetTester"  
module_type: "PSet01"
```

14.8 Member Function Templates and their Arguments

Now that you have seen templates, we can introduce some more language that you will need to know. In the above examples, `get<std::string>` and `get<int>` are *member functions* of the class `ParameterSet`.

On its own, `get` is called a *member function template*; this means that `get` is a set of rules to write a member function. The member function can only be written once the template's argument has been specified. In the future, when we refer to `get`, we will call it either by its proper name:

```
ParameterSet::get<T>
```

or by the abbreviation `get<T>`. In the notation `<T>`, the angle brackets indicate that `get` is a template and the capital letter `T` is a dummy argument that indicates that if you want to use the template, you must supply one template argument. The choice of the letter `T` as the name of the dummy argument is a mnemonic for *Type*, indicating that the template argument is usually the name of a type.¹



If you are familiar with template meta-programming you can find the source for the class `fhicl::ParameterSet` in the files:

```
$FHICL_CPP_INC/fhiclcpp/ParameterSet.h
$FHICL_CPP_DIR/source/fhiclcpp/ParameterSet.cc
```

In particular, this is where you can find the source for `ParameterSet::get<T>`.

14.8.1 Types Known to `ParameterSet::get<T>`

This section describes the different types that can be used as the template argument for `ParameterSet::get<T>`. If you use `ParameterSet::get<T>` “out of the box,” it supports the following types.

- For a parameter that has a simple value, `get<T>` supports: `bool` and `std::string`; any C++ built-in integral type, such as `int`, `unsigned` or `short`; any C++ built-in floating point type, such as `float` or `double`;

¹ Much later in the workbook, you will see one case in which it is something other than the name of a type. **FIXME:** Refer to this section once it is written

- 1 ○ For a parameter whose value is another parameter set, T must be `fhicl::ParameterSet`.
- 2 ○ For a parameter with a value that is a sequence of items, all items in the sequence
- 3 must be of the same type and `get<T>` allows T to be `std::vector<S>`, where
- 4 the template argument S is any of the types given in the previous two bullets.

5 14.8.2 User-Defined Types

6 You can write helper functions that will allow the type T to be almost any type that you

7 might want. How to do this is beyond the scope of this chapter. For an example, see the

8 files `ParameterSetHelpers.h` and `ParameterSetHelpers.cc` under

9 `$TOYEXPERIMENT_DIR/source/toyExperiment/Utilities/`

10 These files allow you to define a FHiCL parameter as:

```
11 zaxis : [ 0., 0., 1.]
```

12 and to read it as

```
13 auto zaxis = pset.get<CLHEP::Hep3Vector>("zaxis");
```



14 14.9 Exceptions (as in “Errors”)

15 14.9.1 Error Conditions

16 There are two sorts of error conditions that may occur when reading parameters from a

17 parameter set:

- 18 1. The requested parameter is not present in the parameter set.
- 19 2. The requested parameter is present but cannot be converted into the requested type.

20 To give an example of the second sort, suppose that on line 6 of Listing 14.1 you change the

21 FHiCL definition of the parameter `c` from `3.14159` to the string `"test"`. Now consider

22 what happens when you try to read this parameter as a double, as is done on line 6 of

23 Listing 14.2:

```
24 double c=pset.get<double>("c");
```

- 1 The code will correctly find that parameter `c` exists but it will produce an error when it
- 2 tries to convert the string "test" to a double.
- 3 In both cases, the code inside `pset` will tell `art` to stop processing events and to perform
- 4 an orderly shutdown, which will be described in Section [14.9.2](#).

5 14.9.2 Error Handling

6 From time to time code within `art` will discover that, because of some error condition,
7 it cannot continue to process events. When this happens `art` can be configured to stop
8 processing events and then to do one of several different things:

- 9 1. It can attempt an orderly shutdown, described below.
- 10 2. It can write the offending event to a separate output file and continue normally with
- 11 the next event.
- 12 3. It can skip the module in which the problem occurred and continue normal process-
- 13 ing with the next module.

14 There are several other options that cannot be described here because the necessary back-
15 ground information has not yet been established.

16 **FIXME:** *When more material on exception handling is in place, add a link to it from*
17 *here.*

18 When `art` attempts an orderly shutdown, it will:

- 19 1. write a message to the log file that describes what happened
- 20 2. record the error condition that stopped processing; this information will be written to
- 21 all output event-data files **FIXME:** *Check with Jim, Marc r Chris; I am not certain*
- 22 *what gets written to the event-data file.*
- 23 3. call the `endSubRun` member function of every module
- 24 4. call the `endRun` member function of every module
- 25 5. call the `endJob` member function of every module
- 26 6. properly flush and close all output and log files


1 7. perform a few other clean-up and shutdown actions for parts of *art* that have not yet
2 been discussed

3 8. return a nonzero status code to the parent process (the status code is the number that
4 appears on the last line of your *art* output, beginning with "Art has completed ...")

5 For most sorts of errors, the orderly shutdown will be successful and your work up to the
6 error will be preserved. However, there are circumstances for which the orderly shutdown
7 will fail, for example when there is no disk space to hold more output.

8 For all exception cases but one, *art*'s default behavior is to attempt an orderly shutdown.
9 The inability of *art* to find a requested data product is the nonstandard case; when this
10 occurs *art* simply continues with the next module.

11 These default behaviors can be changed in the FHiCL file. *When the section that describes*
12 *how to do this is written, a link to that section will be added here.*

13 The technology that *art* uses to interrupt event processing and to take one of the possible
14 follow-on actions is a feature of C++ called *exceptions*. When *art* stops event processing
15 and takes the appropriate follow-up action it is said to *throw an exception*; this phrase will 
16 be used throughout the Workbook. The topic of exceptions is complex and beyond the
17 scope of this chapter. *A chapter yet to be written will describe how to use exceptions in*
18 *your own code to tell art to interrupt processing.*

19 14.9.3 Suggested Exercises

20 In `pset01.fcl`, remove the definition of the parameter `b`. Rerun *art*. You should see
21 an error message like that shown in Listing 14.7. Read the error message and understand
22 what it is telling you. It is important to recognize the error message in case you make this
23 mistake in the future.

Listing 14.7: Output from `PSet01` with `pset01.fcl` (parameter `b` removed)

```
24 %MSG-s ArtException: PSet01:psetTester@Construction (date time)
25 ModuleConstruction
26 cet::exception caught in art
27 ---- Can't_find_key_BEGIN
28   _b
29 ----_Can't find key END
30 %MSG
31 Art has completed and will exit with status 8001.
```

1 Note, too, that the completion status is nonzero.

2 In `pset01.fcl`, restore the definition of `b` and change the definition of `c` to `"test"`.

3 Rerun `art`. You should see an error message like that shown in Listing 14.8. Again, read
4 the error message and understand it.

Listing 14.8: Output from `PSet01` with `pset01.fcl` (parameter `c` misdefined)

```

5 %MSG-s ArtException: PSet01:psetTester@Construction (date time)
6 ModuleConstruction
7 cet::exception caught in art
8 ---- Type mismatch BEGIN
9   c
10  ---- Type mismatch BEGIN
11     error in float string:
12     test
13     at or before:
14  ---- Type mismatch END
15 ---- Type mismatch END
16 %MSG
17 Art has completed and will exit with status 8001.
```

18 14.10 Parameters and Data Members

19 Information from the parameter set is often needed in a member function of the module
20 class. This information is propagated from the parameter set to the member function by
21 storing the values of these parameters as data members of the module class. This is illus-
22 trated in the two files `PSet02_module.cc` and `pset02.fcl`. Open these files with
23 an editor and follow along with the description below.

24 If you need to refamiliarize yourself with the concept of data members of a class, refer to
25 Section 6.7.2.

26 There are three things to notice in `PSet02_module.cc`.

- 27 1. The class declares three data members named `b_`, `c_`, and `f_`. These are declared
28 in the private section so that only the module itself can see them.
- 29 2. In the constructor, these three data members are initialized to values extracted from
30 the module's parameter set.
- 31 3. In the `analyze` member function all three data members are printed out.

1 If you need to refamiliarize yourself with the colon initializer syntax, refer to Section 6.7.5,
 2 or with the conventions about underscore characters in the names of data members, refer
 3 to Section 6.7.7.2 covers

4 To run this example, enter

```
5 art -c fcl/ParameterSets/pset02.fcl >& output/pset02.log
```

6 The expected output from this is given in Listing 14.9.

Listing 14.9: Output from PSet02 with pset02.fcl

```
7 Event number: run: 1 subRun: 0 event: 1 b: 42 c: 3.14159 f: a:4 b:5  

8 Event number: run: 1 subRun: 0 event: 2 b: 42 c: 3.14159 f: a:4 b:5  

9 Event number: run: 1 subRun: 0 event: 3 b: 42 c: 3.14159 f: a:4 b:5
```

10 This example is only relevant when parameters are actually used in member functions. If a
 11 parameter is used only inside the constructor, do not store it as a data member; instead you
 12 should store it as a local variable of the constructor. This brings up a *best practice*: always
 13 declare a variable in the narrowest scope that works.



14.11 Optional Parameters with Default Values

15 It is sometimes convenient to provide a default value for a parameter. Default values may
 16 be provided in the source code that reads the parameter set. This mechanism is illustrated
 17 by the files PSet03_module.cc and pset03.fcl. Open these files with an editor
 18 and follow along with the description below.

19 You have already seen that the member function template `ParameterSet::get<T>`
 20 takes one function argument, the name of the parameter. For example,

```
21 1 int b = pset.get<int>("b");
```

22 It also takes an optional second function argument, a default value for the parameter. For
 23 example,

```
24 1 int b = pset.get<int>("b", 0);
```

25 If the second argument is present, there two cases:

```
et:pset03:output
```

Listing 14.10: Parameter-related portion of output from PSet03 with `pset03.fcl`

```
debug level: 0
g: 1 1 1 1 1
```

- 1 1. If the parameter is not defined in the FHiCL file, then the second argument is re-
- 2 returned as the value of the call to `get`.
- 3 2. If the parameter is defined in the FHiCL file, then the second argument is ignored
- 4 and the value read from the FHiCL file is returned as the value of the call to `get`.

5 When reading the code in this example you will encounter the expression:

```
6 1 std::vector<double>(5, 1.0);
```

7 This tells the compiler to instantiate an object of type `std::vector<double>`, set its
 8 size to 5 and initialize elements 0 through 4 to have the value 1.0. If you are not familiar
 9 with this syntax, you can read about it in the documentation for the C++ Standard Library
 10 (see Section [6.9](#) sec.cpp:references).

11 This expression appears as the second argument of the second call to the member function
 12 `pset.get<T>`. Therefore the compiler will create an unnamed temporary object (the
 13 vector of doubles) and pass that object to the member function `get<std::vector<double>>`
 14 as its the second argument; the compiler ensures that, once function call has completed,
 15 the temporary object is deleted.

16 With the above explanations, the source code for this example should be reasonably self-
 17 explanatory; it looks for two parameters named `debugLevel` and `g` and supplies default
 18 values for each of them. Look at the file `pset03.fcl`; you will see that the parameters
 19 `debugLevel` and `g` are not present in the `testPSet` parameter set; therefore printout
 20 will show the default values.

21 To run this example,

```
22 art -c fcl/ParameterSets/pset03.fcl >& output/pset03.log
```

23 The expected output from this is given in Listing [14.10](#) lst:PSet:pset03:output.

24 As a suggested exercise, edit `pset03.fcl` and, in the parameter set `testPSet`, provide
 25 definitions for the parameters `debugLevel` and `g`. Make their values different from the

- 1 default values. Rerun *art* and verify that the module has correctly read in and printed out
- 2 the values you defined.

3 **14.11.1 Policies About Optional Parameters**

admonition

4 Allowing *optional* parameters is important for developing, debugging and testing; if all
5 parameters were required all of the time, the complete list of parameters could become
6 unwieldy². On the other hand, the use of optional parameters can make it difficult to audit
7 the physics content of a job. Therefore experiments typically have policies for what sorts
8 of parameters may have defaults and what sorts may not. For example, your experiment
9 may prohibit default values for parameters that define the physics behavior, but allow them
10 for parameters that control printout and other diagnostics.



11 Consult your experiment to learn what policies you should follow.

12 **14.12 Numerical Types: Precision and Canonical Forms**

sets-types

13 FHiCL recognizes numbers in both fixed point and exponential notation, for example
14 123.4 and 1.234e2; the letter *e* that separates the exponent can be written in either
15 upper or lower case.

16 In the preceding exercises you defined some numerical values in a FHiCL file, read them
17 into your code and printed them out; the printed values exactly matched the input values.
18 The values used in those exercises were carefully chosen to avoid a few surprises: there
19 are cases in which the printed value will be an equivalent, but not identical, form. This
20 section discusses those cases and provides some examples.

21 When FHiCL recognizes that a parameter value is a number it converts the number into a
22 *canonical form* and stores the canonical form as a string. The transformation to the canoni-
23 cal form preserves the full precision of the number and involves the following steps:

- 24 1. The canonical form has no insignificant characters:
 - 25 (a) no insignificant trailing zeros

²FHiCL has several features that make it easier to deal with large parameter sets. *This will be explained in a future chapter.*

Table 14.1: Canonical forms of numerical values in FHiCL files

Number	Canonical Form	Number	Canonical Form
2	2	1.234E2	1.234e2
2.	2	1.23456E5	123456
2.0	2	1.23456E6	1.23456e6
2.1E2	210	1234567	1.234567e6
+210	210	0.01	1E-2

- 1 (b) no insignificant trailing decimal point
- 2 (c) no insignificant leading plus sign
- 3 (d) no insignificant leading plus sign in the exponent
- 4 2. If a number is specified in exponential notation and if the number can be represented
- 5 as a integer without loss of precision, and if the resulting integer has 6 or fewer
- 6 digits, then the canonical form is the integer. For example, the canonical form of
- 7 `1.23456E5` is `123456` but the canonical form of `1.23456E6` is `1.23456e6`.
- 8 3. The canonical form of all other floating point numbers is exponential notation with
- 9 a single, non-zero digit to the left of the decimal point.
- 10 4. The canonical form of all strings includes beginning and ending quotes; this is true
- 11 even if the string contains no embedded whitespace or other special characters.

12 Some examples of numbers and their canonical forms are given in Table 14.1.

13 If a numerical value, when expressed as a fixed point number, has no fractional part, your
 14 code may ask for the parameter to be returned as either a floating point type (such as
 15 `double` or `float`) or as an integral type (such as `int`, `short unsigned` or `std::size_t`).
 16 For example, fourth non-blank line in the listing in Figure ?? was written

```
17 1 int b = pset.get<int>("b");
```

18 It might also have been written

```
19 1 double b = pset.get<double>("b");
```

20 which would do the expected thing: given the input from `pset01.fcl`, it would read the
 21 value `42` into a variable of type `double`.

1 On the other hand, if a numerical value, when expressed as a fixed point number, does
2 have a fractional part, you may only ask for the parameter to be returned as a floating point
3 type. If you ask for such a value as an `int`, the `ParameterSet::get<int>` member
4 function will throw an exception; similarly for all other integral types. This behavior may
5 not be intuitive: the authors of *art* could have decided, instead, to discard the fractional
6 part and return the integer part. They chose not to do this because when this situation
7 occurs, it is almost always an error.

8 **14.12.1 Why Have Canonical Forms?**

9 What is the point of having canonical forms for numbers?

10 There are times when it is necessary for *art*, or user of *art*, to ask if two parameter sets are
11 the same. An example is when looking at the processing history of a data product, which
12 includes the parameter set used to configure the module that produced the data product.
13 In this example you might wish to ask if a particular data product was created with the
14 official version of the parameter set or with an unofficial version.

15 When comparing two parameter sets a trivial, but pernicious, complication is when two
16 parameter sets differ only by meaningless differences in the representation of numbers,
17 such as “1” vs “1.0”. If one compares the canonical forms of the two parameter sets, this
18 complication is removed.

19 **14.12.2 Suggested Exercises**

20 **14.12.2.1 Formats**

21 The above ideas are illustrated by the files `PSet04_module.cc` and `pset04.fcl`. To
22 run this example,

```
23 art -c fcl/ParameterSets/pset04.fcl >& output/pset04.log
```

24 The expected output from this is given in Listing 14.11. list:PSet:pset04:output Read the source code and the
25 FHiCL file; then examine the output. The first three lines show three different printed
26 formats of the parameter `a`, with the first being the canonical form. While all forms are
27 equal to the number found in the FHiCL file, they all have different formats. Understand
28 why each line has the format it does.

et:pset04:output

Listing 14.11: Output from PSet04 with pset04.fcl

```

1 parameter a as a string: 1.23456e6
2 parameter a as a double: 1.23456e+06
3 parameter a as int:      1234560
4 parameter b as a string: 3.1415926
5 parameter b as a double: 3.14159
6 parameter b as a double with more significant figures: 3.1415926
7 parameter c as a string: 1
8 parameter c as an int:   1

```

1 Line 4 shows the canonical form of the parameter `b`. Line 5 shows what is printed using
 2 the default C++ settings; the two least significant characters were dropped. The code that
 3 produces line 6 shows how to use the `precision` function from the C++ Standard
 4 Library to tell C++ to print more significant figures.

5 If you modify the precision of `cout`, it will change the format of the printout for the rest
 6 of the job; usually this is a bad thing. To avoid this, `PSet04_module.cc` illustrates how
 7 to save and restore the precision of `cout`.

8 Line 7 shows the canonical form of the parameter `c` and line 8 shows the default C++
 9 printed form of the integer.

10 14.12.2.2 Fractional versus Integral Types

11 For the next exercise, edit `pset04.fcl` and change the value of `c` to something with a
 12 fractional part. Rerun `art`; you should see that it throws an exception because it is illegal
 13 to read a numeric value with a fractional part into a variable of integral type. The error
 14 message from `art` is shown in Listing 14.12. Read the error message and understand what
 15 it is telling you so that you will recognize the error message if you make this mistake
 16 yourself.

17 14.13 Dealing with Invalid Parameter Values

-sets-exceptions

18 **FIXME:** *This section was drafted hastily and has not be revisited.*

19 When you read a parameter from a parameter set it is always a good idea to check that
 20 the value of the parameter is within the allowed set or range of values. If a parameter has an

error:output

Listing 14.12: Output from PSet04 with modified `pset04.fcl` (intentional error)

```
%MSG-s ArtException: PSet04:pset@Construction 28-Jul-2013 23:39:31 CDT
ModuleConstruction
cet::exception caught in art
---- Type mismatch BEGIN
  c
  narrowing conversion
---- Type mismatch END
%MSG
Art has completed and will exit with status 8001.
```

1 invalid value, the *art* team recommends that you immediately tell *art* to initiate an orderly
 2 shutdown; the idea of an orderly shutdown was described in Section 14.9.2.³

3 In order to initiate an orderly shutdown, you throw an exception. An example of how to do
 4 so is found in `PSet08_module.cc` (yes, the numbers are out of order; the document
 5 is still evolving). This module is based on the file `PSet02_module.cc`; it is stripped
 6 down from three parameters to one and code has been added to the constructor to test if
 7 the member datum `weight_` is inside the allowed range.

8 The constructor argument to `art::Exception` must be one of the values defined by
 9 the enum found in the file:

10 `$ART_INC/art/Utilities/Exception.h`.

11 From that file, chose an enum value does a good job of describing the sort of issue that
 12 caused you to throw the exception. You can also append arbitrary text, as is shown in
 13 the example; make sure that the text tells users enough so that they can understand the
 14 problem. You should terminate the text with a new line character, not with `std::endl`;
 15 this is a quirk of *art*'s exception handling system.

16 Instead of using the class `art::Exception` your experiment may advise you to use ei-
 17 ther `cet::Exception` or `std::exception`. Consult with your experiment for more
 18 information; while the big picture is very similar for all three cases, the details differ.

19 To execute this example, use the file `pset08.fcl`. It will produce an error message
 20 saying that the parameter is out of bounds. You will notice that the message text you pro-

³ Others may advise you, instead, to set the value to a “sensible default”; while this will often work, when it does not, it frequently leads to subtle bugs that are difficult to notice and then are difficult to diagnose. The *art* team strongly recommends against this and in favour of an orderly shutdown.

1 vided is surrounded by framing information provided by *art*; in that framing information
2 you should identify both the module class name and module label of the offending mod-
3 ule; also identify the fragment that indicates that the error occurred during the call to the
4 constructor, not some other member function.

5 Edit the value of the `weight` parameter in `pset08.fcl` and rerun. Note the behaviour
6 when `weight` is in range, out of range or exactly at one of the limits.

7 You can read more about exceptions in any standard C++ reference.

8 A variant of this example is in `PSet09_module.cc`, which can be done executed using
9 `pset09.fcl`. In this example, both the access to and the validation of the parameter
10 `weight` are encapsulated into a free function named `validWeight`. In a small module
11 such as this there is little advantage in this encapsulation; but in large modules it can be made
12 much more understandable and much more maintainable by splitting small, self-contained
13 pieces of work into their own functions.

14 The other feature of `PSet09_module.cc` is that the helper function is put into an
15 anonymous namespace. You can read about anonymous namespaces in any standard C++
16 reference. The short version is that this is a way to tell the compiler that this function is
17 only for the use of code found in this file; it also adds some random characters to the name
18 of the function so that the linker will not get confused if someone else decides that they
19 need a function with the same name and different behaviour.

20 14.14 Review

21 **FIXME:** *To be added; list from 'what you will learn' repeated here for reference*

- 22 1. read parameter values from a FHiCL file into a module
- 23 2. require that a particular parameter be present in a parameter set
- 24 3. use data members to communicate information from the constructor to other mem-
25 ber functions of a module
- 26 4. print a parameter set
- 27 5. use the colon initializer syntax

- 1 6. provide a default value for a parameter (if the parameter is absent from a parameter
- 2 set)
- 3 7. modify the precision of the printout of floating point types
- 4 8. recognize the error messages for a missing parameter or for a value that cannot be
- 5 converted to the requested type

6 You will also learn:

- 7 1. that you should find out your experiment's policy about what sorts of parameters are
- 8 allowed to have default values
- 9 2. an extra parameter automatically added by *art*, but only in parameter sets that are
- 10 used to configure modules
- 11 3. the canonical forms of parameters

12 Finally, you will learn a small amount about C++ templates and C++ exceptions, just
13 enough to understand the exercise.

14 14.15 Test Your Understanding

15 14.15.1 Tests

16 **FIXME:** *This section was drafted hastily and has not be revisited.* The answers are inten-
17 tionally placed on a new page (remember to try before you read further!).

18 Run each of the three files `bug01.fcl`, `bug02.fcl` and `bug03.fcl`; in each case
19 there is a bug in the FHiCL. To diagnose the bug you may need to look at the module
20 source in order to determine what is expected. Understand the error message and fix the
21 bug.

22 Copy the file `PSet05_module.cc.nobuild` to `PSet05_module.cc`. In your
23 build window, run `buildtool`. This will produce an error; find and fix the error. After the
24 code builds, run it using `pset05.fcl`; verify that it behaves as expected.

25 Repeat this for the other two files ending in `.nobuild`.

1 14.15.2 Answers

2 In `bug01.fcl` the parameter `g` is expected to be a FHiCL sequence, which is defined
3 using square brackets. The error is that the FHiCL uses parentheses instead of square
4 brackets.

5 In `bug02.fcl` the parameter `debugLevel` is defined as a sequence of integers when
6 the code expects a single integer.

7 In `bug03.fcl` the error is in the definition of the parameter `f`. Reference to `PSet02_module.cc`
8 shows that the code expects `f` to be a table with two definitions, `a` and `b`. The error is
9 that the colons are missing in the definitions of `a` and `b`.

10 In `PSet05_module.cc`, the name of the data member `debuglevel_` is missing its
11 trailing underscore when it is used in the initializer list.

12 In `PSet06_module.cc`, the type `string` should be `std::string`; there are two
13 places that need to be fixed.

14 In `PSet07_module.cc`, the call to `pset.get` for the member datum `key_` has no
15 template argument; it should be C++ `<std::string>`.

15 Exercise 5: Making Multiple Instances of a Module

15.1 Introduction

In a typical HEP experiment it is often necessary to repeat one analysis several times, with each version differing only in the values of some cuts; this is frequently done to tune cuts or to study systematic errors. Very often it is both convenient and efficient to run all of the variants of the analysis in a single job.

A powerful feature of *art* is that it permits you to run an *art* job in which you define and run many instances of the same module; when you do this, each instance of the module gets its own parameter set. In this chapter you will learn how to use this feature of *art*.

15.2 Prerequisites

The prerequisite for this chapter is all of the material in Part I (Introduction) and the material in Part II (Workbook) up to and including Chapter 14, but excluding Chapter 13.

15.3 What You Will Learn

In this chapter you will learn how to run an *art* job in which you run the same module more than once. This exercise will make it clear why *art* needs to distinguish the two ideas of *module label* and `module_type`.

15.4 Setting up to Run this Exercise

Follow the instructions in Chapter 11 if you are logging in after having closed an earlier session. If you are continuing on directly from the previous exercise, keep both your source and build windows open. In your

source window, look at the contents of the directory for this exercise, called `ModuleInstances`:

```
ls art-workbook/ModuleInstances
```

```
CMakeLists.txt  magic.fcl  MagicNumber_module.cc
```

In your build window, just make sure that you are in your build directory. All the code for this exercise is already built; this happened the first time that you ran `buildtool`.

The source code for the first module you will run is `MagicNumber_module.cc` and the FHiCL file to run it is `magic.fcl`. The file `CMakeLists.txt` is identical that used by the previous two exercises.

15.5 The Source File `Magic_module.cc`

The source code for this exercise is found in the file `Magic_module.cc`. Look at this file and you should see the following features, all of which you have seen before.

1. The file declares and defines a class named `MagicNumber` that follows the rules to be an *art* analyzer module.
2. The class has a constructor and an `analyze` method.
3. The class has a data member named `magicNumber_`, of type `int`.
4. The class initializes `magicNumber_` by reading a value from its parameter `set`; the name of the parameter is `magicNumber` (without the underscore).

- 1 5. The parameter `magicNumber` is a required parameter.
- 2 6. Both the constructor and the `analyze` method print an informational message that
- 3 includes the value of `magicNumber_`.

4 15.6 The FHiCL File `magic.fcl`

labels-fcl

5 The FHiCL file used to run this exercise is `magic.fcl`. Look at this file and you should
6 see the following features:

- 7 1. Compared to previous exercises, The FHiCL names `process_name`, `source`
8 and `services` have no important differences.
- 9 2. In the `analyzers` parameter set, inside the `physics` parameter set, you will see
10 the definition of four module labels, `boomboom`, `rocket`, `flower` and `bigbird`¹.
11 The value of each definition is a parameter set.
- 12 3. The first three of these parameter sets tell `art` to run the module `MagicNumber` and
13 each provides a value for the required `magicNumber` parameter²
- 14 4. The last parameter set tells `art` to run the module `First`, the source for which was
15 discussed in Chapter 10 and is listed in Listing 10.2; this module does not need any
16 additional parameters.
- 17 5. The path `e1` contains the names of all of the module labels from the `analyzers`
18 parameter set.

19 15.7 Running the Exercise

ls-running

20 In your build directory, run the following command

```
21 art -c fcl/ModuleInstances/magic.fcl >& output/magic.log
```

¹ All of these are nicknames of ice hockey players who played for the Montreal Canadiens ice hockey team; all of them have had their sweater number retired

²In each case the magic number is the sweater number of the hockey player whose nickname is the module label.

e-labels-output

Listing 15.1: Output using `magic.fcl`

```

MagicNumber::constructor: magic number: 9
MagicNumber::constructor: magic number: 5
Hello from First::constructor.
MagicNumber::constructor: magic number: 10
MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 9
MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 5
Hello from First::analyze. Event id: run: 1 subRun: 0 event: 1
MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 10
MagicNumber::analyze: event: run: 1 subRun: 0 event: 2 magic number: 9
MagicNumber::analyze: event: run: 1 subRun: 0 event: 2 magic number: 5
Hello from First::analyze. Event id: run: 1 subRun: 0 event: 2
MagicNumber::analyze: event: run: 1 subRun: 0 event: 2 magic number: 10
MagicNumber::analyze: event: run: 1 subRun: 0 event: 3 magic number: 9
MagicNumber::analyze: event: run: 1 subRun: 0 event: 3 magic number: 5
Hello from First::analyze. Event id: run: 1 subRun: 0 event: 3

```

- 1 The expected output from this command is shown in Listing 15.1; for clarity, the printout
 2 made by *art* has been elided. Compare this printout to the printout from your run; it should
 3 be exactly the same. Inspect the printout and the files `MagicNumber_module.cc`
 4 and `../FirstModule/First_module.cc`; understand why the printout is what it
 5 is.

6 15.8 Discussion

7 15.8.1 Order of Analyzer Modules is not Important



8 As it happens, *art* runs the four analyzer modules in the order specified in the path def-
 9 inition `e1`. But you must not count on this behaviour! Two of the design rules of *art*
 10 are:

- 11 1. Modules may only communicate with each other by putting information into, and
 12 reading information from, the `art::Event`.
- 13 2. Analyzer modules may not put information into the `art::Event`.

14 Therefore *art* is free to run analyzer modules in any order.

15 For producer modules, which may add information to the event, the order of execution is

1 often very important. When you reach the exercises that run producer modules, you will
2 be told how to specify the order of execution.

3 You may wish to review some of the other ideas about *art* paths that are described in
4 Section 9.8.8.

5 **15.8.2 Two Meanings of *Module Label***

6 In the preceding discussion, the name *module label* was used in two subtly different ways,
7 as is illustrated by the module label `rocket`:

- 8 1. `rocket` identifies a parameter set that is used to configure an instance of the module
9 `MagicNumber`.
- 10 2. `rocket` is also used as the name of the module instance that is configured using
11 this parameter set; the elements in the path `e1` are all the names of module instances.

12 Clearly these two meanings are very closely related, which is why the same name, *module*
13 *label*, is used for both ideas. Throughout the remainder of this document suite the name
14 *module label* will be used for both meanings; the authors believe it will be clear from the
15 context which meaning is intended. This is standard usage within the *art* community.

16 **15.9 Review**

17 After working through this exercise, you should:

- 18 1. Understand the difference between a *module label* and `module_type`.
- 19 2. Know how to run multiple instances of the same module within one *art* job.
- 20 3. Understand that *art* does not guarantee the order in which analyzer modules will be
21 run.
- 22 4. Understand the two senses in which the name *module label* is used: as the name of
23 a parameter set and as the name of the corresponding instance of a module.

1 15.10 Test Your Understanding

2 15.10.1 Tests

3 Remember to switch from your source window to your build window as you edit files
4 and run *art*! (Need a refresher? Review the procedure under Section 10.12.) Answers are
5 provided in Section 15.10.2.

6 Edit `magic.fcl` and do the following:

- 7 1. Add a new analyzer module label that configures an instance of the module `Optional`
8 from Chapter 13.
- 9 2. Add the new module label to `e1`.

10 Then re-run

```
11 art -c fcl/ModuleInstances/magic.fcl
```

12 Do you see the expected additional printout?

13 Now run the configuration with the intentional errors.

```
14 art -c fcl/ModuleInstances/bug01.fcl
```

15 This will run but produce an error. There are two bugs. When you fix the first one and
16 rerun, a new one will pop up. Your job is to figure out each problem, in turn, and fix
17 it.

18 The answers are intentionally placed on a new page (remember to try before you read
19 further!).

15.10.2 Answers

For the first activity, you must add a new analyzer module label to `magic.fcl`, within the `analyzers` block, similar to:

```
opt : {
  module_type : Optional
}
```

After you run the `art` command again on `magic.fcl`, the output should contain sets of lines like the following; note in particular `Hello from Optional::analyze...`:

```
MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 9
MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 5
Hello from First::analyze. Event id: run: 1 subRun: 0 event: 1
MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 10
Hello from Optional::analyze. Event id: run: 1 subRun: 0 event: 1
```

Running `art` on `bug01.fcl` the first time will complete and exit with status 0, but the output will show:

```
%MSG-i DeactivatedPath: art 24-Jul-2014 09:35:52 CDT JobSetup
Detected end path "e1" which was not found in
parameter "physics.end_paths". Path will be ignored.
```

The `e1` in `end_paths` was mistyped as `e1` (lower-case L). Once corrected and rerun, the output now shows:

```
%MSG-s ArtException:
      MagicNumber:flower@Construction 24-Jul-2014 09:51:08 CDT
ModuleConstruction
cet::exception caught in art
---- Can't find key BEGIN
  magicNumber
---- Can't find key END
%MSG
Art has completed and will exit with status 8001.
```

Notice the `MagicNumber:flower@Construction` in the top line. This says that the problem is in the parameter set `flower` that is used to make an instance of the module `MagicNumber`, and that the error occurs in the constructor. Again we have a typo in `bug01.fcl`; this time the parameter name (`magicnumber`) has a lower case letter N, which should be upper case.

16 Exercise 6: Accessing Data Products

16.1 Introduction

Section 10.6.3.6 described the class `art::Event` as an `art::EventID` plus a collection of data products. The concept of a data product was described in Section 3.6.4. You have already done several exercises that made use of the `art::EventID` and in this chapter you will do your first exercises that use a data product.

16.2 Prerequisites

Prerequisites for this chapter include all of the material in Part I (Introduction) and the material in Part II (Workbook) up to and including Chapter 14.

You must also be familiar with the toy experiment described in Section 3.7.

This exercise will use class templates and member function templates in several places. The use of templates was introduced in Section 14.6. Recall that a class template is a set of rules for creating a class and that a member function template is a set of rules for creating a member function. You need to know how to use templates but you do not need to know how to write one. You will need a minimal understanding of the class template `std::vector`, which is part of the C++ Standard Library. If you understand the following four points, then you understand enough about `std::vector` for this exercise. If `t` is an object of type `std::vector<T>`, then:

1. `t` behaves much like an array of objects of type `T`. The main difference is that capacity of the array automatically grows to be large enough to hold all of the elements in the array.

- 1 2. The identifier inside the angle brackets is called a *template argument* and it is usually
2 the name of a C++ type.¹
- 3 3. The dynamic sizing occurs in the middle of a running program; not at compile time.
4
- 5 4. This expression sets `nEntries` to the number of entries in `t`:
6 `std::size_t nEntries = t.size();`
- 7 **FIXME:** Consider providing a 1 or 2 page crash course on `std::vector<T>`. Put it some-
8 where else and refer to it from here.

9 16.3 What You Will Learn

what-learn

10 In this exercise you will learn about:

- 11 1. the data type `tex::GenParticleCollection`
- 12 2. the four-part name of an *art* data product
- 13 3. the class `art::InputTag`
- 14 4. the class template `art::Handle`
- 15 5. the class template `art::ValidHandle`
- 16 6. the member function templates of `art::Event`:
17
 - `getByLabel(art::InputTag, art::Handle<T>) const;`
 - 18 ◦ `getValidHandle<T>(art::InputTag) const;`

19 16.4 Background Information for this Exercise

background

20 The input files used for the *art* workbook contain data products created by a workflow that
21 simulates the response of the toy detector to a generated event, described in Section 3.7.2.
22 The first step in this workflow is to use an event generator to create a collection of gen-
23 erated particles, which is stored in the `art::Event` as a data product. That is, there is

¹You will only see one case in the entire workbook in which it is something other than a the name of a C++ type; and this will be during a short side trip to discuss the **cetlib** utility library.

1 a single data product that holds a collection of generated particles; there is not one data
2 product per generated particle.

3 In this exercise you will retrieve this data product and print the number of generated par-
4 ticles in each event. *A future chapter will look at the properties of individual generated*
5 *particles.*

6 **16.4.1 The Data Type** `GenParticleCollection`

7 Each generated particle in the simulated event is described by an object of type
8 `tex::GenParticle`. All of the generated particles in a given event are stored in
9 an object of type `tex::GenParticleCollection`. This object is written to the
10 `art::Event` as a data product.

11 The header files that describe these two classes, `GenParticle.h` and
12 `GenParticleCollection.h`, are found under:

13 `$TOYEXPERIMENT_INC/toyExperiment/MCDataProducts/`
14


15 The content of `GenParticleCollection.h` is shown in Listing 16.1; the
16 include guards and comments have been omitted. This header uses a typedef
17 to declare that the name `tex::GenParticleCollection`; is a synonym for
18 `std::vector<tex::GenParticle>`.

```
19 1 #include "toyExperiment/MCDataProducts/GenParticle.h"
20
21 3 #include <vector>
22
23 5 namespace tex {
24
25 7     typedef std::vector<GenParticle> GenParticleCollection;
26 8 }
```

Listing 16.1: Contents of `GenParticleCollection.h`

27 Why did the authors of the workbook decide to use a typedef and not simply ask you
28 to code `std::vector<GenParticle>` when needed? The reason is future-proofing.
29 Suppose that down the road the authors find that they need to change the definition of

1 `tex::GenParticleCollection`; if you used the typedef, it is much more likely
 2 that your code will continue to compile and work correctly as is. If, on the other hand, you
 3 used `std::vector<GenParticle>`, then you would need to identify and edit every
 4 instance.

5 Please use the typedef `GenParticleCollection` in your own code and do not hand-
 6 substitute its definition. 

7 Why did the authors of the workbook decide to call this typedef
 8 `GenParticleCollection` and not, for example, `GenParticleVector`?
 9 The answer is a different sort of future-proofing. The C++ standard library provides
 10 class templates other than vectors that are collections of objects, and one can imagine a
 11 scenario in which it would make sense to change `GenParticleCollection` to use
 12 a collection type, such as `std::deque` for example. In such a scenario, the following
 13 definition would make perfect sense to the C++ compiler but would be misleading to
 14 human readers:

```
15 typedef std::deque<GenParticle> GenParticleVector
```

16 The generic name *Collection* avoids this problem.

17 16.4.2 Data Product Names

g-dp-names

18 Each *art* data product has a name that is a text string with four fields, delimited by un-
 19 der-score characters (`_`) that represent, in order, the data type, module label, instance name
 20 and process name, e.g.,:

```
21 MyDataType_MyModuleLabel_MyInstanceName_MyProcessName
```

22 Each data product name must be unique within an *art* event-data file. The fields in the data
 23 product name may only contain the following characters²:

24 ◦ `a...z`

25 ◦ `A...Z`

26 ◦ `0...9`

²Experts may want to know that in an *art* event-data file, each data product is stored as a `TBranch` of a `TTree` named `Event`. Only these characters are legal in a `TBranch` name. The name of the `TBranch` is the name of the data product, hence the restriction.

1 ○ :: (double colon)

2 In particular, periods, dashes, commas, underscores, semicolons, white space and single
3 colons are not allowed; underscores are only allowed as the field separator, not within a
4 field.

5 About each field:

6 1. The data type field is the so-called *friendly name* name of the data type for the data
7 product; friendly names are discussed below.

8 2. The module label field is the label of the module that created the data product. Note
9 that it is the module *label* as specified in the FHiCL file, not the `module_type`.

10 3. A given module instance in a given *art* process may make many data products of
11 the same type. These are distinguished by giving each a unique instance name. An
12 empty string is a valid instance name and in fact is the default. The other three fields
13 must be non-empty strings.

14 4. The process name field holds the value of the `process_name` parameter from the
15 FHiCL file for the *art* job that created the data product.

16 The friendly name of a data type is a concept that *art* inherited from the CMS software
17 suite. You will never need to write friendly names but you will need to recognize them.

18 Knowing the following rules will be sufficient in most cases:

19 1. If a type is not a collection type, then its friendly name is the fully qualified name of
20 the class.

21 2. If a type is `std::vector<T>`, its friendly name is `Ts`; the mnemonic is that
22 adding the letter "s" makes it plural.

23 3. If a type is `std::vector<std::vector<T> >`, its friendly name is `Tss`. And
24 so on.

25 4. If a type is `cet::map_vector<T>`, its friendly name is `Tmv`.

26 The full set of rules is given in the Users' Guide. **FIXME:** *ref to chapter/section when*
27 *available*

28 Corollaries of the above discussion include:

- 1 ○ None of the four fields in a product name may contain an underscore character:
2 otherwise the parsing of the name into its four fields is ambiguous.
- 3 ○ If an *art* event-data file is populated by running several *art* jobs, each of which
4 adds some data products, then each *art* job in the sequence must have a unique
5 process_name.

6 16.4.3 Specifying a Data Product

s-exercise

7 To identify a data product, *art* requires that you specify the data type, module label and
8 instance name fields (an empty string is a valid instance name). If the event contains ex-
9 actly one data product that matches this specification, then *art* allows a wild card match
10 on the process name field. If the event contains more than one data product that matches
11 this specification, then *art* requires that you also specify the process name. I.e., *art* allows
12 a wild card match only on the process name field, not on the others.

13 To tell *art* which data type you want, you use a template argument. To specify the other
14 three fields, module label, instance name and process name, you use an object of type
15 `art::InputTag`. Why is the data-type field treated differently than the others? This
16 method allows *art* to look after the translation of the data type to its friendly name. Users
17 of *art* never need to learn how to do this translation.

18 The header for `art::InputTag` is found in the file

19 `$ART_INC/art/Utilities/InputTag.h`.

20 You can construct an input tag by passing it a string with the three fields separated by
21 colons, e.g.,:

```
22  art::InputTag tag("MyModuleLabel:MyInstanceName:MyProcessName");
```

23 For this exercise, the full specification of the input tag includes only the module label and
24 the process name:

```
25  art::InputTag tag("evtgen::exampleInput");
```

26 The double colon indicates that the instance name (which would come between the colons)
27 is an empty string. The process name rarely needs to be specified, and in fact it is not
28 needed in this exercise. It will be sufficient to specify the input tag as

```
29  art::InputTag tag("evtgen");
```



There are other constructors for `art::InputTag` and there are accessor methods that provide access to the individual fields. You can learn about these by looking at the header file but you will not use these features in this exercise.

16.4.4 The Data Product used in this Exercise

The input files used for this exercise contain data products, one of which this exercise will use. This data product has the following attributes:

- it has a data type of `tex::GenParticleCollection`
- it is produced by a module with the label `evtgen`
- its instance name is an empty string
- it is produced by an *art* job with the process name `exampleInput`.

16.5 Setting up to Run this Exercise

Follow the instructions in Chapter 11 if you are logging in after having closed an earlier session. If you are continuing on directly from the previous exercise, keep both your source and build windows open.

In your source window, look at the contents of the directory for this exercise, called `ReadGenParticles`:

```
ls art-workbook/ReadGenParticles
```

```
CMakeLists.txt
readGens1.fcl   ReadGens1_module.cc
readGens2.fcl   ReadGens2_module.cc
readGens3.fcl   ReadGens3_module.cc
```

In your build window, just make sure that you are in your build directory.

All the code for this exercise is already built; this happened the first time that you ran buildtool.

1

2 In this exercise you will run three modules that differ in only a few lines. The three source
3 files use different syntax to accomplish the same thing. Most of the subsequent exercises in
4 the workbook will use the syntax shown in the third version, `ReadGens3_module.cc`,
5 and we recommend using this syntax in most cases. A description of the first two here
6 serves as a pedagogical progression. You will likely see all three types of syntax in your
7 experiment's code.

8 16.6 Running the Exercise

dp-running

9 You will run the exercise from your build directory in your build window. To run this
10 exercise, `cd` to your build directory and type the command:

```
11 art -c fcl/ReadGenParticles/readGens1.fcl >& output/readGens1.log
```

12 This will make the usual *art* output, interspersed with the output made by
13 `readGens1.fcl`. The output from this module is shown in Listing 16.2. For each event
14 it prints the event number and the number of `GenParticles` in that event.

:output:vl

Listing 16.2: Output using `readGens1.fcl`

```
15 ReadGens1::analyze event: 1 GenParticles: 7  
16 ReadGens1::analyze event: 2 GenParticles: 3  
17 ReadGens1::analyze event: 3 GenParticles: 3  
18 ReadGens1::analyze event: 4 GenParticles: 3  
19 ReadGens1::analyze event: 5 GenParticles: 5
```

20 16.7 Understanding the First Version, `ReadGens1`

g-dp-first

21 16.7.1 The Source File `ReadGens1_module.cc`

adgens1-cc

22 The module `ReadGens1_module.cc` contains a new `include` statement for the
23 `GenParticleCollection.h` header file. Here is the set of `include` statements at the
24 top of the file:

```

1 1 #include "toyExperiment/MCDataProducts/GenParticleCollection.h"
2
3 3 #include "art/Framework/Core/EDAnalyzer.h"
4 4 #include "art/Framework/Core/ModuleMacros.h"
5 5 #include "art/Framework/Principal/Event.h"
6
7 7 #include <iostream>
8 8 #include <string>

```

Listing 16.3: Include statements in `ReadGens1_module.cc`

9 In the next portion of the file, notice the new data member `gensTag_` on line 17, which
10 is initialized in the constructor using a string value that is taken from the parameter
11 set:

```

12 9 namespace tex {
13 10     class ReadGens1 : public art::EDAnalyzer {
14
15 12         public:
16 13             explicit ReadGens1(fhicl::ParameterSet const& );
17 14             void analyze(art::Event const& event) override;
18
19 16         private:
20 17             art::InputTag gensTag_;
21 18     };
22 19 }

```

23 Notice two things in the remainder of the file, below: Lines 26-27 introduce the concept
24 of a *handle*(γ), setting `gens` as a handle to the requested `GenParticleCollection`.
25 Lines 29-33 print out the number of entries in the data product — the same number as the
26 number of generated particles in the event.

```

27 20 tex::ReadGens1::ReadGens1(fhicl::ParameterSet const& pset):
28 21     art::EDAnalyzer(pset),
29 22     gensTag_(pset.get<std::string>("genParticlesInputTag")) {
30 23 }

```

```
1 24 void tex::ReadGens1::analyze(art::Event const& event ){
2
3 26   art::Handle<GenParticleCollection> gens;
4 27   event.getByLabel(gensTag_, gens);
5
6 29   std::cout << "ReadGens1::analyze event: "
7 30             << event.id().event()
8 31             << " GenParticles: "
9 32             << gens->size()
10 33            << std::endl;
11 34 }
12
13 36 DEFINE_ART_MODULE(tex::ReadGens1)
```

14 As you work through the *art* workbook you will encounter several types of handles. All of
15 the handle types behave like pointers with additional features:

- 16 1. They have safety features that make it impossible for your code to look at a pointee
17 that is either not valid or not available.
- 18 2. They may also have an interface that lets you access metadata that describes the
19 pointee.

20 The handle is an example of a broader idea sometimes called a *safe pointer*(γ) and some-
21 times called a *smart pointer*(γ).

22 The header for the class template `art::Handle` is found in the file
23 `$ART_INC/art/Framework/Principal/Handle.h`. This file is automati-
24 cally included by the include for `Event.h`.

25 The `art::Handle` line tells the compiler to default construct an object of
26 type: `art::Handle<GenParticleCollection>`. The name of the default-
27 constructed object is `gens`. A default-constructed handle does not point at anything and,
28 if you try to use it as a pointer, it will throw an exception. A handle in this state is said to
29 be invalid.

30 The following line calls `getByLabel`, which uses its first argument (`gensTag_`) to
31 learn three of the four elements of the name of the requested data product. It can deduce

1 the fourth element, the data type, from the *type* of its second argument (*gens*): that is, it
2 knows that it must look for a data product of type `tex::GenParticleCollection`.
3 *art* has tools to compute the friendly name from the full class name, which is why you will
4 never need to write a friendly name.

5 When this line is executed, the event object looks to see if it contains a data product that
6 matches the request. There are three possible outcomes:

- 7 1. the event contains exactly one product that matches
- 8 2. the event contains no product that matches
- 9 3. the event contains more than one product that matches

10 In the first case, the event object will give the handle a pointer to the requested
11 `tex::GenParticleCollection`; the handle *gens* can then be used as a pointer,
12 as is done in the second line of the `std::cout` section. When the handle has received the
13 pointer, it is said to be in a valid state. In the second and third cases, the event object will
14 leave the handle in its default-constructed state and, if you try to use it as a pointer, it will
15 throw an exception.

16 If the event object finds exactly one match, it will also add two pieces of metadata to
17 the handle. One is a pointer to an object of type `art::Provenance`, which contains
18 information about the processing history of the data product. The second is an object of
19 type `art::ProductID`; this is essentially a synonym for the four-field string form of
20 the product name. Both of these will be illustrated in future exercises. **FIXME:** *References*
21 *to them.*

22 The third case bears one more comment: the developers of *art* made a careful decision that,
23 except for the process name field, `getByLabel` will not have a notion of “best match”.
24 When you use `getByLabel` you must unambiguously specify the data product you want
25 or *art* will leave the handle in its default-constructed state.

26 If the `getByLabel` member function does not find the requested data product, e.g., if
27 you run it on a different input file or if you misspell any of the fields in the input tag, the
28 handle will be left in its default-constructed state. In this case, the `gens->size()` call
29 will know that the handle is invalid and will throw an exception.

30 In all cases but one, *art*'s response to an exception is to attempt a graceful shutdown. The

1 one unusual case is *ProductNotFound*, which is the exception thrown by an invalid handle
2 when you try to use it as a pointer. In this case *art* will print a warning message, skip this
3 module and attempt to run the remaining modules in the trigger paths and end paths.

4 It is possible to test the state of *gens* by using the member function `gens.isValid()`,
5 which returns a `bool`. This not illusrated in the example because in most cases we rec-
6 ommend that you let *art* deal with this for you.

7 In the preceding discussion we did not mention that `getByLabel` is actually a *member*
8 *function template*. There is no explicit template argument in the `event.getByLabel`
9 line because the C++ template mechanism is able to deduce the template argument from
10 the type of the second argument.



11 The `art::Event` object supports several other ways to request data products from the
12 event, including a way to get handles to all data products that match a partial specification.
13 This material is beyond the scope of this exercise. **FIXME:** *Reference to where they are*
14 *discussed.*

15 16.7.2 Adding a Link Library to `CMakeLists.txt`

cmakelists

16 **FIXME:** *Marc to change ‘link’ to ‘dynamic’ library? FIXME: MFP: I think we should*
17 *leave the name “link library”, because this is what the CMakeLists.txt file calls this kind*
18 *of thing. I believe the name is intended to point to the difference in use: this is a dynamic*
19 *library that code is linked to, rather than a dynamic library that is explicitly loaded by the*
20 *moduling loading facilities of the art framework.*

21 `ReadGens1_module.so` requires linking to a dynamic library that was not needed by
22 previous exercises, namely

23 `$TOYEXPERIMENT_LIB/libtoyExperiment_MCDataProducts.so`.

24 It is different from the previous dynamic libraries that are explicitly

25 This library contains the object code for the classes and functions defined in the `MCDataProducts`
26 subdirectory of the `toyExperiment UPS` product. In particular it contains object code
27 needed by the data product `tex::GenParticleCollection`.

28 Here we will call it a *link* rather than *dynamic* library since that’s what the `CMakeLists.txt`
29 file calls this kind of library that code is *linked to* rather than one that is explicitly loaded
30 by *art*.

- 1 Adding this library to the link list required a one-line modification to `CMakeLists.txt`.
 2 If you compare this file to the corresponding file for the previous exercise, you will see
 3 that `CMakeLists.txt` for this exercise contains one additional line:

```
4     ${TOYEXPERIMENT_MCDATAPRODUCTS}
```

- 5 The string `TOYEXPERIMENT_MCDATAPRODUCTS` is a `cmake` variable that was defined
 6 when you first ran the `buildtool` command. The translated value of this variable is the name
 7 of the required link library.

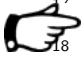
8 16.7.3 The FHiCL File `readGens1.fcl`


- 9 There is only one fragment of `readGens1.fcl` that contains any new ideas. It is the
 10 fragment that configures the module label `read`, reproduced in Listing 16.4

Listing 16.4: Configuring the module label `read` in `readGens1.fcl`

```
11     read : {
12         module_type           : ReadGens1
13         genParticlesInputTag : "evtgen"
14     }
```

- 15 On line 4 of this fragment, the parameter `genParticlesInputTag` specifies the input
 16 tag that identifies the data product to be read by this exercise.

 17 We recommend that you always initialize input tags using parameters from the parameter
 18 set and that you never initialize them using strings defined within the code. This will allow
 19 you run the same module on data products with different input tags; this is a widely used
 20 feature.

 21 We further recommend that you not provide a default value in the call to get the parameter
 22 value from the parameter set. This derives from a general recommendation that parameters
 23 affecting physics output should never have default values; the only parameters with default
 24 values should be those that control debugging and diagnostics.

25 16.8 The Second Version, `ReadGens2`

- 26 Version 2 of this exercise consists of the files `ReadGen2_module.cc` and
 27 `readGen2.fcl`. To run this version, `cd` to your build directory and type the com-

1 mand:

```
2 art -c fcl/ReadGenParticles/readGens2.fcl >& output/readGens2.log
```

3 It will produce the same output as the previous two versions.

4 The only significant change from version 1 to version 2 is that lines

```
5 1 art::Handle<GenParticleCollection> gens;  
6 2 event.getByLabel(gensTag_, gens);
```

7 have been replaced by the single (long) line:

```
8 1 art::ValidHandle<GenParticleCollection> gens =  
9 2     event.getValidHandle<GenParticleCollection>(gensTag_);
```

10 This version is a little verbose but that aspect will be addressed in version 3. Note
11 that the class template `art::Handle` has been replaced by a new class template
12 `art::ValidHandle`. Both class templates are defined in the same header file,
13 `$ART_INC/art/Framework/Principal/Handle.h`.

14 The above line has functionality very similar to that of the two lines from version 1: the
15 net result is that `gens` can be used as a pointer to the requested data product. It also has
16 an interface to access the `art::Provenance` and the `art::ProductID`.

17 However, there are several significant differences between `art::Handle<T>` and
18 `art::ValidHandle<T>`:

- 19 1. Unlike an `art::Handle<T>`, which may be either valid or invalid, an
20 `art::ValidHandle<T>` is guaranteed to be valid. It cannot be default-
21 constructed.
- 22 2. A call to `getValidHandle<T>` will either return a properly constructed
23 `art::ValidHandle<T>` or it will throw a `ProductNotFound` exception.
- 24 3. `art::ValidHandle` does not have an `isValid()` method.
- 25 4. Everytime that you use an `art::Handle<T>` as a pointer, it first checks that the
26 pointer is valid. On the other hand, when you use an `art::ValidHandle<T>` as
27 a pointer, no check is necessary; using an `art::ValidHandle<T>` as a pointer
28 is as fast as using a bare pointer or a reference.

16.9 The Third Version, ReadGens3

Version 3 of this exercise consists of the files `ReadGen3_module.cc` and `readGen3.fcl`. To run this version, `cd` to your build directory and type the command:

```
art -c fcl/ReadGenParticles/readGens3.fcl >& output/readGens3.log
```

It will produce the same output as the previous two versions.

The only change from version 2 is that the call to `getValidHandle` has a slightly different syntax that provides the same behavior but is less verbose (shown here on two lines):

```
1 auto gens =
2     event.getValidHandle<GenParticleCollection>(gensTag_);
```

This version uses a feature of C++ that is new in C++-11, the keyword `auto`. This keyword tells the C++ compiler to automatically determine the correct type for `gens`.

When you call the member function `getValidHandle<T>` the return type will always be `art::ValidHandle<T>`.

Version 3 is the version that we recommend you use but you can use any of the three. We introduced the version using the keyword `auto` as the last version because it is a handy shorthand when you know how to determine the correct type but it is very confusing if you do not know how to do so.

In future exercises we will use the pattern of version 3 regularly.

16.10 Suggested Activities

Edit `readGens3.fcl` and supply the full input tag:

```
genParticlesInputTag : "evtgen::exampleInput"
```

Run `art` and observe that it works correctly.

Edit `readGens3.fcl` and misspell the the requested module label, for example

misspelled

Listing 16.5: Warning message for misspelled module label of data product

```
%MSG
%MSG-w FailModule:  ReadGens3:read 15-Jun-2014 10:00:24 CDT
                                run: 1 subRun: 0 event: 5
Module failed due to an exception
---- ProductNotFound BEGIN
  getByLabel: Found zero products matching all criteria
  Looking for type: std::vector<tex::GenParticle>
  Looking for module label: genevent
  Looking for productInstanceName:
---- ProductNotFound END
```

```
1 genParticlesInputTag : "genevent"
```

2 Run *art* and observe the warning messages, which should look like the message in Listing 16.5. This tells you that:

- 4 1. The error occurred while processing the module with the module label `read`, which is an instance of the `module_type` (`ReadGens3`). You learn this from line 2 of the listing.
- 7 2. The error occurred while calling the member function `getByLabel`.
- 8 3. The data product that was requested had a type of `std::vector<tex::GenParticle>` and was created by the module with the label `genevent`; the requested data product had an instance name of an empty string.
- 12 4. *art* found zero data product matching this request.

13 You would also get an error if *art* had found more than one data product matching this request.

15 **FIXME:** *Rob to Anne: the process name the module type are different - check the capitalization of the first letter. ... (Anne:) readGens3proc vs readGens3mod comes to mind as a possibility; maybe unwieldy. It's just that there are so many avenues for confusion.*

18 Observe that for each event *art* prints the warning message and continues with the next event.

1 Look at the last line of the *art* output and observe that *art* completed with status 0. **FIXME:**
2 *We should include this in the listing. Am trying to rerun, am getting different output. Re-*
3 *visit. AH* This is because *art* treats `ProductNotFound` as a warning, not as an error that
4 will initiate a shutdown.

5 You can reconfigure *art* so that a `ProductNotFound` exception will cause *art* to shutdown
6 gracefully. To do this, edit your modified `readGens3.fcl` and add the following line
7 inside the `services` parameter set:

8

```
9 scheduler : { defaultExceptions : false }
```

10 This line tells *art* that its response to all exceptions should be to attempt a graceful shut-
11 down. When you rerun *art* you should see output like that shown in Listing 16.6. [list:readGen1:shutdown](#)

1:shutdown

Listing 16.6: Exception message for ProductNotFound, default Exceptions disabled

```

1 %MSG-s ArtException: PostCloseFile 15-Jun-2014 10:32:55 CDT PostEndRun
2 cet::exception caught in art
3 ---- EventProcessorFailure BEGIN
4   An exception occurred during current event processing
5   ---- EventProcessorFailure BEGIN
6     An exception occurred during current event processing
7     ---- ScheduleExecutionFailure BEGIN
8       ProcessingStopped.
9
10    ---- ProductNotFound BEGIN
11      getByLabel: Found zero products matching all criteria
12      Looking for type: std::vector<tex::GenParticle>
13      Looking for module label: genevent
14      Looking for productInstanceName:
15
16      cet::exception going through module ReadGens3/read
17                                     run: 1 subRun: 0 event: 1
18    ---- ProductNotFound END
19      Exception going through path end_path
20    ---- ScheduleExecutionFailure END
21    ---- EventProcessorFailure END
22    cet::exception caught in EventProcessor and rethrown
23 ---- EventProcessorFailure END
24 %MSG

```

16.11 Review

In this chapter you have learned:

1. the type `tex::GenParticleCollection`
2. the four-part identifier of data product and the class `art::InputTag`
3. the class templates `art::Handle` and `art::ValidHandle`
4. how to get a handle to a data product, given its type and input tag
5. how to use a handle as a pointer to the requested data product
6. how to recognize a `ProductNotFound` warning message

- 1 7. how to tell *art* to treat the `ProductNotFound` exception as a hard error that will
2 initiate a graceful shutdown.

3 16.12 Test Your Understanding

4 16.12.1 Tests

5 Two source files and one FHiCL file are provided that contain intentional bugs. Your job, of
6 course, is to find the bugs and fix them. The answers are provided in Section 16.12.2.

7 The source files listed below fail to build. Follow essentially the same procedure as out-
8 lined in Section 10.12 on each of these files. To run them in *art*, modify an existing FHiCL
9 file, as needed, or create a new one.

10 `ReadGenParticles/ReadGensBug01_module.cc.nobuild`

11 `ReadGenParticles/ReadGensBug02_module.cc.nobuild`

12 What's with the `.nobuild` suffix? The build system looks at each file in each subdirec-
13 tory and figures out what to do with the file based on the file type of the file; the file type is
14 that last part of the filename, the part that comes after the dot. It has no information on
15 what to do with files that end in `.nobuild` so it simply ignores them. When you rename
16 the file to `.cc` the next run of `buildtool` will discover the file and attempt to build
17 it.

18 **FIXME:** *Rob to Anne: we have not supplied FHiCL files for the two .nobuild exercises.*
19 *They need to write their own. Anne: but the other fcl files work fine... let's discuss. 7/30*
20 *AH: see my text above; ok I think*

21 The FHiCL file `ReadGenParticles/bug03.fcl` runs to completion, printing lots
22 of warning messages. However, it does not make any of the expected printout from the
23 `analyze` member function. Run it, find the problem, and fix it.

24 16.12.2 Answers

25 For the file `ReadGensBug01_module.cc`, when you build you should find the follow-
26 ing error message:

```

1 /home/ahavey/workbook/art-workbook/art-
2 workbook/ReadGenParticles/ReadGensBug01_module.cc:43:57: error: 'class
3 art::ValidHandle<std::vector<tex::GenParticle>_>' has no member named 'size'
4     << "_GenParticles:_ "          << gens.size()

```

5 As the error message indicates, the error is indeed in the line:

```

6     << "_GenParticles:_ "          << gens.size()

```

7 Remember that `gens` is a handle to `GenParticleCollection`, and does not it-
8 self have a member function named `size`. The entity that does have a member function
9 named `size` is the collection itself. Therefore the solution is that `gens.size()` should
10 be `gens->size()`.

11 More generally, the syntax `gens->foo()` says “go find the thing to which points, find
12 its function `foo()` and call it.” The (incorrect) `gens.foo()` says to call the function
13 `foo` of the handle itself, which doesn’t exist.

14 For the file `ReadGensBug02_module.cc`, you will find the error:

```

15 /home/ahavey/workbook/art-workbook/art-
16 workbook/ReadGenParticles/ReadGensBug02_module.cc:34:57:
17 error: no matching function for call to
18 'art::EDAnalyzer::EDAnalyzer()'
19     gensTag_(pset.get<std::string>("genParticlesInputTag")) {

```

20 The problem here is that the constructor of an analyzer module must pass the parameter set
21 to the constructor of `art::EDAnalyzer`, and it’s missing in this file. This was described
22 in Sections [10.6.3.3](#) and [??](#)

23 To make the build work, fix the constructor to look like

```

24 1 tex::ReadGensBug01::ReadGensBug01(fhicl::ParameterSet
25 2     const& pset ) :
26 3     art::EDAnalyzer(pset) ,
27 4     gensTag_(pset.get<std::string>("genParticlesInputTag")) {
28 5 }

```

29 Now it should work fine.

30 Finally, run `art` on `bug03.fcl`. It will produce the error:

```

1 ---- ProductNotFound BEGIN
2   getByLabel: Found zero products matching all criteria
3   Looking for type: std::vector<tex::GenParticle>
4   Looking for module label: Evtgen
5   Looking for productInstanceName:
6
7 ---- ProductNotFound END

```

8 The parameter value for `genParticlesInputTag` has an upper-case E in `Evtgen`; it
9 should be `evtgen`.

10 **FIXME:** *This exercise makes you understand `evtgen`, but doesn't test the use of tem-*
11 *plates*

12 **FIXME:** *Rob to Anne: we can make another `.nobuild` in which the template argument in*
13 *the call to `getByLabel` is misspelled. I will commit it tonight.*

14 **FIXME:** *Useful info for me; do we want to include generally? Section ?? talks about*
15 *`genParticlesInputTag` in `fcl` module with label `evtgen` produces the data product used in*
16 *exercise ?? 16.4.3 input tag (or `art::InputTag`) `mytype:mymod:myinst:myprocess sec ??`*
17 *`art::InputTag` tag("evtgen"); in 16.4.3 16.7.3 shows `evtgen` in `fcl` file ??*

18 **FIXME:** *Rob to Anne: I am not sure what you are asking here. Are you suggesting that*
19 *in the review section we add links to the sections that discuss the material in each review*
20 *bullet point? If so that's a good idea. If not find me tomorrow.*

17 Exercise 7: Making a Histogram

17.1 Introduction

One of the workhorse tools of HEP data analysis is ROOT. Among its many features are tools for data analysis, visualization, presentation and persistency. As was discussed in Section 3.6.9, *art* uses ROOT as a tool for persistency of event-data¹. In the code base of a typical HEP experiment there are many modules that use ROOT to create histograms, graphs, ntuples and trees, all of which are objects used for data analysis, visualization and presentation.

This exercise will show you how to use ROOT in the *art* environment using the ROOT class TH1D — one of many — as an example. Using this class you will create, fill and present 1-dimensional histograms. You can follow the model presented here if you wish to use related ROOT classes, such as the other histogram classes and the classes for graphs, ntuples and trees. If you are not familiar with graphs, ntuples and trees, examples will be given in future exercises. **FIXME:** *Add refs to future exercises when they are available*

Detailed information about ROOT is available from its website, <http://root.cern.ch/drupal>.

Most of the modules that get run in a typical *art* job — plus *art* itself — use ROOT. Due to the way ROOT and *art* interact (a topic beyond the present scope), *art* needs to provide a mechanism to ensure that your module's use of ROOT will not interfere with the use of ROOT by *art* or by other modules running in the same job. The mechanism is an *art* service called TFileService, which does the necessary organizational work.

¹ You may have already guessed this, having seen the module types `RootInput` and `RootOutput` in previous exercises.

- 1 This chapter will introduce you to *art* services in general and to the `TFileService` in
2 particular.



3 All user interactions with ROOT should happen via this service.



6 Note that it is possible to use ROOT as an event-processing framework, e.g., the AliRoot
7 framework used by the ALICE Collaboration. But if you are using *art*, then *art* is always
the event-processing framework and ROOT is used as a toolkit. The AliRoot documenta-
tion is at <http://aliweb.cern.ch/Offline/AliRoot/Manual.html>.

8 17.2 Prerequisites

9 Prerequisites for this chapter include all of the material in Part I (Introduction) and the
10 material in Part II (Workbook) up to and including Chapter 16.

11 17.3 What You Will Learn

12 In this exercise you will learn:

- 13 1. What the `art::TFileService` is and what it does for you.
- 14 2. How to configure the `art::TFileService`.
- 15 3. What an `art::ServiceHandle` is and what it does for you.
- 16 4. How to access ROOT via the `art::TFileService`.
- 17 5. How to create and fill a ROOT TH1D histogram.
- 18 6. How to use the interactive ROOT browser to view the histogram.
- 19 7. How to run a CINT script to view the histogram and to write the histogram to a PDF
20 file.
- 21 8. The naming convention used by the Workbook to distinguish event-data ROOT files
22 from ROOT files containing histograms, ntuples, and so on. This convention is
23 specific to the Workbook and it may differ from what your experiment uses.

17.4 Setting up to Run this Exercise

Follow the instructions in Chapter 11 if you are logging in after having closed an earlier session. If you are continuing on directly from the previous exercise, keep both your source and build windows open.

In your source window, look at the contents of the directory for this exercise, called `FirstHistogram`:

```
ls art-workbook/FirstHistogram
```

```
CMakeLists.txt      FirstHist1_module.cc  
drawHist1.C         firstHist1.fcl
```

In your build window, just make sure that you are in your build directory. All the code for this exercise is already built; this happened the first time that you ran `buildtool`.

The module `FirstHist1_module.cc` is very much like the module `ReadGens3_module.cc` from the previous exercise. The main difference is that it does not create any printout but rather, it fills a histogram displaying the number of generated particles in each event.

The FHiCL file `firstHist1.fcl` is very much like the file `readGens3.fcl` from the previous exercise. The important difference here is that `firstHist1.fcl` configures the `TFileService`.

The file `drawHist1.C`, discussed in Section 17.9.2, is a script written in a ROOT-defined language called CINT. This script contains the commands to open a ROOT file, draw a histogram and write it to a PDF file.

The file `CMakeLists.txt` plays its usual role telling the build system what to do. Compared to the corresponding file for the previous exercise, it has two additional link libraries and contains an explicit directive that `drawHist1.C` should not be built. The meaning of this will become clear in the full discussion of `CMakeLists.txt`.

17.5 The Source File `FirstHist1_module.cc`

1

`FirstHist1-module`

2 The C++ source code for this exercise is found in the file `FirstHist1_module.cc`.
 3 Open the file in your source window to see it as a whole. Listing 17.1 contains a frag-
 4 ment of this file, showing the included headers and the declaration of the module class
 5 `FirstHist1`. Compared to the file `ReadGens3_module.cc` from the previous exer-
 6 cise, four new lines have been added; they appear in the listing as:

- 7 1. line 6, which includes the header for the *art* `TFileService`
- 8 2. line 8, which includes the header for the ROOT class `TH1D`
- 9 3. line 21, which declares the member function `beginJob`
- 10 4. line 28, which declares a new member datum, named `hNGens_`, of type *pointer* to
 11 an object of type `TH1D`.

12 The name `hNGens_` was chosen because this pointer will eventually point at a histogram
 13 object that contains a histogram of the number of generated particles per event.



14 The *art* workbook has adopted the style that all names for pointers to histograms begin
 15 with the lower case letter “h”.

16 The two new headers can be found at:

17 `$ART_INC/art/Framework/Services/Optional/TFileService.h`
 18 `$ROOT_INC/TH1D.h`

19 The conventions for including header files from ROOT differ from those for including
 20 header files from *art* and from *toyExperiment*. To remind you, the conventions for *art* and
 21 the *toyExperiment* UPS product are:

- 22 1. The names of all classes and functions are inside a namespace, *art* or *tex*, respec-
 23 tively.
- 24 2. In the header file `#include` lines, the name of the package to which the header
 25 belongs is always the first element of the path.

26 When ROOT was developed, namespaces were not supported robustly by many C++ com-
 27 pilers. Therefore a different set of conventions were adopted – and remain – for ROOT:

t1:declare

```
1 #include "toyExperiment/MCDataProducts/GenParticleCollection.h"
3 #include "art/Framework/Core/EDAnalyzer.h"
4 #include "art/Framework/Core/ModuleMacros.h"
5 #include "art/Framework/Principal/Event.h"
6 #include "art/Framework/Services/Optional/TFileService.h"
8 #include "TH1D.h"
10 #include <iostream>
11 #include <string>
13 namespace tex {
15     class FirstHist1 : public art::EDAnalyzer {
17     public:
19         explicit FirstHist1(fhicl::ParameterSet const& );
21         void beginJob() override;
22         void analyze(art::Event const& event) override;
24     private:
26         art::InputTag gensTag_;
28         TH1D* hNGens_;
30     };
32 }
```

Listing 17.1: Top portion of `FirstHist1_module.cc`

- 1 1. The names of all ROOT classes and functions are in the *global* namespace, i.e., they
2 are not part of a namespace defined by ROOT.
- 3 2. The names of all ROOT classes begin with a capital letter T followed by an upper
4 case letter (this serves as a weak substitute for using a namespace).
- 5 3. The syntax to include a file from ROOT is to give the filename without any leading
6 path elements. The clue that the file is a ROOT header file comes from the leading
7 capital T.

8 Listing 17.2 shows the implementation section of the file `FirstHist1_module.cc`.

9 The new features in this listing are:

- 10 1. line 5, which initializes `hNGens_` to have the value of a null pointer
- 11 2. lines 8 through 14, which create an empty histogram
- 12 3. line 20, which fills the histogram with the number of generated particles in the cur-
13 rent event

14 The identifier `nullptr`, used in line 5, was added to the C++ core language in the 2011
15 Standard. It is the value of a pointer that points to nothing; in practice it has a value of zero.
16 You will very likely encounter code written prior to the 2011 Standard. In such code you
17 will see the equivalent of line 5 written in one of the following two ways: `hNGens_(0)`
18 or `hNGens_(NULL)`. In the second form, the value `NULL` is a C-Preprocessor **MACRO**
19 variable that is defined to have a value of 0.



20 We strongly recommend, first, that you use `nullptr` for this purpose, and second that
21 you never use the C-Preprocessor `NULL`.

22 17.5.1 Introducing `art::ServiceHandle`

23 Section 3.6.5 discussed the idea of *art* services. These are classes that provide some func-
24 tionality (i.e., a service) that can be used by any module or by other *art* services. In this
25 exercise you will see your first example of an *art* service, the `art::TFileService`,
26 which provides a bookkeeping layer to ensure that your use of ROOT does not interfere
27 with other uses of ROOT within the same *art* job.

Hist1:impl

```
1 tex::FirstHist1::FirstHist1(fhicl::ParameterSet const& pset):
2   art::EDAnalyzer(pset),
3   gensTag_(pset.get<std::string>("genParticlesInputTag")),
4   hNGens_(nullptr) {
5 }
6
7 void tex::FirstHist1::beginJob() {
8
9   art::ServiceHandle<art::TFileService> tfs;
10  hNGens_ = tfs->make<TH1D>( "hNGens",
11    "Number of generated particles per event",
12    20, 0., 20.);
13
14 }
15
16 void tex::FirstHist1::analyze(art::Event const& event ) {
17
18   auto gens =
19     event.getValidHandle<GenParticleCollection>(gensTag_);
20
21   hNGens_->Fill(gens->size());
22
23 }
```

Listing 17.2: Implementation of the class FirstHist1

1 In a similar way that access to data products is provided by the class tem-
 2 plates `art::Handle` and `art::ValidHandle`, access to services is pro-
 3 vided by the class template `art::ServiceHandle`. Line 10 in List-
 4 ing 17.2 tells the compiler to default construct an object, named `tfs`, of type
 5 `art::ServiceHandle<art::TFileService>`. The constructor of `tfs` will
 6 contact the internals of `art` and ask `art` to find a service of type `art::TFileService`.
 7 If `art` can find such a service, it will give the service handle a pointer to the service. If
 8 not, it will throw an exception and attempt a graceful shutdown.

9 Once a service handle has been constructed, the downstream code can use the service
 10 handle as a pointer to the pointee, i.e., to `art::TFileService`.

11 The header file for `art::ServiceHandle` is found at:

12 `$ART_INC/art/Framework/Services/Registry/ServiceHandle.h`

13 It is automatically included by one of the files that are already included in
 14 `FirstHist1_module.cc`.

15 17.5.2 Creating a Histogram

16 Lines 9 through 12 of Listing 17.2,


```
17 9   art::ServiceHandle<art::TFileService> tfs;
18 10   hNGens_ = tfs->make<TH1D>( "hNGens",
19 11       "Number of generated particles per event",
20 12       20, 0., 20.);
```

21 use `art::TFileService` to create a new histogram object of type `TH1D`. In the call
 22 to the member function template `tfs->make` (lines 11 and 12), the type of object to be
 23 created is specified using a template argument. The function arguments, listed below, are
 24 the arguments needed by a constructor of that type of object. You do not need to understand
 25 why things are done this way or how it all works. You just need to follow the pattern. The
 26 return value of the call to `tfs->make` is a pointer to the newly created histogram object
 27 and this value is assigned the member datum `hNGens_`.

28 In the case of creating a `TH1D`, the five function arguments are:

- 29 1. the name by which `ROOT` will know this histogram; the `art` workbook has adopted
 30 the convention that this name will always be the name of the corresponding member

- 1 datum, excluding the underscore (in this case `hNGens`)
- 2 2. the title that will be displayed when the histogram is drawn (given on line 12 of
3 Listing 17.2)
- 4 3. the number of bins in the histogram (20)
- 5 4. the lower edge of the lowest bin of the histogram (0.)
- 6 5. the upper edge of the uppermost bin of the histogram (20.)

7 ROOT defines that the low edge of a bin is within that bin, while the upper edge of a bin is
8 part of the next bin up. Therefore the lower edge of the lowest bin is inside the histogram
9 but the upper edge of the uppermost bin is outside of the histogram. 

10 If you would like to learn more about the `TH1D` class you can look at its header file or you
11 can read about it on the ROOT web site: <http://root.cern.ch/root/html534/TH1D.html>.

12 Where is the histogram created? The histogram is created in memory that is owned and
13 managed by ROOT. ROOT also knows that when the job is finished, it should write the
14 histogram to a ROOT output file that you can inspect at a later time. The name of the output
15 file is specified in the FHiCL file for the *art* job; more on that later. We will call this file the
16 *histogram output file* or *histogram file*. Although histogram files often contain much more
17 than just histograms, the name is in fairly common usage among the experiments that use
18 *art*. Histogram files do not contain *art* data products.

19 Just as file systems have the notion of directories and subdirectories (or folders and
20 subfolders if you prefer), a ROOT file has the notion of directories and subdirectories
21 that are internal to the ROOT file. If a module makes at least one histogram, then the
22 `TFileService` will first create a new top-level directory in the histogram file. The name
23 of this top-level directory is the name of the module label of the module that created the
24 histogram. All ROOT objects that are created by that module will be created within this top
25 level directory. When the contents of ROOT-managed memory are written to the histogram
26 file, this directory structure is preserved.

27 Recall that within a given *art* job each module label must be unique. This ensures that, for
28 every module instance that uses the `TFileService`, a uniquely named top-level direc-
29 tory will be created in the histogram file. It is this strategy that ensures that the histogram
30 names of my module will never collide with the histogram names of your module.

17.5.3 Filling a Histogram

Line 21 of Listing 17.2

```
hNGens_ -> Fill(gens -> size());
```

fills the histogram pointed to by `hNGens_` with the number of generated particles for this event.

If you look up the function prototype for `TH1D::Fill` you will see that it expects an argument that is a double. On the other hand, `gens->size()` returns an unsigned integer. One of the features of C++ is that it can automatically convert the unsigned integer to a double and pass that to the function. For details consult the standard C++ documentation that is listed in Section 6.9.

17.5.4 A Few Last Comments

All of the comments above about management of ROOT directories and writing histograms to files are also true for most other sorts of ROOT objects. In particular they are true for `TTrees` and `TNtuples`.

If you think carefully about `FirstHist1_module.cc` you might wonder why there is no `endJob` member function containing a call to delete the histogram that was created in the `beginJob` member function. The answer is that when you create a histogram that is controlled by ROOT, then ROOT is responsible for calling `delete` at the right time.

If you talk to an HEP old-timer about creating histograms, he or she will probably call it “booking a histogram.” This is language left over from a precursor to ROOT named HBOOK.

Both the histogram files (output) and the *art* event-data files (input, in this example) are written using ROOT. Even though both are ROOT files, the two types of files are structured very differently and are not in any way interchangeable or interoperable.

In order to make it clear which ROOT files are of which type, the *art* Workbook has adopted the convention that *art* event-data files always end in `.art`. All other files ending in `.root` are histogram files.

1 Some experiments, Mu2e, for example, have adopted the same convention. Other experi-
 2 ments have chosen that event-data files always end in `_data.root`, while all other files
 3 ending in `.root` are histogram files.² Yet other experiments have chosen the opposite
 4 convention: files ending in `_hist.root` are histogram files and all other files ending in
 5 `.root` are *art* event-data files. Check with one of your colleagues to learn which conven-
 6 tion your experiment has adopted.

7 17.6 The Configuration File `firstHist1.fcl`

8 The file `firstHist1.fcl`, shown in Listing 17.3, is very much like the file `readGens3.fcl`
 9 from the previous exercise.

10 The most important new feature is at line 12,

```
12 TFileService : { fileName : "output/firstHist1.root" }
```

12 which configures the `TFileService`. This service has one required parameter, which
 13 is the name of the histogram file that contains the histograms, trees, and so on that are
 14 created by the *art* job.

15 If this parameter is missing, or if the configuration for the `TFileService` is missing
 16 entirely, then the first attempt to get a service handle to the `TFileService` will throw
 17 an exception, and *art* will attempt a graceful shutdown.

18 Unlike in the previous excercises, the FHiCL file runs on the large input event-data file,
 19 `inputFiles/input04.art`, which contains 1,000 events.

20 17.7 The file `CMakeLists.txt`

21 The `CMakeLists.txt` file used for this exercise is shown in Listing 17.4. Compared to
 22 the corresponding file for the previous exercise, there are two new features.

- 23 1. Two link libraries have been added.
- 24 2. There is a directive indicating that `cmake` should not build files ending in `.C`.

²Earlier versions of the workbook used this convention

`FirstHist1.fcl`**Listing 17.3:** firstHist1.fcl

```
1 #include "fcl/minimalMessageService.fcl"
2
3 process_name : firstHist1
4
5 source : {
6   module_type : RootInput
7   fileNames   : [ "inputFiles/input04.art" ]
8 }
9
10 services : {
11   message : @local::default_message
12   TFileService : { fileName : "output/firstHist1.root" }
13 }
14
15 physics :{
16   analyzers: {
17     hist1 : {
18       module_type      : FirstHist1
19       genParticlesInputTag : "evtgen"
20     }
21   }
22
23   e1      : [ hist1 ]
24   end_paths : [ e1 ]
25
26 }
```

1 The two new link libraries are specified by
2 `${ART_FRAMEWORK_SERVICES_OPTIONAL_TFILESERVICE_SERVICE}`
3 and `${ROOT_HIST}` (lines 12 and 15). These items are both `cmake` variables that
4 were defined for you when established the development environment. The first variable
5 translates to
6 `$ART_LIB/libart_Framework_Services_Optional_TFileService_service.so`
7 and the second translates to `$ROOTSYS/lib/libHist.so`, which contains the object
8 code for the class `TH1D`, among others.

9 Many projects use the convention that files ending in `.C` contain code written in the `C` pro-
10 gramming language. By default `cmake` will assume that files ending in `.C` follow this con-
11 vention and, therefore, it will try to compile and link them. We have already encountered a
12 `CINT` script that ends in `.C`. The `cmake` needs to ignore `CINT` files. `CMakeLists.txt`
13 includes code to effect this.

14 You do not need to understand the details of how the `CMakeLists.txt` excludes
15 `drawHist1.C` from the build. For those who wish too look up the details, the high level
16 explanation follows. Lines 1 and 2 in the listing of `CMakeLists.txt` tell `cmake` to
17 define a new `cmake` variable named `ROOT_MACROS_DO_NOT_BUILD`. This variable is
18 the set of all filenames that end in `.C` from the same directory as the `CMakeLists.txt`
19 file. Line 4 in `CMakeLists.txt` tells `cmake` that it should do nothing for all files that
20 appear in the translation of this variable.



21 17.8 Running the Exercise

22 To run this exercise, `cd` to your build directory and run `art`:

```
cd $ART_WORKBOOK_WORKING_BASE/<username>/\  
workbook/build-prof  
  
art -c fcl/FirstHistogram/firstHist1.fcl
```

23
24 This module does not make any of its own printout. You should see the standard printout

hist1:cmakelists

Listing 17.4: CMakeLists.tex in the directory FirstHistogram

```

1 file( GLOB ROOT_MACROS_DO_NOT_BUILD
2     RELATIVE ${CMAKE_CURRENT_SOURCE_DIR} *.C )
3 art_make(
4     EXCLUDE ${ROOT_MACROS_DO_NOT_BUILD}
5     MODULE_LIBRARIES
6     ${TOYEXPERIMENT_MCDATAPRODUCTS}
7     ${ART_FRAMEWORK_CORE}
8     ${ART_FRAMEWORK_PRINCIPAL}
9     ${ART_PERSISTENCY_COMMON}
10    ${ART_FRAMEWORK_SERVICES_REGISTRY}
11    ${ART_FRAMEWORK_SERVICES_OPTIONAL}
12    ${ART_FRAMEWORK_SERVICES_OPTIONAL_TFILESERVICE_SERVICE}
13    ${FHICL_CPP}
14    ${CETLIB}
15    ${ROOT_HIST}
16 )

```

1 from *art*, including the final line saying that *art* will exit with status 0. Remember to
 2 add `>& output/<filename>.log` to the end of the command to send the printout to a
 3 file. The Workbook will not always show this in subsequent exercises, but it is always
 4 recommended.

5 You should see that the *art* job created the file `output/firstHist1.root`. This is
 6 the histogram file.

7 17.9 Inspecting the Histogram File

inspect-histogram

8 In this section you will inspect the file `output/firstHist1.root`.

9 First, look again at `fcl/FirstHistogram/firstHist1.fcl`. Note that the mod-
 10 ule label of the `FirstHist1` module is `hist1`.

11 To inspect the histogram you will remain in your build directory and you will run the
 12 interactive `ROOT` program, using the command `root`. This command was put into your
 13 path when you established your build environment. To perform this exercise:

1. Enter:

```
root -l output/firstHist1.root
```

The command line option is a lower case letter L. In your build window, some output and a new prompt will appear:

```
root [0]\\
Attaching file output/firstHist1.root as \_file0...\
root [1]}
```

2. At the root prompt, type the command

```
TBrowser* b = new TBrowser("Browser", _file0);
```

This will open a new window on your display; a screen capture of this window is shown in Figure 17.1. We will refer to this window as the TBrowser window.

3. In the left hand panel of the TBrowser window, you will see a ROOT file icon followed by the name of the file `output/firstHist1.root`. Double click on this line.
4. This will create a new line in the left hand panel of the TBrowser window. The line contains a folder icon followed by the name of a folder, `hist1;1`. Double click on this line.
5. This will create another new line in the left hand panel of the TBrowser window. This line contains a blue histogram icon and the name of a histogram `hNGens;1`. Double click on this line.
6. The histogram will appear in the right hand panel of the TBrowser window. Figure 17.2 shows a screen capture of the window with the histogram drawn.

7. To exit root, return to the build window and, at the root prompt, type `.q` (a period followed by a lower case letter Q).

`.q`

8. Another way to quit root is to click on the “Browser” pull-down menu on the top bar of the TBrowser window. From the menu select “Quit ROOT”.

1

2 In step 4 you should have recognized the name of the folder, `hist1;1`. Ignoring the
 3 trailing `;1`, it is the name of the module label used in `firstHist1.fcl`. In step 5 you
 4 should have recognized the name of the histogram, `hNGens`; ignoring `;1`, it is the name
 5 that you gave the histogram when you created it.

6 About the `;1` that ROOT has stitched onto `hist1` and `hNGens`: ROOT calls these *cycle*
 7 *numbers*. They are part of a checkpointing mechanism that is beyond the scope of this
 8 exercise; if you ever see more than one cycle number for a ROOT object, the highest
 9 number is the one that you want. Consult the ROOT documentation for more details.

10 Now look at the histogram in the right hand panel of the TBrowser window. In the statistics
 11 box on the upper right you should see that it has 1000 entries, one for each event in the
 12 input file. You should also notice that only the odd bins are populated: this is because the
 13 generated events always contains three signal particles, plus a random number of pairs
 14 of background particles ($3 + 2n$). The three signal particles are φ meson and the two
 15 kaons into which it decays. You should also recognize the title and the name that you set
 16 when you created the histogram. Finally you should recognize that the binning matches
 17 the binning you requested when you created the histogram.

18 17.9.1 A Short Cut: the `browse` command

firstHist1-browse

19 The above description for viewing a ROOT file interactively requires a lot of tedious typing
 20 at step 2. The toyExperiment UPS product provides a command named `browse` that does
 21 the typing for you. To use this command:

22 `browse output/firstHist1.root`

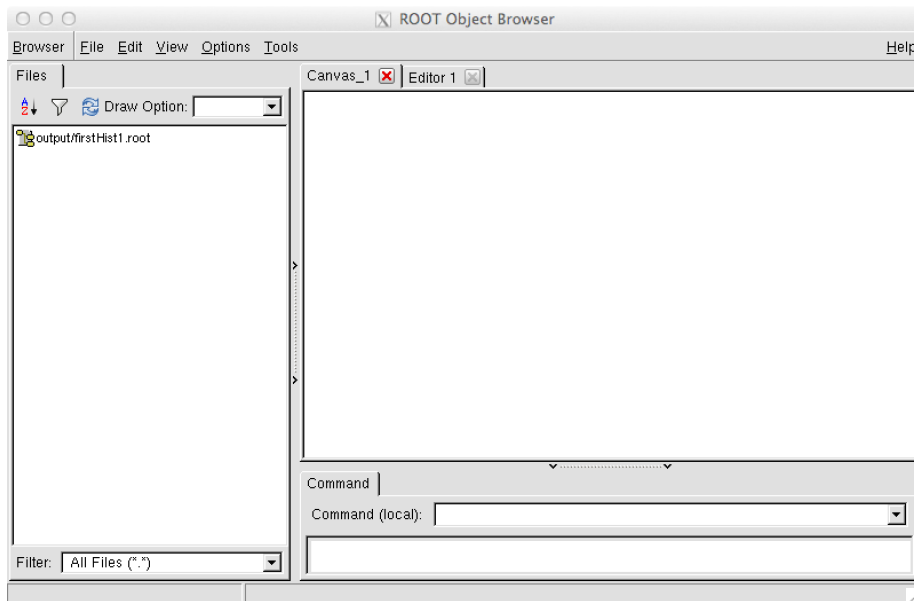


Figure 17.1: The TBrowser window immediately after opening `output/firstHist1.root`.

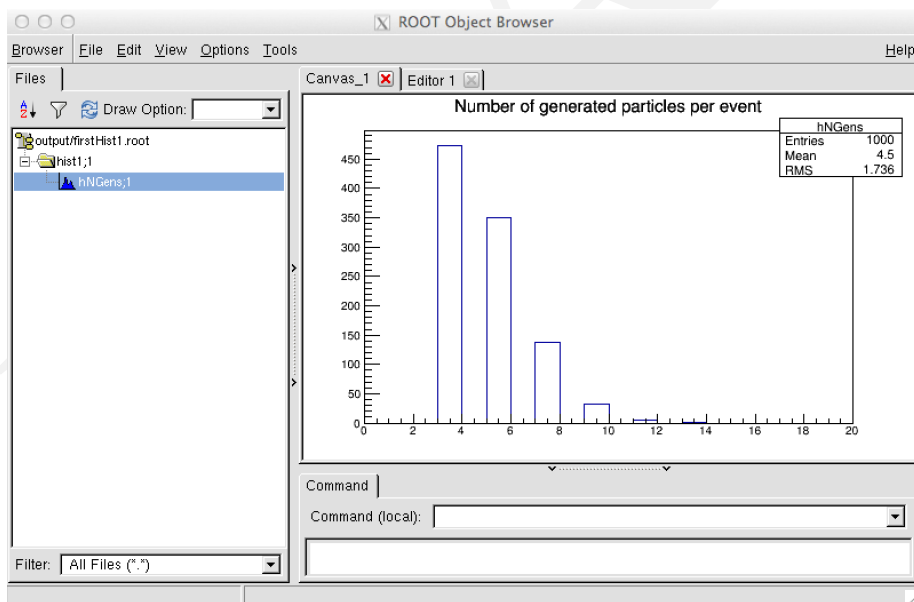


Figure 17.2: TBrowser window after displaying the histogram `hNGens;1`.

1 Then follow the instructions from the previous section, starting at step 3.

2 When you created your *art* build environment, the toyExperiment UPS product put the
3 command `browse` into your path. This command is implemented as a bash script and you
4 can find its definition using the following bash command:

5 `type browse`

6 `browse is <...>/scripts/browse`

7 where `<...>` will change from one site to the next; it will have the value of
8 `$TOYEXPERIMENT_DIR` as defined at that site.

9 **FIXME:** *Mine says 'hashed':*

10 `[aheavey@cluck build-prof]$ type browse`

11 `browse is hashed (/home/kutschke/products//toyExperiment/v0_00_15/scri`

12 `[aheavey@cluck build-prof]$ echo $TOYEXPERIMENT_DIR`

13 `/home/kutschke/products//toyExperiment/v0_00_15`

14 `[aheavey@cluck build-prof]$`

15 **FIXME:** *end fixme; Rob should check this*

16 17.9.2 Using CINT Scripts

17 When you type a command at the root prompt, you are typing commands in a ROOT-
18 defined language called CINT. Scripts in this language are sometimes called ROOT scripts
19 and other times CINT scripts.³ As with many interpreted languages, you may write CINT
20 commands in a file and execute that file as script.

21 It is a common convention that files that contain CINT scripts have a file type of `.C`. This
22 convention is followed throughout the *art* workbook.

23 This exercise provides an example of a CINT script, `drawHist1.C`, in Listing 17.5.

³*CINT* is a (somewhat misleading) acronym for C++ INTERpreter. While CINT will correctly execute a lot of C++ code, there is legal C++ code that is not legal CINT and there is legal CINT that is not legal C++. There is also code that is legal in both C++ and CINT but that does subtly different things in the two environments. Therefore it is more correct to say that the CINT language shares a lot of syntax with C++, not that it *is* C++.

rawHist1.C

Listing 17.5: Sample CINT file DrawHist1.C

```


1 float
2 //
3 // Root script to draw the histogram made by FirstHist1_module.cc
4 //
5 {
6
7 // With this you can reinvoke the script without exiting root
8 // and restarting.
9 gROOT->Reset();
10
11 // Get rid of grey background (ugly for printing).
12 gROOT->SetStyle("Plain");
13
14 // Recommended content of statistics box:
15 // Number of Entries, Mean, Rms, Underflows, Overflows,
16 // Integral within limits
17 gStyle->SetOptStat("emruoi");
18
19 // Open the input file that contains histogram.
20 TFile* file = new TFile("output/firstHist1.root");
21
22 // Get pointer to the histogram.
23 TH1D* hNGens; file->GetObject("hist1/hNGens", hNGens);
24
25 // Open a new canvas on the screen.
26 TCanvas *canvas = new TCanvas("canvas", "Plots_from_Firsthist1.root");
27
28 // "H9": draw outline histogram ("H") in high resolution mode (9)
29 hNGens->Draw("H9");
30
31 canvas->Update();
32 canvas->Print("output/NumberGenerated.pdf");
33
34 }

```

36 To use this script, run:

```
37 root -l drawHist1.C
```

38 This will open a window on your display, draw the histogram in that window and save
39 the window to the PDF file output/NumberGenerated.pdf, which is shown as



`fig:drawHist1`

Figure 17.3: Figure made by running the CINT script `drawHist1.C`.

1 `fig:drawHist1` Figure 17.3. When the script is complete, it returns control to the root prompt in your
 2 build window. At this prompt you can issue more CINT commands. To exit ROOT, type
 3 `.q` at the root prompt.

4 `ch:ug:viewing:printing` Appendix D has instructions on how to view the PDF file interactively. For those of you at
 5 Fermilab, it also has instructions on how to print the PDF file.

6 `fig:firstHist1` If you compare this figure to the histogram in Figure 17.2 you will see that there is differ-
 7 ence in the statistics box in the upper right. In this figure, the name of the histogram is not
 8 shown but three new fields are: the number of entries below the lower limit (Underflow),
 9 the number of entries above the upper limit (Overflow) and the number of entries between
 10 the limits (Integral). The field named “Entries” is the sum of Integral plus Underflow plus
 11 Overflow.

12 `lst:drawHist1:C` It is beyond the scope of this writeup to describe all of the features in Listing 17.5, but we
 13 will describe some of the code.

14 Regarding lines 9 and 12, we will let comments in the code be a guide, and give two
 15 additional hints. The object `gROOT` is a pointer to an instance of the class `TROOT` and the
 16 object `gStyle` is a pointer to an instance of the class `TStyle`. To find the documentation

1 for these classes, see Section [17.10](#). sec: first-histo-firstHist1-cint-rootdocs

2 Line 17 tells ROOT what to draw in the statistics box in the upper right of every histogram.
3 The comments in the code describe the mnemonics of the letter codes. The full set of letter
4 codes is described on the ROOT web site: <http://root.cern.ch/root/html534/TStyle.html> as
5 part of the documentation for the member function `SetOptStat`.

6 Line 20 opens the input file. The ROOT type `TFile` is ROOT's interface to ROOT infor-
7 mation that lives in a disk file; it allows a ROOT program to write ROOT objects to the file
8 and read ROOT objects from the file. You can learn more about the class `TFile` from the
9 ROOT web site <http://root.cern.ch/root/html534/TFile.html>.

10 Line 23 has two statements on it. The first declares `hNGens` as an object of type pointer to
11 `TH1D` and initializes it to 0 (ROOT does not support `nullptr`). The second asks the `TFile`
12 to find a ROOT object named "hist1/hNGens", copy it from the file into memory and to
13 set the pointer `hNGens` to point to that object. If ROOT cannot find the requested object,
14 or if the type of the requested object does not match the type of the pointer, ROOT will
15 leave `hNGens` with a value of 0. This is reminiscent of asking an `art::Event` to fill an
16 `art::Handle`.

17 Line 26 tells ROOT to open a new window on your display. The first argument is an
18 arbitrary name that must be unique within the job; ROOT uses it internally to differentiate
19 multiple canvases. The second argument is the title that will be drawn on the title bar of
20 the window.

21 Line 29 tells ROOT to draw the histogram on the canvas. If, at line 21, ROOT was unable
22 to properly set the pointer, then this line will produce an error message and return control
23 to the root prompt in the build window.

24 Line 31 tells ROOT to flush its internal buffers and make sure that everything that is in the
25 queue to be drawn on the canvas is actually drawn.

26 Line 33 tells ROOT to save the canvas by writing it to the file specified as the function
27 argument. The format in which the file will be written is governed by the file type field
28 in the filename, `.pdf` in this case. Many other formats are supported and a full list is
29 available at: <http://root.cern.ch/root/html534/TPad.html#TPad:Print>

17.10 Finding ROOT Documentation

The main ROOT web site is <http://root.cern.ch/drupal>. On the top navigation bar there is a title labeled *Documentation*. Hover over this and a pull-down menu will appear. From this menu you can find links to Tutorials, How To's FAQs, a User Guide, a Reference Manual and more. Some of the documentation is version-dependent. The version of ROOT used by this version of the workbook is v5.34/30.

One possible starting point for learning ROOT is the ROOT Primer. You can find it by first going to the User's Guide page or you can follow the direct link:
<http://root.cern.ch/drupal/content/users-guide#primer>

One of the most useful parts of the ROOT documentation suite is the Reference Guide, which can be reached from the pull down menu. The direct link to this page is:
<http://root.cern.ch/root/html534/ClassIndex.html>
This section has a description of all of the members in each ROOT class.

17.10.1 Overwriting Histogram Files

Suppose that you run *art* and make a histogram file. If you run *art* again, what happens? The answer is that it will overwrite the existing histogram file and replace it with the one created in the second job.

To illustrate this, rerun the exercise but tell *art* to only do 500 events:

```
art -c fcl/FirstHistogram/firstHist1.fcl -n 500
```

Inspect the histogram file and you will see that the histogram now has only 500 entries.

It is your responsibility not to overwrite files that you wish to keep. One way to keep a file that is valuable is to use the unix `chmod` command to change the protections on the file so that it is readonly:

```
chmod -w <filename>
```

1 To restore the file to a writeable state the unix command is:

2 `chmod o+w <filename>`

3 17.10.2 Changing the Name of the Histogram File

4 You can change the name of the histogram file by editing the FHiCL file but you can also
5 do so from the *art* command line by using the `--TFileName` option; the short form of
6 this option is `-T`.

7 Run the two following commands, both shown here on two lines (note that in both cases a space is required before the backslash since a space exists in each command line at these points):

```
art -c fcl/FirstHistogram/firstHist1.fcl --TFileName \  
output/anotherName.root -n 400
```

```
art -c fcl/FirstHistogram/firstHist1.fcl -T \  
output/yetAnotherName.root -n 750
```

8 After each run, inspect the output file and verify that the number of entries in the histogram matches the number of events requested on the command line.

9 17.10.3 Changing the Module Label

10 In `firstHist1.fcl`, change the name of the module label, `hist1`. Rerun the job and
11 browse the histogram file. You should see that the name of the directory in the histogram
12 file has changed to match the new module label.

13 17.10.4 Printing From the TBrowser

14 You can use the `browse` command to open the histogram file and view the histogram.

1 In the TBrowser window, click on the “File” button. This will open a pull-down menu.

2 Click on the line “Save As ...”. This will open a dialog window that will let you save the histogram displayed on the canvas in a variety of formats, including .png, .gif, .jpg and .pdf.

3 17.11 Review

4 In this exercise you have learned:

- 5 1. How to configure the TFileService.
- 6 2. How to use an `art::ServiceHandle` to access the TFileService.
- 7 3. How to use the TFileService to create a histogram that will automatically be
8 written to the histogram file.
- 9 4. Three different ways to view the contents of the histogram file: by launching a
10 TBrowser by hand, by using browse command and by running a CINT script.
- 11 5. The convention used by the *art* Workbook to differentiate histogram files from *art*
12 event data files.

13 **FIXME:** *Would be nice to have targeted tests for these things ... true for each chapter;*
14 *just thinking as I go.*

15 17.12 Test Your Understanding

16 17.12.1 Tests

17 Two source files and one FHiCL file are provided that contain intentional bugs. Your job, of
18 course, is to find the bugs and fix them. The answers are provided in Section 17.12.2.

1 The source files listed below fail to build. Follow essentially the same procedure as out-
2 lined in Section [10.12](#) on each of these files. To run them in *art*, use the FHiCL file that
3 matches the number 1 or 2, respectively.

4 `FirstHistogram/FirstHistBug01_module.cc.nobuild`

5 `FirstHistogram/FirstHistBug02_module.cc.nobuild`

6 Next, run *art* on this FHiCL file that has a problem; then diagnose and fix it:

7 `FirstHistogram/bug03.fcl`

8 The answers are intentionally placed on a new page (remember to try before you read
9 further!).

DRAFT

17.12.2 Answers

1

ex7:test-answers

2 When you run `buildtool` after you have made the file `FirstHistBug01_module.cc`,
3 the output includes the following error message (lines split here for readability):

```
4 /home/aheavey/workbook/art-workbook/art-workbook/
5     FirstHistogram/FirstHistBug01_module.cc: In member function
6     'virtual_void_tex::FirstHistBug01::analyze(const_art::Event&)':
7
8 /home/aheavey/workbook/art-workbook/art-workbook/
9     FirstHistogram/FirstHistBug01_module.cc:56:10: error:
10    base operand of '->' has non-pointer type 'TH1D'
11
12    hNGens_->Fill(gens->size());
```

13 The cause of the problem is that `hNGens_` was declared as type `TH1D`, not as a *pointer* to
14 type `TH1D`. The solution is to add an asterisk to the type declaration so that it reads

```
15     private:
16         ...
17     TH1D* hNGens_;
```

18 **FIXME:** *Would be nice to say why it was declared that way, ...*

19 When you run `buildtool` after you have made the file `FirstHistBug02_module.cc`,
20 the output includes the following error message (lines split here for readability):

```
21 /home/aheavey/workbook/art-workbook/art-workbook/FirstHistogram/
22     FirstHistBug02_module.cc:
23 In member function
24     'virtual_void_tex::FirstHistBug02::analyze(const_art::Event&)':
25
26 /home/aheavey/workbook/art-workbook/art-workbook/
27     FirstHistogram/FirstHistBug02_module.cc:56:11:
28 error: request for member 'Fill' in
29     '((tex::FirstHistBug02*)this)->tex::FirstHistBug02::hNGens_',
30 which is of pointer type 'TH1D*' (maybe you meant to use '->'?)
```

31 The solution is to replace the dot in the `Fill` line with a `->` in order to dereference the
32 pointer `hNGens_`:

```
33     hNGens_->Fill(gens->size());
```

1 hNGens_ is not the histogram class (TH1D) that has a Fill member function; it is a
2 pointer to that class. **FIXME:** *Please correct/embellish description*

3 When you try to run:

```
4 art -c fcl/FirstHistogram/bug03.fcl
```

5 you will obtain the following error message:

```
6 RootOutput cannot ascertain a unique temporary filename for output  
7 based on stem "outputs/TFileService": No such file or directory.
```

8 To fix bug03.fcl, correct the spelling of outputs to output.

DRAFT

18 Exercise 8: Looping Over Collections

ing-genparticles

18.1 Introduction

In Chapter 16 you learned about the data product that describes all of the particles in a generated event. It has a type of `tex::GenParticleCollection` and an input tag of `evtgen`. In this data product each generated particle is represented by an object of type `tex::GenParticle`. This chapter will introduce you to most of the properties of a `tex::GenParticle` and to looping over all of the `tex::GenParticle` objects in the data product. The remaining properties of a `tex::GenParticle` object will be discussed in Chapter 20.

18.2 Prerequisites

Prerequisites for this chapter include all of the material in Part I (Introduction) and the material in Part II (Workbook) up to and including Chapter 17.

18.3 What You Will Learn

In this exercise you will learn:

1. the following properties of a `tex::GenParticle`
 - (a) the particle data group ID code, `PDGCode::type`
 - (b) the position of the particle when it was created
 - (c) the 4-momentum of the particle when it was created

- 1 2. the recommended way to write a loop over a `GenParticle` collection
- 2 3. seven other ways to write the same loop; you will encounter many of these in code
- 3 written by others
- 4 4. about the classes `CLHEP::Hep3Vector` and `CLHEP::HepLorentzVector`
- 5 and how to use them

6 18.4 Setting Up to Run Exercise

setting:up

Follow the instructions in Chapter 11 if you are logging in after having closed an earlier session. If you are continuing on directly from the previous exercise, keep both your source and build windows open.

In your source window, look at the contents of the directory for this exercise, called `LoopOverGens`:

```
ls art-workbook/LoopOverGens
```

```
CMakeLists.txt          LoopGens3_module.cc
loopGens1.C             loopGens4.C
loopGens1.fcl           loopGens4.fcl
LoopGens1_module.cc    LoopGens4_module.cc
loopGens2.C             LoopGens5_module.cc.nobuild
loopGens2.fcl           LoopGens6_module.cc.nobuild
LoopGens2_module.cc    LoopGens7_module.cc.nobuild
loopGens3.C             loopGens8.fcl
loopGens3.fcl           LoopGens8_module.cc
```

In your build window, just make sure that you are in your build directory.

- 7
- 8 The main body of the exercise consists of the files `LoopGens1_module.cc`,
- 9 `loopGens1.fcl` and `loopGens1.C`. The files with the numerals 2 and 3 in their
- 10 names show alternate ways to write the loop over the generated particles. The files with

1 the numeral 4 in their names collectively provide the solution to an exercise that will be
 2 assigned towards the end of this chapter. The files with the numerals 5, 6, 7, 8 in their
 3 names contain deliberate errors that you will be asked to find and fix.

4 18.5 The Class `GenParticle`

5 This chapter discusses the data product `GenParticleCollection`, which you learned
 6 in Section 16.4.1 is a `std::vector<GenParticle>`. This section will describe the
 7 class `GenParticle`, which contains information about one generated particle.

8 The header file and the source file can be found at:

```
9 $TOYEXPERIMENT_INC/toyExperiment/MCDataProducts/GenParticle.h
10 $TOYEXPERIMENT_DIR/source/toyExperiment/MCDataProducts/GenParticle.cc
```

11 Open the files to follow along with the discussion.

12 **FIXME:** *Add url to code in the LXR browser.*

13 18.5.1 The Included Header Files

14 The header file `GenParticle.h` includes six header files that have not yet been used in
 15 the Workbook:

```
1b #include "toyExperiment/DataProducts/PDGCode.h"
1c #include "art/Persistency/Common/Ptr.h"
1d #include "art/Persistency/Common/PtrVector.h"
1e #include "CLHEP/Vector/LorentzVector.h"
1f #include "CLHEP/Vector/ThreeVector.h"
2g #include <iosfwd>
```

22 The PDG in the name of first header file refers to the *Particle Data Group*, which has
 23 defined a standard set of about 400 integer codes to identify particles in simulations of HEP
 24 events. This scheme provides identifiers for individual quarks and leptons, for their anti-
 25 partners, for the gauge bosons, for the Higgs, for the graviton, for hadrons and for many
 26 proposed particles that are beyond the Standard Model. This scheme has been broadly
 27 adopted by the HEP community and is described on the PDG web site:
 28 <http://pdg.lbl.gov/2007/reviews/montecarlohpp.pdf>

1 The Workbook includes a class that declares and initializes a small subset of the PDG
2 codes, only those that are of interest to the toy experiment. It does so by way of an *enum*,
3 which is a standard C++ type that sets a restricted range of possible values to a variable.
4 This class is defined in:

```
5 $TOYEXPERIMENT_INC/toyExperiment/DataProducts/PDGCode.h
```

6 The only content of this class is the enum that defines mnemonic names for each code.
7 By using the mnemonic names instead of the integer codes, your code will be under-
8 standable by people who do not have all of the codes committed to memory. To use
9 one of the codes in the Workbook, you need to give its full name; for example, use
10 `PDGCode::pi_minus` to specify the code for a π^- . If you are not familiar with enums,
11 consult any standard C++ text.

12 The class templates from the next header files, `Ptr.h` and `PtrVector.h`, will be dis-
13 cussed in Chapter 20. For this exercise you need to know only that these classes provide a
14 type of smart pointer that can be written to disk and read back. In the `GenParticle` class
15 they are used to record the parent-child relationships among the generated particles.

16 **FIXME:** *Rob to Anne: answering the above. When you write an object out and read it
17 back it, it will in general live at a different memory address. So something needs to figure
18 out what the new memory address is, track down the pointers that point to it and update
19 them with the new address. This will be discussed in the referenced chapter. The problem of
20 how to make pointers that survive being written to a file is really hard and there is no fully
21 satisfactory solution. The ones that “just do everything” are too slow for practical use
22 so we need to find the right set of compromises. Every framework in HEP has a different
23 solution that reflects the priorities set by the experiments using that framework. AH:ok.
24 Am keeping this fixme for my info*

25 The header files `ThreeVector.h` and `LorentzVector.h` are part of a package
26 named CLHEP, which is a library of classes widely used in HEP. If you are not familiar
27 with CLHEP, there is general discussion in Appendix E. The first file contains the decla-
28 ration of the class `CLHEP::Hep3Vector` that represents 3-vectors in 3-dimensional
29 space. The second file contains the delaration of the class `CLHEP::HepLorentzVector`
30 that represents a 4-vector in 4-dimensional space-time. These classes are used to describe,
31 respectively, the position and the 4-momentum of the `GenParticle` at the time that it
32 was created. The header files can be found at:

```
1 $CLHEP_INC/CLHEP/Vector/ThreeVector.h
2 $CLHEP_INC/CLHEP/Vector/LorentzVector.h
```

3 You can find specific discussions about `Hep3Vector` in Section E.6.1 and about `LorentzVector` in Section E.6.2.

5 The last of the six newly introduced headers is `<iosfwd>` which is part of the C++ standard library. It contains forward declarations (a standard C++ technique that allows the compiler to work more efficiently) for classes, functions and objects that are declared in the headers `<istream>`, `<ostream>` and `<iostream>`. By using `<iosfwd>` here instead of one of the other three headers, some code that includes `GenParticle.h` will compile more quickly.

11 18.5.2 Particle Parent-Child Relationships

12 One of the important notions in many HEP event generators is that some particles are created by a primary interaction; these are called primary particles. Other particles are created by the decay of one of the primary particles or by the interaction of one of the primary particles with material in the experiment; these are called secondary particles.

16 Therefore one of the ideas encoded in a `GenParticle` is the identity of its parent particle and the identities of its children, if any. A common convention, one that is used in the Workbook, is that primary particles have no parent.

19 18.5.3 The Public Interface for the Class `GenParticle`

20 The first element of the class declaration for `GenParticle` in `GenParticle.h` is a typedef:

```
22 typedef art::PtrVector<GenParticle> children_type;
```

23 This type is a collection of smart pointers that allows you to find the children of a given `GenParticle` and it will be discussed in Chapter 20.

25 The `GenParticle` class has two constructors:

```
26 GenParticle();
```

27 which is the default, and another that declares a few data members:


```

1     GenParticle( PDGCode::type           pdgId,
2                 art::Ptr<GenParticle> const& parent,
3                 CLHEP::Hep3Vector const& position,
4                 CLHEP::HepLorentzVector const& momentum );

```

5 It also has some private member functions with access to them provided by (public) acces-
6 sor functions; accessors are discussed in Section 6.7.7. The accessor functions are given
7 in Table 18.1.

8 The public interface has one non-const member function that is used to add a child particle
9 to a GenParticle:

```

10    void addChild( art::Ptr<GenParticle> const& child );

```

11 This member function will not be discussed here but it will be discussed *in a future*
12 *chapter that discusses the EventGenerator module in the toy experiment UPS package.* It
13 is not needed when you are using GenParticles that already exist, as in this exercise, only
14 when making new ones. For this exercise, you only need to concern yourself with the eight
15 accessors. the remainder of this section will discuss each of them. **FIXME:** *Reference to*
16 *future chapter when available*

17 Finally there is a do-nothing definition of operator less-than; this is required to work
18 around a bug in a tool named `genreflex` that is part of ROOT's persistency system.
19 This operator should never be called because it does not do anything meaningful. Once the
20 bug is fixed, we will remove this operator.

21 The member function `position` returns the position at which the particle was gener-
22 ated. The return type is `CLHEP::Hep3Vector` which was described in Section 18.5.1;
23 additional details can be found in Section E.6.1.

24 The member function `momentum` returns the 4-momentum of the particle at the place of
25 its creation. The return type is `CLHEP::HepLorentzVector` which was described in
26 Section 18.5.1; additional details can be found in in Section E.6.2.

27 As was described in Table 3.2, the position is given in mm and the 4-momentum is given
28 in MeV.

29 The member function `pdgId` returns the value of the `PDGCode`; the next exercise, in
30 Chapter 19, will describe the type `PDGCode::type` in more detail.

article:accessors

Listing 18.1: The const accessors of class `GenParticle`. The order has been changed from that found in the header file to match the order of discussion in the text. The white space has also been altered to make the code fit better in the narrow page of this document.

```
1
2 CLHEP::Hep3Vector      const& position() const { return _position; }
3 CLHEP::HepLorentzVector const& momentum() const { return _momentum; }
4
5 PDGCode::type pdgId() const { return _pdgId;    }
6
7 children_type const& children() const {
8     return _children;
9 }
10
11 bool hasParent() const { return _parent.isNonnull(); }
12 bool hasChildren() const { return !_children.empty(); }
13
14 art::Ptr<GenParticle> const& parent() const {
15     return _parent;
16 }
17
18 art::Ptr<GenParticle> const& child( size_t i ) const {
19     return _children.at(i);
20 }
```

1 The member function `children` returns a `const` reference to the collection of smart
2 pointers to the children of this `GenParticle`. For this exercise you only need to know
3 one thing: `art::PtrVector` has a member function:

```
4 std::size_t size() const;
```

5 that returns the size of the collection, which corresponds to the number of children of the
6 `GenParticle`.

7 The member function `hasParent` returns `true` if the `GenParticle` has a parent. In
8 the `toyExperiment`, this will only be true for the daughters of the ϕ meson; all others are
9 primary particles.

10 The member function `hasChildren` returns `true` if the `GenParticle` has children. In
11 the `toyExperiment`, this will only be true for the ϕ meson.

12 You will notice that the implementations of the member functions `hasParent` and
13 `hasChildren` are simple, almost trivial, logic on data members. So why bother to pro-
14 vide them? The reason is that coding

```
15 if ( gen.hasParent() ) { ... }
```

16 makes the programmer's intent manifest, while coding

```
17 if ( gen.parent().isNonnull() ) { ... }
```

18 requires that the reader draw some context-dependent inferences to understand the pro-
19 grammer's intent.

20 The remaining two member functions have to do with parent-child navigation and will be
21 described in Chapters 19 and 20.

22 In `GenParticle`, all accessors follow the recommendation that accessors for large ob-
23 jects return their information by `const` reference while accessors for small objects return
24 their information by value. With very few exceptions, the data product classes in the `toy-`
25 `Experiment` and in the `Workbook` follow this pattern.¹

26 **FIXME:** *At some point we will have a best practices section that describes this recom-*
27 *mendation. When it is written, add a reference here..*

¹The boundary line between large and small is usually taken to be the size of a pointer, 8 bytes on most computers on which `art` will run; a pointer is included in the set of small objects.

1 The final part of the public interface comes after the declaration of the class.² It is the free
2 function, called the stream insertion operator:

```
3 std::ostream& operator<<(std::ostream& ost, const tex::GenParticle& genp );
```

4 If `gen` is an object of type `GenParticle`, this allows you to print it using the syn-
5 tax:

```
6 std::cout << gen << std::endl;
```

7 18.5.4 Conditionally Excluded Sections of Header File

8 ROOT has a tool called `genreflex` that is used to generate ROOT dictionaries. These
9 are needed for all classes that are part of a data product but *art* does not need a complete
10 dictionary: it only needs to know about the data members, the default constructor and, if
11 present, the destructor. To reduce both the time needed to create the dictionaries and the
12 memory footprint of the dictionaries, it makes sense to hide code from `genreflex` that
13 it doesn't need. In addition, `genreflex` does not understand C++11 syntax; therefore
14 C++11 features in a data product class must be hidden from `genreflex`. To enable this
15 exclusion, `genreflex` sets a C-preprocessor macro `__GCCXML__` that can be used to
16 hide portions of code from `genreflex`. Two sections of `GenParticle.h` are condi-
17 tionally marked for exclusion from `genreflex` in this way:

```
18 #ifndef __GCCXML__  
19 // Code to be excluded  
20 #endif
```

21 *A future chapter will describe a use case for which it is useful to generate complete
22 dictionaries and will describe how to do so.*

23 **FIXME:** *Add reference to the future section after it is written.*

24 18.6 The Module `LoopGens1`

25 This module has some new features.

²Why is it there? It is not part of the class so it must be declared outside of the class declaration (i.e. after the closing `};`). There are many different sorts of “helper” free functions that are not part of the class but are closely associated with it, and which must appear in the same namespace. It is a common convention to declare these functions in the same header file as the class.

- 1 1. The class declaration declares two new histogram pointers, `hP_` and `hcz_`.
 - 2 (a) `hP_` will point to a histogram of the magnitude of the 3-momentum of each
 - 3 generated particle.
 - 4 (b) `hcz_` will point to a histogram of the cosine of the polar angle of the 3-
 - 5 momentum of each generated particle.
- 6 2. The constructor initializes these pointers to `nullptr`.
- 7 3. The `beginJob` member function creates the two histograms and sets `hP_` and
- 8 `hcz_` to point to them.
- 9 4. The `analyze` member function loops over the generated particles and fills these
- 10 histograms with one entry per particle.

11 Listing 18.2 shows the loop over the generated particles, found in the `analyze` member
 12 function of `LoopGens1_module.cc`. This loop uses a feature of C++ that intro-
 13 duced in C++-11 called the *range-based for loop*. You can read about range-based for
 14 loops in a standard C++ text that is up-to-date with C++-11 features.

15 The simple picture of this range-based for loop is that the body of the loop will be executed
 16 once for each element in the collection specified by the expression to the right of the colon.
 17 On the first pass, `gen` will be a const reference to the first element in the collection. On
 18 the second pass it will be a const reference to the second element in the collection and
 19 so on.

20 Had the first line of the loop been written,

```
21 for ( GenParticle gen : *gens ) {
```

Listing 18.2: The loop over the generated particles, from `LoopGens1_module.cc`

```
1 for ( GenParticle const& gen : *gens ) {
2
3   CLHEP::HepLorentzVector const& p4 = gen.momentum();
4   CLHEP::Hep3Vector const& p      = p4;
5
6   hP_ -> Fill( p.mag() );
7   hcz_ -> Fill( p.cosTheta() );
8
9 }
```

1 The code would have compiled and it would have produced the correct results. However,
 2 on each pass of the loop, `gen` would be a copy of the corresponding element of the col-
 3 lection. This copy is unnecessary and will cause the code to be slower.

4 Why is there a `*` in front of `gens`? In the opening line of a range-based for loop, the
 5 expression on the right-hand side of the colon must be a collection type. Recall that the
 6 object named `gens` is a handle to a `GenParticleCollection`. Therefore it is not
 7 itself a collection type. On the other hand, the expression `*gens` dereferences the han-
 8 dle and returns a `const` reference to the pointee of the handle, which in this case is a
 9 `GenParticleCollection`. This *is* a collection type and therefore it can be put to the
 10 right of the colon. If you are unsure of how this works, the key is to know that, used in this
 11 context, the character “`*`” is known as the *dereferencing operator*. You can look this up in
 12 any standard C++ text. Then you can consult the header for `art::ValidHandle` to see
 13 that it obeys the standard conventions for the dereferencing operator.

14 What sort of collection types may be on the right-hand side of the colon? Any standard
 15 library collection type is allowed. User-written collection types may also be used.

16 The first two lines of the body of the loop are

```
17 CLHEP::HepLorentzVector const& p4 = gen.momentum();
18 CLHEP::Hep3Vector const& p      = p4;
```

19 Recall that a reference behaves like a compile-time alias; therefore there is no run-time
 20 cost for these lines to be present. The code is written this way because the author thinks
 21 that it makes the loop body easier to read.

22 The last two lines of the loop body fill the histograms with the requested information. If
 23 you are not familiar with the accessor member functions of `Hep3Vector` consult Sec-
 24 tion E.6.1.

25 As a last comment, the line that defines the variable `p` uses a somewhat obscure fea-
 26 ture of `CLHEP`: a `HepLorentzVector` vector has an implicit conversion operator
 27 to a `Hep3Vector`. This operator returns a `const` reference to the space part of the
 28 `HepLorentzVector`. One might have considered writing this line differently, for ex-
 29 ample

```
30 CLHEP::Hep3Vector p = p4.vect();
```

- 1 While this version makes the intent more explicit, and perhaps makes it a little easier to
- 2 understand, it also forces an unnecessary copy. So the authors of the Workbook elected to
- 3 choose the more efficient, albeit more obscure, variant.

4 **18.7** CMakeLists.txt

cmakelists

- 5 The file CMakeLists.txt is unchanged from the previous exercise.

6 **18.8** Running the Exercise

ng:running

7

To run this exercise, cd to your build directory and run *art*:

```
cd $ART_WORKBOOK_WORKING_BASE/<username>/\  
workbook/build-prof
```

```
art -c fcl/LoopOverGens/loopGens1.fcl >& output/loopGens1.log
```

art should complete with status 0. It does not make any interesting printout. The next step is to view the histograms.

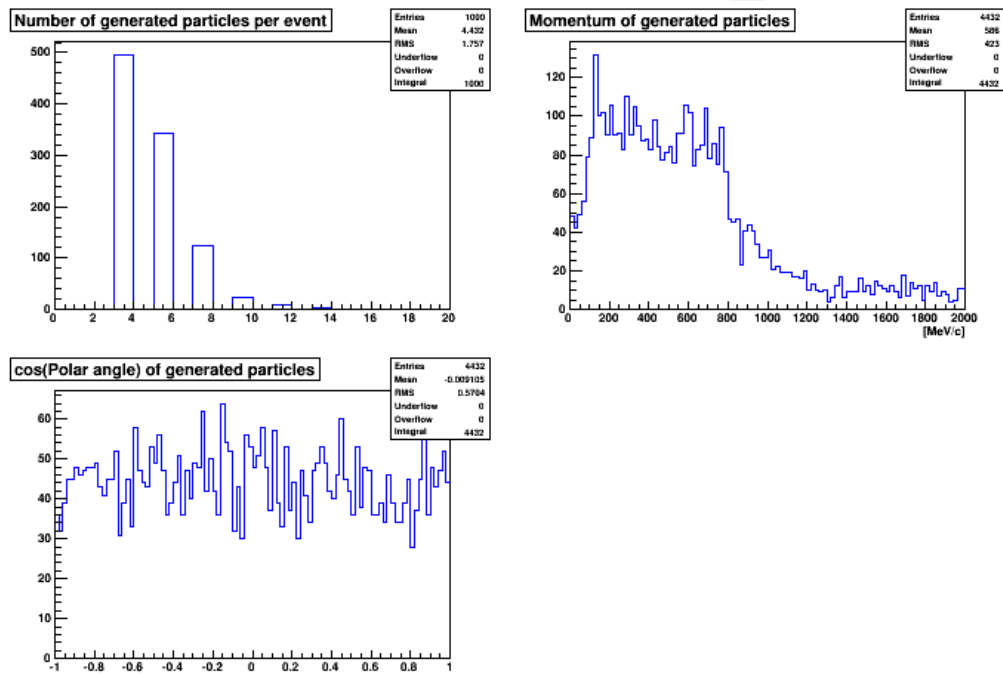
```
root -l fcl/LoopOverGens/loopGens1.C
```

8

- 9 The last step will make a figure like that shown in Figure 18.1. It will also write the figure
- 10 to the file `output/loopGens1.png`. Appendix D has instructions on how to view the
- 11 `.png` file interactively.

- 12 For those of you at Fermilab, it also has instructions on how to print the `.png` file.



`fig:loopGens1`**Figure 18.1:** Histograms made by `loopGens1.fcl`

18.9 Variations on the Exercise

18.9.1 LoopGens2_module.cc

The file `LoopGens2_module.cc` shows another way to write the loop over generated particles. Compared to the first version, there are three differences, the first of which is the most important:

1. The body of the loop has been moved into a new member function, `fillParticleHistograms` that is in the private section of the class declaration.
2. In the range-based for loop, the keyword `auto` was used to tell the compiler to automatically figure out the correct type for `gens`.
3. Inside `fillParticleHistograms`, the computation of `p` has been reduced from two lines to one.

The main reason for showing this variation is to advocate for writing code in more small functions rather than in fewer larger functions. In this example the body of the loop is short and there is little benefit from abstracting it into a function. However, if the body of the loop is long, then there is a real benefit in doing so — it allows the reader of the `analyze` member function to see the entire function at once.

The loop body should be abstracted into a function if it meets any of the following criteria:

1. The body of the loop may be repeated in several places in the code.
2. The body of the loop, implemented as a separate function, allows it to be tested independently of other code.
3. The details of the body of the loop distract from understanding the tasks performed by `analyze`.

To run this variation of the exercise, do the following in your build directory:

```
art -c fcl/LoopOverGens/loopGens2.fcl >& output/loopGens2.log
```

```
root -l fcl/LoopOverGens/loopGens2.C
```

1 This will produce a page of histograms that should be identical to
 2 those made by `loopGens1.C`. It will also write the figure to the file
 3 `output/loopGens2.png`.

4 Regarding the second item on the list, when a type is as simple as `GenParticle`, there
 5 is little benefit to using `auto`; it may even obfuscate things a little. The benefit of `auto`
 6 arises when the correct type is a long, complicated name that is hard to get right.

7 Regarding the third item on the list, the Workbook authors felt it would be confusing to
 8 use this form before showing you the first form. However, from now on we will use this
 9 form to get a `const` reference to its 3-momentum from a `GenParticle`.

10 18.9.2 `LoopGens3_module.cc`

ations:LoopGens3

11 The file `LoopGens3_module.cc` shows six other ways to write the loop over generated
 12 particles. These are shown because you will likely see them in code written by others.
 13 There is no need, at this time, to become proficient in their use; just know where to look
 14 them up if you encounter them.

15 This module makes seven histograms of the momentum spectrum of the generated parti-
 16 cles. The first is made using the familiar range-based `for` loop and is provided for refer-
 17 ence.

18 The second method uses operator square brackets:

```
19   for ( size_t i=0; i!=gens->size(); ++i ){
20       CLHEP::Hep3Vector const& p = (*gens)[i].momentum();
21       hP2_->Fill( p.mag() );
22   }
```

23 Recall that the type of `*gens` is a `GenParticleCollection`, which is a typedef for
 24 `std::vector<GenParticle>`. This example shows that a `std::vector` can be
 25 used as an array; as for an array, `(*gens)[i]` returns a reference to the i^{th} element in
 26 the collection.

27 The third method uses the `at` member function:

```
28   for ( size_t i=0; i!=gens->size(); ++i ){
29       CLHEP::Hep3Vector const& p = gens->at(i).momentum();
30       hP3_->Fill( p.mag() );
31   }
```

1 The member function `at(i)` is a run-time bounds-checked version of `[i]`. That is,
 2 `at(i)` first checks that the value of `i` is within the bounds of the `std::vector`. If
 3 it is, then `at(i)` returns a reference to the i^{th} element in the collection. If it is not, then
 4 `at(i)` will throw an exception; `art` will catch the exception and will attempt a graceful
 5 shutdown. The check for a valid value of `i` is performed at run-time during each call to
 6 `at(i)`; therefore the check will safely follow changes to the size of the collection while
 7 the program is running.

8 If you use the operator `[i]` with a value of `i` that is out of bounds, then the calling
 9 code will receive a reference to undefined memory. If you are lucky your code will crash
 10 quickly. If you are unlucky you will get subtly incorrect results. If you are really unlucky
 11 your Mars lander might join Romana in E-Space.

12 We strongly recommend that, when possible, you use a *range-based for* loop. If you need
 13 to use one of the indexed versions, we strongly recommend `at(i)` over `[i]`; only when
 14 code is very well tested and is proven to be time-critical should you switch to `[i]`.

15 The fourth method uses iterators:

```
16 GenParticleCollection::const_iterator i=gens->begin();
17 GenParticleCollection::const_iterator end=gens->end();
18 for ( ; i!=end; ++i ){
19     CLHEP::Hep3Vector const& p = i->momentum();
20     hP4_->Fill( p.mag() );
21 }
```

22 An iterator behaves like a pointer to the selected element. If you wish to learn about itera-
 23 tors, consult any standard C++ reference.

24 The fifth method also uses iterators:

```
25 for ( auto j=gens->begin(); j!=gens->end(); ++j ){
26     CLHEP::Hep3Vector const& p = j->momentum();
27     hP5_->Fill( p.mag() );
28 }
```

29 Here we have used the `auto` keyword to avoid typing:

```
30 GenParticleCollection::const_iterator
```

31 This is the first time in the Workbook that using `auto` has really paid off. All that we need
 32 to know about this iterator is that `j->` will do the right thing; we don't need to spend a lot
 33 of time spelling its type correctly.

1 We recommend the fifth method over the fourth method. In the fifth method, the scope
 2 of the loop iterator `j` is strictly inside the loop. In the fourth method two iterators, `i` and
 3 `end` are defined outside the scope of the loop. Their scope extends to the closing brace
 4 of the member function `analyze`. After the body of the loop has ended, this pollutes the
 5 remaining scope of `analyze` with identifiers that no longer serve any purpose.

6 The sixth method introduces the comma operator:

```
7   for ( auto j=gens->begin(), jend=gens->end(); j!=jend; ++j ){
      CLHEP::Hep3Vector const& p = j->momentum();
      hP6_->Fill( p.mag() );
    }
```

8 This version says that it will evaluate, sequentially from left to right, the two comma-
 9 separated expressions as part of the initializer phase of the `for` loop. Compared to the
 10 previous version, this version needs to evaluate `gens->end()` only once, not once per
 11 loop iteration. Therefore this version is slightly more efficient.

12 Finally, the seventh version uses the `std::begin` and `std::end` free functions:

```
13   for ( auto j=std::begin(*gens), jend=std::end(*gens); j!=jend; ++j ){
14     CLHEP::Hep3Vector const& p = j->momentum();
15     hP7_->Fill( p.mag() );
16   }
```

17 The benefit to using this form only appears if you are writing templates — so whereas you
 18 are unlikely to use it yourself, you may see others using it.

19 The C++ language also supports `while` and `do-while` loops. There are appropriate
 20 places to use these features but writing a simple loop over a `std::vector` is not one
 21 of them. Why? These loop styles require that at least one identifier be defined outside
 22 of the scope of the loop. You can learn about these types of loops in any standard C++
 23 reference.

24 The bottom line is this: if a *range-based for* loop will do, use it. Because this is a new
 25 feature in C++-11, legacy code will contain many loops written in the other ways.

26 To run this variation of the exercise, do the following in your build directory:

```
27 art -c fcl/LoopOverGens/loopGens3.fcl >& output/loopGens3.log
```

```
28 root -l fcl/LoopOverGens/loopGens3.C
```

1 This will make two files of histogram output:

2 `output/loopGens3_1.png` `outout/loopGens3_2.png`

3 Each file will show a 2×2 array of histograms; four on the first page, three on the second.
 4 On the first page, the upper left histogram is the histogram made by method 1, which
 5 should be the same as that from the previous two exercises. The remaining six histograms
 6 show, for each method, the difference between the histogram made by that method and the
 7 histogram made by method 1. All of these histograms should be flat lines at zero.

8 **FIXME:** *AH: (I've left this fixme to remind you to talk to Chris.) I can't get kpdf to work;*
 9 *display works RK: I can ask Chris to put this on cluck. This is going to be a monster ongo-*
 10 *ing problem ... we can control the GPCF machines - everyone else is on their own.*

11 You can inspect the file `fcl/LoopOverGens/loopGens3.C` to learn how to use a
 12 CINT script to subtract two histograms. It will also show you how one CINT script can
 13 produce multiple `png` files.

14 **FIXME:** *Rob's note to self: In this section Do we want to recommend pre-increment over*
 15 *post increment or is that just distracting?*

16 18.9.3 `LoopGens3a_module.cc`

17 `LoopGens3a_module.cc` shows a minor variation on `LoopGens3_module.cc`
 18 that many people prefer. You will certainly see code written this way and you are free
 19 to write it this way if you find it more natural than the earlier version. **FIXME:** *This*
 20 *should include a clear statement of what this section teaches, i.e., This shows how to use*
 21 *const reference the object of type GenParticle instead of pointer to it ... or some such...*
 22

23 In `LoopGens3_module.cc` the first line of the member function `analyze` was:

24 `auto gens = event.getValidHandle<GenParticleCollection>(gensTag_);`

25 In `LoopGens3a_module.cc` this has been changed too:

26 `auto const& gens(*event.getValidHandle<GenParticleCollection>(gensTag_));`

27 You might find it easier to understand the second version if you break it into two steps:

```
1 auto gensHandle = event.getValidHandle<GenParticleCollection>(gensTag_);  
2 GenParticleCollection const& gens(*gensHandle);
```

3 Note that both blocks of code end by defining a variable named `gens`, which is used by
4 the subsequent code. The difference is that, in the first version you use `gens` as a pointer
5 to the `const` data product whereas, in the second version, you use `gens` as a `const`
6 reference to the data product.

7 In the remainder of the member function `analyze` there two sets of changes: `gens->`
8 was changed to `gens.` and `(*gens)` was changed to `gens`.

9 18.10 Review

10 In this exercise you have learned:

- 11 1. how to use most of the accessor member functions of the class `GenParticle`
- 12 2. how to use the stream insertion operator of `GenParticle`
- 13 3. how to use some of the accessors functions of `CLHEP::Hep3Vector` and
14 `CLHEP::HepLorentzVector`
- 15 4. CLHEP's use of `.icc` files to hold inline implementation
- 16 5. eight different ways loop over a `GenParticleCollection`
- 17 6. that, for most purposes, the preferred way to write the loop is to use a ranged for
18 loop
- 19 7. that if the body of a loop is long, or if it will be repeated elsewhere, then put it into
20 a separate function
- 21 8. how to write a CINT script that writes multiple figure files
- 22 9. how to write a CINT script that writes a multi-page PDF file
- 23 10. how to use a CINT script to subtract two histograms

18.11 Test Your Understanding

Test 1 asks you to write a module. For tests 2 and 3, files are provided that contain intentional bugs. Your job, of course, is to find the bugs and fix them. The answers, provided in Section 18.11.4, are intentionally placed on a new page.

18.11.1 Test 1

FIXME: *Anne did not run this test*

In this assignment you will write your own module class and FHiCL file, build your module, run *art* on it and check your results.

Create your new `_module.cc` file and FHiCL file in the source directory for this exercise. Do NOT pick a name of the form `LoopGensN` or `loopGensN`, where `N` is a single digit; there are more of these files to come.

In your new module file, extend `LoopGens1_module.cc` to add six additional histograms and the printout described below. The numbers in parentheses are a suggested histogram binning in the format: (nbins, xlow, xhigh). The six histograms are:

1. The azimuth of the momentum 3-vector of each generated particle (100, $-\pi$, π).
2. The rest mass of each generated particle (100, 0., 1100.)
3. The x , y and z positions of each generated particle, (100, -1., 1.).
4. The number of children of each generated particle (5, 0., 5.)

Hint: the C++ header `<cmath>` defines the symbolic constant `M_PI`, which contains the value of π .

At the time of this writing, all of the generated particles in the example input files are produced at the origin. A more interesting distribution will be added in a future release.

FIXME: *Update the generator code to fix this.*

For each generated particle print to `cout`

1. The event number

1 2. The index of the `GenParticle` within the `GenParticleCollection`

2 3. The full information about the `GenParticle`; see the hint below.

3 Add a parameter set parameter to limit this printout to a number of events specified in the
4 parameter set. The default should be for no printout.

5 **Hint:** To print all of the information about the `GenParticle` named `gen`, you can use
6 the stream insertion operator (operator `<<`):

```
7     std::cout << gen << std::endl;
```

8 When you are ready to build your module, go to your build directory and give the com-
9 mand:

```
10    buildtool
```

11 This will compile your module and link it into a shared library.

12 Prepare a FHiCL file to run this module on the file `inputs/input04_data.fcl`. Be
13 sure to pick a unique name for the histogram file created by the `TFileService`. Keep
14 the FHiCL file in the source directory for this exercise.

15 To run your module, go to your build directory and give the command:

```
16    art -c fcl/LoopOverGens/<your-file-name>.fcl >& output/<your-file-name>.log
```

17 When you are done, you can compare your solution to that given in the files
18 `LoopGens4_module.cc` and `loopGens4.fcl`. To run this module, be sure that you
19 are in your build directory and give the command:

```
20    art -c fcl/LoopOverGens/loopGens4.fcl >& output/loopGens4.fcl
```

21 Compare the printout made by `LoopGens4_module.cc` with your printout. Compare
22 the histograms made by `LoopGens4_module.cc` with your histograms; a discussion
23 of the histograms will follow the discussion of the printout.

24 The first line of printout, when split over two lines, will look like:

```
25    Event:      1 GenParticle:      0 : [ pdg: 211 Position: (0,0,0)
26       4-momentum (229.7,536.378,-81.9304;605.521) parent: none children: none ]
```

27 The printout for each `GenParticle` is enclosed in square brackets and consists of:

28 1. The PDG Id code

29 2. The position at which the particle was created: (x, y, z)

1 3. The 4-momentum at which the particle was created: $(p_x, p_y, p_z; E)$

2 4. An identifier of the parent, or “none” if the particle has no parent

3 5. Identifiers of the children, in parentheses, or “none” if the particle has no children

4 You can inspect the printout to see that only five different PDG Id codes occur: ± 211 ,
5 ± 321 and 333. The next exercise, in Chapter 19, will explain how to interpret these codes.

6 You will learn that the code 333 is for the ϕ meson, that ± 211 are for π^\pm and that ± 321 are
7 for K^\pm . The meaning of the parent/child identifiers will be discussed in Chapter 20.

8 To help compare histograms you can make a multipage PDF file containing all of the ref-
9 erence histograms with the command:

```
10 root -l fcl/LoopOverGens/loopGens4.C
```

11 This will make a file named `output/loopGens4.pdf` that has three pages of his-
12 tograms.

13 Appendix D has instructions on how to view a multipage PDF file interactively.

14 For those of you at Fermilab it also has instructions on how to print it using the Fermilab
15 printers.



16 You can also study the file `fcl/LoopOverGens/loopGens4.C` to learn how to use
17 ROOT to make a multipage PDF file of histograms from a ROOT file.

18 18.11.2 Test 2

19 In the source directory for this exercise there are three files that end in `.cc.nobuild`.
20 Each of these has an error expressly inserted. Your assignment is to find and fix each of the
21 errors, one at a time, rebuilding each time. The answers are given in Section 18.11.4.

22 The source files listed below fail to build. For each in turn, follow essentially the same
23 procedure as outlined in Section 10.12.

24 For `LoopGens5_module.cc` the error message is short.

25 When you have found and fixed the error in `LoopGens5_module.cc`, follow the di-
26 rections in the preceding paragraph for the file `LoopGens6_module.cc.nobuild`.
27 When you build this module it will produce a very long error message. Scroll back to
28 first few lines of the error message because they contain the most useful information. This

1 is generally good advice when C++ produces very long error messages: the most useful
2 information is in the first few lines.

3 If you have trouble finding the error and wish to try one of the other files, rename the `.cc`
4 file back to `.cc.nobuild` and start to work on the other file.

5 ***FIXME:** I (Anne) moved the ‘what’s with the nobuild to accessing data products; the first*
6 *time it’s encountered*

7 **18.11.3 Test 3**

8 The code in `LoopGens8_module.cc` builds and runs but it produces incorrect out-
9 put. To run the code and look at its output, go to your build directory and run the com-
10 mands:

```
11 art -c fcl/LoopOverGens/loopGens8.fcl >& output/loopGens8.log
```

```
12 browse output/loopGens8.root
```

13 Navigate to the histogram `gens/hP` and view it. The title says that it is the momentum
14 spectrum of the generated particles; therefore it should look like the corresponding his-
15 togram in any of the other root files from this exercise. But it looks very different. Figure
16 out why and fix it.

17 ***FIXME:** My histogram is blank. Anne*

anding:LoopGens8

18.11.4 Answers

18.11.4.1 Test 1

FIXME: We should provide a solution

18.11.4.2 Test 2

In `LoopGens5_module.cc` the error is in the range-based for loop

```
for ( GenParticle& gen : *gens ){
```

The declaration to the left of the colon should be a const reference, not a non-const reference. The underlying reason for this is that the `art::Event` only grants const access to data products already in the event. In this case, the error message from the compiler is fairly clear:

```
error: invalid initialization of reference of type 'tex::GenParticle&'
from expression of type 'const tex::GenParticle'
```

This error message is shown here on two lines for readability.

In `LoopGens6_module.cc` the error is again in the range-based for loop

```
for ( GenParticle const& gen : gens ){
```

This time the error is that the `*` is missing from `*gens` to the right of the colon. The reason why it is needed was described on page 332. This time the text of error message is not as explanatory:

```
error: no matching function for call to
'begin(art::ValidHandle<std::vector<tex::GenParticle> >&)'
  for ( GenParticle const& gen : gens ){
                                ^
```

However the caret does correctly indicate where to look for the error. Again, the error message is shown here on two lines for readability.

In `LoopGens7_module.cc` the error is in the body of the range-based for loop:

```
CLHEP::Hep3Vector& p = gen.momentum();
```

1 The error is that the type of `p` should be a const reference, not a non-const reference. The
 2 underlying reason is the same as for `LoopGens5_module.cc`; the immediate reason is
 3 that `gen` is itself a const reference so the code may only call const functions of `gen`. This
 4 time the error message is short and on point:

```
5 error: invalid initialization of reference of type 'CLHEP::Hep3Vector&'
6 from expression of type 'const CLHEP::HepLorentzVector'
7     CLHEP::Hep3Vector& p = gen.momentum();
8                               ^
```

9 An inferior solution would be to remove the ampersand and make `p` a copy of the space
 10 part of the generated momentum, not a const reference to it. This would allow the code to
 11 compile and run correctly but it will execute more slowly than it should.

12 18.11.4.3 Test 3

13 In `LoopGens8_module.cc` the error is in the line

```
14     hP_-> Fill( gen.momentum().mag() );
```

15 The issue is that both `HepLorentzVector` and `Hep3Vector` have a function named
 16 `mag()`. The former returns the invariant mass of a 4-vector and the latter returns the
 17 magnitude of a 3-vector. The line in the above listing very clearly says to call the member
 18 function `mag` from the class `HepLorentzVector`, which was not the intent stated in
 19 the title of the histogram.

20 The preferred solution is write the code the way that `LoopGens2_module.cc` is writ-
 21 ten:

```
22     CLHEP::Hep3Vector const& p = gen.momentum();
23     hP_-> Fill( p.mag() );
```

24 This makes it clear that we want to call the `mag` function of a `Hep3Vector`. Also read
 25 the section that begins “As a last comment” on page 332.

26 **FIXME:** *Rob’s note to self: Is this the right place to talk about single phase constrcution*
 27 *or should we wait until we have more examples?*

19 Exercise 9: Accessing and Using Particle Data

chap:PDT

19.1 Introduction

FIXME: *This chapter is in the middle of a major reorganization and does not match the code.* RK July 27, 2014

FIXME: Introduce PDT first In Chapter ?? you saw that each particle species has a PDG ID code which allows you to look up its properties, e.g., name, charge, and so on. **FIXME:** *I think this is true:* This exercise introduces you to the *particle data table* (abbreviated PDT), which is a FHiCL file that stores this information in the form of a table, with each particle identified by its PDG code. Access to this table is provided by an *art* service called the PDT service. It is a *user-written art* service, meaning that it is not an integral part of *art*, but rather must be supplied — usually as part of the experiment code.

FIXME: Might want to clarify table vs service The PDT table/service used in this exercise is supplied as part of the toyExperiment UPS product, and for simplicity and clarity, it stores only the information required by this package: the mass, electric charge and name of each of the particle species used in the the Workbook exercises, π^\pm , K^\pm and ϕ . Your experiment's PDT will likely include many more properties (e.g., error on mass, width or lifetime and its error, spin, parity and so on) and will include many more particles.

This exercise focuses on *using* the PDT service. You do not need to understand its insides. *A chapter yet to be written will discuss writing your own services and will include the PDT service as a case study.*

FIXME: here introduce activity The first module in this exercise looks at all of the generated particles in the input stream and counts the occurrence of each unique species, i.e.,




Figure 19.1: Number of each type of particle in the input stream; made by `countGens1.C`.

- 1 `PDGCode::type`. Listing 19.1 shows a representative example of the output, which con-
 2 tains three columns: the name of the particle, its PDG ID and the number of times that a
 3 `GenParticle` of that type was found in the input stream.

Listing 19.1: Particle data for this exercise; output from `countGens1.fcl`

```

4 Summary of Particle types found in this job
5 Name: K-   Id: -321   Count: 1000
6 Name: pi-  Id: -211   Count:  716
7 Name: pi+  Id:  211   Count:  716
8 Name: K+   Id:  321   Count: 1000
9 Name: phi  Id:  333   Count: 1000

```

- 10 Figure 19.1 presents the same information as a histogram.

11 **FIXME:** *Not yet sure where this info should go*

- 12 1. To save memory space there is not a copy of this information for each generated
 13 particle, just the PDG ID.
- 14 2. The information in the PDT service is read at the start of the *art* job and is available
 15 to be used by modules in the constructor or at any later time.

- 1 3. **FIXME:** *Reference to this chapter once it is written.*
- 2 4. This exercise also introduces a class that is not a module but is used by the module.

3 19.2 Prerequisites

4 Prerequisites for this chapter include all of the material in Part I (Introduction) and the
5 material in Part II (Workbook) up to and including Chapter ??.

6 19.3 What You Will Learn

7 In this exercise you will learn how to:

- 8 1. use the PDT service; the classes `tex::PDT` and `tex::ParticleInfo`.
- 9 2. loop over `GenParticle` collection and count the occurrence of each particle
10 species.
- 11 3. print an end-of-job summary.
- 12 4. make histograms of individual particle properties and event totals.
- 13 5. write a class that is not a module class and make it available to the module.
- 14 6. maintain the `CMakeLists.txt` file by adding `PDT_service.so` and `UsePDTService.so`
15 to the link list.
- 16 7. use a `std::map`.
- 17 8. define and use an `enum`.
- 18 9. open a text output file and write to it.

19 You will also discover why the *art* developers chose to implement the PDT as a service.
20 **FIXME:** *What are you going to learn here? How to tell whether contents are valid from*
21 *c'tor onwards? A service Handle is valid from module c'tor onwards but its contents may*
22 *not be — you just gotta know. but the PDT is defined at all times that any module is being*
23 *called: c'tor, beginJob, beginJob*

1 19.4 Setting Up to Run Exercise

e-pdt:setting:up

Follow the instructions in Chapter [11](#) if you are logging in after having closed an earlier session. If you are continuing on directly from the previous exercise, keep both your source and build windows open.

In your source window, look at the contents of the directory for this exercise, called `UsePDTSvc`:

```
ls art-workbook/UsePDTSvc
```

```
CMakeLists.txt          countGens2.fcl
countGens1.C            CountGens2_module.cc
countGens1.fcl          IdCounter.cc
CountGens1_module.cc    IdCounter.h
countGens2.C
```

In your build window, just make sure that you are in your build directory. All the code for this exercise is already built; this happened the first time that you ran `buildtool`.

2

- 3 The file `CountGens1_module.cc` is the module that you will study in this exercise.
4 The file `countGens1.fcl` runs the exercise and the `countGens1.C` makes a histogram of the results. The files `IdCounter.h` and `IdCounter.cc` are the source for
5 a class that is used as a “helper class” by `CountGens1_module.cc` to do its job. The
6 files `CountGens2_module.cc`, `countGens2.C`, and `countGens2.fcl` are the
7 solutions to an exercise given in Section [??](#). **FIXME:**
8

19.5 The Particle Data Table Service

19.5.1 The PDT FHiCL File

The particle data table is a FHiCL file that should be provided by the experiment. **FIXME: right?** For the toy experiment the file is:

```
$TOYEXPERIMENT_DIR/databaseFiles/pdt.fcl
```

The contents of the file is reproduced in Listing 19.2.

Listing 19.2: Contents of `pdt.fcl`

```
1 BEGIN_PROLOG
2
3 mPiCharged : 139.57018
4 mKaCharged : 493.677
5 mPhi       : 1019.455
6
7 END_PROLOG
8
9
10 particles : [
11   { name : "pi+" pdgId : 211 mass : @local::mPiCharged charge: 1 },
12   { name : "pi-" pdgId : -211 mass : @local::mPiCharged charge: -1 },
13   { name : "K+"  pdgId : 321 mass : @local::mKaCharged charge: 1 },
14   { name : "K-"  pdgId : -321 mass : @local::mKaCharged charge: -1 },
15   { name : "phi" pdgId : 333 mass : @local::mPhi       charge: 0 }
16 ]
```

The particle data table itself is contained in the FHiCL parameter `particles`, which is a FHiCL sequence of FHiCL tables (parameter sets). Inside the sequence, each FHiCL table defines the information for one particle species. As specified in Table 3.2 **FIXME: ref not found**, the masses are given in MeV and the charges in units of the charge of the proton. This file introduces two new features of FHiCL, the `@local::` syntax and `prologs(γ)`.

The syntax `@local::XXX` tells FHiCL to look earlier in the same or an included file **FIXME: right?** to find a definition for a parameter named XXX. If it can find such a definition, it replaces `@local::XXX` with the value of the parameter XXX. The pion mass `@local::mPiCharged`, for example, will thus be replaced with the value 139.57018 MeV.

1 If FHiCL cannot locate the definition of XXX it will throw an exception and *art* will attempt an orderly shutdown.

3 The FHiCL definitions between the `BEGIN_PROLOG` and `END_PROLOG` line are resources that FHiCL can use when resolving `@local::` requests but these lines are not part of the final FHiCL document. **FIXME:** *What is the 'final fcl document' relative to the file `pdt.fcl`?*

7 Chapter [31](#) contains a complete description of FHiCL and will describe these features further.

9 In summary, when FHiCL has finished processing `pdt.fcl` it will have been reduced to:
10 **FIXME:** *When you run `art`, it will change this file during the processing? What happens for the next `art` job? This isn't clear; something's missing.*

```
12 particles : [
13   { name : "pi+" pdgId : 211 mass : 139.57018 charge: 1 },
14   { name : "pi-" pdgId : -211 mass : 139.57018 charge: -1 },
15   { name : "K+" pdgId : 321 mass : 493.677 charge: 1 },
16   { name : "K-" pdgId : -321 mass : 493.677 charge: -1 },
17   { name : "phi" pdgId : 333 mass : 1019.455 charge: 0 }
18 ]
```

19 19.5.2 The PDT Service

20 When code in the toy experiment or the Workbook needs to access information about the particle data table, it should do so by using the [PDT](#) service. This is an *art* service that is part of the toyExperiment UPS product. For this exercise you do not need to understand how the [PDT](#) service works, just how to use it.

24 User code sees the [PDT](#) service as simply the class [PDT](#). The header file for this class is found in:

```
26 $TOYEXPERIMENT_INC/toyExperiment/PDT/PDT.h
```

27 The source file for this class is found in:

```
28 $TOYEXPERIMENT_INC/toyExperiment/PDT/PDT_service.cc
```

29 You do not need to look at the source file. Notice, however, the structure of the filename, because it is significant; the suffix `_service` tells *art* that this file contains a plug-in that

1 can create an *art* service whose name is the part before the suffix, in this case `PDT`.

2 In general, when an *art* job starts, and before any modules are created, *art* creates each
3 service listed in the specified FHiCL file (not `pdt.fcl`). When the `PDT` service is created,
4 *art* reads the file `pdt.fcl` and stores its **FIXME:** *is this where 'final' form comes in?*
5 contents in the service. Once this step has completed, any code that wishes to access the
6 particle data table information can now do so via the `PDT` service.

7 Look again at the header file, `PDT.h`. Three things require comment. The others are part
8 of the pattern that makes the class `PDT` and *art* service, and *will be discussed in a yet-to-*
9 *be-written chapter on services.*

10 **FIXME:** *Reference the new chapter when it is ready.*

11 The three points are:

- 12 1. `PDT` has a member function:
13

```
ParticleInfo const& getById( PDGCode::type ) const;
```


14 This function returns all of the information about the species of particle specified by
15 the argument. The information is made available as a `const` reference to an object of
16 type `ParticleInfo`, which is discussed in Section 19.5.5.
- 17 2. `PDT` has a member function:
18

```
void print( std::ostream& ) const;
```


19 This function prints the content of the particle data table to the output stream that is
20 given as an argument.
- 21 3. `PDT` has a stream insertion operator (`operator<<`), that provides a second way to
22 print the content of the particle data table.

23 19.5.3 Using the PDT Service

24 To get access to the particle data table, your code needs to include the header file:

```
25 1 #include "toyExperiment/PDT/PDT.h"
```

26 and default construct a service handle to the `PDT` service:

```
27 1 art::ServiceHandle<PDT> pdt;
```

1 This can be done in any code that is executing inside *art*, including in your module class.
 2 Once the object `pdt` has been constructed it can be used as a pointer to the particle
 3 data table. If you need to review the class template `art::ServiceHandle`, see Sec-
 4 tion 17.5.1.

5 19.5.4 Configuring the PDT Service

6 If you plan to use the `PDT` service in your *art* job, then your FHiCL file must provide a
 7 parameter set that *art* can use to configure the service. You do this by adding a FHiCL
 8 definition inside the `services` parameter set, as shown in Listing 19.3.

Listing 19.3: FHiCL fragment to configure the PDT service

```

1 services : {
2
3   // Parameter sets for other art supplied services - elided for clarity
4
5   PDT : { pdtFile : "databaseFiles/pdt.fcl" }
6
7   // Parameter sets for other user supplied services - elided for clarity
8   }
9 }
```

18 Recall that `PDT` is a user-written service; it is not provided by *art* itself. *art* requires
 19 that all user-written services be configured inside of a parameter set named `user` that is
 20 placed inside the `services` parameter set. The configurations for user-written services
 21 are not allowed to be placed directly in the `services` parameter set. This separation
 22 of the configurations for *art*-supplied and user-supplied services is a precaution against
 23 accidental name collisions.

24 The user-defined services you encounter will typically be provided by your experiment,
 25 in this case the toy experiment. But you or any of your colleagues may also provide your
 26 own.

27 The parameter set for a service does not contain an analog of the `module_type` key-
 28 word that is required for a parameter set that configures an *art* module (see Section 9.8.1).
 29 Instead the convention is that the name of the FHiCL definition must be the name of the
 30 class that implements the service, in this case `PDT`. When *art* sees a parameter set for the
 31 `PDT` service, it will search `LD_LIBRARY_PATH` to find a file named

1 `lib*PDT_service.so`

2 It will find it in the dynamic library:

3 `$TOYEXPERIMENT_LIB/libtoyExperiment_PDT_PDT_service.so`

4 If your code uses the `PDT` service you need to add this library to the link list.

5 **FIXME:** *Forward reference to the section service interfaces where I think this rule may*
6 *be a little different.*

7 The `PDT` class determines the parameters that can go inside the parameter set `PDT`. Only
8 the two following parameters are allowed, the first of which is required. They are shown
9 with typical values.

```
1b pdtFile   : "databaseFiles/pdt.fcl"
1d verbosity : 0
```

12 `pdtFile` gives the name of the file that contains the information that will be read into
13 the particle data table. You may recall from Section 19.5.1 that the actual path to the file is
14 the relative path given here prefixed by `$TOYEXPERIMENT_DIR`. In Chapter 31 you will
15 learn how `FHiCL` knew to look in this directory.

16 The second parameter, `verbosity`, determines how much information gets printed out.
17 If it is zero, `PDT` makes no printout; if it is greater than 0 but less than or equal to 2, `PDT` will
18 print out a summary of the information it received from the parameter set; if it is greater
19 than 2, `PDT` will print out the full particle data table.

20 19.5.5 The Class `ParticleInfo`

21 The class `ParticleInfo` holds the particle data information about one species of par-
22 ticle **FIXME:** *at a time?*. The source code for this class is in:

```
23 $TOYEXPERIMENT_INC/toyExperiment/PDT/ParticleInfo.h
24 $TOYEXPERIMENT_INC/toyExperiment/PDT/ParticleInfo.cc
```

25 The object code for this class is found in the dynamic library:

26 `$TOYEXPERIMENT_LIB/libtoyExperiment_PDT.so`

27 If your code uses `ParticleInfo` you will need to add this library to the link list.

1 Look at the header file. The features of `ParticleInfo` that you will need for this exer-
 2 cise are the four accessor member functions:

```
3 1 PDGCode::type      id()      const { return _id;  }
4 2 double            mass()    const { return _mass; }
5 3 double            charge()  const { return _q;   }
6 4 std::string const& name()    const { return _name; }
```

7 If you are not sure what information each of these returns, see Section ?? **FIXME:** .
 8 Similar to the way the `GenParticle` class behaves, this class returns small pieces of
 9 information by value and large pieces by const reference.

10 19.6 The File `CountGen1_module.cc`

11 This file has a few new features, including two new headers:

```
12 1 #include "art-workbook/UsePDTService/IdCounter.h"
13 2 #include <fstream>
```

14 The first is the helper class mentioned in Section 19.4. The second is the header for reading
 15 and writing files.

16 It also includes two new data members:

```
17 1 std::string      outputFilename_;
18 2 IdCounter       counter_;
```

19 The first is the name of a text output file that you can ask *art* to write. It transfers infor-
 20 mation from the module class to `countGens.C`, the file used to generate the histogram.
 21 The second is an instance of the helper class — everything related to the PDT service will
 22 be inside the `IdCounter` class.

23 **FIXME:** *None of the 4 lines above are in this cc file; what's going on?*

24 In the constructor for `CountGen1` two new data members are initialized:

```
25 1 tex::CountGen1::CountGen1(fhicl::ParameterSet const& pset ) :
26 2   art::EDAnalyzer(pset),
27 3   gensTag_(pset.get<std::string>("genParticlesInputTag")),
```

```

1 4   outputFilename_(pset.get<std::string>("outputFilename", "")),
2 5   counter_() {
3 6 }

```

4 **FIXME:** *outputFilename is not in this cc file, nor is counter; gensTag is*

5 C++ does not require putting `counter_` here but we consider it good C++ hygiene.

6 **FIXME:** *what does it do?* Note that the `outputFilename` is optional. If it is absent,
7 the value of the data member `outputFilename_` is an empty string.

8 For each particle in the event, the code increments the counter.

```

9 1   for ( GenParticle const& gen : *gens ) {
10 2     counter_.increment( gen.pdgId() );
11 3   }

```

12 **FIXME:** *this code is not what's there either*

13 What does it mean to increment the counter? Conceptually it is:

- 14 1. Refer to `countGens1.fc1` (Listing 19.1). **FIXME:** *it's wrong; check*
`list:pdt:output:full`
- 15 2. If the PDG ID of this particle is already in the table, then just increment its count.
- 16 3. If it is not yet in the table, then add a new line to the table, set the PDG ID column
17 appropriately and set the count to 1.
- 18 4. At this point in the processing, nothing is done with the particle name; that is filled
19 in at the end of the job.

20 At the end of the job, tell `counter_` to print itself out to the log file:

```

21 1   counter_.print( std::cout );

```

22 **FIXME:** *not there either* It is optionally written to another file. **FIXME:** *in other words,*
23 *you can choose to write it to another file? Why/when would you want to do this?*

```

24 1   if ( outputFilename_.empty() ) return;
25 2   std::ofstream out( outputFilename_.c_str() );
26 3   counter_.print( out, IdCounter::minimal );

```

1 Note that we pass “where we want the printout to go” as an argument! **FIXME:** to
2 `std::ofstream`; where do you set `outputFilename`? In your `fc1` file?

3 `empty()` is a member function of the class `std::string`; it is true if there are no char-
4 acters in the string. If the output filename was not specified then the value of `outputFilename_`
5 will be an empty string.

6 This code snippet performs the following logic: If `outputFilename_` is not empty,
7 then open an output file with the given name, and write a second copy of the `IdCounter`
8 output to it.

9 `std::ofstream` is the name of the class that knows how to write text to files. `out` is an
10 object of this type. The constructor argument is the address of a C-style string that holds
11 the name of the file.

12 What is the `.c_str()`? Inside every `std::string` is a C-style string; a
13 `std::string` contains a lot more information. The designers of the C++ standard li-
14 brary did not want the IO subsystem to depend on the class `std::string`. Therefore the
15 constructor of `std::ofstream` takes an argument that is a pointer to a C-style string,
16 not a `std::string`. The class `std::string` has a member function `c_str()` that
17 returns the address to the C-style string that lives inside it.

18 The second call to `print` has a second argument, `IdCounter::minimal`, that was not
19 present in the first call. You will see that this tells the member function `print` to format
20 the information a little differently. The first format is optimized to be easy for people to
21 understand. The second version is optimized to be easy for the `.C` file to understand. That
22 output looks like:

```
c:output:minimal
```

Listing 19.4: The contents of `output/countGens1.txt`

```
23 K- 1000  
24 pi- 716  
25 pi+ 716  
26 K+ 1000  
27 phi 1000
```


19.7 Running the Exercise

To run this exercise, cd to your build directory and run *art*:

```
cd $ART_WORKBOOK_WORKING_BASE/<username>/\  
workbook/build-prof
```

```
art -c fcl/UsePDTSERVICE/countGens1.fcl >& output/countGens1.log
```

art should complete with status 0. This makes the printout shown in Listing 19.1. It also makes the file `output/countGens1.txt`

```
root -l fcl/LoopOverGens/countGens1.C
```

This makes the histogram shown in Figure 19.1.

About Listing 19.1. It should be self explanatory. The last line says that when we looked at every gen particle in the file, we found 1000 phi mesons. This is good because there were 1000 events in the input file and, by construction, each event has exactly one phi meson. And each decays into a K^+K^- pair; so there should be exactly 1000 of these and there are. The background component of each event is N pairs of $\pi^+\pi^-$ mesons, where N is drawn from Poisson distribution with a mean of 0.75; therefore the number of pi+ and pi- in the listings should be statistically compatible with 750 — and it is (716 is about 1.25 sigma low).

The order of particles in the list is in order of increasing numerical value of the PDG ID code.

About the file `output/countGens1.txt`. Why is it there. The short answer is that we need a way to get the text axis labels from the module to `countGens1.C`. It's easy to store the bin contents has a histogram but there is no way to assign text labels to each bin until after the histogram has been drawn! It is possible to stuff the labels in to the root file but that is really awkward and tedious. It's just easier to write them to a file. Once were are sending the bin labels this way, we might as well send the contents this way too.

1 Another alternative would be to write a script to parse the log file and pull out the infor-
2 mation. But that's a last resort — never do that as a first choice because it there are simply
3 too many things that can go wrong.

4 19.8 The Class `IdCounter`

5 This class is found in the files `IdCounter.h` and `IdCounter.cc`. Look at the header
6 file. The core of this class is the data member:

```
7 1 Counter_type counts_;
```

8 where the type is defined as:

```
9 1 typedef std::map<PDGCode::type,int> Counter_type;
```

10 `std::map<T1,T2>` is another class template from the C++ standard library. It has two
11 template arguments that are the names of two C++ types. You can think of a `std::map<T1,T2>`
12 as a kind of table. In this table each row has two things: a name and content. The name of
13 each row is an object of type `T1` and the content of each row is an object of type `T2`.

14 In more correct language, the name of the row is called its *key* and the content of the row
15 is called its *value*. An `std::map` is a collection of *key-value* pairs.

16 In this class the key is of type `PDGCode::type` and the value is of type `int`, which is
17 the count of how many times that ID has been seen.

18 This exercise touches just a tiny part of the functionality of `std::map`. We will have
19 more to say about `std::map` when we discuss `IdCounter.cc`. You can read about
20 `std::map` in any standard C++ reference.

21 The other new elements of the header file are listed below. There is one new include
22 file:

```
23 1 #include <map>
```

24 This is the header for `std::map`.

25 There is an enum:

```
26 1 enum outputStyle { full, minimal};
```

1 that is defined *inside* the class. This enum is used to select which printout style to use.
2 The value `full` chooses the style of printout shown in Listing 19.1. The value `minimal`
3 chooses the style of printout shown in Listing 19.4.

4 For enums that support a class, always define them inside the class. This gives them scope
5 protection: ie another class can use the same names and they won't collide.

6 The name `outputStyle` is now the name of a type.

7 When you use the identifiers `full` and `minimal` inside the class, you can use them
8 just the way they are. If you use them outside the class you need to say their full names
9 `IdCounter::full` and `IdCounter::minimal`.

10 The public interface has only 5 items:

- 11 1. the enum just discussed
- 12 2. a default constructor that initializes `counts_` as an empty `std::map`.
- 13 3. the number function `increment`, discussed later
- 14 4. The accessor member function named `counts()`.
- 15 5. The member function `print`.

16 The member function `increment` is defined in the class header. Therefore the compiler
17 will automatically consider it as a candidate for inlining. The function is defined as:

```
18 1 void increment( PDGCode::type id ) {  
19 2     ++counts_[id];  
20 3 }
```

21 The return type `void` means that nothing is returned. For an `std::map`, the operator
22 square brackets takes an argument that must be the same type as the *key*. If the key exists
23 in the the map, then the map will return *a reference to* the value associated with that key.
24 That is, you can think of the second line of `increment` as having two steps:

```
25 1 int& count= counts_[id];  
26 2 ++count;
```

27 Because the operator `[]` returns a *reference* the `++` increments the counter in place; no
28 copies are needed.

1 On the other hand, if the key does not exist in the map, then the map will create a new *key-*
2 *value* pair in which the value is appropriately initialized, in this case initialized to zero. It
3 will then return a reference to that value.

4 In the member function `print` must always be provided. It is a reference to an output
5 stream, ie to an object of type `std::ofstream`. This is where the printout will go. Two
6 examples of output streams are `std::cout` and `std::cerr`. In the discussion of the
7 `endJob` member function of `CountGens1` you that you can create your own object of
8 that type and pass it as an argumnet to `print`.

9 The second argument of `print` is optional. If present it must be one of two values de-
10 fined in the enum. If the argument is not present, the compiler will substitute the value
11 `IdCounter::full`.

12 Now look at `IdCounter.cc`. This contains the constructor, which simply calls the de-
13 fault constructor of the map; this creates an empty map. The only other function in this file
14 is the `print` function. In the headfile the second argument was declared optional; here it
15 is not optional. The optional bit was that, when the compiler compiled the caller, it would
16 automagically fill in a missing second argument. In the `.cc` file there is no such thing as an
17 optional argument — it is always there.

18 Finally, we get to the PDT service.

19 19.9 Discussion

- 20 1. Comment on why we chose to put the PDT in a service: don't want singletons be-
21 cause of lack of control over lifetime and lack of control over run time configuration.
- 22 2. A service Handle is valid from module c'tor onwards but its contents may not be —
23 you just gotta know.
- 24 3. PDT is defined at all times that any module is being called: c'tor, beginJob, beginJob
25
- 26 4. If you want a nice exercise for yourself, reorder the printout into decreasing fre-
27 quency.

1 19.10 Suggested Activities

DRAFT

20 The `art::Ptr` Class Template

reading:art::ptr

20.1 Introduction

In Chapter 18 you learned that an object of class `GenParticle` knows about its parent and children particles; you also learned that this relationship is expressed using a type of smart pointer. This chapter will describe that type of smart pointer and how to it.

To motivate why we need a special kind of pointer, consider the following story. Suppose that we have a job in which

1. We create a data product of type `GenParticleCollection`.
2. Other modules in the job read the data product.
3. The data product is discarded at the end of the job; that is, the data product is never written to disk.

In this limited situation we could use the following technique to connect a particle to its parent. The `GenParticle` could have a private data member:

```
14 GenParticle const* parent_;
```

and a matching public accessor:

```
16 GenParticle const* parent() const { return parent_; }
```

With the `GenParticle` class designed in this way one can write code like:

```
18 GenParticle gen; // Initialized somehow.
19 if ( gen.parent() != nullptr ) {
20     std::cout << "The_PDG_ID_of_the_parent_particle_is:__"
21               << gen.parent()->pdgId() << std::cout;
22 }
23 }
```

1 This technique works because, once it has been created, the data product occupies a fixed
 2 place in memory. Therefore the address of each `GenParticle` object is fixed for the
 3 duration of the event. Therefore it is safe for a child particle to hold a pointer to its par-
 4 ent.

5 Now consider what happens when you write the data product to disk and read it in again. In
 6 the general case, the data product will occupy a different range of memory locations than
 7 it did in the job that created it. Therefore the pointer that connects a child to its parent no
 8 longer points the parent `GenParticle`; it points, instead, at a memory location that has
 9 nothing to do with the parent.

10 There is not even enough information present in the data product to figure out how to
 11 recompute the pointers and to reset them to the correct memory addresses. The parent-
 12 child information has been irretrievably lost.

13 There is a second problem with the above example. The user of the pointer must always
 14 remember to check if the pointer is non-null before using it. If the pointer is null and if
 15 the user uses it anyway, this will cause a segmentation fault. *art*'s exception management
 16 system is not able to catch a segmentation fault; therefore *art* will stop execution immedi-
 17 ately, without going through the orderly shutdown process. All of the work in the *art* job
 18 will be lost.

19 To solve both of these problems *art* has created a class template named `art::Ptr`¹. With
 20 this technology, each `GenParticle` has a data member:

```
21 art::Ptr<GenParticle> parent_;
```

22 and a matching public accessor:

```
23 art::Ptr<GenParticle> const& parent() const { return parent_; }
```

24 With this modification, the code fragment from earlier in this section would read:

```
25 GenParticle gen; // Initialized somehow.
26 if ( gen.parent().isNonnull() ){
27     std::cout << "The_PDG_ID_of_the_parent_particle_is:__"
28             << gen.parent()->pdgId() << std::cout;
29 }
```

¹ The symbol `Ptr` is pronounced as two syllables: “pu”-“ter”, with the emphasis on the first syllable; the “u” in the first syllable is pronounced as in “put”.

1 Compare the way that the `pdgId()` function is accessed in this fragment to the way
2 that it is accessed in the earlier fragment: they are exactly the same. In this sense an
3 `Ptr<GenParticle>` can be used exactly as if it were a bare pointer, a `GenParticle`
4 `const*`.

5 When a `Ptr` is created, it is given enough information about the pointee that, after being
6 read from disk, it can recompute a valid pointer to the pointee. Moreover, the `Ptr` itself
7 knows if it has no pointee; for example, a primary particle in a `GenParticleCollection`
8 has no parent. A `Ptr` in this state is said to be invalid or to be null. As illustrated in the
9 above fragment you can ask if a `Ptr` is valid by calling its member function `isNonnull()`.
10 Whenever you follow a `Ptr`, it does an internal validity check. If the `Ptr` is valid, it simply
11 behaves as if it were a bare pointer. If the `Ptr` is invalid, then it will throw an exception.
12 *art* will catch the exception and will attempt a orderly shutdown. Your work to that point
13 will be preserved.

14 In summary, a `Ptr` has two big advantages over a bare pointer:

- 15 1. After being read from disk, it knows how to recompute a valid pointer to its pointee.
- 16 2. If you try to follow an invalid `Ptr`, the `Ptr` will throw an exception and *art* will
17 attempt an orderly shutdown.

18 The above story was told to explain why the class `GenParticle` uses a `Ptr` to spec-
19 ify its parent particle. A similar problem exists for specification of the children of a
20 `GenParticle`. Because a `GenParticle` may have more than one child, a collec-
21 tion type is needed to specify the children and the type used by `GenParticle` is
22 `art::PtrVector`, which is a collection of `Ptrs`, all of which point to objects within
23 the same data product.

24 In other places in the *art* documentation you will see the statement that a `Ptr` is a form
25 of persistable pointer. That simply means that you can write it to disk and it will still be
26 usable after being read back.

27 `Ptr` and `PtrVector` are two of the three class templates that *art* provides to describe
28 relationships between objects in data products. The third class template, `art::Assns`,
29 will be discussed in Chapter 32. ch:wbk:assns

20.2 Prerequisites

Prerequisites for this chapter include all of the material in Part I (Introduction) and the material in Part II (Workbook) up to and including Chapter 19

20.3 What You Will Learn

This exercise shows how to use `Ptrs` that are found in an existing data product. *A future exercise will show how to create `Ptrs`.*


FIXME: Reference to the future chapter once it is written.

In this exercise you will learn about:

1. The class `art::ProductID`, which is a second way to identify a data product.
2. The class template `art::Ptr`
3. The components of a `Ptr`: the product ID and the key.
4. The class template `art::PtrVector`
5. The many different notions of validity
6. How to test a `art::Ptr` for validity.

20.4 Rough Notes

1. Introduce `art::Ptr`.
 - (a) Describe the problem with persisting pointers.
 - (b) `art::Ptr` acts as a persistable pointer.
 - (c) under the covers it contains an `art::ProductID` and an offset
 - (d) Can only point to a second tier object and can only do it if the the container has a random access iterator: `std::vector`, `cet::map_vector` not sure about `std::map<integral_type,T>?`
2. Introduce `PtrVector`.

- 1 3. Distinguish `std::vector<art::Ptr<T> >` from `art::PtrVector<T>`.
- 2 4. Exercise:
 - 3 ○ Find the phi meson in the `GenParticleCollection`; find its daughters.
 - 4 ○ Compute the invariant mass of the two daughters and plot it.
 - 5 ○ Verify that `Ptr` to parent for first generation particles is null; is `isNull` and `is-`
 - 6 `Nonnull`.
- 7 5. User exercise: round trip: that daughters all point to the mother.
- 8 6. Bug finding:
 - 9 ○ Remove the `hasParent` test and watch the code crash
 - 10 ○ Maybe look at daughters 1 and 2 for a particle with no daughters?
- 11 7. No need to `event.get` the pointee data product. `art` does that for you.
- 12 8. Reserved value to indicate null.
- 13 9. Can only write `ptr` to a second tier object.
- 14  10. No support for writing `Ptr`'s into `run` or `subrun` objects; (since we have not yet
- 15 introduced `run` and `subrun` products)
- 16 11. Distinguish the many ideas of validity/invalidity.
 - 17 (a) `offset` is the reserved value.
 - 18 (b) `ProductID` does not exist in the registry.
 - 19 (c) `ProductID` exists but the data product is not in memory and not in the input file.
 - 20 (d) `ProductID` exists and the product is in the input file but not yet in memory.
 - 21 (e) `ProductID` exists and the product is in memory but the index does not exist.
 - 22 (f) `ProductID` exists and the product is in memory and the index does exist.

- 1 **20.5 Running the Exercise**
- 2 **20.6 Discussion**
- 3 **20.7 Test Your Understanding**
- 4 **20.8 Review**

DRAFT

1 21 The Geometry Service

geometryservice

- 2 1. The reason for doing this now is that these are needed for the next exercises.
- 3 2. Nothing new about services. This is just a tour of the tracker geometry and the
4 tracker efficiency.
- 5 3. `tex::Geometry`, `tex::Tracker`, `tex::TrackerComponent`, `tex::Shell`
- 6 4. Geometry is only valid if a run is in memory - so you cannot use it at `c'tor` or
7 `beginJob` time.
- 8 5. print the geometry information
- 9 6. make some histograms of the geometry info
- 10 7. Change some geometry numbers and watch the histograms change.

11 21.1 Prerequisites

12 21.2 What You Will Learn

13 21.3 Running the Exercise

14 21.4 Discussion

15 21.5 Suggested Activities

1 22 Instance Names of Data Products

2 In the input files, the Intersections object come in two instances, one for each detec-
3 tor.

- 4 1. Show how to get a data product that has an instance name.
- 5 2. Refer to future examples for `getMany` and selector functions; don't do them here.
- 6 3. Advise people NOT to use `getMany` or selectors as a substitute for “`getByType`”,
7 which we removed from the API.

1 23 InRun DataProducts

2 The DetectorSimulation module also creates a Run summary object that describes the
3 results of

4 the simulation.

- 5 1. Get the MCRunSummary data product from the event.
- 6 2. Histogram the number of particles with 5 or more intersections.
- 7 3. No example provided for InSubRun but just follow the pattern.
- 8 4. Show that you can also get the Run and SubRun objects from the event; and the Run
9 from the SubRun.
- 10 5. This data product is not a collection type. That is allowed everywhere.

1 24 Provenance

- 2 1. Get Provenance pointer from the handle.
- 3 2. Get full product name, and names of each part
- 4 3. Friendly and non-friendly type name
- 5 4. Get parameter set(s) in the creation chain: GenParticles are in the chain for Intersec-
- 6 tions.

DRAFT

1 25 Listing the Data Products in a File

- 2 1. Tell user to re-read story of the toyExperiment from the intro.
- 3 2. 4-part name of a data product.
- 4 3. How to list what data products are in a file
- 5 4. No code, just a .fcl file this time.
- 6 5. Maybe: Add a 5th example input file which is the same as the 4th one but done in
- 7 two steps so that we can
- 8 illustrate ProductID's better. Maybe defer this until later?

1 26 Producer Module

t:producer

2 Need to extend toyExperiment/RecoDataProducts to have a class to represent a recon-
3 structed

4 phi meson candidate, tex::Candidate

5 has art::Ptr to FittedHelixData Collection, plus mass hypothesis, plus mass and covari-
6 ance.

7 Both read facade and write facade this time.

8 1. Loop over all pairs of FittedHelixData.

9 2. Only consider oppositely charged pairs

10 3. Compute the invariant mass; use RecoTrk.

11 4. Add candidates that pass some loose cuts to the data product?

12 5. Read it back and make histograms from it; compare to histograms in the creator.

13 6. This module will show how to create Ptrs to the fitted helix data.

14 7. This introduces trigger paths.

15 8. Write a companion analyzer module to make some plots; run it both within the
16 writing job and

17 in a separate readback job.

18 9. Use the filedumper to see that it really is there.

1 27 Filter Module

- 2 1. Loop over data product created by the example producer (`tex::Candidate`).
3 Select events that have at least one `tex::Candidate` that passes some tighter cuts than
4 the producer cuts.
- 5 2. We could have written the producer as a filter but we recommend that you keep the
6 two ideas separate so that
7 you can write different workflows.

1 28 Configuring Output Modules

- 2 1. Select events from trigger paths completion status. Mention that you can do boolean
3 logic on multiple paths.
- 4 2. Multiple output files in one job.
- 5 3. Select which data products go into which file.
- 6 4. Use `filedumperoutput` to verify what is in each file.
- 7 5. mention fast cloning; superficial description; when might you want to enable/disable
8 it?
- 9 6. Show the grammar to select events that threw exceptions or otherwise failed.

1 29 Creating Your Own Data Products

- 2 1. Not sure yet what to use as an example; maybe the right answer is to make this
3 orthogonal to the rest of the toyExperiment.
- 4 2. Create a data product library with just this new data product.
- 5 3. You should learn your experiment's policy on whether they want to have many li-
6 braries with few
7 data products each or if they want to bundle related data products into a smaller
8 number of libraries.
- 9 4. classes.h, classes_def.xml
- 10 5. Your data product class must be either movable or swappable.
- 11 6. art::Wrapper - this holds the TObjectness.
- 12 7. Structure of an art event-data file.

1 30 Module Name Collisions

- 2 1. A second HelloWorld module that prints something that is different enough from the
3 HelloWorld module in exercise 2.
- 4 2. Need to delay making the .so until we are at this exercise or else the first module
5 exercise will fail.

DRAFT

31 More FCL Concepts

chap:more:fcl

1. This will be a standalone module with its's own toy .fcl file. I initially thought to use the full MC Chain but that is too much to swallow at once.
2. Explain PROLOG, @local, scoping rules, dot notation
3. When they arrive, @table and @sequence (and @echo?)
4. true, false, infty, +infty, -infty, @nil
5. Check Mu2e tutorial for additional details needed here.
6. To get fully resolved fcl file. ART_DEBUG_CONFIG=1; art
7. Explain how to specialize get_if_present<T>. That is, how to write

```
auto v = pset.get<MyClass>("key");
```

A good example might be a binning class:
myHist1 : nbins : 100 xlow : 0 xhigh : 1.
Maybe this belongs in the User's Guide?

1 32 art::Assns Connecting Hits to Intersec- 2 tions

:wbk:assns

- 3 1. The model is that a tex::Intersection is an MC truth object but a TrkHit is reco object
4 - ie it could come
5 from the real experiment so it must not contain any MCTruth info
- 6 2. Because of inefficiency there are Intersections with no corresponding TrkHit.
- 7 3. When we run on actual experimental data we don't want to load any MC libraries.
8 Therefore we follow the pattern
9 that we segregate MC classes into their own directories. Therefore
10 Classes in RecoDataProducts must not know about classes in MCDataProducts.
11 Classes in MCDataProducts may know
12 about classes in RecoDataProducts. This pattern is common in many experiments.
- 13 4. Simulations/HitMaker_module.cc:
 - 14 (a) Has a model of inefficiency; look at run-time config for this but not the code.
 - 15 (b) produces TrkHitCollection - a single collection for both tracker components -
16 just because.
 - 17 (c) produces TrkHitMatch - connection between Hits and Intersections
- 18 5. The name "God's block" - does anyone still use it?
- 19 6. Example code: shows how to find an intersection given a hit.

- 1 7. Example code: shows how to find a hit given an intersection - it is allowed to not
2 have a match.
- 3 8. Example code: loop over all matches.
- 4 9. Assns< A,B> can be used as Assns< B,A>.
- 5 10. under the cover the data product is just pairs of art::Ptr<A> and art::Ptr
- 6 11. Not illustrated Assns<A,B,D>. Statement about this and a reference to a later ex-
7 ample.
- 8 12. Discussion of when to use Ptr and when to use Assns.
- 9 13. Third option: data products with lock-step indexing.

1 33 Facade Pattern: Fitted Helices

- 2 1. This has a read facade only.
- 3 2. Fitted helix is a geometric object. No knowledge of momentum until you add external information
4
5 (ie the magnetic field).
- 6 3. The data product is FittedHelixData
- 7 4. The fully powerful object is FittedHelix - it is not a data product because it contains
8 non-data-product information,
9 like it's connection to services. Breaking this rule can lead to linkage loops.
- 10 5. Some code to explore the properites of a fitted helix and make histograms.

1 34 art::View

- 2 1. Do we really want to do this without a good example of polymorphic data?
- 3 2. Show the mechanics of how to do it using one of the GenParticle exercises.
- 4 3. State that it is useful for polymorphic data (which we won't have an example of
- 5 here unless we can
- 6 make one using TOF and tracker).

35 Random Numbers

1. Brief discussion: engines vs distributions; seeds vs state;
2. CLHEP weirdness: fire and shoot and global engines.
3. Two modules: one with a single engine and one with two engines
4. The single engine module will print an int from a flat distribution, limited by max-print.
Printout will be event number and value of the random variate.
5. Histograms of flat, gauss, poisson and a few others.
6. Always use RandGausQ and RandPoissonQ
7. Special case of Geant4 and CLHEP global engine
8. RandGeneral - belongs in UserGuide not here.
9. Discussion of seeding strategies; should I put the seed service into the toyExperiment?
10. Save the seed. Write an output file.
11. Upon request art team could provide StdRandomNumberService that works with random engines from the C++11 std library instead of the CLHEP ones. But, for now, G4 controls.

1 36 Repeating Random Numbers

- 2 1. Start at event N and verify that random numbers are properly repeated.
- 3 2. Run 1000 events; then run 500 events and save the state at end of job; continue the
4 job
5 and compare to the full job. Just use the module that prints int's drawn from a flat
6 distribution.
- 7 3. Warning: sometimes G4 hadronics code does not play nice.

1 37 Writing Your Own Service

- 2 1. Discuss legacy service, global service and schedule local services.
- 3 2. Inspect PDT, Geometry and Conditions services in the toyExperiment code.
- 4 3. Assign a project to write a service that duplicates the behaviour of tracer service? Is
- 5 there a better project?

1 38 Running the MC Chain

- 2 1. Copy input04.fcl.
- 3 2. Explain what it does.
- 4 3. add the phi meson candidate producer
- 5 4. Run the MC chain.
- 6 5. Run some analyzer modules on the output.

1 39 Reconstruction on Demand

- 2 1. Describe reco on demand
- 3 2. Show the .fcl file that implements the MC chain as reco on demand.
- 4 3. Use `-tracer` to follow what it is doing.

DRAFT

1 40 Run the Existing 2D EventDisplay

ch:2devtdisp

- 2 1. Describe what the fcl parameters do.
- 3 2. Scan some events.
- 4 3.

1 41 Review of all Modules in toyExperiment

- 2 1. Discuss the code and other files in the toyExperiment package.
- 3 2. Say what each does.
- 4 3. Comment on features found in each that were not covered by the above.

DRAFT

1 42 Running a Grid Job

- 2 1. Cartoon description of the grid
- 3 2. Run the MC chain as a grid job.
- 4 3. Compare histograms
- 5 4. Need to work with FIFE people to define this.

¹ 43 Introducing SAM and dcache

- ² 1. Need to work with FIFE people to define this.

DRAFT

44 3D Event Displays

ch:3devtdisp

44.1 Introduction

3devtdisp-intro

Most high energy physics experiments have some form of an *Event Display* used to visualize physics objects together with the detector geometry. Even simple visualization tools, like the 2D example introduced in Chapter 40, **FIXME: not in the doc yet** are indispensable for developing reconstruction algorithms, validating simulation code, and providing valuable insight when doing analyses. This chapter will show you how to combine the *art* event-processing framework with the visualization framework provided by ROOT's Event Visualization Environment (EVE). Doing so will allow you to easily create sophisticated yet flexible 3D visualization tools for your experiment. Using the **FIXME: existing** HEP-centric visualization classes in EVE will also help you avoid re-inventing the wheel, freeing up valuable time that is better spent on doing physics.

This chapter will describe a 3D event display for visualizing the detector geometry, simulated hits, and simulated tracks of the toyExperiment. It was built using ROOT's GUI and EVE classes in an *art* analyzer module. In addition, it uses an *art* service called `EvtDisplayService`. Aside from the normal operation of stepping forward sequentially through events in an input ROOT file, this service allows you to rewind or go backwards sequentially. It also provides random access to any event in the file. We will begin by describing the user interface of the event display and how to run it. We will then examine the code for the analyzer module in detail to understand how it works.

Detailed information on the ROOT GUI and EVE classes can be found starting from the ROOT Reference Guide page, <http://root.cern.ch/drupal/content/reference-guide>.

Additional information for the ROOT GUI classes can also be found in the ROOT User's guide, <http://root.cern.ch/drupal/content/users-guide>.

1 There is also a ROOT page dedicated to EVE, <http://root.cern.ch/drupal/content/eve>, where
2 there are references to presentations and write-ups on EVE from conference proceed-
3 ings.

4 44.2 Prerequisites

5 Prerequisites for this chapter include all of the material in Part I (Introduction) and all
6 of the chapters in Part II (Workbook) up to and including Chapter 21. *Completion of*
7 *Chapter 40 on the 2D event display is also recommended.*

8 44.3 What You Will Learn

9 In this chapter you will learn:

- 10 ○ how to run the event display for the toyExperiment
- 11 ○ what the various widgets in the EVE browser mean and do
- 12 ○ how to use `art::EvtDisplayService`
- 13 ○ how to create a default EVE browser
- 14 ○ how to extend the EVE browser by adding a navigation panel and multiview ortho-
15 graphic projections
- 16 ○ how to use button and text entry widgets in the navigation panel
- 17 ○ how to connect signals emitted by widgets to receiver slots
- 18 ○ how to use the EVE projection manager for easily creating 2D orthographic views
- 19 ○ how to import detector geometry from a file and draw it
- 20 ○ how to use EVE classes to draw physics objects like hits and tracks

21 44.4 Setting up to Run this Exercise

22 If you are logging in after having closed an earlier session, follow the instructions in
23 Chapter 11. If you are continuing on directly from the previous exercise, keep both your

1 source and build windows open.

2 44.5 Running the Exercise

3 44.5.1 Startup and General Layout

4 To start up the event display, run the following command in your build directory:

5 **FIXME:** *Need to create symlinks before it will work.*

6 `art -c fcl/EventDisplay3D/eventDisplay01.fcl .`

7 This brings up a TEveBrowser which is a customized version of the ROOT TBrowser for
 8 ROOT's Event Visualization Environment (EVE). The first thing you will notice is that
 9 the TEveBrowser used in this exercise, which is shown in Figure 44.1, looks just like a
 10 TBrowser except for some key differences which will be described next. On the left hand
 11 side are three tabbed panes, which will be referred to collectively here as the control panel.
 12 Aside from the usual *Files* pane, this panel includes two new tabbed panes labeled *Eve* and
 13 *Event Nav*. Like a TBrowser, the right hand side is divided into two major areas with the
 14 usual ROOT command console at the bottom. Above the command console is the main
 15 EVE display area with two tabbed panes labeled *Viewer 1* and *Ortho Views*. Technically,
 16 each of these three areas is a ROOT TGTAB widget. Internally, EVE uses the indices
 17 kLeft, kRight, and kBottom to refer to the control panel, main EVE display area,
 18 and ROOT command console, respectively.

19 44.5.2 The Control Panel

20 44.5.2.1 The List-Tree Widget and Context-Sensitive Menus

21 Clicking on the *Eve* tab in the control panel reveals a list-tree with four top-level items
 22 labeled *WindowManager*, *Viewers*, *Scenes*, and *Event* (see Figure 44.2a). To facilitate our
 23 discussion, think of each tabbed pane in the main EVE display area as a window. Each
 24 window can have a single frame, like the first one named *Viewer 1*, or have multiple sub-
 25 frames, like the second one named *Ortho Views*. Each frame or subframe has a viewer
 26 associated with it and each viewer can consist of one or more scenes.

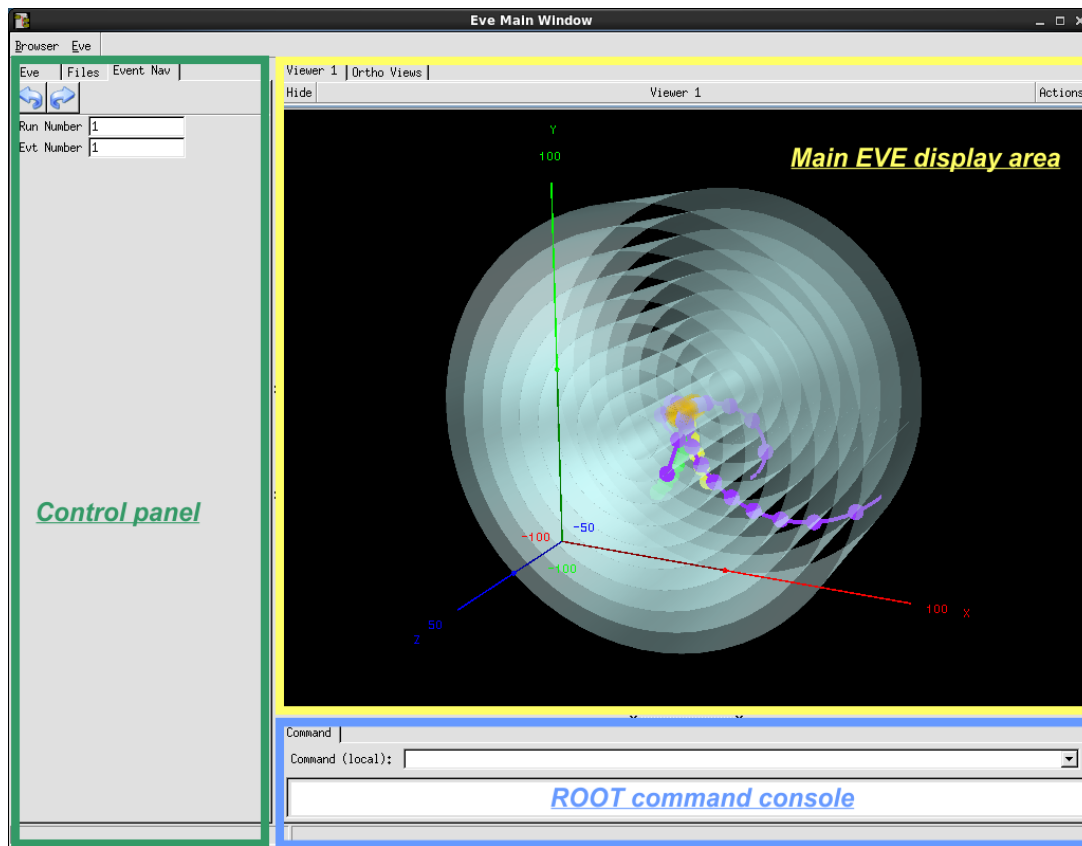


Figure 44.1: The TEveBrowser is a specialization of the ROOT TBrowser for the ROOT EVE Event Visualization Environment. Shown above is the one used in this workbook exercise which is divided into three major regions: 1) a control panel, 2) a main EVE display area, and 3) a ROOT command console.

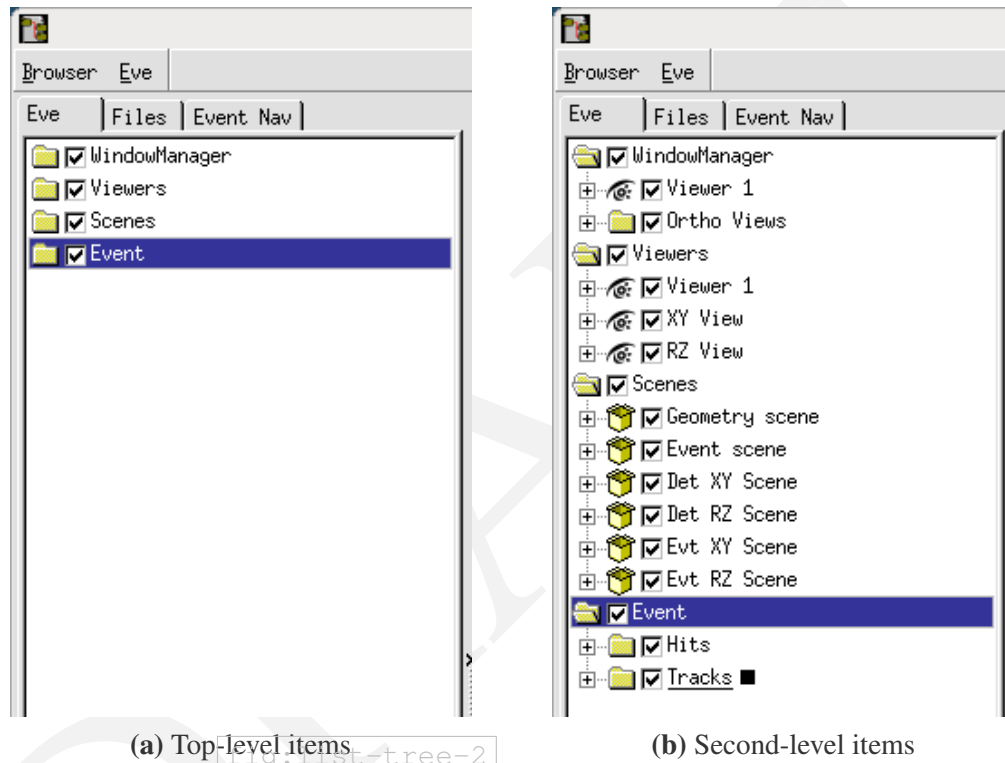


fig:list-tree-1

(a) Top-level items

fig:list-tree-2

(b) Second-level items

Figure 44.2: Shown above are two different views of the list-tree widget, showing the top-level items in (a), and expanded to second-level items in (b).

1 Expanding the first top-level item in the list-tree labeled *WindowManager*, you will find
 2 two children, labeled *Viewer 1* and *Ortho Views*, as shown in Figure 44.2b. These items
 3 represent the windows in the two tabbed panes of the main EVE display area. Associated
 4 with the *Viewer 1* window is a viewer of the same name which presents a 3D perspec-
 5 tive view of the detector and event (we will refer to the representation of the detector and
 6 the event collectively as the *detector-event*). Associated with the *Ortho Views* window are
 7 two viewers named *XY View* and *RZ View*. The *XY View* presents an orthographic view of
 8 the detector-event looking down along the negative *z*-axis. The *RZ View* presents an ortho-
 9 graphic view of the detector-event, looking down along the positive *x*-axis. All three view-
 10 ers, which are C++ objects belonging to the ROOT *TEveViewer* class, show up as chil-
 11 dren when the *Viewers* top-level list-tree item is expanded (refer to Figure 44.2b).

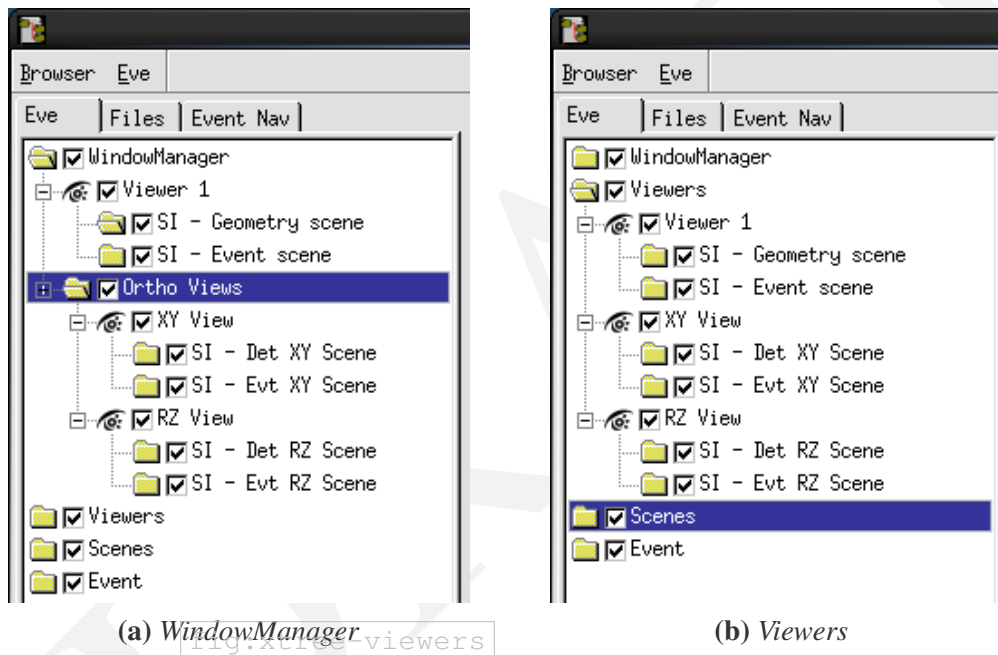
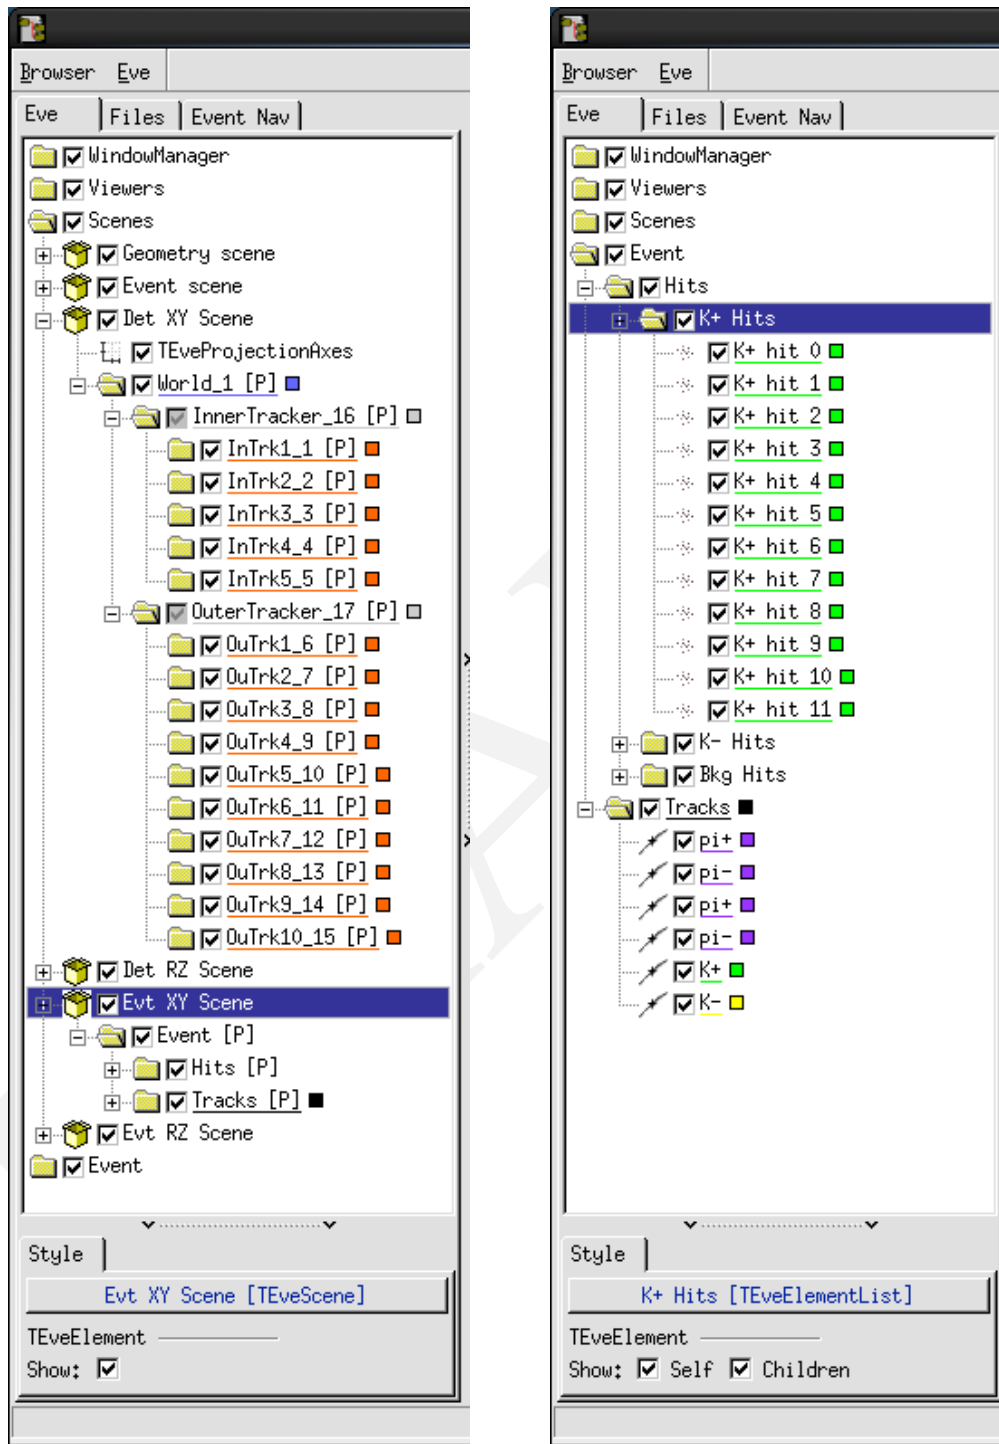


Figure 44.3: Expanded views of the (a) *WindowManager* and (b) *Viewers* list-tree items.

12 As shown in the expanded list-tree widgets in Figures 44.3a and 44.3b, each view contains
 13 two separate scenes (ROOT *TEveScene* class). The first view is associated with the
 14 static components of the detector-event representing the geometry of the detector. The
 15 second view is associated with the dynamic components of the detector-event which, in
 16 this example, are the generated hits and tracks. All six scenes in our example (two per
 17 viewer) are listed as children when the *Scenes* top-level list-tree item is expanded (see



(a) Scenes

(b) Event

Figure 44.4: Expanded views of the (a) Scenes and (b) Event list-tree items.

fig:list-tree-2
1 Figure 44.2b). The two scenes named *Geometry Scene* and *Event Scene* are associated
2 with the viewer named *Viewer 1*. The two named *Det XY Scene* and *Evt XY Scene* are
3 associated with the viewer named *XY View*. Finally, the two scenes named *Det RZ Scene*
4 and *Evt RZ Scene* are associated with the viewer named *RZ View*.

5 Expanding each of the second-level items under *Scenes* reveals a third layer of children.
6 For example, fig:xtree-scenes Figure 44.4a shows expanded views of the *Det XY Scene* and *Evt XY Scene*
7 items in the *Scenes* branch. In general, the static detector scenes have a direct descendant
8 named *World_1*, representing the master volume. Expanding *World_1*, you will find two
9 items, labeled *InnerTracker_16* and *OuterTracker_17*, representing the two subvolumes
10 within *World_1*, which, in turn, contain the inner and outer tracking detectors, respectively.
11 In addition to *World_1*, the static scenes associated with the orthographic views have a
12 second direct descendant labeled `TEveProjectionAxes` (ROOT class), representing
13 the coordinate axes displayed in these views. Dynamic event scenes have a single direct
14 descendant named *Event* which has two children labeled *Hits* and *Tracks* representing
15 the generated tracks and the detector hits produced by these tracks (see bottom part of
16 fig:xtree-scenes Figure 44.4a).

17 The fourth and last top-level item in the list-tree widget is labeled *Event*, representing
18 the dynamic component of a detector-event. In ROOT, this item is associated with an ob-
19 ject belonging to the `TEveEventManager` class. Referring to fig:xtree-event Figure 44.4b, we find
20 that clicking on this item expands it into two children labeled *Hits* and *Tracks*, repre-
21 senting the generated hits and tracks associated with an event. These items are associ-
22 ated with objects belonging to the ROOT `TEveElementList` class. Expanding the *Hits*
23 item expands it into another layer, consisting of three list items (also associated with
24 `TEveElementList` class objects) representing the three categories of generated parti-
25 cles: [1] positive kaons (K^+) and [2] negative Kaons (K^-) coming from the ϕ meson, and
26 [3] everything else not originating from phi-meson (*Bkg*). Clicking on each of these cate-
27 gories expands them into a list of hits which are leaf-type items that terminate the list-tree
28 branches. Going back up two levels in the list-tree hierarchy and clicking on the *Tracks*
29 item expands it into a list of leaf-type items representing each generated track in the event
30 labeled by the name of the particle type associated with the track.

31 One last thing to note about the list-tree widget is the presence of a check box preceding
32 the item label. This box is used to toggle the visibility of the graphical representation asso-
33 ciated with the item in the appropriate scenes shown in the main EVE display area.

- 1 Directly below the list-tree widget, in the same *Eve* panel, is a context-sensitive menu that
 2 presents various options relevant to the type of list-tree item selected (see Figure 44.5). The
 3 ROOT class to which the object associated with the selected item belongs is also displayed
 4 at the top of the menu, just below the tab labels. For example, selecting the *Viewer 1* item
 5 under the *Viewers* top-level list-tree item (*Viewers*→*Viewer 1*) shows us that the selected
 6 item is associated with a standalone GL-viewer object of the ROOT *TGLSAViewer* class.
 7 Four sets of options, grouped under separate tabs, are available for this item. Clicking on
 8 the third tab labeled *Clipping* presents options for creating cutaway views of the detector-
 9 event with either a clipping plane or box that can be interactively manipulated.

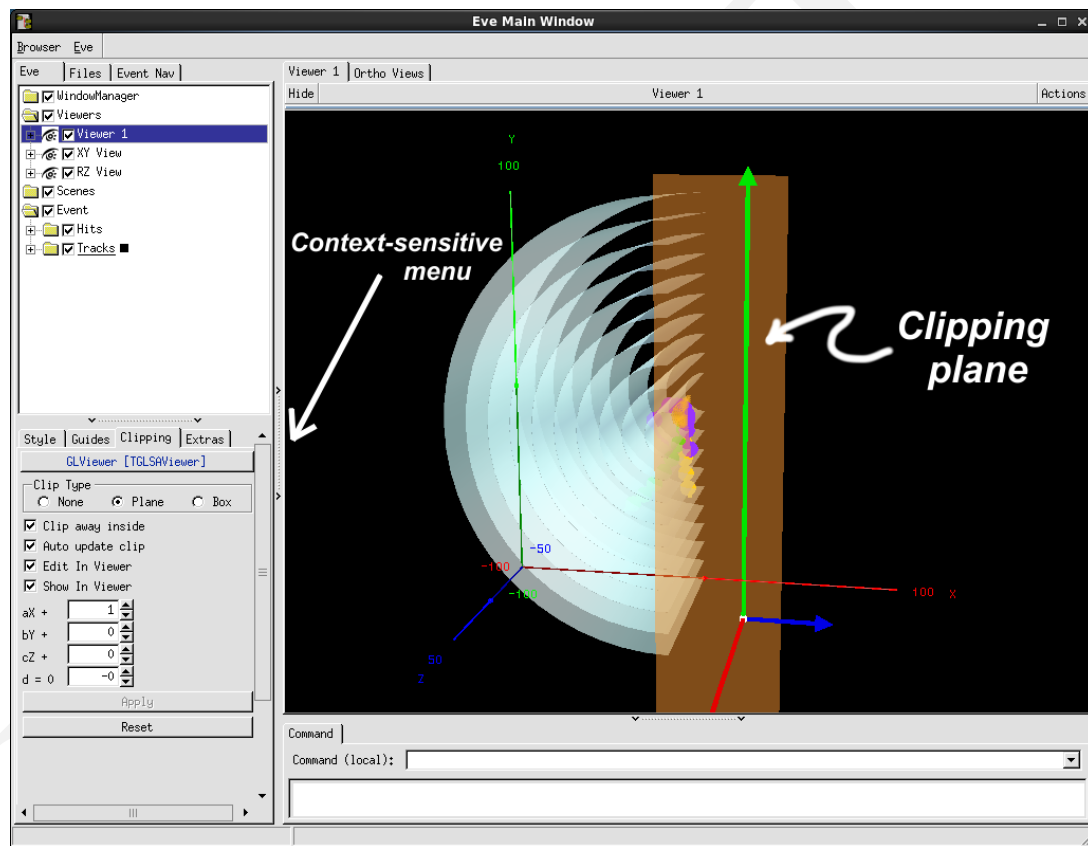


fig:clipping

Figure 44.5: The context-sensitive menu below the list-tree widget changes in response to the selected list-tree item. Shown above is the menu for a viewer-type item. In this example, we have enabled an interactive clipping plane.

- 10 As a second example, select one of the leaf items representing a generated track under the
 11 *Tracks* child of the *Event* top-level list-tree item. The context-sensitive menu tells us that

- 1 the selected item is associated with an object of the ROOT `TEveTrack` class. Included
- 2 under a single tab labeled *Style* are options that allow you to change the color and thickness
- 3 of the track and edit the attributes of its propagator. The context menus are generally
- 4 intuitive and their functionality fairly easy to figure out.

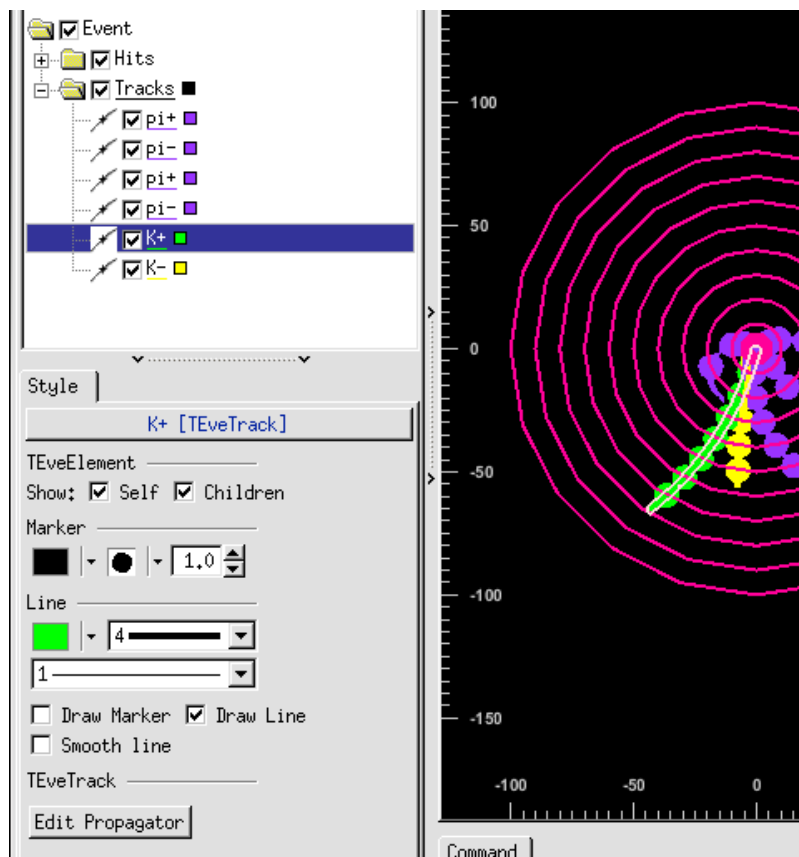


Figure 44.6: Shown above is the context-sensitive menu displayed below the list-tree widget when a track element is selected.

44.5.2.2 The Event-Navigation Pane

- 6 Let us skip the second tabbed pane labeled *Files* in the event control panel since this
- 7 identical to that found in a generic `TBrowser`. Instead, click on the third tab labeled *Event*
- 8 *Nav*. At the top of this pane, shown in Figure 44.7, are two arrow keys. Clicking on the
- 9 left arrow key moves you back one event in the input ROOT file and displays this event.

- 1 Clicking on the right arrow key moves you forward one event and displays this event.
- 2 Underneath the two arrow keys are two text entry widgets labeled *Run Number* and *Event*
- 3 *Number*. Entering valid Run and Event numbers allows you to jump directly to an event in
- 4 the input ROOT file associated with the specified run and event numbers and displays this
- 5 event.



Figure 44.7: The *Event Nav* pane on the control panel.

6 44.5.3 Main EVE Display Area

- 7 Having covered the event-navigation panel in some detail, let us now move on to the
- 8 right hand side of our TEveBrowser. Skipping the ROOT command console in the bottom,
- 9 which is identical to that in a generic TBrowser, let us focus our attention on the main EVE
- 10 display area which was briefly introduced in Section 44.5.1. As described previously, this
- 11 area has two tabbed panes labeled *Viewer 1* and *Ortho Views*, with windows showing a 3D
- 12 perspective shown in Figure 44.1 and two different 2D Orthographic views shown in Fig-
- 13 ure 44.8. Positioning the mouse cursor within a viewport allows you to rotate, pan across,
- 14 or zoom in and out of a scene. For simplicity, we use the term *zoom* here interchangeably
- 15 to mean modifying a scene by either changing the focal length of the camera or dollying
- 16 the camera. With the cursor positioned in the viewport, clicking on the left mouse but-
- 17 ton, while moving the mouse, rotates the scene using a virtual trackball centered on the
- 18 reference point. Clicking on the middle button, while moving the mouse, pans across the
- 19 scene. Clicking on the right button while moving the mouse zooms in and out of the scene.
- 20 Note that rotation of the scene with the right button is only possible in the perspective
- 21 view. For mice with a scroll middle button, zooming in and out of the scene can also be
- 22 accomplished by turning the wheel.

- 23 In the viewports of a tabbed pane, you will also notice a title bar at the top with the words
- 24 *Hide* and *Actions* on each end and the name of the associated viewer in the middle. If you

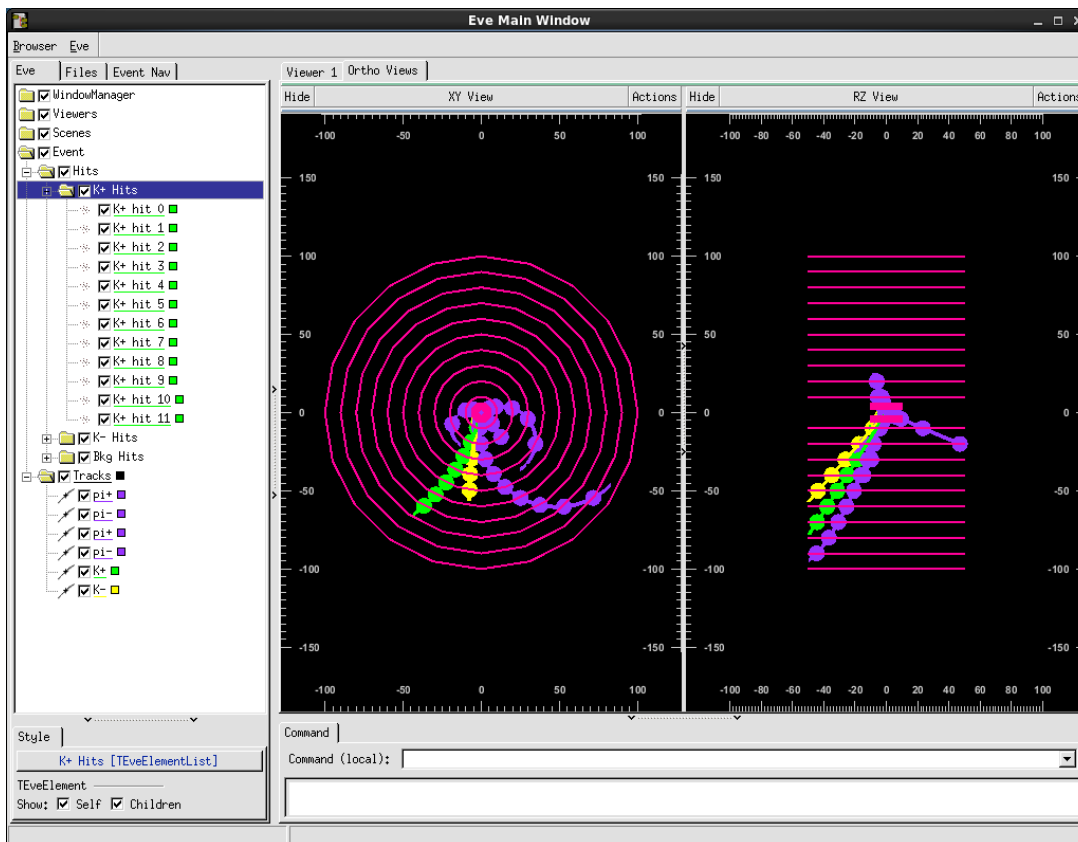


Figure 44.8: The orthographic XY and RZ views in the *Ortho Views* tabbed pane of the main EVE display panel.

- 1 hover the mouse cursor over the lower edge of this title bar, a pull-down menu bar, like the
 2 one shown in Figure 44.9, will appear underneath it. Selecting *Save* or *Save As* in the *File*
 3 pull-down menu allows you to save the scene displayed in the viewport to several output
 4 file format options which can then be sent to the printer. You can learn more about other
 5 available options and features by using the *Help* pull-down menu on the right.

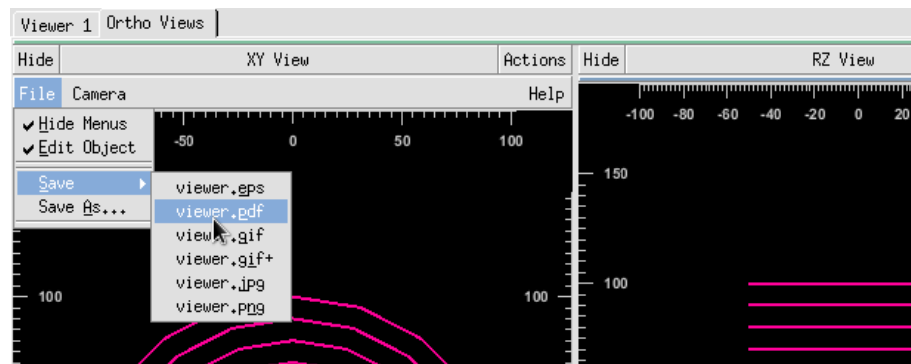


fig:menubar

Figure 44.9: Hovering the mouse cursor over the lower edge of the title bar in a viewport reveals a pull-down menu bar with more options.

- 6 You may have noticed that hovering your mouse cursor over an element in the dynamic
 7 event-type scenes (as opposed to static detector-type scenes), such as one representing
 8 a generated track and hit, highlights that particular element. It also pops up a text box,
 9 known as *tooltips*, like the ones shown in Figure 44.10, with information relevant to that
 10 element. In the case of generated track elements, for example, this information includes
 11 the PDGID, production vertex, and momenta at this vertex of the particle associated with
 12 the track. For the generated hits, this information includes the category this hit belongs to
 13 (K^+ , K^- , or Bkg), the hit index, and the ID of the detector element whose interaction with
 14 the generated particle produced the hit.

15 44.6 Understanding How the 3D Event Display Module Works

- 16
 17 The 3D event display module in this workbook exercise is an *art* analyzer module, which
 18 was introduced in Section 10.6.3. It uses *art* services, which were introduced in Sec-
 19 tion 3.6.3 and in Chapter 17. Aside from the geometry, conditions, and particle data table
 20 services used by other example modules discussed in this workbook, it also requires the




Figure 44.10: Tooltips with relevant information show up when the mouse cursor is hovered over (a) track and (b) hit elements

1 *art* service named `EvtDisplayService` that provides functions useful for event dis-
 2 plays. This service starts up a ROOT `TApplication` and provides sequential (forward
 3 and backward) or direct access of events in the input ROOT file.

4 The event display service used in this exercise is a much simplified version of the one used
 5 in the NOvANO ν A experiment.

6 With *art* providing the event-processing framework, the 3D event display module uses
 7 ROOT's Event Visualization Environment (EVE) to create a visual representation of the
 8 event-data and the detector geometry and to allow user interaction with it. In this section,
 9 we will describe the source file for the module, `EventDisplay3D_module.cc`, in
 10 some detail to demonstrate how to use ROOT's EVE to create event displays that work
 11 within the *art* framework.

12 44.6.1 Overview of the Source Code File `EventDisplay3D_module.cc`

13 For our discussion of the source code for the event display module, refer to the following
 14 file in the *art-workbook* source directory:

15 `art-workbook/EventDisplay3D/EventDisplay3D_module.cc`

1 At the top of this file, you will find the included headers arranged according to package.
 2 Following the includes for `toyExperiment` and `art`, are the ROOT includes, further
 3 broken down by the ROOT library they are associated with. The last included header,
 4 immediately after the ROOT headers, is for the `EvtDisplayUtils` class which will be
 5 described in more detail later.

6 Coming after the section containing the included headers, is an anonymous namespace
 7 with helper functions for (1) setting the transparency and color of drawn detector com-
 8 ponents, and (2) drawing generated hits. This is followed by the declaration of the
 9 `EventDisplay3d` class in the namespace `tex` and its constructor. Below this is the
 10 source code for the `makeNavPanel()` new member function used to create a GUI panel
 11 for event navigation.

12 Starting at around the middle of the file, is the source code for the `art`-defined member
 13 functions `beginJob()` and `analyze()`. Tasks that only need to be done once or are as-
 14 sociated with things that remain constant from event-to-event, such as starting up the EVE
 15 application manager, creating the GUI, initializing the graphics environment, and draw-
 16 ing detector components, are done at the beginning of the `art` job in the `beginJob()`
 17 member function. On the other hand, things that change from event-to-event, like drawing
 18 generated hits and tracks, are done in the `analyze()` member function.

19 44.6.2 Class Declaration and Constructor

20 Let us now take a closer look at the code, beginning with the class declaration shown in
 21 Listing 44.1. The constructor and `art`-defined methods are declared in lines 7-9. Lines 13-
 22 22 declare the data members that are initialized to parameter set values. As discussed in
 23 Section 16.7.1, the first member, `gensTag_`, specifies part of the requested data product
 24 name. The other data members in these lines are described below:

- 25 1. `drawGenTracks_` turns drawing of generated tracks on or off.
- 26 2. `drawGenHits_` turns drawing of generated hits on or off.
- 27 3. `hitMarkerSize_` determines the size of the spheres used to represent a generated
 28 hit.
- 29 4. `trkMaxR_` sets the maximum radial extent to which generated tracks are drawn.

Listing 44.1: The declaration of the class `EventDisplay3D` from `EventDisplay3D_module.cc`

```

1 namespace tex {
2
3   class EventDisplay3D : public art::EDAnalyzer {
4
5   public:
6     explicit EventDisplay3D(fhicl::ParameterSet const& pset);
7     void beginJob() override;
8     void endJob() override;
9     void analyze(const art::Event& event) override;
10
11  private:
12    // Set by parameter set variables.
13    art::InputTag gensTag_;
14    bool          drawGenTracks_;
15    bool          drawHits_;
16    Double_t      hitMarkerSize_;
17    Double_t      trkMaxR_;
18    Double_t      trkMaxZ_;
19    Double_t      trkMaxStepSize_;
20    Double_t      camRotateCenterH_;
21    Double_t      camRotateCenterV_;
22    Double_t      camDollyDelta_;
23
24    art::ServiceHandle<Geometry> geom_;
25    art::ServiceHandle<PDT>      pdt_;
26
27    std::unique_ptr<tex::EvtDisplayUtils> visutil_;
28    TEveGeoShape* fSimpleGeom;
29
30    TEveViewer *fXYView,*fRZView;
31    TEveProjectionManager *fXYMgr,*fRZMgr;
32    TEveScene *fDetXYScene,*fDetRZScene;
33    TEveScene *fEvtXYScene,*fEvtRZScene;
34
35    TGTextEntry *fTeRun,*fTeEvt;
36    TGLabel     *fTlRun,*fTlEvt;
37
38    TEveTrackList *fTrackList;
39    TEveElementList *fHitsList;
40
41    void makeNavPanel();
42  };
43 }

```

- 1 5. `trkMaxZ_` sets the maximum extent in the z direction to which generated tracks
2 are drawn.
- 3 6. `trkMaxStepSize_` sets the upper limit for the step size used by the propagator
4 to draw the track. Smaller step sizes produce smoother tracks.
- 5 7. `camRotateCenterH_` specifies the camera elevation angle in radians above (neg-
6 ative) or below (positive) the $x - z$ plane.
- 7 8. `camRotateCenterV_` specifies the camera's azimuth or polar angle in radians in
8 the $x - z$ plane, with positive angles corresponding to a clockwise rotation.
- 9 9. `camDollyDelta_` specifies the amount, in mouse pixel units, by which to move
10 the camera towards (positive) or away from (negative) the reference point without
11 changing the camera's focal length.

12 On lines 24-25, we declare the `geom_` and `pdt_` service handles (see Section 17.5.1)
13 for the geometry and particle data table services, respectively. A `unique_ptr`, a type of
14 smart pointer provided by the C++ Standard Library, is declared on line 27 to manage a raw
15 pointer to an object of type `EvtDisplayUtils`. As we will see in more detail later, this
16 class provides the communication link between the ROOT EVE GUI and the event dis-
17 play service. Without going into too much detail, using a smart pointer for a dynamically
18 created object, like `visutil_`, helps avoid problems with dangling pointers and memory
19 leaks, by automatically freeing up resources when they are no longer needed.

20 Pointers to the viewers associated with the 2D XY and RZ orthographic views, described
21 in 44.5.2.1, are declared on line 30. Pointers to the projection managers associated with
22 each of these viewers are declared on line 31. The pointers, declared on lines 32 and 33,
23 refer, respectively, to the static detector geometry and per-event scenes associated with
24 each viewer. The declarations on lines 35 and 36 are for pointers to the text entry widgets
25 and their labels, shown in Figure 44.7. On lines 38 and 39, we declare pointers to contain-
26 ers for `TEveTrack` and `TEvePointSet` type objects, respectively, which are used to
27 draw generated tracks and hits. Finally, on line 41, we declare a new member function that
28 creates the event-navigation pane described in 44.5.2.2.

29 Shown in listing 44.2 is the code implementing the constructor of the `EventDisplay3D`
30 class, which makes use of the colon initializer syntax described in Section 6.7.5. At the top
31 of the initializer list on line 2, is the required constructor for the base class of our analyzer

Listing 44.2: Implementation of the constructor for the EventDisplay3D class from EventDisplay3D_module.cc

```

1 tex::EventDisplay3D::EventDisplay3D(fhicl::ParameterSet const& pset):
2   art::EDAnalyzer(pset),
3   gensTag_      ( pset.get<std::string>("genParticleTag") ),
4   drawGenTracks_ ( pset.get<bool>      ("drawGenTracks", true) ),
5   drawHits_     ( pset.get<bool>      ("drawHits", true) ),
6   hitMarkerSize_ ( pset.get<Double_t>  ("hitMarkerSize", 2.) ),
7   trkMaxR_      ( pset.get<Double_t>  ("trkMaxR", 100.) ),
8   trkMaxZ_      ( pset.get<Double_t>  ("trkMaxZ", 50.) ),
9   trkMaxStepSize_ ( pset.get<Double_t>  ("trkMaxStepSize", 1.) ),
10  camRotateCenterH_ ( pset.get<Double_t>  ("camRotateCenterH",-0.26) ),
11  camRotateCenterV_ ( pset.get<Double_t>  ("camRotateCenterV",-2.
12  ) ),
13  camDollyDelta_   ( pset.get<Double_t>  ("camDollyDelta",500.) ),
14  geom_(),pdt_(),
15  visutil_(new tex::EvtDisplayUtils()),
16  fSimpleGeom(nullptr),
17  fXYView(nullptr),fRZView(nullptr),
18  fXYMgr(nullptr),fRZMgr(nullptr),
19  fDetXYScene(nullptr),fDetRZScene(nullptr),
20  fEvtXYScene(nullptr),fEvtRZScene(nullptr),
21  fTeRun(nullptr),fTeVt(nullptr),
22  fTlRun(nullptr),fTlEvt(nullptr),
23  fTrackList(nullptr),fHitsList(nullptr){
24  if ( trkMaxStepSize_ < 0.1 )trkMaxStepSize_ = 0.1;
25
26 }

```

1 module (see Section 10.6.3.3). Lines 3 through 12 initialize the data members that can
2 be configured through the FHiCL file to the parameter set values. The service handles for
3 the geometry and particle data table services are initialized on line 13, and resources are
4 dynamically allocated for the EvtDisplayUtils-type member object on line 14. Lines
5 15 through 22 initialize the pointers, for the member objects that belong to EVE classes,
6 to the value `nullptr` (see Section 17.5). The only work that is actually done in the body
7 of the constructor on line 24, is to guard against having the track propagator take step sizes
8 that are too small.

1 **44.6.3 Creating the GUI and Drawing the Static Detector Components in the** 2 `beginJob()` **Member Function**

3 In the code overview of Section 44.6.1, we mentioned that one-time tasks, such as starting
 4 up the EVE manager and drawing detector components, are done in the `beginJob()`
 5 member function of our analyzer module. The full implementation of this function is
 6 shown in Listings 44.3 through 44.5.

7 **44.6.3.1 The Default GUI**

8 Referring to Listing 44.3, the very first thing done, on line 5, is to start up the EVE central
 9 application manager. More precisely, this line checks whether the pointer, `gEve`, is set. If
 10 not, it constructs an EVE manager object of type `TEveManager` and points `gEve` to it.
 11 Since no arguments are passed to the static `Create()` member function, the following
 12 defaults are used:

```
13 static TEveManager* Create(Bool_t map_window=kTRUE,Option_t* opt="FIV");
```

14 These values are passed to the constructor for `TEveManager`, which, among other things,
 15 starts up a `TEveBrowser` window. Having `map_window` set to its default, means that
 16 the EVE browser window will be made visible. After this, a list-tree and editor widget is
 17 created in the left vertical area of the EVE browser. This widget is actually a composite,
 18 consisting of the list-tree widget and context-sensitive menu widget described in Section
 19 44.5.2.1.

20 The number of "V"s in the options string tells the manager how many GL view-
 21 ers (one in our case) to spawn in the right main area. After replacing the "V" with
 22 an empty string, the remainder of the options string, "FI", is passed on to the
 23 `TEveBrowser::InitPlugins()` member function. The "F" in the options string
 24 instructs the function to create a `TGFileBrowser` object in the left vertical area. The
 25 "I" in the options string instructs the `TRootBrowser::InitPlugins()` member
 26 function of the EVE browser's base class to create a command line console in the right
 27 bottom area.

28 Since one "V" was counted in the options string earlier, a single GL viewer is created and
 29 embedded in the right main area. Two scenes—a *Global* scene and an *Events* scene—are
 30 also created and added to the viewer. The *Global* scene is meant to hold objects, like

1 those representing detector geometry, which remain "resident" or constant as one navigates
 2 through events. The *Events* scene is meant for objects, like detector hits and reconstructed
 3 tracks, that are unique for every event.

4 44.6.3.2 Adding the Global Elements

5 From the discussion above, we see that line 5 of Listing 44.3 actually does quite a bit
 6 of work for us by providing us with a browser with an initial set of widgets that in-
 7 clude object and file browsers, a 3D viewer, and a command line console. With all of
 8 these tasks taken care of by EVE, we can immediately focus our attention on visualizing
 9 the detector components. This is a straightforward task which, essentially, only requires
 10 providing EVE with a description of the detector geometry through the EVE manager's
 11 `AddGlobalElement()` member function.

12 44.6.3.3 Customizing the GUI

13 Having created the default GUI and drawn the detector components, we will now look at
 14 extending the event display by adding multiview orthographic projections and an event-
 15 navigation panel.

16 The two views we will create are:

- 17 1. an *XY* view with a projection plane that is perpendicular to the *z*-axis and the positive
 18 *z*-axis pointing out of the plane
- 19 2. an *RZ* view with a projection plane that is perpendicular to the *x*-axis and the positive
 20 *x*-axis pointing into the plane.

21 Referring to Listing 44.4, we begin by creating two new scenes for each of these views in
 22 lines 7 through 10. The two arguments passed to the `SpawnNewScene` function are the
 23 name and title of the scene, respectively. The name is used for list-tree item for the scene
 24 in the left vertical area of the EVE browser. The title is used for the tooltips that pop up
 25 when the mouse cursor is hovered over the list-tree item for the scene. In our example,
 26 we use an empty string for the title so no tooltips show up for the scene items. The two
 27 scenes having names starting with "`Det_`" are *Global* scenes used for the static detector
 28 geometry. The two scenes with names starting with "`Evt_`" are *Event* scenes used for the
 29 displaying the generated hits and tracks associated with an event.

Listing 44.3: Implementation of the `beginJob()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued on listing 44.4).

```

1 void tex::EventDisplay3D::beginJob() {
2
3     // Initialize global Eve application manager (return gEve)
4     // ~~~~~
5     TEveManager::Create();
6
7     // Import simplified extracted toy detector geometry for ortho views
8     // ~~~~~
9     TFile* geom = TFile::Open("toyDetector-extract.root");
10    TEveGeoShapeExtract* gse =
11        (TEveGeoShapeExtract*) geom->Get("ToyDetector");
12    fSimpleGeom = TEveGeoShape::ImportShapeExtract(gse, 0);
13    geom->Close();
14    delete geom;
15    gEve->AddGlobalElement(fSimpleGeom);
16    // ... Turn off rendering of simplified geometry so it does not appear
17    //     in main 3d window
18    gEve->GetGlobalScene()->FindChild("World_1")->SetRnrState(kFALSE);
19
20    // Import toy detector geometry from root file
21    // ~~~~~
22    gGeoManager = gEve->GetGeometry("toyDetector.root");
23
24    // ... Get top volume and find inner/outer tracker assemblies
25    //     underneath it
26    TGeoVolume* topvol = gGeoManager->GetTopVolume();
27    TEveGeoTopNode* intrackernode =
28        new TEveGeoTopNode(gGeoManager, topvol->FindNode("InnerTracker_16"));
29    TEveGeoTopNode* outrackernode =
30        new TEveGeoTopNode(gGeoManager, topvol->FindNode("OuterTracker_17"));
31
32    // ... Use helper to recursively make inner/outer tracker descendants
33    //     transparent & set custom colors
34    setRecursiveColorTransp(intrackernode->GetNode()->GetVolume(),
35        kOrange, 60);
36    setRecursiveColorTransp(outrackernode->GetNode()->GetVolume(),
37        kCyan-10, 60);
38
39    // ... Add static detector geometry to global scene
40    gEve->AddGlobalElement(intrackernode);
41    gEve->AddGlobalElement(outrackernode);
42
43        .
44        .
45        .

```


Listing 44.4: Implementation of the `beginJob()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued from listing 44.3 and continued on listing 44.5).

```

1           .
2           .
3           .
4
5         // Create detector and event scenes for ortho views
6         // ~~~~~
7         fDetXYScene = gEve->SpawnNewScene("Det_XY_Scene", "");
8         fDetRZScene = gEve->SpawnNewScene("Det_RZ_Scene", "");
9         fEvtXYScene = gEve->SpawnNewScene("Evt_XY_Scene", "");
10        fEvtRZScene = gEve->SpawnNewScene("Evt_RZ_Scene", "");
11
12        // Create XY/RZ proj mgrs, draw proj axes, & add them to scenes
13        // ~~~~~
14        fXYMgr = new TEveProjectionManager(TEveProjection::kPT_RPhi);
15        TEveProjectionAxes* axes_xy = new TEveProjectionAxes(fXYMgr);
16        fDetXYScene->AddElement(axes_xy);
17
18        fRZMgr = new TEveProjectionManager(TEveProjection::kPT_RhoZ);
19        TEveProjectionAxes* axes_rz = new TEveProjectionAxes(fRZMgr);
20        fDetRZScene->AddElement(axes_rz);
21
22        // Create adjacent ortho XY & RZ views in new tab & add det/evt scenes
23        // ~~~~~
24        TEveWindowSlot *slot = 0;
25        TEveWindowPack *pack = 0;
26
27        slot=TEveWindow::CreateWindowInTab(gEve->GetBrowser()->GetTabRight());
28        pack = slot->MakePack();
29        pack->SetElementName("Ortho_Views");
30        pack->SetHorizontal();
31        pack->SetShowTitleBar(kFALSE);
32
33        pack->NewSlot()->MakeCurrent();
34        fXYView = gEve->SpawnNewViewer("XY_View", "");
35        fXYView->GetGLViewer()->SetCurrentCamera(TGLViewer::kCameraOrthoXOY);
36        fXYView->AddScene(fDetXYScene);
37        fXYView->AddScene(fEvtXYScene);
38
39        pack->NewSlot()->MakeCurrent();
40        fRZView = gEve->SpawnNewViewer("RZ_View", "");
41        fRZView->GetGLViewer()->SetCurrentCamera(TGLViewer::kCameraOrthoXOY);
42        fRZView->AddScene(fDetRZScene);
43        fRZView->AddScene(fEvtRZScene);
44           .
45           .

```

1 Let us now look at lines 14 through 16. The first line creates a projection manager of type
 2 `TEveProjectionManager`. A projection manager takes care of applying the appro-
 3 priate projection to the elements in a scene. On this line, we pass an `enum` type to the
 4 constructor to specify the desired type of projection. In this case, `kPT_RPhi` specifies the
 5 projection for the *XY* view described above. On line 15, we create projected coordinate
 6 axes for our *XY* view by passing the pointer to the projection manager to the constructor
 7 of the `TEveProjectionAxes` object. We add the projected axes to the *XY* scene on
 8 line 16. These steps are repeated in lines 18-20 for the case of the *RZ* view.

9 On lines 24 through 43, we create viewers for the orthographic scenes and projection
 10 managers constructed above. The first two lines declare and initialize pointers used in the
 11 code. On the next line, we use a static helper function which creates an empty window
 12 slot in the tab widget in the right main area, and returns a pointer to it. This tab widget is
 13 specified by using the EVE browser's `GetTabRight()` member function. There are two
 14 other similar functions, namely `GetTabLeft()` and `GetTabBottom()`, for accessing
 15 the left vertical area and lower command line area, respectively. The new slot will be on a
 16 pane underneath the existing one, with a little tab at the top for raising it. Next, we make
 17 the empty window slot a *pack* container which is a vertical or horizontal stack of frames for
 18 holding widgets. Since the default orientation is vertical, we explicitly specify a horizontal
 19 one in which new frames will be inserted from left to right. Setting the element name to
 20 "Ortho Views" puts this string on the little tab at the top of the new pane. The last line
 21 suppresses the title bar for the top-level frame holding our stack of frames..

22 On line 33, we insert the first frame in our horizontal stack of frames and select it as the
 23 current window. The next line then gets the current window as an empty slot, creates a GL
 24 viewer, and embeds it in the window slot. The first argument passed to `SpawnNewViewer`
 25 specifies the name of the viewer which will show up on the title bar of the viewer window
 26 and on the list-tree item associated with it. The second item specifies the text displayed
 27 in the tooltips popups when the mouse cursor is hovered over the list-tree item associated
 28 with the viewer. After creating the viewer, we specify the properties of the camera associ-
 29 ated with it. For both the *XY* and *RZ* orthographic views, we should specify the camera,
 30 as shown in the code, to `kCameraOrthoXOY`. The projection managers defined earlier
 31 will apply the correct projections on the elements. Finally, we add the *Global* and *Event*
 32 *XY* scenes created earlier to our viewer. The steps we just described in this paragraph
 33 are repeated on lines 39 through 43 for the case of the *RZ* view. This view shows up in
 34 a second window to the right of that for the *XY* view in the tabbed pane labeled *Ortho*

1 *Views.*

2 Moving on now to Listing 44.5, on lines 6 and 7, we use the projection manager ob-
3 ject's `ImportElements()` member function to import the detector geometry into the
4 orthographic *XY* and *RZ* scenes we set up above for displaying static components. Each
5 projection manager applies its own projection to all the imported objects. The first argu-
6 ment to the function is a pointer to the objects to import and the second argument is a
7 pointer to the scene to import them into. On line 13, we bring the first pane (index 0, with
8 the default 3D GL viewer) of the tab widget on the right main area to the front. Line 17
9 calls the `makeNavPanel()` member function of our `EventDisplay3D` class to cre-
10 ate add an event-navigation pane to the tab widget in the left vertical area. We will discuss
11 this function in more detail in a Section 44.6.3.4 below. On line 23, we create a new EVE
12 event manager object of type `TEveEventManager` and pass it as an argument to the
13 `AddEvent()` member function of the EVE manager. The EVE event manager contains
14 a list of the objects in our event. The `AddEvent` function adds these objects to the *Event*
15 scenes we set up earlier and to a top-level item in the list-tree widget in the left vertical
16 area. The name of this list-tree item is the first argument ("Event") passed to the con-
17 structor of the EVE event manager. The second argument ("Toy Detector Event")
18 passed to the constructor is the text in the tooltips that pops up when the mouse cursor is
19 hovered over the list-tree item.

20 The last few lines in `beginJob`, starting at line 26, set up the appearance of the coordinate
21 axes and reference point, and the initial orientation of the camera for the main 3D view in
22 the right main area. The `GetDefaultViewer()` member function of the EVE manager
23 object returns a pointer to the default viewer which is the first one listed under the *View-*
24 *ers* item in the list-tree widget. With this pointer, we then use the `SetGuideState()`
25 member function of the viewer object to configure the axes and reference point. The first
26 argument to the this function, `kAxesEdge`, specifies the three axes to be drawn along
27 the edges of the bounding-box containing the drawn components. Two other options,
28 `kAxesNone` and `kAxesOrigin`, specify no axes and axes that go through the origin,
29 respectively. The second argument to the function turns depth-testing on for hidden-line
30 removal when drawing the axes. The third argument turns drawing of the reference point
31 off. The fourth argument specifies the position of the reference point if it were drawn. The
32 last two lines (28 and 29) set the initial orientation of the camera for our default 3D viewer

Listing 44.5: Implementation of the `beginJob()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued from listing 44.4).

```

1
2
3
4
5 // ... Import simplified geom into ortho views and apply projections
6 fXYMgr->ImportElements(fSimpleGeom, fDetXYScene);
7 fRZMgr->ImportElements(fSimpleGeom, fDetRZScene);
8
9 // ... Turn rendering of simplified geom ON in ortho views
10 fDetXYScene->FindChild("World_1_[P]")->SetRnrState(kTRUE);
11 fDetRZScene->FindChild("World_1_[P]")->SetRnrState(kTRUE);
12
13 gEve->GetBrowser()->GetTabRight()->SetTab(0);
14
15 // Create navigation panel
16 // ~~~~~
17 makeNavPanel();
18
19 // Add new Eve event into "Event" scene and make it the current event
20 // ~~~~~
21 // (Subsequent elements added using "AddElements" will be added to
22 // this event)
23 gEve->AddEvent(new TEveEventManager("Event", "Toy_Detector_Event"));
24
25 // ... Set up initial camera orientation in main 3D view
26 TGLViewer *glv = gEve->GetDefaultGLViewer();
27 glv->SetGuideState(TGLUtil::kAxesEdge, kTRUE, kFALSE, 0);
28 glv->CurrentCamera().RotateRad(camRotateCenterH_, camRotateCenterV_);
29 glv->CurrentCamera().Dolly(camDollyDelta_, kFALSE, kFALSE);
30
31 }

```

1 using the parameters described in Section 44.6.2.

2 44.6.3.4 Adding the Navigation Pane

3 The `EventDisplay3d::makeNavPanel()` member function was briefly mentioned
 4 in Section 44.6.3.3. The code for this function is shown in Listing 44.6, where ROOT's
 5 GUI classes are used to construct the event-navigation pane shown in Figure 44.7. We
 6 begin in Line 7 by getting a pointer to the EVE browser. We then use the browser object's
 7 `StartEmbedding()` member function to specify that a new external frame should be
 8 embedded as a new pane of the tab widget, in the left vertical area (identified by `kLeft`)
 9 of the browser. Behind the scenes, the current root window is being set to a newly created
 10 pane in the tab widget. This makes the top-level frame, created on Line 10, a child of
 11 the new pane and embeds it in the pane. Passing `kDeepCleanup` as an argument to
 12 `SetCleanup` in Line 12 turns on the automatic hierarchical cleanup of the composite
 13 main frame and its children in the destructor.

14 On the next two lines, we create a horizontally aligned and a vertically aligned compos-
 15 ite sub-frame having the top-level frame (`frmMain`) as parent. On Line 17, we specify
 16 `$ROOTSYS/icons` as the location of the `.gif` files containing pictures of the forward
 17 and backward arrows to use on our event-navigation buttons. The back-arrow picture but-
 18 ton is created on Line 22 with the horizontally aligned frame (`navframe`) as its par-
 19 ent and the back-arrow to be used as the picture. Next we use `AddFrame` to add this
 20 button into the horizontally aligned frame's list of children widgets. Then we establish
 21 a connection between the `Clicked()` signal, emitted when the button is pressed, and
 22 the `PrevEvent()` member function of the `EvtDisplayUtils` class introduced in
 23 Section 44.6.2. The `PrevEvent()` function acts as a receiver slot which is executed
 24 whenever it detects the `Clicked()` signal. The names of the emitted signal and the slot
 25 receiving it are passed as the first and fourth arguments to `Connect`, respectively. The
 26 second argument specifies the name of the class the receiver slot function belongs to. The
 27 third argument is a pointer to an object belonging to that class. These steps for creating the
 28 back-arrow button are repeated for the forward-arrow button on Lines 29 through 32.

29 The buttons described above allow us to step through events sequentially, one at a time,

Listing 44.6: Implementation of the `makeNavPanel()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued on listing 44.7).

```

1 void tex::EventDisplay3D::makeNavPanel()
2 {
3     // Create control panel for event navigation
4     // ~~~~~
5
6     // ... Insert nav frame as new tab in left vertical area
7     TEveBrowser* browser = gEve->GetBrowser();
8     browser->StartEmbedding(TRootBrowser::kLeft);
9
10    TGMainFrame* frmMain = new TGMainFrame(gClient->GetRoot(), 1000, 600);
11    frmMain->SetWindowName("EVT_NAV");
12    frmMain->SetCleanup(kDeepCleanup);
13
14    TGHorizontalFrame* navFrame = new TGHorizontalFrame(frmMain);
15    TGVerticalFrame* evtidFrame = new TGVerticalFrame(frmMain);
16    {
17        TString icondir(TString::Format("%s/icons/",
18                                     gSystem->Getenv("ROOTSYS")));
19        TGPictureButton* b = 0;
20
21        // ... Create back button & connect to "PrevEvent" slot
22        b = new TGPictureButton(navFrame, gClient->GetPicture(icondir +
23                                                         "GoBack.gif"));
24        navFrame->AddFrame(b);
25        b->Connect("Clicked()", "tex::EvtDisplayUtils", visutil_.get(),
26                "PrevEvent()");
27
28        // ... Create forward button & connect to "NextEvent" slot
29        b = new TGPictureButton(navFrame, gClient->GetPicture(icondir +
30                                                         "GoForward.gif"));
31        navFrame->AddFrame(b);
32        b->Connect("Clicked()", "tex::EvtDisplayUtils", visutil_.get(),
33                "NextEvent()");
34
35        // ... Create run num text entry widget and connect
36        //     to "GotoEvent" slot
37        TGHorizontalFrame* runoFrame = new TGHorizontalFrame(evtidFrame);
38        fTlRun = new TGLabel(runoFrame, "Run_Number");
39        fTlRun->SetTextJustify(kTextLeft);
40        fTlRun->SetMargins(5, 5, 5, 0);
41        runoFrame->AddFrame(fTlRun);
42        .
43        .
44        .

```

1 in the forward or backward direction. The next few lines of the code we will discuss
2 create text entry widgets that allow us to navigate directly to an event by entering its run
3 and event numbers. We start by creating a horizontally aligned frame for the run number
4 text entry widget on line 37 of Listing 44.6. We wish to label this widget with the text
5 *Run Number* on its left. So the next line creates a label widget, passing a pointer to the
6 horizontal frame as its first argument to specify the parent widget, and a string for the label
7 text as its second argument. We specify that the text in the label should be left-justified.
8 The margin width around the text is set to 5 pixels on the left, top, and right, and to 0 pixels
9 on the bottom. This widget is added as the first entry in the horizontally aligned frame's
10 list of children.

11 Moving on to Listing 44.7, the text entry widget is created on Line 5. The first argument to
12 the constructor is a pointer to the horizontally aligned frame serving as the parent widget.
13 We create a 5-character text buffer to hold the string for the value of the run number and
14 point the `fTbRun` data member of the `EvtDisplayUtils` class to it. This is passed
15 as the second argument of the constructor. The buffer object will be adopted by the text
16 entry widget which will be responsible for its deletion. On Line 7, we set the initial text
17 in the buffer by inserting a "1" at the 0'th position. Next, we establish a connection be-
18 tween the `ReturnPressed()` signal, emitted when the return key is hit to signal com-
19 pletion of text entry, and the `EvtDisplayUtils::GotoEvent()` member function.
20 The text entry widget is added into the horizontally aligned frame's list of children on
21 Line 25. In addition, a `TGLLayoutHints` object is also passed as a second argument to
22 `AddFrame`. This tells the layout manager of the parent frame how to position the child
23 widget within the frame. Layout managers (`TGLLayoutManager`) are associated with
24 composite frames like the horizontally aligned frame containing our text entry widget.
25 The `kLHintsExpandX` layout hint we use tells the layout manager to expand the child
26 text entry widget horizontally up to the free space available to the child.

27 The steps described above for the run number text entry widget are repeated on
28 Lines 14 through 25 for the event number text entry widget. On Lines 28 and 29, we
29 stack the run and event number widgets up on top of each other by making them chil-
30 dren of the vertically aligned frame created earlier (`evtidFrame`). This frame, to-
31 gether with the one containing the forward and backward buttons, are added to the
32 top-level frame's (`frmMain`) list of children on Lines 32 through 35. A horizontal
33 line (`TGHHorizontal3DLine`) is also added between the two widgets to separate
34 them.

Listing 44.7: Implementation of the `makeNavPanel()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued from listing 44.6).

```

1
2
3
4
5   fTeRun = new TGTextEntry(runoFrame,
6       visutil_>fTbRun = new TGTextBuffer(5));
7   visutil_>fTbRun->AddText(0, "1");
8   fTeRun->Connect("ReturnPressed()", "tex::EvtDisplayUtils",
9       visutil_.get(), "GotoEvent()");
10  runoFrame->AddFrame(fTeRun, new TGLayoutHints(kLHintsExpandX));
11
12  // ... Create evt num text entry widget and connect
13  //     to "GotoEvent" slot
14  TGHorizontalFrame* evnoFrame = new TGHorizontalFrame(evtidFrame);
15  fTlEvt = new TGLLabel(evnoFrame, "Evt_Number");
16  fTlEvt->SetTextJustify(kTextLeft);
17  fTlEvt->SetMargins(5, 5, 5, 0);
18  evnoFrame->AddFrame(fTlEvt);
19
20  fTeEvt = new TGTextEntry(evnoFrame,
21       visutil_>fTbEvt = new TGTextBuffer(5));
22  visutil_>fTbEvt->AddText(0, "1");
23  fTeEvt->Connect("ReturnPressed()", "tex::EvtDisplayUtils",
24       visutil_.get(), "GotoEvent()");
25  evnoFrame->AddFrame(fTeEvt, new TGLayoutHints(kLHintsExpandX));
26
27  // ... Add horiz run & event number subframes to vert evtidFrame
28  evtidFrame->AddFrame(runoFrame, new TGLayoutHints(kLHintsExpandX));
29  evtidFrame->AddFrame(evnoFrame, new TGLayoutHints(kLHintsExpandX));
30
31  // ... Add navFrame and evtidFrame to MainFrame
32  frmMain->AddFrame(navFrame);
33  TGHorizontal3DLine *separator = new TGHorizontal3DLine(frmMain);
34  frmMain->AddFrame(separator, new TGLayoutHints(kLHintsExpandX));
35  frmMain->AddFrame(evtidFrame);
36
37  frmMain->MapSubwindows();
38  frmMain->Resize();
39  frmMain->MapWindow();
40
41  browser->StopEmbedding();
42  browser->SetTabTitle("Event_Nav", 0);
43  }
44  }

```


1 The call to `MapSubWindows()` on Line 37 goes through all sub-frames contained within
2 the top-level frame and turns their visibility on. The next line resizes the top-level window
3 to a default value and applies the layout specifications for the contained widgets. After all
4 of this is completed, the top-level frame is made visible, causing its children to also appear
5 on the screen. On Line 41, the root window is reset to the default, ending the embedding of
6 external or top-level frames in the navigation pane temporarily set to be the root window
7 earlier in the code. On the very last line of the `makeNav()` member function, we set the
8 label that appears on the tab at the top of the event-navigation panel to *Event Nav*. The
9 second argument of 0 to `SetTabTitle` specifies the tab widget in the left vertical area
10 of the browser. The third argument is used to select a particular tab or pane in the widget.
11 Since none is provided here, the current one is selected.

12 **44.6.4 Drawing the Generated Hits and Tracks in the `analyze()` Member** 13 **Function**

14 The previous sections dealt with setting up the static components of the event display.
15 This section will deal with the components associated with an event and which change
16 from event to event. These components, which include run and event numbers, and gener-
17 ated hits and tracks, are created or updated in the `analyze()` member function of our
18 `analyzer` module. The complete code for this function is shown in Listings 44.8 through
19 44.12. Referring to Listing 44.8, we start by updating the text entry widgets in our event-
20 navigation pane to display the correct values for the current event. We first create an output
21 string stream object and use the insertion operator to send the run number to this object.
22 After clearing the text buffer for the text entry widget, it is updated with the current run
23 number in the `ostringstream` object by passing a pointer to the C-string representing
24 it to the `AddText` method of the buffer. The host graphics system is then informed that
25 the text entry widget for the run number needs to be redrawn to display the contents of the
26 updated buffer. After emptying the string in the `ostringstream` object, the steps above
27 are repeated for the event-number text entry widget.

28 Before we begin drawing the generated hits and tracks, we first delete any non-global
29 elements associated with the previous event, such as the generated tracks and hits, and
30 interactively added annotations. Line 19 gets the list of all the GL viewers associated with
31 the current EVE manager and deletes all the overlay elements in each viewer that are
32 annotations. The next line deletes all the non-global elements that were previously added

display3D-analyze1

Listing 44.8: Implementation of the `analyze()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued on listing 44.8).

```

1
2 void tex::EventDisplay3D::analyze(const art::Event& event ){
3
4     // ... Update the run and event numbers in the TGTEntry widgets in
5     //     the Navigation panel
6     std::ostringstream sstr;
7     sstr << event.id().run();
8     visutil_>fTbRun->Clear();
9     visutil_>fTbRun->AddText(0,sstr.str().c_str());
10    gClient->NeedRedraw(fTeRun);
11
12    sstr.str("");
13    sstr << event.id().event();
14    visutil_>fTbEvt->Clear();
15    visutil_>fTbEvt->AddText(0,sstr.str().c_str());
16    gClient->NeedRedraw(fTeEvt);
17
18    // ... Delete visualization structures associated with previous event
19    gEve->GetViewers()->DeleteAnnotations();
20    gEve->GetCurrentEvent()->DestroyElements();
21
22    .
23    .

```

lin:anal-del

lin:anal-del

1 through the EVE manager's `AddElement` method. As we will see later in the code, these
2 elements are actually not destroyed, but simply removed from the list-trees displayed in
3 the left vertical area of the browser.

4 With our viewers and list-trees cleared of old structures, we can now draw the generated
5 hits for the current event. In Listing 44.9, after setting up a handle pointing to the *art*
6 data product for generated hits, we check to see if the `fHitsList` pointer to the list of
7 objects representing hits is set. If not, an EVE element list named "Hits" is constructed
8 and `fHitsList` is set to point to it. The *deny-destroy* counter for this element is also
9 incremented to 1, protecting it from deletion when the EVE manager is asked to destroy
10 the non-global elements from the previous event in Line 20 of Listing 44.8. If, on the other
11 hand, `fHitsList` is set, all its children are destroyed.

12 On Lines 19-21, we create separate containers for hits originating from K^+ , K^- , or *Bkg*
13 particles. On Lines 23-34, we loop through all the hits in the collection, drawing them as
14 spheres of a particular color, and assigning them to the appropriate container, based on the
15 particle identity of the generated track producing the hit.

16 The code to draw the hits is implemented in the `drawHit()` helper function shown in
17 Listing 44.10. In this function, we begin by creating strings for the text displayed in the
18 tooltips associated with the hits. These tooltips show the hit index and the detector layer
19 the hit belongs to. On Line 10, we create an `TEvePointSet` object to represent our hit.
20 Even though such an object can hold a collection of hits, we use it to represent a single
21 hit so that each hit shows up as a separate item in our list-tree. The `Form()` function
22 used here returns a pointer to a C-string constructed in `printf`-style fashion from its
23 arguments. This string determines the name of the list-tree item for the hit. In the next line,
24 the `Form()` function is used again to construct the text that is displayed in the tooltips
25 for the hit. In the next three lines, we set the coordinates, color, and size of the spherical
26 marker representing the hit. The hit is then added into list for the type of particle that
27 generated it. Returning now to Listing 44.9 after looping through all the hits, we add the
28 lists for each particle type into the main list for all hits and then add this main list into the
29 EVE *Event* scene.

30 Let us now look at how generated tracks are drawn in Listing 44.11. As for the generated
31 hits, we begin by setting up a handle to point to the *art* data product for the generated
32 tracks. In a similar fashion, we also set up the `fTracksList` pointer to the list of objects
33 representing generated tracks. The difference here is that, instead of `TEveElementList`,

Listing 44.9: Implementation of the `analyze()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued from listing 44.8 and continued on listing 44.11).

```

1
2
3
4
5 // Draw the detector hits
6 // ~~~~~
7 if (drawHits_) {
8     std::vector<art::Handle<IntersectionCollection>> hitsHandles;
9     event.getManyByType(hitsHandles);
10
11     if (fHitsList == 0) {
12         fHitsList = new TEveElementList("Hits");
13         fHitsList->IncDenyDestroy(); // protect element against destruction
14     }
15     else {
16         fHitsList->DestroyElements(); // destroy children of the element
17     }
18
19     TEveElementList* KpHitsList = new TEveElementList("K+_Hits");
20     TEveElementList* KmHitsList = new TEveElementList("K-_Hits");
21     TEveElementList* BkgHitsList = new TEveElementList("Bkg_Hits");
22
23     int ikp=0, ikm=0, ibkg=0;
24     for ( auto const& handle: hitsHandles ){
25         for ( auto const& hit: *handle ){
26             if ( hit.genTrack()->pdgId() == PDGCode::K_plus ){
27                 drawHit("K+", kGreen, hitMarkerSize_, ikp++, hit, KpHitsList);
28             } else if ( hit.genTrack()->pdgId() == PDGCode::K_minus ){
29                 drawHit("K-", kYellow, hitMarkerSize_, ikm++, hit, KmHitsList);
30             } else{
31                 drawHit("Bkg", kViolet+1, hitMarkerSize_, ibkg++, hit, BkgHitsList);
32             }
33         }
34     }
35     fHitsList->AddElement(KpHitsList);
36     fHitsList->AddElement(KmHitsList);
37     fHitsList->AddElement(BkgHitsList);
38     gEve->AddElement(fHitsList);
39 }
40
41
42

```

3D-drawhit

Listing 44.10: The `drawHit()` helper function from `EventDisplay3D_module.cc`.

```

1 void drawHit(const std::string &pstr, Int_t mColor, Int_t mSize,
2             Int_t n, const tex::Intersection &hit,
3             TEveElementList *list)
4 {
5     std::string hstr="_hit_%d";
6     std::string dstr="_hit#_%d\nLayer:_%d";
7     std::string strlst=pstr+hstr;
8     std::string strlab=pstr+dstr;
9
10    TEvePointSet* h = new TEvePointSet(Form(strlst.c_str(),n));
11    h->SetTitle(Form(strlab.c_str(),n,hit.shell()));
12    h->SetNextPoint(hit.position().x()*0.1,
13                  hit.position().y()*0.1,
14                  hit.position().z()*0.1);
15    h->SetMarkerColor(mColor);
16    h->SetMarkerSize(mSize);
17    list->AddElement(h);
18 }

```

-drawhit

1 we use `TEveTrackList`, which is a specialized EVE element list for tracks and includes
 2 the former class as one of its three base classes. This type of list offers convenient features
 3 such as setting common track attributes and allowing selection based on track parameters.
 4 Such a feature is used on line 12, where we set the line width attribute for tracks in this
 5 list.

6 On Line 19, we get the track propagator object associated with the track list. This is used to
 7 calculate the path of the track for a given magnetic field. In the following lines, we specify
 8 the magnetic field, the maximum extents of the track in the radial and z directions, and the
 9 maximum step size to use for track ropagation. The negative sign for the argument passed
 10 to the `SetMagField()` method is necessary because of the inverted convention used in
 11 EVE for the magnetic field direction.

12 Once the propagator is set up, we loop over the generated tracks in our collection, skip-
 13 ping those that have decayed. We first construct a `TParticle` type object, setting its
 14 momentum, production vertex, and particle identity from the information in the *art* data
 15 product. Then, we create a `TEveTrack` type object with the form of the constructor for a
 16 `TParticle` and using the track's index in the *art* collection for the EVE simulation la-
 17 bel. After the constructor, we set the EVE reconstruction label to 0. The `SetStdTitle()`

1 method sets the text to display in the tooltips to a default set. In the next line, we apply the
2 line width attribute set in Line 12 to the track.

3 Continuing with the track loop on line 5 in Listing 44.12, we assign a color for each track
4 based on its particle identity. At the end of each iteration through the loop, we add the
5 track to the list of tracks pointed to by `fTrackList`. When we exit the loop, we use
6 the `MakeTracks()` method to generate the trajectory for each track in the list using the
7 propagator specified earlier. The list is then added into the EVE *Event* scene.

8 At the very end of the `analyze()` member function, we retrieve the elements added
9 above into the EVE *Event* scene. After clearing the 2D orthographic views of structures
10 associated with the previous event, the retrieved elements are imported into these views
11 with the appropriate projections applied.

Listing 44.11: Implementation of the `analyze()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued from listing 44.9 and continued on listing 44.12).

```

1           .
2           .
3           .
4
5 // Draw the generated tracks as helices in a uniform axial field
6 // ~~~~~
7 if (drawGenTracks_) {
8     auto gens = event.getValidHandle<GenParticleCollection>(gensTag_);
9
10    if (fTrackList == 0) {
11        fTrackList = new TEveTrackList("Tracks");
12        fTrackList->SetLineWidth(4);
13        fTrackList->IncDenyDestroy(); // protect against destruction
14    }
15    else {
16        fTrackList->DestroyElements(); // destroy children of the element
17    }
18
19    TEveTrackPropagator* trkProp = fTrackList->GetPropagator();
20    trkProp->SetMagField(-geom_->bz()*1000.);
21    trkProp->SetMaxR(trkMaxR_);
22    trkProp->SetMaxZ(trkMaxZ_);
23    trkProp->SetMaxStep(trkMaxStepSize_);
24
25    int mcindex=-1;
26    for ( auto const& gen: *gens){
27        mcindex++;
28        // ... Skip tracks decayed in the generator.
29        if ( gen.hasChildren() ) continue;
30        TParticle mcpart;
31        mcpart.SetMomentum(gen.momentum().px(), gen.momentum().py(),
32                            gen.momentum().pz(), gen.momentum().e());
33        mcpart.SetProductionVertex(gen.position().x()*0.1,
34                                   gen.position().y()*0.1, gen.position().z()*0.1, 0.);
35        mcpart.SetPdgCode(gen.pdgId());
36        TEveTrack* track = new TEveTrack(&mcpart, mcindex, trkProp);
37        track->SetIndex(0);
38        track->SetTitle();
39        track->SetAttLineAttMarker(fTrackList);
40
41        .
42        .

```

Listing 44.12: Implementation of the `analyze()` member function of the `EventDisplay3D` class from `EventDisplay3D_module.cc` (continued from listing 44.11).

```

1
2           •
3           •
4
5   if ( gen.pdgId() == PDGCode::K_plus ){
6       track->SetMainColor(kGreen);
7   } else if ( gen.pdgId() == PDGCode::K_minus ){
8       track->SetMainColor(kYellow);
9   } else {
10      track->SetMainColor(kViolet+1);
11  }
12  fTrackList->AddElement(track);
13  }
14  fTrackList->MakeTracks();
15  gEve->AddElement(fTrackList);
16  }
17
18  // Import event into ortho views and apply projections
19  // ~~~~~
20  TEveElement* currentv = gEve->GetCurrentEvent();
21
22  fEvtXYScene->DestroyElements();
23  fXYMgr->ImportElements(currentv, fEvtXYScene);
24
25  fEvtRZScene->DestroyElements();
26  fRZMgr->ImportElements(currentv, fEvtRZScene);
27
28 } // end tex::EventDisplay3D::analyze

```


¹ 45 Live Histogram Update

- ² 1. Run MC events and update live histograms.

DRAFT

1 46 Checklist

2 Make sure that all of the following have a home above:

3 1. Add suggested exercises (and solutions) to each chapter; some are already there but
4 it is incomplete.

5 Some exercises should be examples of code that does not compile or which produces
6 incorrect results; the task

7 is to fix the problem. Others should be of the form: write code to do something. In
8 both cases we need to

9 provide verified solutions.

10 2. What is the best way to illustrate `art::Assns<A,B,D>`? One option is MCTruth
11 matching of Reco tracks; here D

12 is the chisquared of the match or something like that. Will I

13 get enough multiple matches to make this example work? Maybe reduce the resolu-
14 tion of the tracker until I do get

15 a few dozen cases of multiple matches in 1000 events?e A second option is a barrel
16 TOF system in which the D object

17 is something like the inferred time and particle ID probabilities. Use the particle ID
18 probabilities when

19 making the $K+K^-$ mass plot.

20 3. Changes to the `toyExperiment` product

- 1 ○ Follow the guidelines so that only member data, c'tor and d'tor are exposed
2 to genreflex
- 3 ○ Add version numbers to the classes following Chris Green's model. A path for
4 future expansion
5 is to make a change to a class that requires a handwritten streamer?
- 6 ○ Add the extra genparticles and modify downstream code; maybe exponentially
7 falling
- 8 ○ Find a way to naturally introduce map_vector
- 9 ○ Find a way to introduce EnumToString - maybe in GenParticle? Or is that too
10 much too soon.
- 11 4. admonition to always start a new login session for new projects - how many different
12 places?
- 13 5. run and subrun data products
- 14 6. scrub .fcl discussion for inappropriate use of "keyword".
- 15 7. Event::getManyByType and Selector functions.
- 16 8. PtrVector<T> and std::vector< art :: Ptr <T>>;
- 17 9. map_vector
- 18 10. Making subdirectories in the TFileService
- 19 11. Polymorphic data and Views. Maybe the right model for a view is to have HOTs and
20 hits. If a Hit subclasses off of hit, then it is a good model
21 to excercsie View.
- 22 12. ParameterSet part 2:
23 (a) include, prolog, @local, redefining parameters
24 (b) ART_DEBUG_CONFIG=1 mu2e -c Ex03/ex03.fcl
25 (c) FHICL_FILE_PATH
- 26 13. Grid/SAM

- 1 (a) Generate and Reco MC
- 2 (b) Probe job
- 3 (c) Stage in and stage out files.
- 4 14. Check to see which of the CLHEP things I really need in the dictionary for Dat-
- 5 aProducts.
- 6 15. Modify toyExperiment to add more tracks; update the description and figures in the
- 7 document.
- 8 16. Timebomb module to illustrate exceptions.
- 9 17. Memory leak module to illustrate the crude memory leak checker.
- 10 18. Show how to send events that give an exception to a separate output file.
- 11 19. Parse a separate fcl file.
- 12 20. Message logger system
- 13 21. Services calling services
- 14 22. How to ensure that service B called before service A?
- 15 23. Drop on output; drop on input
- 16 24. write your own TNtuple and TTree; for TTree need to discuss provision of dictio-
- 17 naries.
- 18 25. enum matched to string class template.
- 19 26. validity checking of psets
- 20 27. examples to illustrate how to use gdb, totalview etc
- 21 28. Example of how to use the timing service; do we need to slow modules for this to
- 22 work well?
- 23 Maybe a producer module that finds all prime numbers less than N - choose N from
- 24 RandGauss?
- 25 29. VERBOSE flag on buildtool - probably belongs in the Build and Run the first mod-
- 26 ule section?

- 1 30. TFileDirectory and making directories within the TFileService.
- 2 31. std::vector hygiene: understand that it has a capacity and that it
- 3 auto-reallocates as needed. Use reserve. Understand what it means that
- 4 push_back (emplace_back) invalidates all iterators. Know that insertion
- 5 on a list or deque does not invalidate iterators.
- 6 32. scrub the text to remove instances of “variable” - prefer to use object in most cases.
- 7 There may be
- 8 a few cases in which variable is what we want.
- 9 33. scrub the text to remove “method” in favor of “member function” - may be complete
- 10 34. scrub the text for definitions of site. It needs to be expanded to include “your laptop”.
- 11 A particular
- 12 experiment+institution may have several sets of computing hardware - their own
- 13 cluster, their desktops, their
- 14 laptops, the cluster belonging to their friend from which they borrow cycles...
- 15 35. an exercise to illustrate cet::map_vector
- 16 36. Does it make sense to teach some git stuff as follows:
- 17 (a) clone the code from our repository
- 18 (b) make a new repository under your name on redmine (or elsewhere)
- 19 (c) push all of the code to their repository
- 20 (d) commit changes as needed
- 21 I don't see how to deal with updates.

1 47 Classes in the toyExperiment

2 This is meant to be a manifest of the toyExperiment data product. Fill in any details that
3 are not well

4 covered in the above.

5 In DataProducts:

6 **EnumToString** Def

7 **EnumToStringSparse**

8 **PDGCode**

9 In MCDataProducts

10 **GenParticle**

11 **GenParticleCollection**

12 **Intersection**

13 **IntersectionCollection**

14 **MCRunSummary**

15 **TrkHitMatch**

16 In RecoDataProducts

17 **DetectorStatus**

18 **DetectorStatusRecord**

- 1 **FittedHelixData**
- 2 **FittedHelixDataCollection**
- 3 **Helix**
- 4 **RecoTrk**
- 5 **RecoTrkCollection**
- 6 **TrkHit**
- 7 **TrkHitCollection**

- 8 Producer modules:
- 9 **DetectorSimulation_module**
- 10 **EventGenerator_module**
- 11 **HitMaker_module**
- 12 **FindAndFitHelix_module**

- 13 Analyzer modules:
- 14 **EventDisplay_module**
- 15 **HelloWorld_module**
- 16 **InspectFittedHelix_module**
- 17 **InspectGenParticles_module**
- 18 **InspectIntersections_module**
- 19 **InspectTrkHits_module**
- 20 **MassPlot_module**

- 21 Analyzer modules:
- 22 **Conditions/Conditions_service**
- 23 **Geometry/Geometry_service.**

- 1 **Conditions/Conditions**
- 2 **Geometry/Geometry**
- 3 Other source files:
- 4 **Analyzers/PlotOrientation**
- 5 **Conditions/ParticleInfo**
- 6 **Conditions/PDT**
- 7 **Conditions/ShellConditions**
- 8 **Geometry/IntersectionFinder**
- 9 **Geometry/Shell**
- 10 **Geometry/Tracker**
- 11 **Geometry/TrackerComponent**
- 12 **Reconstruction/FittedHelix**
- 13 **Utilities/Decay2Body**
- 14 **Utilities/dphi**
- 15 **Utilities/eventIDToString**
- 16 **Utilities/ParameterSetFromFile**
- 17 **Utilities/phi_norm**
- 18 **Utilities/phi_small**
- 19 **Utilities/polar3Vector**
- 20 **Utilities/RandomUnitSphere**
- 21 **Utilities/safeSqrt**
- 22 **Utilities/sqrtOrThrow**
- 23 **Utilities/TwoBodyKinematics**

- 1 Directories:
- 2 **Analyzers**
- 3 **Conditions**
- 4 **DataProducts**
- 5 **Geometry**
- 6 **HelloWorldScripts**
- 7 **MCDataProducts**
- 8 **RecoDataProducts**
- 9 **Reconstruction**
- 10 **Simulations**
- 11 **Utilities**
- 12 **bin**
- 13 **databaseFiles**
- 14 **fcl**
- 15 **inputFiles**

DRAFT

1 48 Troubleshooting

2 48.1 Updating Workbook Code

3 If the remote machine that you log onto to run the Workbook exercises runs into problems
4 during the setup procedure, it's possible that the admin for that machine has not installed
5 the most recent versions of the Workbook code, or some dependent code. Contact the
6 administrator.

7 48.2 XWindows (xterm and Other XWindows Products)

8 9 48.2.1 Mac OSX 10.9

10 The XWindows products, xterm, xclock and so on, likely reside in the directory
11 /opt/x11/bin/. You will need to add this to your PATH. When your machine is con-
12 nected to a second monitor, the XWindows products may not position properly on the
13 screen. You may need to contact a Macintosh support person to configure the X11 setup
14 properly so that it works with the multiple screen configuration.



5 At Fermilab, open a service desk ticket at [https://fermi.service-now.com/](https://fermi.service-now.com/navpage.doService-Now)
6 navpage.doService-Now.

17 48.3 Trouble Building

18 If the buildtool doesn't seem to find your module, check that it can see it:

1 `ls lib`

2 Make sure it's not some other module that's causing the build failure.

3 **48.4 *art* Won't Run**

4 Make sure you're in your build window

5 Make sure the path to the FHiCL file is correct, e.g.,

6 `art -c fcl/<dir-for-exercise>/<filename>.fcl`

DRAFT

1

Part III

2

User's Guide

1 49 git

ch:ug-git

2 The source code for the exercises in the *art* workbook is stored in a source code manage-
3 ment system called *git* and maintained in a repository managed by Fermilab. Think of *git*
4 as an enhanced *svn* or (a VERY enhanced) *cvs* system. The repository is located at . You
5 will be shown how to access it with *git*.

6 If you want some background on *git*, we suggest the Git Reference.

7 You will need to know how to install *git*, download the workbook exercise files initially to
8 your system and how to download updates. You will not be checking in any code.

9 **FIXME:** *Revisit once Rob extracts the git info from /ds50/app/ds50/ds50.sh and/or /ds50/app/ds50/setup_workbook*
10 *Will student run 'setup git' or the 'git clone' command?*

11 To install *git* on a Mac:

```
12 $ http://git-scm.com/download/mac
```

13 This will automatically download a disk image. Open the disk image and click on the .pkg
14 file.

15 In your home directory, edit the file `.bash_profile` and add the line:

```
16 $ export PATH=/usr/local/git/bin/:${PATH}
```

```
17 $ git clone ssh://p-art-workbook@cdcvs.fnal.gov/cvs/projects/art-workbook
```

18 and how to download updates as the developers make them:

```
19 $ git pull
```

1 49.1 Aside: More Details about git

date:git:details

2 **FIXME:** *I propose putting this in the users guide and copying in here only a minimum. To*
 3 *do. AH*

4 To bring your working copy of the workbook code up to date, you need to use `git`. Before
 5 describing the required `git` commands, we need to explain a little more about how `git`
 6 works and how the art-workbook team have chosen to use it. If you are familiar with `git`
 7 you can skip this section.

8 49.1.1 Central Repository, Local Repository and Working Directory

9 At any given time, there are three copies of the code that you need to be aware of, the
 10 central repository, your local clone of the central repository and the working copy of the
 11 code in your source directory.

- 12 1. The central git repository that contains all of the versions of the workbook is hosted
 13 by a machine named `cdcvs.fnal.gov`¹ The art-workbook team updates this
 14 repository as it develops and maintains the exercises.
- 15 2. In section [10.4.1](#), in step 5b) you used the `git clone` command to make a copy of
 16 the central repository in your source directory. This clone contains a complete his-
 17 tory of the development of art-workbook *as it existed at the time that you made the*
 18 *clone*. The local clone is found in the `.git` subdirectory of your source directory.
- 19 3. In section [10.4.1](#), in step 5d), you used the `git checkout` command to choose
 20 one of the tagged versions of art-workbook. This command looked into your local
 21 clone of the central repository, found all of the files in the requested version and
 22 put copies of them in the correct spot in the directory tree rooted at your source
 23 directory.



26 There are two other source code management systems that are widely used in HEP, `cvs`
 27 and `svn`. If you are familiar with either of these, `git` has an extra level: the concept of
 a local clone of the central repository does not exist in those systems. That is, when are
 using `cvs` or `svn` and you want to switch to another version of the code, you need to

¹ Originally this machine hosted only `cvs` repositories, hence its name. It now hosts `cvs`, `svn` and `git` repositories.

1 contact the central repository but, when you are using `git`, you need only to contact your
2 local clone of the central repository.

3 To bring your working code up to date you need to do two steps:

- 4 1. Update your local clone of the central repository.
- 5 2. Checkout the new version from the local clone.

6 The discussion of the checkout has several cases. Each is discussed in one of the following
7 sub-sections. It is possible for all four of these cases to occur on any given checkout.

8 **49.1.1.1 Files that you have Added**

ails:added

9 When you worked on Exercise 2, you added some files to your working directories; for
10 example you added the files `Second_module.cc` and `second.fcl`. When you do the
11 checkout of the new version, these files will remain in your working directory and will not
12 be modified; however the checkout command will generate some informational messages
13 telling you that your working directories contain files that are not part of the checked out
14 version.

15 **FIXME:** *Add example*

16 You do not need to take any action; just be aware of the situation.

17 **49.1.1.2 Files that you have Modified**

t:Modified

18 Another case occurs for files that have the following properties:

- 19 1. They were part of the old version.
- 20 2. You have modified them.
- 21 3. They have not been modified in the central repository since you cloned the reposi-
22 tory.

23 For example, suppose that you modified `first.fcl`; it is very unlikely that this file
24 would have been modified in the central repository after you cloned the repository.

1 In this case, the checkout command will issue a warning message to let you know that your
2 working version contains changes that are not part of the release you checked out.

3 **FIXME:** *Add example*

4 You do not need to take any action; just be aware of the situation.

5 **49.1.1.3 Files with Resolvable Conflicts**

6 Another case occurs for files that have the first two properties from the list in Section 49.1.1.2
7 but which have been modified in the central repository since you cloned the repository.
8 This will happen from time to time when we update exercises based on suggestions
9 from users.

10 When this happens there are two cases, one of which is discussed here, while the other is
11 discussed in the next sub-section.

12 If the two sets of changes (yours and those in the repository) are on different lines of the
13 file `git`, will usually successfully merge these changes; `git` will then issue an warning
14 message telling you what it has done.

15 **FIXME:** *Add example*

16 It is your responsibility to identify these cases, understand the changes made in the repos-
17 itory and understand if `git` did the merge correctly.

18 **49.1.1.4 Files with Unresolvable Conflicts**

19 The final case is a variant of the previous case; it occurs when `git` is unable to automat-
20 ically merge conflicting changes. This will happen when the changes you made and the
21 changes made in the repository affect the same line, or lines, of code. When `git` does not
22 know how to merge the changes it will give up, add markup to the offending files to mark
23 the conflict and issue an error message. This leaves the offending files in an unusable state
24 and you must correct the conflicts, by hand, before continuing.

25 The art-workbook has been designed so that this should happen very, very rarely. Most
26 readers should bookmark this spot for future reference and only read it when they need
27 to.

1 **FIXME:** *Finish this section.*

2 49.1.2 git Branches

3 git supports a concept known as *branches*. This is a very powerful feature that simplifies
 4 the task of having many developers collaboratively working on a single code base. More-
 5 over, different experiments can choose to use branches in different ways; therefore a full
 6 description of branches is very open ended topic.

7 Fortunately, to use the workbook you do not need to know very much about branches; all
 8 that you need to know is summarized in Figure 49.1, which shows a simplified view of the
 way that the art-workbook team uses git branches. In Figure 49.1 time starts at the bottom




Figure 49.1: A figure to illustrate the idea of git branches, as used in the Workbook; the figure is described in the text.

9
 10 of the figure and runs upward. The art-workbook team has adopted the convention that the

1 most up to date version of the art-workbook code will always be found by checking out a
2 branch named *develop*. In Figure 49.1 the vertical line represents the develop branch.

3 At the earliest time represented in Figure 49.1, the develop branch existed in some state
4 that the art-workbook team liked. So they tagged the develop branch with the name *v_1*,
5 for version 1. Shortly afterwards, the art team needed to add some improvements. To do
6 this they did:

- 7 1. Use the `git pull` command to make sure that their local copy of the repository
8 is up to date.
- 9 2. Check out the develop branch.
- 10 3. Start a new branch; in this example the new branch has the name *f1*, for “feature
11 number 1”.
- 12 4. Do all the development work on this branch. When they change files and commit
13 their changes, the changes stay local to the branch.
- 14 5. Once the new code has been tested it is merged back into the develop branch.
- 15 6. Use the `git pull` command to make sure that their local copy of the repository
16 is up to date; in this example, no one else has made.
- 17 7. Finally the developer must push their local copy of the repository to the central
18 repository.

19 This is illustrated in Figure 49.1 by the blue line labelled *f1*. While the developer is
20 working on *f1*, he can change back and forth between the develop branch and the *f1*
21 branch.

22 In Figure 49.1 the red line labelled *f2* represents a second feature that is added to the code
23 base following the same pattern as the first.

24 In this example, the development team decided to tag the develop branch after the *f2*
25 branch was merged back in; the tag was given the name *v2*; this is represented by the
26 green filled circle in the figure.

27 The next items on the figure are the branches named *f3* and *f4*. In this example, someone
28 started with the develop branch and began work on the feature *f3*. A little later someone
29 else (or maybe the same person) started with the develop branch and began work on the

1 feature `f4`. The person starting work on `f4` did so before the changes from `f3` were
2 merged back into the `develop` branch; therefore the two branches `f3` and `f4` both start
3 from the same place, the `v2` tag of `develop`. In this example, the next item on the timeline
4 is that the developer of `f3` commits their changes back to the `develop` branch. Sometime
5 after that the developer of `f4` merges their changes back. At this time the developer of `f4`
6 has the responsibility to check for conflicts that occurred during the merge and fix them;
7 this may or may not require consultation with the author of `f3`.

8 After this, the `develop` branch is again tagged, this time with a version named `v_3`.

9 The next items on the timeline are the branches named `f5` and `f6`. This example was
10 included to show that it is legal for `f6` both start and end during the time that `f5` is
11 alive.

12 Finally, the `develop` branch is tagged one more time, this time with the name `v_4`.

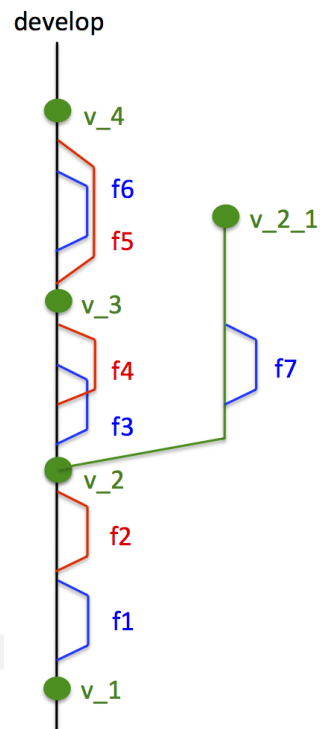
13 In Figure 49.1 consider a time when both branches `f5` and `f6` are active.

14 Fortunately, to use the workbook you do not need to know very much about branches. You
15 really need to know only two things.

- 16 1. In the `art-workbook` repository, there is a branch named `develop`; this branch is the
17 head of the project.
- 18 2. When the `art-workbook` team decides that a new stable version of the code is avail-
19 able, they `checkout` the `develop` branch and then start a new branch. This new
20 branch, called a release branch has the same name as the version number of the re-
21 lease. For example, the code for version `v0_00_13` is found in branch `v0_00_13`,
22 and so on. New work, towards the next release, continues on the `develop` branch.

23 This is a bit of simplification but it captures the big ideas. Users of `art-workbook` should
24 always work with one of the release branches; and they should always consult the docu-
25 mentaion to learn which version of the code is matched to that version of the documenta-
26 tion.

27 Users of the `art-workbook` should never work in the `develop` branch; at any given time that
28 branch may contain code that is still under development.



branches: feature

Figure 49.2: A figure to illustrate the idea of git branches, as used in the Workbook; the figure is described in the text. **FIXME:** *Not referenced in text*

1 49.1.3 Seeing which Files you have Modified or Added

```
2
```

```
3 At any time you can check to see which files you have modified and which you have  
4 added. To do this, cd to your source directory and issue the git status command.
```

```
5 Suppose that you have checked out version v0_00_13, modified first.fcl and added  
6 second.fcl. The git status command will produce the following output:
```

```
7 $ git status
```

```
8
```

```
9 # On branch v0_00_13
```

```
10 # Changes not staged for commit:
```

```
11 #   (use "git add <file>..." to update what will be committed)
```

```
12 #   (use "git checkout -- <file>..." to discard changes in  
13 working directory)
```

```
14 #
```

```
15 # modified:   first.fcl
```

```
16 #
```

```
17 # Untracked files:
```

```
18 #   (use "git add <file>..." to include in what will be committed)
```

```
19 #
```

```
20 # second.fcl
```

```
21 no changes added to commit (use "git add" and/or "git commit -a")
```

22 You should not issue the `git add` or `git commit` commands that are suggested above.

23 In the rare case that you have neither modified nor added any files, the output of `git status` will be:

```
24 $ git status
```

```
25 # On branch v0_00_13
```

1 50 *art* Run-time and Development Environ- 2 ments

3 50.1 The *art* Run-time Environment

4 Your *art* run-time environment consists of:

- 5 ○ your current working directory
- 6 ○ all of the directories that you can see and that contain relevant files, including system
7 directories, project directories, product directories, and so on
- 8 ○ the files in the above directories
- 9 ○ the environment variables in your environment (not sure how to say this nicely)
- 10 ○ any aliases or shell functions that are defined

11 fig:runtimeenv3time-env4
12 Figures 50.1, 50.2 and 50.3 show the elements of the run-time environment in various
13 scenarios, and a general direction of information flow for job execution.

14 When you are running *art*, there are three environment variables that are particularly im-
15 portant:

- 15 ○ PATH
- 16 ○ LD_LIBRARY_PATH
- 17 ○ FHICL_FILE_PATH

18 They are colon-separated lists of directory names. When you type a command at the com-
19 mand prompt, or in a shell script, the (bash) shell splits the line using whitespace and the

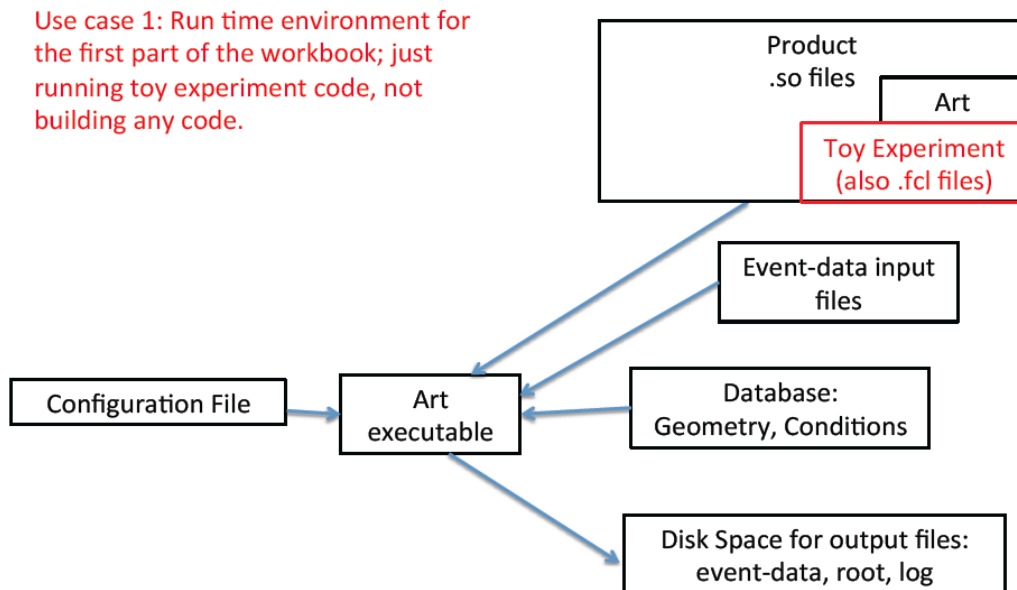


Figure 50.1: Elements of the *art* run-time environment, just for running the Toy Experiment code for the Workbook exercises

1 first element is taken as the name of a command. It looks in three places to figure out what
 2 you want it to do. In order of precedence:

- 3 1. it first looks at any aliases that are defined
- 4 2. secondly, it looks for shell keywords in your environment with the command name
 5 you provide
- 6 3. thirdly, it looks for shell functions in your environment with that name
- 7 4. then it looks for shell built-ins in your environment with that name
- 8 5. finally, it looks in the first directory defined in `PATH` and looks for a file with that
 9 name; if it does not find a match, it continues with the next directory, and so on,
 10 followed by the paths defined in the other two variables.

11 Some parts of the run-time environment will be established at login time by your login
 12 scripts. This is highly site-dependent. We will describe what happens at Fermilab - consult
 13 your site experts to find out if anything is provided for you at your remote site.

Use case 3: Run time environment for someone at their terminal running the experiments code on a job using prebuilt everything

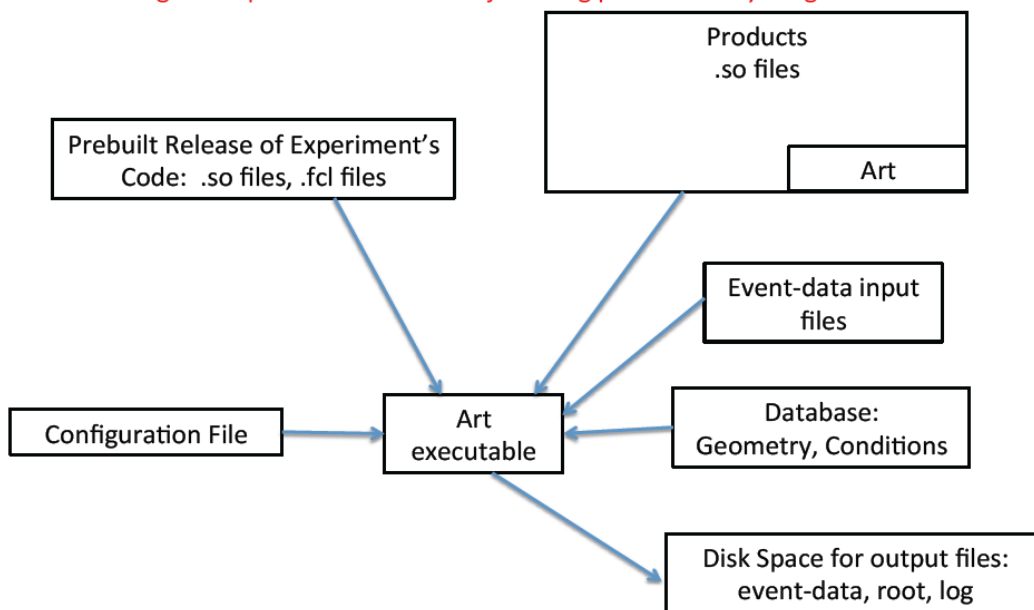
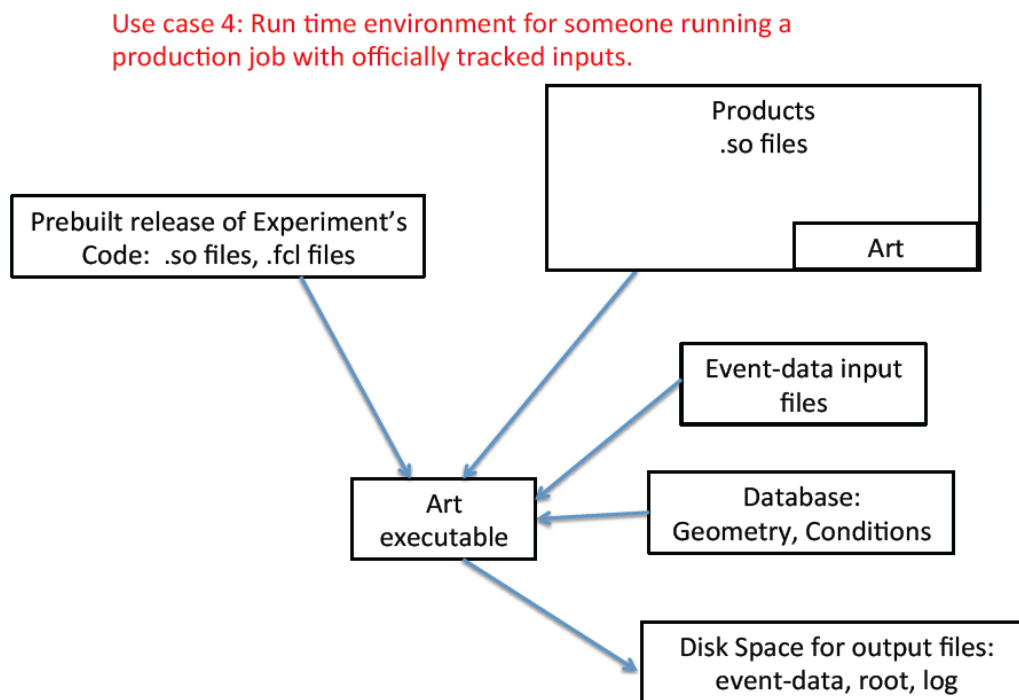


Fig:runtime-env3

Figure 50.2: Elements of the *art* run-time environment for running an experiment's code (everything pre-built)



runtime-env4

Figure 50.3: Elements of the *art* run-time environment for a production job with officially tracked inputs

1 When running the workbook, the interesting parts of your environment are established in
2 two steps:

- 3 ○ source a site-specific setup script
- 4 ○ source a project-specific setup script

5 The Workbook, and the software suites for most IF experiments, are designed so that all
6 site dependence is encoded in the site-specific setup script; that script adds information to
7 your environment so that the project-specific scripts can be written to work properly on
8 any site.

9 50.2 The *art* Development Environment

sec:ug-dev-env

10 The development environment includes the run-time environment in Section 50.1 plus the
11 following.

sec:ug-run-env

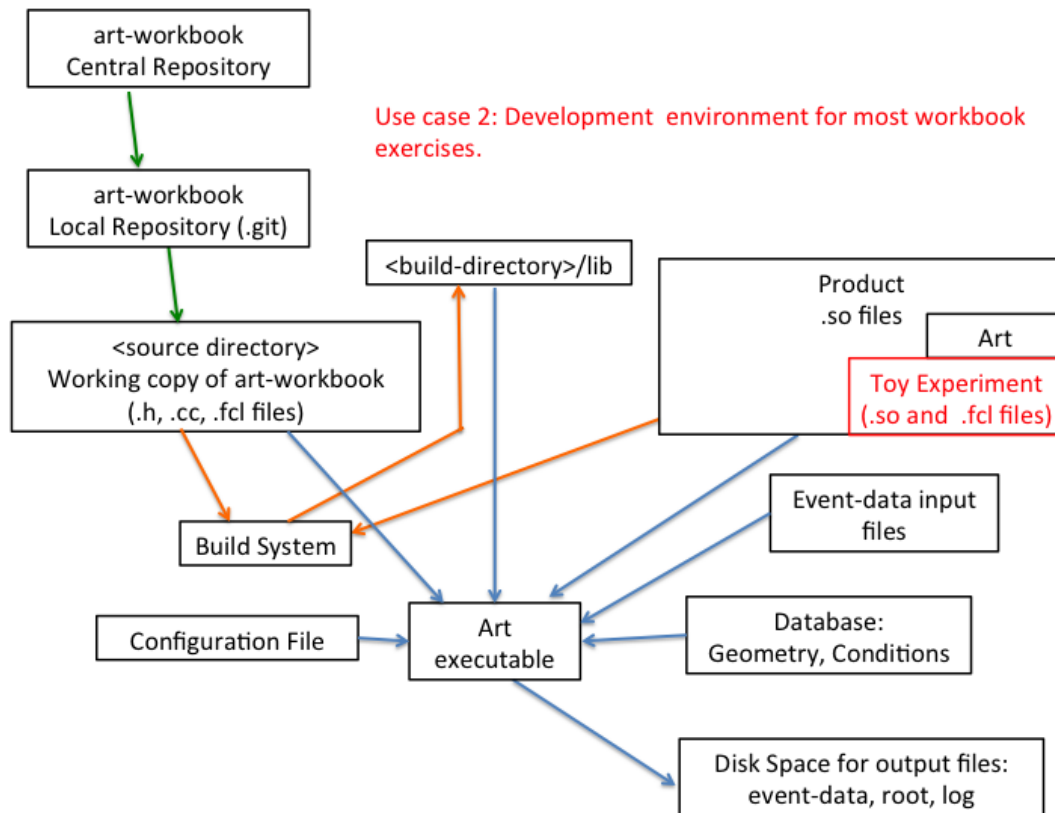
- 12 ○ the source code repository
- 13 ○ the build tools (these are the tools that know how to turn `.h` and `.cc` files in to `.so`
14 files)
- 15 ○ additional environment variables and `PATH` elements that simplify the use of the
16 above

17 Figures 50.4, 50.5 and 50.6 illustrate the development environment for various scenar-
18 ios.

19 In some experiments the run-time and development environments are identical.

20 It turns out that there is no perfect solution for the job that build tools do. **FIXME:** *Means*
21 *'there is no perfect build tool'*? As a result, several different tools are widely used. Every
22 tool has some pain associated with it. You never get to avoid pain entirely but you do get
23 to pick where you will take your pain.

24 The workbook uses a build tool named **cetbuildtools**. Other projects have chosen `make`,
25 `cmake`, `scons` and Software Release Tools (SRT). Here is something to watch out for:
26 “build tools” written as two words refers generically to the above set of tools; but “build-
27 tools” written as one word is the name of the executable that runs the build for **cetbuild-**
28 **tools**.



g:dev-env2

Figure 50.4: Elements of the *art* development environment as used in most of the Workbook exercises

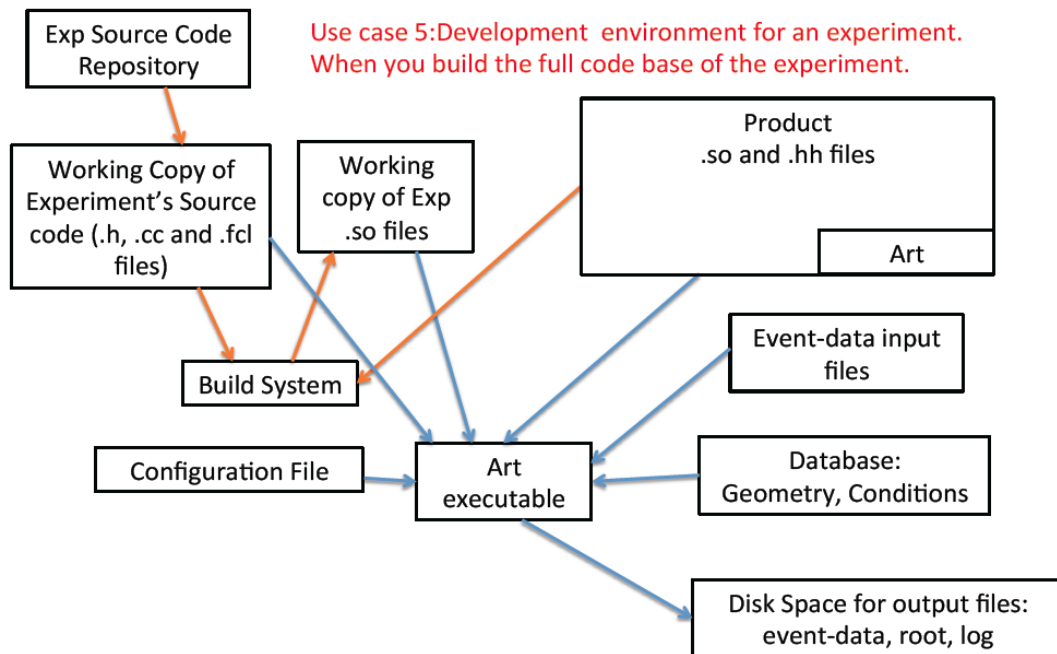


fig:dev-env5

Figure 50.5: Elements of the *art* development environment for building the full code base of an experiment

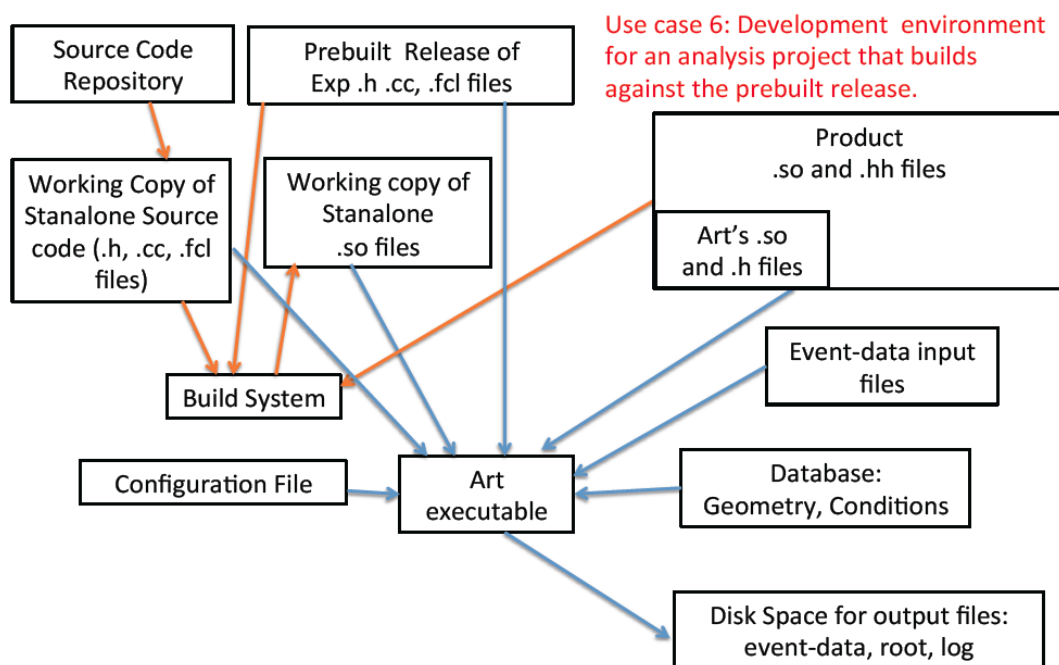


Figure 50.6: Elements of the *art* development environment for an analysis project that builds against prebuilt release

g: dev-env6

1 51 *art* Framework Parameters

2 **FIXME:** *from https://cdcvs.fnal.gov/redmine/projects/art/wiki/ART_framework_parameters.*
3 *I don't think it contains adequate info yet. AH 3/22*

4 This chapter describes all the parameters currently understood by the *art* framework, in-
5 cluding by framework-provided services and modules. The parameters are organized by
6 category (module, service or miscellaneous), and **FIXME:** *move the fcl intro* preceded
7 by a general introduction to the expected overall structure of an *art* FHiCL configuration
8 document.

9 51.1 Parameter Types

10 The parameters are described in tables for each module. **FIXME:** *prev sent not clear* The
11 type of a defined parameter may be:

- 12 ○ TABLE: A nested parameter set, e.g., `set: { par1: 3 }`
- 13 ○ SEQUENCE: A homogeneous sequence of items,
14 e.g., `list: [1, 1, 2, 3, 5, 8]`
- 15 ○ STRING: A string (enclosing double quotes not required when the string matches
16 `[A-Za-z_][A-Za-z0-9_]*`). (Note: Identifiers when quoted do not function as special
17 identifiers.) E.g.,

```
18     simpleString: g27  
19     harderString: "a-1"  
20     sneakystring1: "nil"  
21     sneakystring2: "true"  
22     sneakystring3: "false"
```

- 1 ○ COMPLEX: A complex number; e.g., `cnum: (3, 5)`
- 2 ○ NUMBER: A scalar (integer or floating point), e.g., `num: 2.79E-8`
- 3 ○ BOOL: A boolean, e.g.,
 - 4 `tbool: true`
 - 5 `fbool: false`
 - 6

7 **51.2 Structure of *art* Configuration Files**

8 **FIXME:** *I think all this is now covered in if-ug-fcl; if not, it should be there not here; 4/25*
9 *AEH*

10 The expected structure of an *art* configuration file

11 Note, any parameter set is optional, although certain parameters or sets are expected to be
12 in particular locations if defined.

13 **FIXME:** *add prolog info to if-ug-fcl. Add: what 'is' the prolog, what's it for? And putting*
14 *this example in before 9.3 may be helpful.*

```
15 # Prolog (as many as desired, but they must all be contiguous with only
16 # whitespace or comments inbetween.
17 BEGIN_PROLOG
18 pset:
19 {
20   nested_pset:
21   {
22     v1: [ a, b, "c-d" ]
23     b1: false
24     c1: 29
25   }
26 }
27 END_PROLOG
28
29 # Defaulted if missing: you should define it in most cases.
```

```
1 process_name: PNAME
2
3 # Descriptions of service and general configuration.
4 services:
5 {
6   # Parameter sets for known, built-in services here.
7   # ...
8
9   # Parameter sets for user-provided services here.
10
11  # General configuration options here.
12  scheduler:
13  {
14  }
15 }
16
17 # Define what you actually want to do here.
18 physics:
19 {
20   # Parameter sets for modules inheriting from EDProducer.
21   producers:
22   {
23     myProducer:
24     {
25       module_type: MyProducer
26       nested_pset: @local::pset.nested_pset
27     }
28   }
29
30   # Parameter sets for modules inheriting from EDFilter.
31   filters:
32   {
33     myFilter: { module_type: SomeFilter }
34   }
35
```



```
1 # Parameter sets for modules inheriting from EDAnalyzer.
2 analyzers:
3 {
4 }
5
6 # Define parameters which are lists of names of module sets for
7 # inclusion in end_paths and trigger_paths.
8
9 p1: [ myProdroducer, myFilter ]
10 e1: [ myAnalyzer, myOutput ]
11
12 # Compulsory for now: will be computed automatically in a future
13 # version of ART.
14
15 trigger_paths: [ p1 ]
16 end_paths: [ e1 ]
17 }
18
19 # The primary source of data: expects one and only one input source
20 parameter set.
21 source:
22 {
23 }
24
25 # Parameter sets for output modules should go here.
26 outputs:
27 {
28
29 }
```

1 **51.3 Services**

2 **51.3.1 System Services**

3 **FIXME:** *Define, list, describe, something!* These services are always loaded regardless of
4 whether a configuration is specified.

5 **51.3.2 FloatingPointControl**

6 These parameters control the behavior of floating point exceptions in different modules.

7 **51.3.3 Message Parameters**

8 These parameters configure the behavior of the message logger (this is a pseudo-service –
9 not accessible via ServiceHandle).

10 **51.3.4 Optional Services**

11 These services are only loaded if a configuration is specified (although it may be empty).

12 **51.3.5 Sources**

13 **51.3.6 Modules**

14 Output modules

Table 51.1: *art* Floating Point Parameters

Enclosing Table Name	Parameter Name	Type	Default	Notes
services	floating_point_control	TABLE	{}	Top-level parameter set for the service
floating_point_control	setPrecisionDouble	BOOL	false	
	reportSettings moduleNames	BOOL SEQUENCE	false []	Each module name listed should also have its own parameter set within floating_point_control. One may also specify a module name of, "default" to provide default settings for the following items:
<module-name>	enableDivByZeroEx	BOOL	false	
	enableInvalidEx	BOOL	false	
	enableOverFlowEx	BOOL	false	
	enableUnderFlowEx	BOOL	false	

Table 51.2: *art* Message Parameters

Enclosing Table Name	Parameter Name	Type	Default	Notes
services	message	TABLE		Top-level parameter set for the service
message				

1 52 Job Configuration in *art*: FHiCL

ch:fcl

2 Run-time configuration for *art* is written in the Fermilab Hierarchical Configuration Lan-
3 guage (FHiCL, pronounced “fickle”), a language that was developed at Fermilab to support
4 run-time configuration for several projects, including *art*. For this reason, this chapter will
5 need to discuss FHiCL both as a standalone language and as used by *art*.

6 By convention, the names of FHiCL files end in `.fcl`. Job execution is performed by
7 running *art* on a FHiCL configuration file, which is specified via an argument for the `-c`
8 option:

```
9 $ art -c run-time-configuration-file.fcl
```

10 See Figure 50.1 in Section 50.1 to see how the configuration file fits into the run-time
11 environment.

12 The FHiCL concept of *sequence*, as listed in brackets [], maps onto the C++ concept of
13 `std::vector`, which is a sequence container representing an array that can change in size.
14 Similarly, the FHiCL idea of *table*, as listed in curly brackets {}, maps onto the idea of
15 `fhiCL::ParameterSet`. **FIXME:** *need to define 'table' within fcl*. Note that `ParameterSet` is
16 not part of *art*; it is part of a utility library used by *art*, `FHiCL-CPP`, which is the C++
17 toolkit **FIXME:** *add ref* used to read FHiCL documents within *art*. FHiCL files provide
18 the parameter sets to the C++ code, specified via module labels and paths, that is to be
19 executed. **FIXME:** *I'm going to have to read this pgraph over later to see if it makes*
20 *sense!*

52.1 Basics of FHiCL Syntax

52.1.1 Specifying Names and Values

A FHiCL file contains a collection of definitions of the form

```
name : value
```

where “name” is a parameter that is assigned the value “value.” Many types of values are possible, from simple atomic values (a number, string, etc., with no internal whitespace) to highly structured table-like values; a value may also be a reference to a previously defined value. The white space on either side of the colon is optional. However, to include whitespace within a string, the string must be quoted (single or double quotes are equivalent in this case).

The fragment below will be used to illustrate some of the basics of FHiCL syntax:

```
# A comment.
// Also a comment.

name0 : 123           # A numeric value. Trailing comments
                    # work, too.
_name0 : 123         # Names can begin with underscores

name00 : "A quoted comment prefix, # or //, is just part of a
           # quoted string, not a comment"

name1:456.           # Another numeric value; whitespace is
                    # not important within a definition

name2 : -1.e-6

name3 : true         # A boolean value
NAME3 : false       # Other boolean value; names are case-
                    # sensitive.

name4 : red          # Simple strings need not be quoted
name5 : "a quoted string"
name6 : 'another quoted string'
```

```

1
2 name7 : 1 name8 : 2      # Two definitions on a line, separated by
3                          # whitespace.
4 name9
5 :                        # Same as name9:3 ; newlines are just
6 3                        # whitespace, which is not important.
7
8
9 namea : [ abc, def, ghi, 123 ] # A sequence of atomic values.
10                                     # FHiCL allows heterogeneous
11                                     # sequences, which are not,
12                                     # however, usable via the C++ API.
13
14 nameb :                    # A table of definitions; tables may nest.
15 {
16     name0: 456
17     name1: [7, 8, 9, 10 ]
18     name2:
19     {
20         name0: 789
21     }
22 }
23
24 namec : [ name0:{ a:1 b:2 } name1:{ a:3 c:4 } ]
25         # A sequence of tables.
26
27 named : []                # An empty sequence
28 namee : {}                # An empty table
29
30 namef : nil                # An atomic value that is undefined.
31
32 abc : 1                    # If a definition is repeated twice within
33 abc : 2                    # the same scope, the second definition
34 def : [ 1, 2, 3 ]         # will win (e.g., "abc" will be 2 and
35 def : [ 4, 5, 6 ]         # "def" will be [4,5,6])

```

```

1 name : {
2   abc : 1
3   abc : 2
4 }
5
6 cont1:{x: 1.0 y: 2.0 z: 3.0} # Hierarchical (compound) names denote
7 cont1.x : 5                 # levels of scope; here set x in cont1 to 5.
8 OR
9 cont2:[1, 2, 3]
10 cont2[0] : 1               # Here, redefine the first (atomic) value
11                             # for cont2, assign it the value 1. I.e., her
12                             # no action. Indices of PHiCL sequences
13                             # begin with 0. \fixme{right?}
14
15 name0:{ a:1 b:2 }
16 x : @local::name0.a       # Using reference notation "@local," this assigns
17                             # to x the value of a in table name0, in the
18                             # line above, this value is 1.

```

19 52.1.2 FHiCL-reserved Characters and Identifiers

20 Several identifiers, characters and strings are *reserved to* FHiCL. What does this mean?
 21 Whenever FHiCL encounters a *reserved* string, FHiCL will interpret it according to the
 22 *reserved* meaning. Nothing prevents you from using these reserved strings in a name or
 23 value, but if you do, it is likely to confuse FHiCL. FHiCL may produce an error or warn-
 24 ing, or it may silently do something different than what you intended. Bottom line: don't
 25 use reserved strings or symbols in the FHiCL environment for other than their intended
 26 uses.

27 The following characters, including the two-character sequence ::, are reserved to FHiCL:

28
 29 `, : :: @ [] { } ()`

30 The following strings have special meaning to FHiCL. They can be used as parameter
 31 values to pass to classes, e.g., to initialize a variable within a program, but their uses will

1 not be fully described here because of subtleties and variations. As you work with C++
2 and FHiCL, the way to use them will become clearer.

3 **true, false** These values convert to a boolean

4 **nil** This value is associated with no data type. E.g. if `a : nil`, then `a` can't be converted
5 to any data type, and it must be redefined before use

6 **infinity, +infinity, -infinity** These values initialize a variable to positive (the first two) or
7 negative (the third) infinity

8 **BEGIN_PROLOG, END_PROLOG** (FIXME: *I have no notes on this*)

9 The first six strings (three lines) above function as identifiers reserved to *art* only when
10 entered as lower case and unquoted; the last two strings (the last line) are reserved to *art*
11 only when they are in upper case, unquoted and at the start of a line. Otherwise these
12 are just strings. You may include any of the above reserved characters and identifiers in a
13 “quoted” string to prevent them from being recognized as reserved to *art*.

14 52.2 FHiCL Identifiers Reserved to *art*

15 FHiCL supports run-time configuration for several projects, not only for *art*. *art* reserves
16 certain FHiCL names as identifiers that it uses in well-defined ways. (Other projects may
17 use FHiCL names differently.) Within FHiCL files used by *art*, these FHiCL names obey
18 scoping rules similar to C++. These identifiers appear in the FHiCL file with a scope,
19 i.e.,

20 **FIXME:** *These defs are two narrow; need more work*

```
21 identifier : {
22 ...
23 }
```

24 if they define a list of modules or a processing block, or with square brackets

```
25 identifier :[
26 ...
27 ]
```

28 if they define a list of paths.

1 The following is a list of the identifiers reserved to *art* and their meanings. In the outermost
2 scope within a FHiCL file, the following can appear:

3 **process_name** A user-given name to identify the configuration defined by the FHiCL file
4 (it is recommended to make it similar to the FHiCL file name). This must appear at
5 the top of the file. It may not contain the underscore character (`_`).

6 **source** Identifies the data source, e.g., a file in ROOT format containing HEP events.

7 **services** Identifies ...**FIXME:** *describe*

8 **physics** Identifies the block of code that configures the scientific work to be done on every
9 event (as contrasted with the “bookkeeping” portions).

10 **outputs** List of output modules. **FIXME:** *that define the output format, structure and so*
11 *on?*

12 The following may appear within the `physics` scope:

13 **producers** Specifies the configurations of producer modules **FIXME:** ; *see Chapter 54* [|ch:ug-producer-modules](#)

14 **analyzers** Specifies the configuration of analyzer modules **FIXME:** ; *see Chapter 55* [|ch:ug-analyzer-modules](#)

15 **filters** Specifies the configuration of filter modules **FIXME:** ; *see Chapter 56* [|ch:ug-filter-modules](#)

16 **trigger_paths** Sequence of pathnames; the paths named here may contain only producer
17 and/or filter modules.

18 **end_paths** Sequence of pathnames; the paths named here may contain only analyzer and
19 output modules.

20 The last two elements specify which of the modules will be executed in an *art* job. (It
21 is legal for a module to be configured but not to be executed.) To understand order of
22 execution see Sections 52.4 and 52.7. [|sec:fcl:sectar:fcl:scheduling](#)

23 The identifier `process_name` is really only reserved to *art* within the outermost scope
24 (but it would seem to be needlessly confusing to use `process_name` as the name of
25 a parameter within some other scope). The names `trigger_paths` and `end_paths`
26 are artifacts of the first use of the CMS framework, to simulate the several hundred par-
27 allel paths within the CMS trigger; their meaning should be come clear after reading the
28 remainder of this page. **FIXME:** *check this*

1 **FIXME:** Why called 'trigger' paths – do they need to contain all producer and filter
2 module paths used in the fcl file? Similar question for 'end' paths...

3 **52.3 Structure of a FHiCL Run-time Configuration** 4 **File for *art***

5 Here is a sample FHiCL file called `ex01.fcl` that will do a physics analysis using
6 the code in the *art* module `Ex01_module.so` (the object file of the C++ source file
7 `Ex01_module.cc`). In this configuration, *art* will operate sequentially on the first three
8 events contained in the source file `inputFiles/input01.art`.

```
9 #include "fcl/minimalMessageService.fcl"  
10  
11 process_name : ex01  
12  
13 source : {  
14     module_type : RootInput  
15     fileNames   : [ "inputFiles/input01.art" ]  
16     maxEvents   : 3  
17 }  
18  
19 services : {  
20     message : @local::default_message  
21 }  
22  
23 physics : {  
24     analyzers : {  
25         hello : {  
26             module_type : Ex01  
27         }  
28     }  
29  
30     e1 : [ hello ]  
31     end_paths : [ e1 ]
```

1 }

2 Let's look at it step-by-step.

```
3 #include "fcl/minimalMessageService.fcl"
```

4 Similar to C++ syntax, this effectively replaces the '#include' line with the contents
5 of the named file. This particular file sets up a messaging service.

```
6 process_name : ex01
```

7 The value of the parameter `process_name` (`ex01`, here, the same as the FHiCL
8 file name) identifies this *art* job. It is used as part of the identifier for data products
9 produced in this job. For this reason, the value that you assign may not contain
10 underscore (`_`) characters. If the `process_name` is absent, *art* substitutes a default
11 value of "DUMMY."

```
12 source : {  
13   module_type : RootInput  
14   fileNames   : [ "inputFiles/input01.art" ]  
15   maxEvents   : 3  
16 }
```

17 This `source` parameter describes where events come from. There may be at most
18 one source module declared in an *art* configuration. At present there are two options
19 for choosing a source module:

20 **module_type : RootInput** *art*::Events will be read from an input file or from a list of
21 input files; files are specified by giving their pathname within the file system.

22 **module_type : EmptyEvent** Internally *art* will start the processing of each event by in-
23 crementing the event number and creating an empty *art*::Event. Subsequent modules
24 then populate the *art*::Event. This is the normal procedure for generating simulated
25 events.

26 Here `RootInput` is used; the data input file, in ROOT format, is assigned to the vari-
27 able `fileNames`. The `maxEvents` parameter says: Look at only the first three
28 events in this file. (A value of -1 here would mean "read them all.")

29 Note that if no source parameter set is present, *art* substitutes a default parameter set
30 of:

1 **FIXME:** *Want to indent this:*

```
2 source : {
3   module_type : EmptyEvent
4   maxEvents : 1
5 }
```

6 **FIXME:** *To do later* See the web page about configuring input and output modules for
7 details about what other parameters may be supplied to these parameter sets.

```
8 services : {
9   message : @local::default_message
10 }
```

11 Before starting processing, this puts the message logger in the recommended con-
12 figuration.

```
13 physics :{
14   analyzers: {
15     hello : {
16       module_type : Ex01
17     }
18   }
```

19 In *art*, `physics` is the label for a portion of the run-time configuration of a job. It
20 contains the “meat” of the configuration, i.e., the scientific processing instructions,
21 in contrast to the more administrative or bookkeeping information. The `physics`
22 block of code may contain up to five sections, each labeled with a reserved identi-
23 fier (that together form a parameter set within the FHiCL language); the strings are
24 *analyzers*, *producers*, *filters*, *trigger_paths* and *end_paths*. In our example it’s set to
25 *analyzers*.

26 The `analyzers` identifier takes values that are FHiCL tables of parameter sets
27 (this is true also for `filters` and `producers`). Here it takes the value `hello`,
28 which is defined as a table with one parameter, namely `module_type`, set to the
29 value `Ex01`. The setup defined a variable called `LD_LIBRARY_PATH`; *art* knows
30 to match the value defined by the name `module_type` to a C++ object file with the
31 name `Ex01_module.so` somewhere in the path defined by `LD_LIBRARY_PATH`.

1 **FIXME:** *Ex01_module.so must be an EDAnalyzer?*

2 We will expand on the `physics` portion of the FHiCL configuration in Section [52.5](#). sec:fcl-physics


```
3 e1          : [ hello ]
4 end_paths  : [ e1 ]
```

5 **FIXME:** *I don't know what this does yet. I know that trigger path are different from*
6 *end paths, they can contain different types of modules; event gets frozen after trigger*
7 *path.*

8 52.4 Order of Elements in a FHiCL Run-time Config- 9 uration File for *art*

sec:fcl-order

10 In FHiCL files there are very, very few places in which order is important. Here are the
11 places where it matters:

- 12 ○ A `#include` must come before lines that use names found inside the `#include`.
- 13 ○ A later definition of a name overrides an earlier definition of the same name.
- 14 ○ The definition of a name resolved using `@local` needs to be earlier in the file than
15 the place(s) where it is used. 
- 16 ○ Within a trigger path, the order of module labels is important. **FIXME:** *Can we say*
17 *'the order of module labels sets the order of execution'?*



18 Here is a list of *a few places* (of many) where order does not matter. This list is by no
19 means exhaustive.

- 20 ○ Inside the `physics` scope, the order in which modules are defined does NOT matter
21 for filters and analyzers blocks. These blocks define the run-time configurations of
22 *instances* of modules.
- 23 ○ The five *art*-reserved words that appear in the outermost scope of a FHiCL file can
24 be in any order. You could put `outputs` first and `process_name` last, as far as FHiCL
25 cares. It may be more difficult for humans to follow, however.
- 26 ○ Within the `services` block, the services may appear in any order.

1 Regarding `trigger_paths` and `end_paths`, the following is a conceptual description of how
2 *art* processes the FHiCL file:

- 3 1. *art* looks at the `trigger_paths` sequence. It expands each trigger path in the sequence,
4 removes duplicate entries and turns the result into an ordered list of module labels.
5 The final list has to obey the order of each contributing trigger path, but there are no
6 other ordering constraints.
- 7 2. It does the same for the `end_paths` sequence but there is no constraint on order.
- 8 3. It makes one big sequence that contains everything in 1 followed by everything in 2.
- 9 4. It looks throughout the file to find parameter sets to match to each module label in
10 the big list in 3.
- 11 5. It gives warning messages if there are left over parameter set definitions not matched
12 to any module label in 3.
- 13 6. It then parses the rest of the physics block to make a “dictionary” that matches
14 module labels to their configuration.

15 A conceptual description for the processing of services is as follows:

- 16 1. *art* first makes a list of all services, sorted alphabetically.
- 17  2. It makes a dictionary that matches service names to their parameter sets. A corollary
18 is that service names must be unique within an *art* job.
- 19 3. *art* has some “magic” services that it knows about internally. It loads the `.so` file
20 for each of them and constructs the services.
- 21 4. It loads the `.so` files for all of the services and calls their constructors, passing each
22 service its proper parameter set.
- 23 5. It works through its list of modules in 5 **FIXME:** ? - it loads the `.so` and calls the
24 constructor, passing the constructor the right parameter set.
- 25  6. It gives warning messages if there are left-over parameter set definitions not matched
26 to any module label in 3.

27 When one service relies on another, things get a bit more complicated. If service A requires
28 that service B be constructed first, then the constructor of service A must ask *art* for a

1 handle to service B. When this happens, *art* will start to construct service A since it is
2 alphabetically first. When the constructor of A asks for a handle to B, *art* will interrupt the
3 construction of service A, construct service B, and return to finish service A. Next, *art* will
4 see that the next thing in the list is B, but it will notice that B has already been constructed
5 and will skip to the next one.

6 Got that? Whew!

7 **52.5 The *physics* Portion of the FHiCL Configuration**

8
9 *art* looks for the experiment code in *art* modules. These must be referenced in the FHiCL
10 file via *module labels*, which are just variable names that take particular values, as this
11 section will describe. The structure of the FHiCL file – or a portion thereof – therefore
12 defines the event loop for *art* to execute. The event loop, as defined in the FHiCL file, is
13 collected into a scope labeled *physics*.

14 For a module label you may choose any name, as long as it is unique within a job, con-
15 tains no underscore (_) characters and is not one of the names reserved to *art*. In the
16 sample *physics* scope code below, we define *aProducer*, *bProducer*, *checkAll*,
17 *selectMode0* and *selectMode1* as module labels.

```
18 physics: {  
19  
20   producers : {  
21     aProducer: { module_type: MakeA }  
22     bProducer: { module_type: MakeB }  
23   }  
24  
25   analyzers : {  
26     checkAll: { module_type: CheckAll }  
27   }  
28  
29   filter : {  
30     selectMode0: {
```

```

1     module_type: Filter1
2     mode: 0
3   }
4   selectModel: {
5     module_type: Filter1
6     mode: 1
7   }
8 }

```

9 The minimum configuration of a module is:

```
10 <moduleLabel> : { module_type : <ClassName> }
```

11 for example, in our code above:

```
12 aProducer: { module_type: MakeA }
```

13 `aProducer` is the module label and `MakeA` corresponds to a module of experiment code (i.e., an *art* module) named `MakeA_module.so`, which in turn was built from `MakeA_module.cc`. Since it falls within the scope `producers`, it must be a module of type `EDProducer`.

17 Let's take this a step farther, and assume that this `EDProducer`-type module `MakeA` accepts four arguments that we want to provide to *art*. The configuration may look like this:

```

19 moduleLabel : {
20   module_type : MakeA
21   pname0 : 1234.
22   pname1 : [ abc, def]
23   pname2 : {
24     name0: {}
25   }
26 }

```

27 This list under `module_type : MakeA` represents parameters that will be formed into a `fhicl::ParameterSet` object and passed to the module `MakeA` as an argument in its constructor. `pname0` is a double, `pname1` is a sequence of two atomic character values, `pname2` consists of a single table named `name0` with undefined contents.

1 **FIXME:** *where to put?* Note that *paths* are lists of module labels, while the two reserved
2 names, `trigger_paths` and `end_paths` are lists of paths.

3 52.6 Choosing and Using Module Labels and Path 4 Names

5 For a module label or a path name, you may choose any name so long as it is unique within
6 a job, contains no underscore (`_`) characters and is not one of the names reserved to *art*
7 (see Section [52.2](#), sec:art-identifiers).

8 Any name that is a top-level name inside of the `physics` parameter set is either a reserved
9 name or the name of a path.

10 It is important to recognize which identifiers are module labels and which are path names
11 in a FHiCL file. It is also important to distinguish between a class that is a module and
12 instances of that module class, each uniquely identified by a module label. **FIXME:** *But*
13 *does this belong here? FHiCL files will only call modules, right?*

14 *art* has several rules that were recommended practices in the old framework but which
15 were not strictly enforced by that framework. *art* enforces some of these rules and will,
16 soon, enforce all of them: **FIXME:** *current status?*

- 17 ○ A path may go into either the `trigger_paths` list or into the `end_paths` list,
18 but not both.
- 19 ○ A path that is in the `trigger_paths` list may only contain the module labels of
20 producer modules and filter modules.
- 21 ○ A path that is in the `end_paths` list may only contain the module labels of analyzer
22 modules and output modules.

23 Analyzer modules and the output modules may be separated into different paths; that might
24 be convenient at some times but it is not necessary. On the other hand, keeping trigger
25 paths separate has real meaning. **FIXME:** *because the order of producer and filters is*
26 *important?*

52.7 Scheduling Strategy in *art*

A set of scheduling rules is enforced in *art*. (Some of the details are remnants of compromises and conflicting interests with CMS.) One of the top-level rules in the scheduler is that all producers and filters must be run first, using the ordering rules specified below. After that, all analyzer and output modules will be run. Recall that analyzer modules and output modules may not modify the event, nor may they produce side effects that influence the behavior of other analyzer or output modules. Therefore, *art* is free to run analyzer and output modules in any order.

The full description of the scheduler strategy is given below:

- If a module name appears in the definition of a path name but it is not found among the the list of defined module labels, FHiCL will issue an error.
- One each event, before executing any of the paths, execute the source module.
- On each event, execute all of the paths listed in the `trigger_paths`.
 - Within one path, the order of modules listed in the path is followed strictly; at present there is one exception to this: see the discussion about the remaining issues
 - *art* can identify module labels that are in common to several `trigger_paths` and will execute them only once per event. In the above example, `aProducer` and `bProducer` are executed only once per event.
 - The various paths within the `trigger_paths` may be executed in any order, subject to the above constraints.
 - If a path contains a filter, and if the filter return false, then the remainder of the path is skipped.
 - The module name of a filter can be negated in path using, `!moduleLabel`; **FIXME:** *clarify, e.g., !myModule?* in this case the path will continue if the filter returns false and will be aborted if the filter returns true.
 - If the module label of a filter appears in two paths, negated in one path and not negated in the other, *art* will only run the instance of filter module once and will use the result in both places.

- 1 – If a module in a trigger path throws, the default behaviour of *art* is to stop
2 all processing and to shut down the job as gracefully as possible. *Art* can be
3 configured, at run time, so that, for selected exceptions, it behaves differently.
4 For example it can be configured to continue with the current trigger path, skip
5 to the next trigger path, skip to the next event, and so on.
- 6 ○ On each event, execute all of the paths listed in the `end_paths`.
 - 7 – The module labels listed in `end_paths` are executed exactly once per event,
8 regardless of how many paths there are in the `trigger_paths` and regardless of
9 any filters that failed.
 - 10 – If a module label appears multiple times among the end paths, it is executed
11 only once. No warning message is given.
 - 12 – Even if all `trigger_paths` have filters that fail, all module labels in the end path
13 will be run.
 - 14 – `End_path` is free to execute the modules in the `end_path` in any order.
 - 15 – If a module in the `end_path` throws, the default response of *art* is to make a best
16 effort to complete all other modules in the end path and then to shutdown the
17 job in an orderly fashion. This behaviour can be changed at run-time by adding
18 the appropriate parameter set to the top level `.fcl` file.
- 19 ○ One can ask that an output module be run only for events that pass a given trig-
20 ger_path; this is done using the `SelectEvents` parameter set, **FIXME:** *as illustrated*
21 *above. FIXME: I believe, but am not certain, that SelectEvents allows some simple*
22 *boolean logic on the pass/fail status of several paths; I have not found the documen-*
23 *tation for this.*
- 24 ○ At present there is no syntax to ask that an analyzer module be run only for events
25 that pass or fail some of the trigger paths. A planned improvement to *art* is to give
26 analyzer modules a `SelectEvents` parameter that behaves as it does for output mod-
27 ules. **FIXME:** *status?*
- 28 ○ If a path appears in neither the `trigger_paths` nor the `end_paths`, there is no warning
29 given.
- 30 ○ If a module label appears in no path, a warning will be given.

1 **FIXME:** *Do we want to include this? What is the status?* In the above there is a lot of
2 focus on which groups of modules are free to be run in an arbitrary order. This is laying
3 the groundwork for module-parallel execution: *art* is capable of identifying which modules
4 may be run in parallel and, on a multi-core machine, *art* could start separate threads for
5 each module. At present both ROOT and G4 are not thread-safe so this is not of immediate
6 interest. But there are efforts underway to make both of these thread-safe and we may one
7 day care about module-parallel execution; our interest in this will depend a great deal on
8 the future evolution of the relative costs of memory and CPU.

9 For simple cases, in which there is one trigger path with only a few modules in the path,
10 and one end path with only a few modules in the path, the extra level of bookkeeping is just
11 extra typing with no obvious benefit. The benefit comes when many work groups wish to
12 run their modules on the same events during one *art* job; perhaps this is a job skimming off
13 many different calibration samples or perhaps it is a job selecting many different streams of
14 interesting Monte Carlo events. In such a case, each work group needs only to define their
15 own trigger path and their own end path, without regard for the requirements of other work
16 groups; each work group also needs to ensure that their paths are added to the `end_paths`
17 and `trigger_paths` variables. *Art* will then automatically, and correctly, schedule the work
18 without redoing any work twice and without skipping work that must be done. This feature
19 came for free with *art* and, while it imposes a small burden for novice users doing simple
20 jobs, it provides an enormously powerful feature for advanced users. Therefore it was
21 retained in *art* when some other features were removed.

22 **FIXME:** *What about the 'some remaining issues' – is that current, too?*

23 52.8 Scheduled Reconstruction using Trigger Paths

24 Consider the following problem. You wish to run a job that has:

- 25 ○ Two producers `MakeA_module.cc` and `MakeB_module.cc`. You want to run both
26 producers on all events.
- 27 ○ One analyzer module that you want to run on all events, `CheckAll_module.cc`.
- 28 ○ You have a filter module, `Filter1_module.cc` that has two modes, 0 and 1; the mode
29 can be selected at run time via the parameter set.

- 1 ◦ You wish to write all events that pass mode 0 of the filter to the file file0.art and you
- 2 wish to write all events that pass mode 1 of the filter to file1.art

3 Here is code that would accomplish this:

```
4 process_name: filter1
5
6 source: {
7   # Configure some source here.
8 }
9
10 physics: {
11
12   producers : {
13     aProducer: { module_type: MakeA }
14     bProducer: { module_type: MakeB }
15   }
16
17   analyzers : {
18     checkAll: { module_type: CheckAll }
19   }
20
21   filter : {
22     selectMode0: {
23       module_type: Filter1
24       mode: 0
25     }
26     selectMode1: {
27       module_type: Filter1
28       mode: 1
29     }
30   }
31
32   mode0: [ aProducer, bProducer, selectMode0 ]
33   mode1: [ aProducer, bProducer, selectMode1 ]
34   analyzermodes: [ checkAll ]
```

```
1  outputFiles:  [ out0, out1 ]
2
3  trigger_paths : [ mode0, mode1 ]
4  end_paths : [ analyzermods, outputFiles ]
5  }
6
7  outputs: {
8    out0: {
9      module_type: RootOutput
10     fileName: "file0.art"
11     SelectEvents: { SelectEvents: [ mode0 ] }
12   }
13
14   out1: {
15     module_type: RootOutput
16     fileName: "file1.art"
17     SelectEvents: { SelectEvents: [ mode1 ] }
18   }
19
20 }
```

21 Recall that the names `process_name`, `source`, `physics`, `producers`, `analyzers`, `filters`, `trigger_paths`, `end_paths` and `outputs` are reserved to *art*. **FIXME:** *can any of these use the underscore? or only process_name restricted?* The names `aProducer`, `bProducer`, `check-`
22 `All`, `selectMode0`, `selectMode1`, `out0` and `out1` are module labels, and these are names of
23 paths: `mode0`, `mode1`, `outputFiles`, `analyzermods`.

26 52.9 Reconstruction On-Demand

27 **FIXME:** *This note needs to be extended to include the use of filters when reconstruction*
28 *on demand is enabled. It also needs to talk about the meaning of "keep (with asterisks)"*
29 *in an output module when reconstruction on demand is enabled: this will trigger running*
30 *all of the registered producers and all of their data products will be written to the output*
31 *file. For a discussion about the keep/drop syntax when using Scheduled Reconstruction,*
32 *see the discussion of configuring output modules.*

1 **52.10 Bits and Pieces **FIXME:****

2 What variables are known to *art*? physics (which has the five reserved identifiers: fil-
3 ters, analyzers, producers, trigger paths and end paths), what else? input file type RootIn-
4 put

5 I know that trigger path are // different from end paths, they can contain different types of
6 modules; // event gets frozen after trigger path.

7 *art* knows to match the value defi

8 ned by the name 'module_name' to a C++ object fi

9 le with the name module_name_module.so" somewhere in the path defi

10 ned by LD LIBRARY PATH.

11 Further information on the FHiCL language and usage can be found at **FIXME:** *will*
12 *change* the mu2e FHiCL page.

53 Data Products

ug-dataproducts

2 **FIXME:** *From web page*

3 **FIXME:** *About duplication of material here vs Chapter 16 and Section 16.4.2. It's a good*
 4 *question. The material in those other places incomplete. It is just enough information*
 5 *to get through the early exercises. This chapter will contain the core dump about data*
 6 *products - there is a lot more information to come. One of my boundary conditions is that*
 7 *Section 16.4.2 must be complete by itself. Therefore there is necessarily some duplication.*
 8 *When we write this chapter we may be able to do one of two things. (1) we may be able to*
 9 *build this chapter in a way that it has sections that are standalone and can be refered from*
 10 *the earlier chapters. Or (2), we may be able to write the common fragment in a separate*
 11 *file that is \input'ed into both places. I say "may" be able to because we won't know until*
 12 *we understand everything that needs to be in this section - the sections serve two different*
 13 *audiences so there is no guarantee that one of these two strategies will work. Another thing*
 14 *that I have said few times but may have been forgotten: 100the Users' Guide is just a*
 15 *random collection of stuff copied from various web sites. I have given little thought to*
 16 *appropriate organization of anything in the Users's Guide and I have checked little of it*
 17 *for correctness (it was all correct at one time).*

53.1 Overview

19 A *data product* is anything that you can add to an event or see in an event. Examples
 20 include the generated particles, the simulated particles produced by Geant4 **FIXME:** *ref,*
 21 the hits produced by Geant4, tracks found by the reconstruction algorithms, clusters found
 22 in the calorimeters and so on.

53.2 The Full Name of a Data Product

Each data product within an event is uniquely identified by a four-part identifier that includes all namespace information. The four parts are separated by underscores:

```
DataType_ModuleLabel_InstanceName_ProcessName
```

DataType identifies the data type that is stored in the product. It is a “friendly” identifier in the way that its syntax has been standardized to deal with *collection* types, as follows:

- If a product is of type T **FIXME:** *non-collection?*, then the friendly name is “T”.
- If a product is of type `mu2e::T`, then the friendly name is “mu2e::T”.
- If a product is of type `std::vector<mu2e::T>`, then the friendly name is “mu2e::Ts”.
- If a product is of type `std::vector<std::vector< mu2e::T > >`, then the friendly name is “mu2e::Tss”.
- If a product is of type `cet::map_vector<mu2e::T>`, then the friendly name is “mu2e::Tmv”. See below **FIXME:** *ref clearly* for a discussion about where underscores may not be used; this example is safe because of the substitution of “mv” for `map_vector`.

ModuleLabel identifies the module that created the product; this is the module label **FIXME:** *ref where defined*, which distinguishes multiple instances of the same module within a produces **FIXME:** *'product' or 'producer'? If 'produces' then need explanation of this*. It is *not* the class name of the module.

InstanceName is a label for the data product that distinguishes two or more data products of the same type that were produced by the same module, in the same process **FIXME:** *'same process' means 'same run of art?'*. If a data product is already unique within this scope, it is legal to leave this field blank **FIXME:** *still need two underscores?*. The instance label is the optional argument of the call to “produces” **FIXME:** *?* in the constructor of the module (xxxx below) **FIXME:** *with quotes?*:

```
produces<T> ("xxxx") ;
```

ProcessName is the name of the process that created this product. It is specified in the FHiCL file that specifies the run-time configuration for the job (shown as *ReadBack02*

1 below):

2 `process_name : ReadBack02`

3 Because the full name of the product uses the underscore character to delimit fields, it
4 is forbidden to use underscores in any of the names of the fields. Therefore, none of the
5 following items may contain underscores:

- 6 ○ the class name of a class that is a data product; the exception is the `cet::map_vector`
7 template; when creating the friendly name, *art* internally recognizes this case and
8 protects against it
- 9 ○ the namespace in which a data product class lives
- 10 ○ module labels
- 11 ○ data product instance names
- 12 ○ process names

13 It is important to know which names need to match each other; see Section 59.1. **FIXME:**
14 *but this is just 'module' names*

15 **FIXME:** *get from web page*

[|sec:modname-rules](#)

1 54 Producer Modules

er-modules

DRAFT

55 Analyzer Modules

analyzer-modules

2 **FIXME:** *From oink tutorial 'analyzer interface' slide*

3 Analyzer modules request data products, do not create new ones; make histograms, etc.

4 **FIXME:** *need links to data product.*

5 An analyzer interface **FIXME:** *why 'interface'?* looks like the following. **FIXME:** *Need*
6 *explanation*

```

7 class EAnalyzer {
8     // explicit EAnalyzer(ParameterSet const&)
9
10    virtual void analyze(Event const&) = 0
11    virtual void reconfigure(ParameterSet const&)
12
13    virtual void beginJob()
14    virtual void endJob()
15    virtual bool beginRun(Run const &)
16    virtual bool endRun(Run const &)
17    virtual bool beginSubRun(SubRun const &)
18    virtual bool endSubRun(SubRun const &)
19
20    virtual void respondToOpenInputFile(FileBlock const& fb)
21    virtual void respondToCloseInputFile(FileBlock const& fb)
22    virtual void respondToOpenOutputFiles(FileBlock const& fb)
23    virtual void respondToCloseOutputFiles(FileBlock const& fb)
24 }
```

56 Filter Modules

er-modules

2 **FIXME:** *From oink tutorial 'filter interface' slide*

3 Filter modules request data products and can alter further processing using return values

4 **FIXME:** *need links to data product and return values.*

5 A filter interface **FIXME:** *why 'interface'?* looks like the following. **FIXME:** *Need explanation*

```
7 class EDFilter {
8     // explicit EDFilter(ParameterSet const&)
9
10    virtual bool filter(Event&) = 0
11    virtual void reconfigure(ParameterSet const&)
12
13    virtual void beginJob()
14    virtual void endJob()
15    virtual bool beginRun(Run &)
16    virtual bool endRun(Run &)
17    virtual bool beginSubRun(SubRun &)
18    virtual bool endSubRun(SubRun &)
19
20    virtual void respondToOpenInputFile(FileBlock const& fb)
21    virtual void respondToCloseInputFile(FileBlock const& fb)
22    virtual void respondToOpenOutputFiles(FileBlock const& fb)
23    virtual void respondToCloseOutputFiles(FileBlock const& fb)
24 }
```

1 57 *art* Services

ch:ug-services

2 **FIXME:** *I WILL COME BACK TO THIS LATER WHEN I UNDERSTAND IT BETTER.*
 3 ANNE 4/28/14 **FIXME:** *There are diagrams called firstdotfcl-diagram and how-services-*
 4 *fit-in that I made in graffle; in artdoc-local*

5 **FIXME:** *From Chris G's wiki page*

6 This chapter describes what services are, what types of services are recognized, how
 7 they are used, how they are constructed. **FIXME:** *I think that's what we want to put in*
 8 *here!*

9 57.1 About Services

10 i


11 *art* Services are introduced in Section 3.6.5. A service in the context of *art* is a class for
 12 which an instance is configured at runtime by a FHiCL configuration and made available
 13 for use by modules and other services, and whose lifetime and configuration are managed
 14 by *art*. A service provides modules access to resources or information that is outside the
 15 purview of the modules, and that is valid for some aggregation of events, subRuns or runs,
 16 or over some time interval. Services may also be used to provide certain types of utility
 17 functions.

18 **FIXME:** *Figure caption* How services fit into *art* during job processing. *art* first pulls
 19 in the parameter set from the user's FHiCL file, instantiates the services called for, calls
 20 the constructor for each module, then begins the event loop. **FIXME:** *Show how user-*
 21 *defined service gets pulled in; can I differentiate between art services and non-internal*
 22 *ones?*

1 Services can be provided by three different sources. Some services are provided by *art* to
 2 implement internal functions used by both *art* itself and experiment code, e.g., the mes-
 3 sage service **FIXME:** *ref it*. Other services **FIXME:** *Where do these come from?* imple-
 4 ment functions that are not internal to *art* but are used by most experiments, for exam-
 5 ple `TFileService`, which manages a secondary ROOT-format file for histograms, and
 6 **FIXME:** *name* and **FIXME:** *name*, which provide interfaces to geometry and calibration
 7 information, respectively. Users (experiments and/or individual experimenters) can add
 8 services to *art*, too.

9 Services can be publicly available, i.e., callable by user-defined classes, or callable by
 10 other services or by *art* itself. Services that are publicly available typically have a header
 11 (.h) file, whereas those intended to be called by *art* itself do not. A service may *register*
 12 functions with *art* **FIXME:** *I don't understand this concept*; a registered function can
 13 then be called at appropriate moments directly by the *art* framework, e.g., before or after
 14 `beginRun`, `endSubRun` and so forth.

15 Services can be made available for environments that support parallelism. More in Sec-
 16 tion 57.6.1. `sec:declare:service`

17 A service should *not* provide “backdoor” communication between modules of physics 
 18 data or to/from anything that ought to be stored in an event, `subRun` or `run`, and/or anything
 19 that requires provenance tracking, since the provenance would not be captured.

20 To access a given service, the calling entity requests a *service handle* for it; this is a type
 21 of smart pointer that **FIXME:** *I don't understand well enough yet...* Depending on what a
 22 service does and what external system(s), if any, it needs to access, a *service interface* may
 23 be required. **FIXME:** *Add something like: this will be discussed in section whatever.*

24 **FIXME:** *need pic: art module, two sep services outside it, service handles to each. One,*
 25 *say to data file, has no interface; the other, say to SAM, has interface; then also a non*

26 Several types of standard *art* services exist: **FIXME:** *but these are not art-provided,*
 27 *right?*

- 28 ○ `TFile`: Controls the ROOT directories (one per module) and manages the his-
 29 togram file.
- 30 ○ `Timing`: Tracks CPU and wall clock time for each module for each event

- 1 o Memory: Tracks increases in overall program memory on each module invocation
- 2 o FloatingPointControl: Allows configuration of FPU hardware “exception”
- 3 processing
- 4 o RandomNumberService: Manages the state of a random number stream for each
- 5 interested module
- 6 o srcv MessageFacility: Routes user-emitted messages from modules based on type
- 7 and severity to destinations

8 57.2 Service Handles

9 To use a service, your code must point to it. Your FHiCL file must point to the service(s)
10 you need and include its configuration, e.g., for TFileService:

```
11 #include "fcl/minimalMessageService.fcl"
12 process_name : ex06
13 ...
14 services : {
15     ...
16     TFileService : { fileName : "output/ex06.root" }
17 }
18 ...
```

19 The .h or .cc file’s declaration of your module (class) under “private” must declare a service
20 handle to the service. This shows part of the code in `Ex06_module.cc` in which a
21 service handle to the TFileService is declared with the name `tfs_`. The **FIXME:** *code*
22 *for the handle?* is in the namespace `art`. **FIXME:** *so TFileService is not provided by art,*
23 *but it’s placed in this namespace so that it’s available to all art modules? Why not under*
24 *art:: in fcl file?:*

```
25 class Ex06 : public art::EDAnalyzer {
26 public:
27     explicit Ex06(fhicl::ParameterSet const& );
28     ...
29 private:
30     ...
31     art::ServiceHandle<art::TFileService> tfs_;
```


1 The class's constructor calls/declares? it as a member function of the class, with no arguments:
2

```

1 tex::Ex06::Ex06(fhicl::ParameterSet const& pset ):
2     gensTag_(pset.get<std::string>("genParticlesInputTag")),
3     tfs_(),
4     hNGen_(nullptr) {
5 }

```

8 57.3 Implementing Simple Services

9 A class that is intended as a service must be declared as a service to the *art* framework and
10 its implementation defined. The way to do this depends on whether the service is provided
11 by *art* or not, and whether it requires an interface.

12 An *art*-provided service such as Timing or SimpleMemoryCheck should be configured for use in this way. In the FHiCL file

```
14 services.SimpleMemoryCheck: { }
```

15 Services not provided by *art* must be configured to be nested under `services.user`:

```
16 services.user.MyService: { }
```

17 Services implementing an interface require yet another syntax for configuration. In `services`
18 or `services.user` **FIXME:** *what are these objects?* there must be the line (with the appropriate ServiceName):

```
20 InterfaceName: { service_provider: ServiceName }
```

21 For example:

```

22 services.user.CatalogServiceInterface: {
23     service_provider: IFCatalogInterface
24 }

```

25 57.4 Configuring a Service

26 First, a service needs to be configured for use. In a FHiCL file, this would look like (shown
27 for the service TFileService):

```

1     services:
2     {
3         TFileService:
4         {
5             fileName: "tfile_output.root"
6         }
7
8         # experiment- or user-defined plugin services
9         ...
10    }

```

11 57.5 Accessing a Service

12 Services are accessed via an instantiation of the `art::ServiceHandle` template. **FIXME:**
 13 *Why is this a 'template'?* E.g., to access the `TFileService` from your module, use this
 14 syntax:

```
15 art::ServiceHandle<art::TFileService> h;
```

16 Treat the handle so created as any other smart pointer. **FIXME:** *Add a x-reference once*
 17 *we have one*

18 If you intend to access one service from a second service, you should obtain a `ServiceHandle`
 19 for the second service from the constructor of the first, even if you don't plan to use it there.
 20 This ensures the correct order of construction (and destruction) of the services, which oth-
 21 erwise would be alphabetical by class name.

22 Any library containing code that uses a service should link explicitly to that service.

23 Note that obtaining an `art::ServiceHandle` is relatively expensive in terms of per-
 24 formance, and should not be done in tight loops. It is recommended to cache the han-
 25 dle in your module class scope, at the class or module entry point (i.e., `analyze()`,
 26 `filter()`). **FIXME:** *Could use short example here*



28 To avoid problems when parallelism is in operation, do not cache a bare pointer to a service
 29 at a greater-than-module-entry-point scope. **FIXME:** *What's a 'bare' pointer relative to a*
handle?

1 De-referencing the handle, on the other hand, while not zero-cost, is relatively cheap. As
 2 for certain types of services, *art* runs a check on the handle to verify that the service
 3 referenced is appropriate for the context in which the handle will be used. If necessary, the
 4 handle is “re-seated” for the new context. **FIXME:** *What does that mean?*

5 57.6 Writing a Service

6 ServiceMacros.h **FIXME:** *I don’t think we’ve defined “service macro” or “DECLARE”*
 7 *or “DEFINE” macros. What is it, what’s it used for, and why is it a header file.*

8 The code in the .cc file implementing a service must begin with a call to the *art service*
 9 *macro*: What does ServiceMacros do?

```
10 #include "art/Framework/Services/Registry/ServiceMacros.h"
```

11 If the service implements an interface, the interface definition takes the form of a header
 12 file containing a class definition (declaration?) and interface declaration. Clarify ‘define’
 13 vs ‘declare’ vs ‘implement’; I don’t think we’re using them consistently (sec 6.6.4 says in
 14 ptest.cc a line declares the variable p0 even though ptest.h exists)

15 An access interface **FIXME:** *or a ‘service’ interface?* looks like the following (shown for
 16 the service TFileService):

```
17 #include "art/Framework/Services/Optional/TFileService.h"
18     ...
19     art::ServiceHandle<art::TFileService> tfs;
20     fFinalVtxX = tfs->make<TH1F>("fFinalVtxX",
21                               "Circe Vertex X; Xfit-Xmc (cm); Events",
22                               200, -50.0, 50.0);
```

23 An interface implementation must inherit publicly from the interface, and declare and
 24 define itself as an implementation of that interface. **FIXME:** *A full example would be*
 25 *great here.*

```
26 #include "art/Framework/Services/??/???.h"
```

1 57.6.1 Declaring and Defining Services

declare:service

2 In general, DECLARE macro invocations should be in the **FIXME:** *in this particular?*
 3 header file, if it exists. DEFINE macro invocations should be in the implementation.
 4 **FIXME:** *Add: DECLARE macros do such and such, whereas DEFINE macros do blah...*

5 A **FIXME:** *standard?* service may have one **FIXME:** *or MUST be declared with one and*
 6 *only one?* of three scope indicators:

7 **LEGACY** All currently defined services are “legacy” in scope, including *art*-defined ser-
 8 vices. **FIXME:** *Then what does it mean to be ‘legacy’?* You should (for now) define
 9 all your services to be “LEGACY”. The presence of even one LEGACY service in
 10 the configuration precludes the simultaneous processing of multiple events (“multi-
 11 schedule” operation).

12 **GLOBAL** Defining a service as “global” in scope gives a guarantee of thread safety: the
 13 service may be called simultaneously from anywhere at any time and continue to
 14 behave correctly.

15 **PER_SCHEDULE** These services may be instantiated multiple times if the *art* frame-
 16 work is configured for multi-schedule operation.

17 For example, this *declaration* would be included in `ServiceMacros.h`:

18 `DECLARE_ART_SERVICE(myexpt::MyService, LEGACY)`

19 whereas this *definition* would be in the implementation file (after the include statement for
 20 the header):

21 `DEFINE_ART_SERVICE(myexpt::MyService)`

22 Next you declare an interface; follow that by declaring then defining a service that will
 23 implement the interface: **FIXME:** *declare in header, define in implementation file?*

24 `DECLARE_ART_SERVICE_INTERFACE(myexpt::Interface, LEGACY)`

25 `DECLARE_ART_SERVICE_INTERFACE_IMPL(myexpt::MyService,`
 26 `myexpt::Interface, LEGACY`

27 `DEFINE_ART_SERVICE_INTERFACE_IMPL(myexpt::MyService,`
 28 `myexpt::Interface)`

29 Next, include the constructor:

```
1 MyService(fhicl::ParameterSet const &, art::ActivityRegistry &);
```

2 Via **FIXME:** *some mechanism that we've undoubtedly set into motion above - please specify,* the service is provided at construction time with access to its configuration parameters

3 **FIXME:** *i.e its fcl file?* and to the `source:art/Framework/Services/Registry/ActivityRegistry` for the purposes of registering for callbacks at any of a large number of synchronization

4 points in the execution of the *art* framework.

5 To register for callbacks, invoke a `watch()` function on the signal data member of

6 `ActivityRegistry` that corresponds to the point in the execution of the *art* framework at which your function is to be called. For instance:

```
7 aReg.sPreBeginRun.watch(&MyService::preBeginRun, *this);
```

8 **FIXME:** *What's aReg and does user know about sPreBeginRun, etc.?*

9 The function you provide must have the signature specified for that signal. **FIXME:** *signature of a physics signal? Can you give an example?* The first template parameter **FIXME:** *Where do I see a 'template parameter'? Has this been defined?* of the signal specifies the calling order of registered callbacks; the second is the return type of the callback function; any subsequent parameters specify the ordered parameters of the callback function. The provided function may be any of:

- 10 ○ a free function, bound or unbound
- 11 ○ a `const` or non-`const` member function of your service or of any other class, bound or (if you provide an object with which to bind it) unbound
- 12 ○ a function whose operator `()` signature matches the required signature

13 A service plugin is compiled into a plugin recognizable to the *art* framework by making a library whose name is `lib[dir1_dir2_]ClassName_service.so`.

14 **FIXME:** *This section should present a set of files that show all the pieces (the header, the .cc, etc.) with all service-related declarations and definitions shown.*

25 57.7 Service Interfaces

26 `<Interface>` is an interface that can point to service `MyService` and `YourService` and `TheirService`. 'something' is a name that a user module can use to access any of these services. Is

- 1 it an implementation? How does that work? What's the syntax for that?
- 2 A service interface defines a base class that is registered to *art* as such, thereby allowing
- 3 a service to inherit from this base class/interface and declare itself to the *art* framework.
- 4 Subsequently, *art* can be configured at runtime via FHiCL to use a given service that
- 5 implements a particular interface. External users of that interface (i.e., other services, or
- 6 on occasion the *art* framework itself) access the given service via the interface without
- 7 ever knowing (or needing to know) exactly which (other) service implements it.
- 8 Art can be configured and/or a user module can be configured to use an interface?
- 9 A service that requires an interface must have a header? Or only if used by user module?
- 10 And the header must include (what in the parens is mandatory?):
- 11 `DECLARE_ART_SERVICE_INTERFACE(myexpt::Interface,LEGACY)`
- 12 **FIXME:** *provide example and maybe diagram*

58 *art* Input and Output

-in-output

58.1 Input Modules

58.1.1 Configuring Input Modules to Read from Files

When reading from an existing file, *art* allows you to select the input files, the starting event, the number of events to read, etc., either from the command line or from the FHiCL file. If a particular quantity is controlled from both the command line and the FHiCL file, the value on the command line takes precedence.

The following code fragment tells *art* to read event data from the file of type “ROOT,” named “file01.art” and to start at the beginning of the file. A value of “-1” for `maxEvents` tells *art* to read events until the end of file is reached:

To tell *art* to read 100 events, or until the end of file, whichever comes first, change the parameter `maxEvents` to 100. This also shows how to specify a list of input files:

The number of files in the list of input files is arbitrary. The following fragment tells *art* to skip the first two events (and thus start with the third):

The fragment below shows some other parameters that can be included in the source pa-

source-sample

Listing 58.1: Reading in a ROOT data file

```
1 source :{
2   module_type : RootInput
3   fileNames   : [ "file01.art" ]
4   maxEvents   : -1
5 }
```

od-source-sample

Listing 58.2: Reading in a ROOT data file

```

1 source : {
2   module_type : RootInput
3   fileNames   : [ "file01.art", "file02.art", "file03.art" ]
4   maxEvents   : 100
5 }

```

od-source-sample

Listing 58.3: Reading in a ROOT data file

```

1 source : {
2   module_type : RootInput
3   fileNames   : [ "file01.art", "file02.art", "file03.art" ]
4   maxEvents   : 100
5   skipEvents  : 2
6 }

```

1 parameter set:

2 The parameters whose names start with *first* specify that the first event to be processed will
 3 be the first event that has an EventID greater than or equal to the specified event. If one of
 4 the *first** parameters is not specified, it takes a default value of -1 and is excluded from the
 5 comparison.

6 If a file of unsorted events is read in, *art* will, by default, present the events for processing
 7 in order of increasing event number. As a corollary to this, the output file will contain the
 8 events in sorted order. This sorting occurs one input file at a time; *art* does not sort across
 9 file boundaries in a list of input files. If the *noEventSort* parameter is set to *true*, the sorting
 10 is disabled, which will, in most cases, yield a minor performance improvement.

11 **FIXME:** I have not yet learned the precise meaning of the *skipBadFiles* and the *fileMatch-*

od-source-sample

Listing 58.4: Reading in a ROOT data file

```

1   firstRun           : 0
2   firstSubRun       : 0
3   firstEvent        : 0
4   noEventSort       : false
5   skipBadFiles      : false
6   fileMatchMode     : "permissive"
7   inputCommands     : ""

```


source-sample

Listing 58.5: Reading in a ROOT data file

```
1 source :{  
2   module_type : EmptyEvent  
3   maxEvents   : 200  
4 }
```

1 Mode parameters.

2 The `inputCommands` parameter tells *art* to delete certain data products from the copy of the event in memory after reading the input file. In other words, the input file itself is
3 not modified but data products are removed from the copy of the event in memory before
4 any modules are called. The syntax of this language is the same as for `outputCommands`,
5 described **FIXME:** *ref*.

7 **FIXME:** ? In the pre-*art* versions of the framework, there were methods to select ranges
8 of events or ranges of `SubRuns`. This is not yet working in *art*; the *art* developers will add
9 this feature back once we decided exactly what we mean by "ranges of events".

10 **FIXME:** *I don't think this is needed (AH)* Specifying Many Input Files In the pre-*art*,
11 python based, configuration language, the standard syntax to initialize a list of input files
12 was limited to 255 files, after which an alternate syntax was required. This is no longer
13 necessary; the length of a `fhicl` list is limited only by available memory.

14 Empty Source

15 In many simulation applications one wishes to start with an empty event, run one or more
16 event generators, pass the generated particles through the `Geant4`, and so on. In *art* the
17 first step in this chain is accomplished using a source module named `EmptySource`, as
18 follows:

19 Instead of reading event-data from a file, the empty source increments the event number
20 and presents an empty event to the modules that will do the work. One may configure
21 `EmptySource` to specify the `EventId` of the first event, to specify the maximum number of
22 events in a `SubRun` or `SubRuns` in a run.

23 The last option tells *art* to reset event numbers to start at 1 whenever *art* starts a new
24 `SubRun` begins; this is the default behaviour and is opposite to the behaviour we inherited
25 from `CMS`.

od-source-sample

Listing 58.6: Reading in a ROOT data file

```

1 source :{
2   module_type      : EmptyEvent
3   firstRun        : 2
4   firstSubRun     : 1
5   firstEvent      : 1
6   numberEventsInRun : 1000
7   numberEventsInSubRun : 100
8   maxEvents       : 200
9   resetEventOnSubRun : true
10 }

```

1 58.2 Output Filtering

2 **FIXME:** *From oink tutorial 'output filtering' slide*

3 Any output module can be configured to write out only those events passing a given trigger
4 path.

5 The parameter set that configures the output module uses a parameter `SelectEvents` to
6 control the output, as shown in the example below:

```

7   # this is only a fragment of a full configuration ...
8   physics:
9     {
10      pathA: [ ... ] # producers and filters are put in this path
11      pathB: [ ... ] # other producers, other filters are put in this
12
13      outA: [ passWriter ] # output modules and analyzers are put in t
14      outB: [ failWriter ] # output modules and analyzers are put in t
15      outC: [ exceptWriter ] # output modules and analyzers are put in
16
17      trigger_paths: [ pathA, pathB ] # declare that these are "trigge
18      end_paths: [ outA outB outC ] # declare these are "end paths"
19    }
20
21   outputs:
22   {

```

```
1     passWriter:
2     {
3         module_type: RootOutput
4         fileName: "pathA_passes.art"
5         # Write all the events for which pathA ended with 'true' from filter
6         # Events which caused an exception throw will not be written.
7         SelectEvents: { SelectEvents: [ "pathA&noexception" ] }
8     }
9     failWriter:
10    {
11        module_type: RootOutput
12        fileName: "pathA_failures.art"
13        # Write all the events for which pathA ended with 'false' from filter
14        # Events which caused an exception throw will not be written.
15        SelectEvents: { SelectEvents: [ "!pathA&noexception" ] }
16    }
17    exceptWriter:
18    {
19        module_type: RootOutput
20        fileName: "pathA_exceptions.art"
21        # Write all the events for which pathA or pathB ended because an exception
22        SelectEvents: { SelectEvents: [ "exception@pathA", "exception@pathB" ] }
23    }
24 }
```

25 58.3 Configuring Output Modules

1 59 *art* Misc Topics that Will Find Home

2 59.0.1 The Bookkeeping Structure and Event Sequencing Imposed by *art*

3
4 In almost all HEP experiments, the core idea underlying all bookkeeping is the *event*. In a
5 triggered experiment, an event is defined as all of the information associated with a single
6 trigger; in an untriggered spill-oriented experiment, an event is defined as all of the in-
7 formation associated with a single spill of the beam from the accelerator. Another way of
8 saying this is that an event contains all of the information associated with some time inter-
9 val, but the precise definition of the time interval changes from one experiment to another.
10 Typically these time intervals are a few nano-seconds to a few tens of mirco-seconds. The
11 information within an event includes both the raw data read from the Data Acquisition
12 System (DAQ) and all information that is derived from that raw data by the reconstruc-
13 tion and analysis algorithms. An event is smallest unit of data that *art* can process at one
14 time.

15 In a typical HEP experiment, the trigger or DAQ system assigns an event identifier (event
16 ID) to each event; this ID uniquely identifies each event. The simplest event ID is a mono-
17 tonically increasing integer. A more common practice is to define a multi-part ID.

18 *art* has chosen to use a three-part ID. In *art*, the parts are named

- 19 ○ run number
- 20 ○ subRun number
- 21 ○ event number

22 In a typical experiment the event number will be incremented every event. When some
23 condition occurs, the event number will be reset to 1 and the subRun number will be

1 incremented, keeping the run number unchanged. This cycle will repeat until some other
2 condition occurs, at which time the event number will be reset to 1, the subRun number
3 will be reset to 0 and the run number will be incremented.



4 *art* does not define what conditions cause these transitions; those decisions are left to each
5 experiment. Typically, experiments will choose to start new runs or new subRuns when
6 any of the following happen:

- 7 ○ a preset number of events have been acquired
- 8 ○ a preset time interval has expired
- 9 ○ a disk file holding the output has reached a preset size
- 10 ○ certain running conditions change

11 *art* requires only that a subRun contain zero or more events and that a run contain zero or
12 more subRuns.

13 As runs are collections of subRuns, and subRuns are collections of events, events in turn
14 are collections of *data products*. A data product is the smallest unit of data that can be
15 added to or retrieved from a given event. Each experiment defines types (classes and
16 structs) for its own data products. These include types that describe the raw data, and types
17 to define the reconstructed data and the information produced by simulations. *art* knows
18 nothing about the internals of any experiment's data products; for *art*, the data product is
19 a “fundamental particle.”

20 At the outside shell of the Russian doll that is the bookkeeping structure in *art*, runs are
21 collected into the *event-data*, defined as all of the data products in an experiment's files;
22 plus the metadata that accompanies them.

23 When an experiment takes data, events read from Data Acquisition System (DAQ) are
24 typically written to disk files, with copies made on tape. *art* imposes only weak constraints
25 on the event sequence within a file. The events in a single subRun may be spread over
26 several files; conversely a single file may contain many runs, each of which contains many
27 subRuns.

28 A critical feature of *art*'s design is that each event must be uniquely identifiable by its event
29 ID. This requirement also applies to simulated events.

59.1 Rules for Module Names

Within any experiment's software, sometimes names of files, classes, libraries, etc., must follow certain rules. Other times, conventions are just conventions. This section is concerned with actual rules only.

Consider a class named `MyClass` that you wish to make into an *art* module. First, your class must inherit from one of the module base classes, `EDAnalyzer`, `EDProducer` or `EDFilter`. Secondly, it must obey the following rules, all of which are case-sensitive.

1. it must be in a file named `MyClass_module.cc`
The build system will make this into a file named `lib/libMyClass_module.so`.
2. the module source file must look like Listing 59.1 (where your experiment's namespace replaces `xxxx`):

This example is for an analyzer. To create a producer or a filter module, you must inherit from either `art::EDProducer` or `art::EDFilter`, respectively. The last line (`DEFINE_ART_MODULE`) invokes a macro that inserts additional code into the `.so` file.



For the experts: it inserts a factory method to produce an instance of the class and it inserts an auto-registration object that registers the factory method with *art*'s module registry.

To declare this module to the framework you need to have a fragment like the following in your FHiCL file:

```

1
2     physics :
3     {
4         analyzers:
5         {
6             looseCuts : { module_type : MyClass }
7
8             // Other analyzer modules listed here ...
9         }
10    }
```

where the string `looseCuts` is called a *module label* and is defined below. **FIXME:**

3. the previous item was for the case that your module is an analyzer. If it is a producer or filter, then the label *analyzers* needs to be either *producers* or *filters*.

rce-sample


Listing 59.1: Module source sample

```
1 namespace xxxx {
2
3     class MyClass : public art::EDAnalyzer {
4
5     public:
6         explicit MyClass(fhicl::ParameterSet const& pset);
7         // Compiler generated destructor is OK.
8
9         void analyze( art::Event const& event );
10
11     };
12
13     MyClass::MyClass(fhicl::ParameterSet const& pset){
14         // Body of the constructor. You can access information
15         in the parameter set here.
16     }
17
18     void MyClass::analyze(art::Event const& event){
19         mf::LogVerbosim("test")
20         << "Hello, _world._From_analyze._"
21         << event.id();
22     }
23
24 } // end namespace xxxx
25
26 using xxxx::MyClass;
27 DEFINE_ART_MODULE(MyClass);
```

- 1 4. When you put a data product into an event, the data provenance system records the
2 module label of the module that did the “put.”

3 59.2 Data Products and the Event Data Model

4 The part of *art* that deals with the bookkeeping of the data products is called the *Event*
5 *Data Model*, which concerns itself with the following ideas:

- 6 1. what a data product looks like when it is in the memory of a running program
- 7 2. what it looks like on disk
- 8 3. how it moves between memory and disk
- 9 4. how a data product refers to another piece of event-data within the same event
- 10 5. how a given piece of experiment code accesses a data product
- 11 6. how the experiment code adds a new data product to the event
- 12 7. metadata that describes, for each data product,
 - 13 ○ what piece of code was used to create it
 - 14 ○ what is the run-time configuration of that code
 - 15 ○ what data products were read in by this experiment code
- 16  8. The mechanism by which the metadata is “married” to the data

17 One of the core principles of *art* is that experiment code modules may communicate with
18 each other only via the event.

19 59.3 Basic *art* Rules

20 *art* prescribes that your classes (i.e., your *art modules*) always contain a *member function*
21 that has a particular name, takes a particular set of arguments, and operates on every event;
22 *art* will call this member function for every event read from the data source (input). If no

1 member function with these attributes exists, then at execution time *art* will print an error
2 message and stop execution.¹

3 If your module provides any optional functions, then *art* requires a name and a set of
4 arguments for each. For each of these that is present in a given class, *art* will make sure
5 that it is called at the right time.

6 The details of the *art* rules will be discussed in **FIXME**: *section ??*.

7 **59.4 Compiling, Linking, Loading and Executing C++** 8 **Classes and *art* Modules**

9 When you write code to be executed by *art*, you provide it to *art* as a group of C++
10 functions. To make this group of functions visible to *art*, you write a C++ class that obeys
11 a set of rules defined by *art* (summarized in Section ??). Such a class is called an *art*
12 *module*, or just *module* in this documentation (this should not be confused with the notion
13 of a *module* as defined more generally in the programming world). The container source
14 code file for an *art* module gets compiled into a dynamic library that can be loaded at
15 run-time by *art*.

16 The *experiment's dynamic code libraries* in Figures ?? and ?? may include libraries con-
17 taining standard C++ classes as well as *art* modules.

18 Experiments typically have many, many C++ classes for offline processing, and physicists
19 add to them all the time. Classes from many files can be linked into a single library, as
20 shown in Figure 59.1. The dynamic libraries may have one-way dependencies on each
21 other; i.e. if library 'a' depends on library 'b', then the reverse cannot be true.

22 *art* modules, as mentioned above, follow a special structure, illustrated in Figure 59.2.
23 They do not use header (.h) files (everything for a module is contained within a single .cc
24 file), a single module builds a single dynamic library, and the name (as recognized by *art*)
25 for each file in the build chain must end in `_module`, e.g., `MyCoolMod_module.cc`.
26 Moreover, *art* recognizes `MyCoolMod_module.cc` as the source for `libxxx_MyCoolMod_module.so`.
27 (Discussion of the *xxx* will be deferred.) **FIXME**: *I should put it in UG and refer to*
28 *it*.

¹Actually the loader that loads the dynamic library, rather than *art* itself, will figure this out.

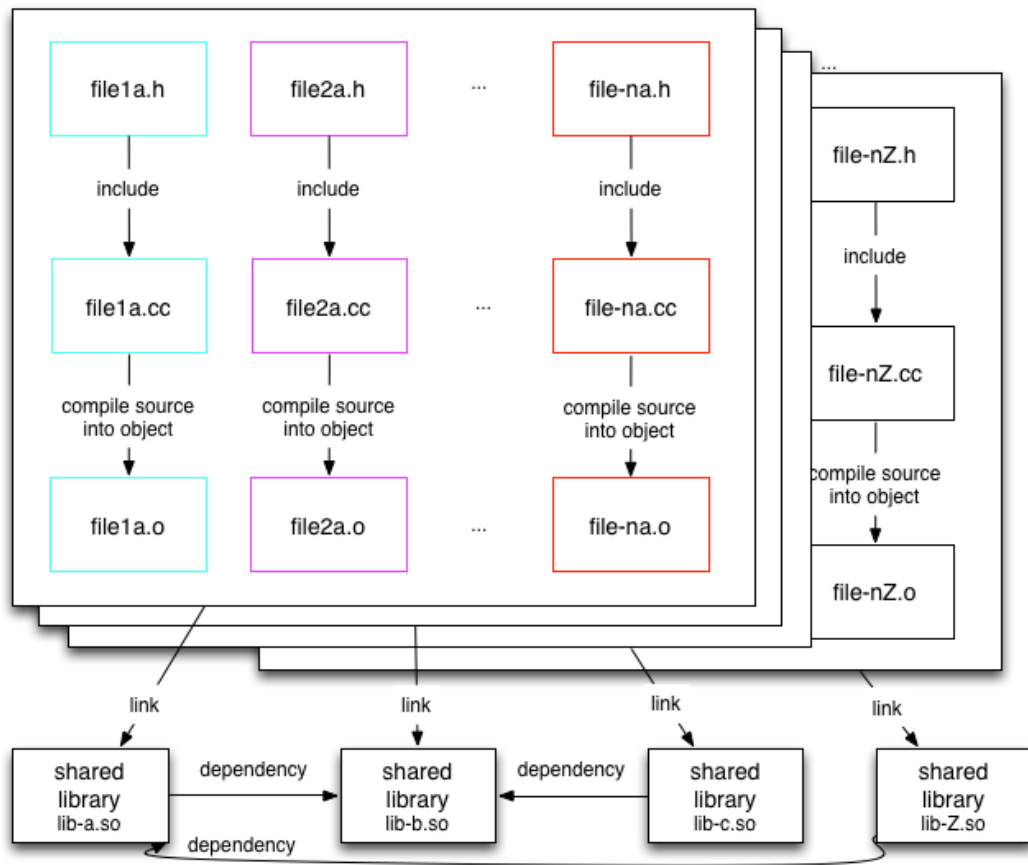


Figure 59.1: Illustration of compiled, linked “regular” C++ classes (not *art* modules) that can be used within the *art* framework. Many classes can be linked into a single dynamic library.

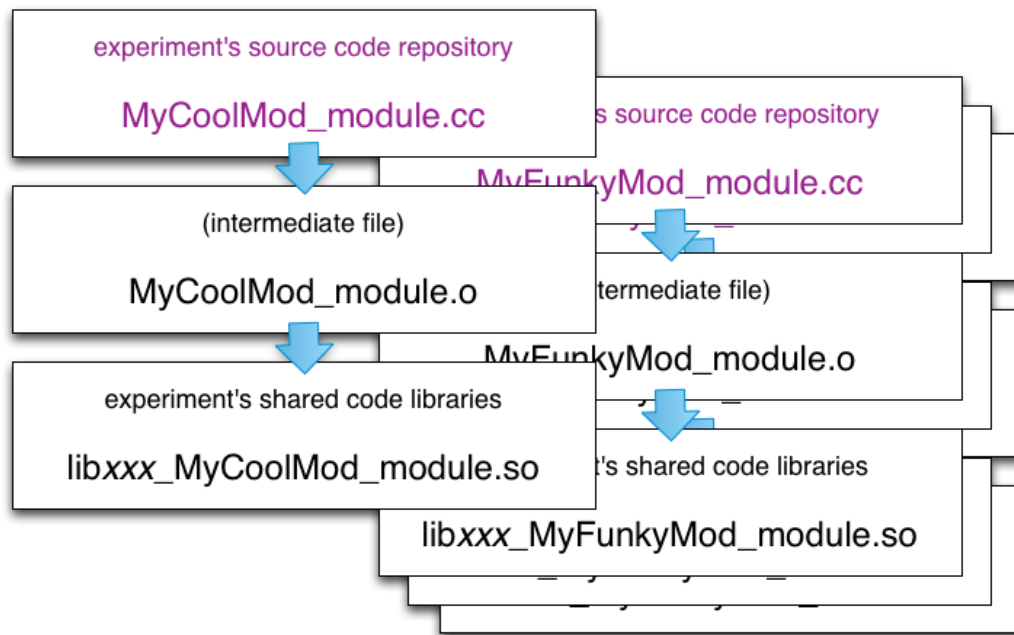


Figure 59.2: Illustration of compiled, linked *art* modules; each module is built into a single dynamic library for use by *art*

59.5 Dynamic Libraries and *art*

When you execute code within the *art* framework, the main executable is provided by *art*, not by your experiment. Your experiment provides its code to the executable in the form of dynamic libraries that *art* loads at run time; these libraries are also called *dynamic load libraries*, *shared libraries* or *plugins*.

Your experiment will likely have many “regular” C++ classes (as distinct from the C++ classes that are modules, aka “*art* modules”). These “regular” classes get built into a set of dynamic libraries, where each library contains object code for multiple classes.

Your experiment will likely have many modules, too. In fact you will likely be writing some for your own analyses. A module must be compiled into its own dynamic library, i.e., there is a one-to-one correspondance between the `.cc` file and the `.so` file for a given module. When the configuration file tells *art* to run a particular module, *art* finds the corresponding `.so` file, loads it, and calls the appropriate member function at each stage of the event loop.

FIXME: *Would be nice to say why some things can be regular classes and others should be modules.*

59.6 Namespaces, *art* and the Workbook

A *namespace* is a prefix that is used to keep different subsets of code distinguishable from one another; i.e., if the same identifier (variable name or type name) is used within multiple namespaces, each will remain distinguishable via its namespace prefix. The otherwise ambiguous identifier should be written as

```
<namespace> :: <identifier>
```

The notion of *namespace* is related to that of *scope*: Within a C++ source file (`.cc` files) a scope is designated by a set of curly braces (`{ ... }`). Once a namespace is defined within a given scope, any identifiers within that scope that “belong to” that namespace no longer need to be written with the prefix. E.g., the following fragment uses the `analyze` defined in the namespace `tex` (i.e., `tex :: analyze`):

```
namespace tex {
```

```
1     class First : public art :: EDAnalyzer {
2     public :
3         explicit First ( fhicl :: ParameterSet const & );
4         void analyze ( art :: Event const & event ) override ;
5     };
6 }
```

7 Note that `EDAnalyzer` is defined in the namespace `art`, as is `Event`, and `ParameterSet`
8 is in `fhicl`.

9 Note also that namespaces are often associated with UPS product code, although the prod-
10 uct and the namespace names may not always be identical. E.g., code associated with the
11 UPS product *fhiclcpp* is in the namespace `fhicl`.

12 All of the code in the `toyExperiment` UPS product was written in a namespace named `tex`;
13 the name `tex` is an acronym-like shorthand for the `toyExperiment` (ToyEXperiment) UPS
14 product. Because all of the `Workbook` code builds on top of the `toyExperiment` code, this
15 code has been placed in the same namespace. The `tex` namespace has no special meaning
16 to *art*, it is just a convenience. (Note that the *art* code itself is in a separate namespace
17 called `art`.)

18 If you need more information about the C++ notion of *namespaces*, see a standard C++
19 reference.

20 59.7 Orphans **FIXME:**

21 A best practice: define ids in the narrowest scope possible to avoid accidental name colli-
22 sions

23 **FIXME:** *where to put the idea that you can add information to an event during processing*
24 During processing, derived information in the event may be changed, added to or deleted;
25 the raw data is not modified. The *event* is the smallest unit of data that *art* can process at
26 one time. **FIXME:** *format 'art'*

27 How bash shell scripts work

28 If you would like to understand how they work, the following will be useful:

- 1 ○ BASH Programming - Introduction HOW-TO
- 2 <http://tldp.org/HOWTO/Bash-Prog-Intro-HOWTO.html>
- 3 ○ Bash Guide for Beginners
- 4 <http://www.tldp.org/LDP/Bash-Beginners-Guide/html/Bash-Beginners-Guide.html>
- 5 ○ Advanced Bash Scripting Guide
- 6 <http://www.tldp.org/LDP/abs/html/abs-guide.html>
- 7 The first of these is a compact introduction and the second is a more comprehensive intro-
- 8 duction.
- 9 The above guides were all found at the Linux Documentation Project: Workbook:
- 10 ○ <http://www.tldp.org/guides.html>

11 59.8 Inheritance

12 59.8.1 Introduction

13 This section introduces a few of the ideas behind *inheritance* and *polymorphism*. There
14 are many, many different ways to use *inheritance* and *polymorphism* but you only need
15 to understand the small subset that are relevant for the Workbook exercises. You can read
16 about inheritance and polymorphism at the following url:

17 <http://www.cplusplus.com/doc/tutorial/inheritance/>

18 Skip the section on Friendship and start at the section on inheritance. When you get to
19 the bottom of the page, continue to the next page by clicking on the arrow for “Polymor-
20 phism”. You can skip the discussion of protected and private inheritance because you will
21 only need to know about public inheritance.

22 After you have learned this material, return to this section and work through the following
23 example which serves as a test that you have learned the necessary material. This example
24 is motivated by the Polygon example given in the referenced material. In this example
25 there is a base class named `Shape` and three derived classes, `Circle`, `Triangle` and
26 `Rectangle`. The main program that exercises these four classes is `itest.cc`.

1 59.8.2 Homework

e-homework

2 To build and run this example:

3 1. log in and follow the follow the steps in Section [ssec:quick-start-env-second](#)??

4 2. cd to the directory for this exercise

5 \$ cd Inheritance/v1

6 \$ ls

7 build Circle.h Rectangle.cc Shape.cc Triangle.cc

8 Circle.cc itest.cc Rectangle.h Shape.h Triangle.h

9 3. build the exercise

10 \$../build

11 This will create the executable file itest

12 4. run the exercise

13 \$ itest

14 Area of circle c1 is: 3.14159

15 Area of circle c2 is: 12.5664

16 Area of rectangle r1 is: 4

17 Area of triangle t1 is: 0.5

18 Area of triangle t2 is: 2

19 This circle has an area of 3.14159 and a color of undefined

20 This circle has an area of 12.5664 and a color of red

21 Unknown shape has color: blue

22 This triangle has an area of 0.5 and a color of green

23 This triangle has an area of 2 and a color of yellow

24

25 When you run the code, all of the printout should match the above printout exactly.

26 Read the code in the example and apply what you learned from the [cplusplus.com](#)
27 website. Understand why the example prints out what it does.

28 The next subsection contains some discussion about the example. In particular it will dis-
29 cuss the *explicit* and *override* identifiers.

59.8.3 Discussion

The heart of this example is the base class `Shape`, found in `Shape.h` and `Shape.cc`. This class illustrates the following ideas:

1. it has a data member named `color_`, which describes an attribute that is common to all shapes. This data member is `protected` so it is visible to derived classes.
2. the two constructors guarantee that the `color_` data member will be initialized whenever a derived class is instantiated.
3. the class has two virtual functions, one of which is pure virtual. Therefore you cannot instantiate an object of type `Shape`.
4. the class provides an implementation for the virtual method `print`.
5. the class provides an accessor for `color_`.

The one argument constructor of `Shape` is declared `explicit`. Since `Shape` cannot be constructed, we will use an imaginary class named `T` to illustrate.

The derived class `Circle`:

1. has one data member, the radius of the circle.

59.9 Inheritance Relic

FIXME: *Rob moved it to the end since it is now a relic - need to discuss before we delete*
FIXME: *Anne moved from ch 8; needs work – probably change example to 'Point' if nothing else*

The first line of the class `First`'s declaration is:

```
class First : public art::EDAnalyzer {
```

The fragment `(: public art::EDAnalyzer)` tells the C++ compiler that the class `First` *inherits* from the class named `art::EDAnalyzer` via *public inheritance*^{c0}. “In-

^{c0}Inheritance can be either *public* or *private*; the Workbook exercises always use public inheritance.

1 inheritance is a way of creating new classes which extend the facilities of existing classes by
2 including new data and functions. The class which is extended is known as the *base class*
3 and the result of an extension is known as the *derived class*; the derived class inherits the
4 data and function members of the base class^{c0}.” In the current example `art::EDAnalyzer`
5 is a base class and `First` is a derived class.

6 The idea of *inheritance* is a very powerful feature of C++ that has many uses, only a few of
7 which are relevant for *art* modules. This discussion should help you focus on the relevant
8 information if you need to consult C++ references on inheritance.

9 **59.10 Pointers**

10 **FIXME:** *Taken out of glossary – TMI there*

11 C++, like many other computer languages, allows you to define variables that are pointers
12 to information held in other variables. The value of a pointer is the memory address of the
13 information held by the given variable. A native C++ pointer is often referred to as a *bare*
14 *pointer*. While pointers provide great flexibility for producing fast, efficient algorithms,
15 they are also easy to misuse. *art* has been designed so that user code will rarely, if ever,
16 interact with *art* via bare pointers; when pointer-like behaviour is required, *art* will provide
17 that information inside a wrapper that is generically referred to as a *smart pointer* or a *safe*
18 *pointer*; *art* defines different sorts of smart pointers for use in different circumstances. The
19 job of a smart pointer is to recognize misuse and to protect against it. One commonly used
20 type of smart pointer is called a *handle*.

21 **59.11 RootOutput and table of event IDs**

22 When `RootOutput` writes a file, it writes the event information to the file and it also writes
23 a table of event IDs that allows it to random access a single event without needing to read
24 all of the events before it. This table is kept in order of increasing event id. When you open
25 a file and read it, `RootInput` starts reads events in the order found in the table.

^{c0}D.M. Capper’s *Introducing C++ for Scientists, Engineers and Mathematicians*, Springer-Verlag Limited 1994, Chapter 11

1 **59.12 Troubleshooting**

2 (Section ^{[|sec:setup-command](#)}7.3) `setup` returns the error message

3 You are attempting to run ```setup``` which requires administrative
4 privileges, but more information is needed in order to do so.

5 The simplest solution is to log out and log in again.

1

Part IV

2

Appendices

DRAFT

1 A Obtaining Credentials to Access Fermilab 2 Computing Resources

kerb-at-fermilab

3 To request your Fermilab computing account(s) and permissions to log into the your exper-
4 iment's nodes, fill out the form Request for Fermilab Visitor ID and Computer Accounts.
5 Typically, experimenters that are not Fermilab employees are considered *visitors*. You will
6 be required to read the Fermilab Policy on Computing.

7 After you submit the form, an email from the Fermilab Service Desk should arrive within
8 a week (usually more quickly), saying that your Visitor ID (an identifying number), Ker-
9 beros Principal and Services Account have been created. You will need to change the
10 password for both Kerberos and Services.

11 A.1 Kerberos Authentication

12 Your Kerberos Principal is effectively a username for accessing nodes that run Kerberos
13 in what's called the FNAL.GOV *realm* (all non-PC Fermilab machines).^{c0}

14 To change your Kerberos password, first choose one (minimum 10 characters with mixture
15 of upper/lower case letters and numbers and/or symbols such as !, , #, \$, %, &, *, %). From
16 your local machine, log into the **FIXME:** *appropriate* machine using `ssh` or `slogin`
17 and run the `kpasswd` command. Respond to the prompts, as follows:

```
18 $ kpasswd <username>@FNAL.GOV
```

19

```
20 Password for username@FNAL.GOV: <--- type your current password he
```

21

^{c0}The FERMI.WIN.FNAL.GOV realm is available for PCs.

```
1     New password:                <--- type your new password here
2
3     New password (again):        <--- type your new password here for c
4
5     Kerberos password changed.
6 Your Kerberos password will remain valid for 400 days.
```

7 **A.2 Fermilab Services Account**

8 The Services Account enables you to access a number of important applications at Fermi-
9 lab with a single username/password (distinct from your Kerberos username/password).
10 Applications available via the Services Account include SharePoint, Redmine, Service
11 Desk, VPN and others.

12 To get your initial Services Account password, a user must first contact the Service Desk
13 to get issued a first time default password. Once a default password is issued, users can
14 access <http://password-reset.fnal.gov/> to change it.

15 If you are not on-site or connected to the Fermi VPN, call the Service Desk at 630-840-
16 2345. You will be given a one-time password and a link to change it.

1 B Installing Locally

2 This appendix describes how to install all of the software needed to run the *art* workbook
3 on your own laptop, desktop or on your institution's computing resources.

4 The *art* team provides downloadable binary distributions of the required packages. A cheat
5 sheet about downloading just what you need to run the *art* workbook is given in Sec-
6 tion B.1. Links to the full instructions are in Section B.3.

7 The *art* team also provides instructions on how to build binaries from source:
8 https://cdcv.s.fnal.gov/redmine/projects/cet-is-public/wiki/Build_packages_required_by_art
9 No additional information on building from source will be provided in this document.

10 B.1 Install the Binary Distributions: A Cheat Sheet

11 At this writing, binary distributions are available for:

- 12 1. Scientific Linux Fermi, versions 5 (SLF5) and 6 (SLF6).
 - 13 ○ Experience has shown that the SLF binaries work on installations of Scientific
14 Linux CERN or vanilla Scientific Linux.
- 15 2. Mac OSX, versions Mavericks (10.9) and Yosemite (10.10).
 - 16 ○ Experience has shown that the Mac OSX Mavericks binaries work on OSX10.8
17 (Mountain Lion) and some earlier versions of OSX.

18 The binary files are distributed in the format of relocatable UPS packages ^{ch:ups:setup}7. The first step
19 is to choose a directory into which the UPS packages will be installed; the full path to this
20 directory is arbitrary; the directory name is arbitrary but it is traditionally chosen to be

1 products.^{c0} In the following, the full path to the products directory will be denoted by
2 <products>.

3 The following procedure illustrates how to download the binaries for toyExperiment plus
4 all of the binaries for the products on which toyExperiment depends. The example is
5 given for SLF6 for non-neutrino experiments. Other options are discussed after the proce-
6 dure:

1. Create a products directory and a temporary directory:

```
mkdir -P <products>
```

```
cd <products>
```

```
mkdir ../tmp
```

```
cd ../tmp
```

The disk must have at least 10 GB of free space.

2. Download the script that will do the work; make it executable:

```
curl -O http://scisoft.fnal.gov/scisoft/bundles/tools/pullProducts
```

```
chmod u+x pullProducts
```

3. If curl is not available on your computer, you can try:

```
wget http://scisoft.fnal.gov/scisoft/bundles/tools/pullProducts
```

```
chmod u+x pullProducts
```

4. Download the binaries:

```
./pullProducts <products> slf6 toyExperiment-v0_00_29 s14-e7 prof
```

```
./pullProducts <products> slf6 toyExperiment-v0_00_29 s14-e7 debug
```

5. Cleanup:

```
cd ..
```

```
rm -rf tmp
```

6. Options:

- (a) If you are working on one of the neutrino experiments, use `s14-e7-nu` instead of `s14-e7`.
- (b) Instead of `slf6` you may choose one of:
 - o `slf5`
 - o `d13` - OSX 10.9 (Mavericks); known to work for some earlier versions of OSX.

^{c0} For historical reasons, Mu2e has traditionally chosen `artexternals` instead of `prodcuts`.

- d14 - OSX 10.10 (Yosemite)
- (c) You only need to download one of prof or debug; you may download both.
- (d) If you do download both, `pullProducts` will not repeat unnecessary work.

1

2

B.2 Preparing the Site Specific Setup Script

e:specific:setup

3

This section describes how to prepare a bash script that serves as the site specific setup procedure. Everyone who wishes to use the *art* workbook will need to source this at the start of each login session. You may put the script in any directory that you want; there are three common choices:

6

7

- At the top level of the products directory

8

- As a peer to the products directory

9

- In the home directory of an account named after your experiment

The name of the file is arbitrary; for purposes of this example we will call it `setup-site.sh`.

The listing below shows the required contents of this file:

```
1 source <products>/setup
2 export ART_WORKBOOK_OUTPUT_BASE=<path to the user area on a data disk>
3 export ART_WORKBOOK_WORKING_BASE=<path to the user area on a code disk>
4 export ART_WORKBOOK_QUAL=<qualifiers>
```

You should also ensure that this script is not executable:

```
17 chmod -x setup-site.sh
```

This is because users must source, not execute, this script.

The first line of the file initializes the UPS system for the current login session. Be sure to change `<products>` to the appropriate path for your site.

The art workbook instructions presume that the administrators of site may have policies such as: source code and binaries belong on one disk while event-data files, log files and root files belong on another disk. You should define

```
24 ART_WORKBOOK_WORKING_BASE
```


1 so that the directory
2 `$ART_WORKBOOK_WORKING_BASE` /<username>
3 is a directory in which the user specified by <username> can work. It should point to a
4 disk on which the user may put source and binary files.

5 You should define
6 `ART_WORKBOOK_OUTPUT_BASE`
7 so that the directory
8 `$ART_WORKBOOK_OUTPUT_BASE` /<username>
9 is a directory in which the user specified by <username> can work. It should point to a
10 disk on which the user may put event-data, root and log files.

11 For example, on Mu2e machines at Fermilab, the values of these variables are:

```
12 export ART_WORKBOOK_OUTPUT_BASE=/mu2e/data/users  
13 export ART_WORKBOOK_WORKING_BASE=/mu2e/app/users
```

14 It is OK if both environment variables point at the same directory; everything will still
15 work correctly.

16 If constraints from your site policies do not allow you to define these two environment
17 variables then you should tell users to follow self managed procedures that are presented
18 as options in Sections 9.6.1.2 and 10.4.3.2.

19 The environment variable `ART_WORKBOOK_QUAL` should be set to the value of the
20 UPS options string that is required to make the following command work properly:

```
21 setup toyExperiment -v v0_00_29 -q${ART_WORKBOOK_QUAL}:prof
```

22 Note that the prof/debug option is NOT part of `ART_WORKBOOK_QUAL` .

23 In the example of downloading binaries earlier in this session, the option string was
24 `s14-e7` . Therefore the environment variable should be set to:

```
25 export ART_WORKBOOK_QUAL=s14:e7
```

26 Note the change in punctuation.

27 You may wish to add additional local customizations to this script. For example if you get
28 git from your UPS products area you may wish to setup git. In this file you should not
29 setup versions of toyExperiment, *art* or any of the *art* tool chain.

1 **B.3 Links to the Full Instructions**

call:local:links

2 The general instructions for downloading binaries are at:

3 https://cdcvns.fnal.gov/redmine/projects/cet-is-public/wiki/Get_binary_distributions

4 Instructions that are specific to this version of toyExperiment are at: The full instructions
5 are available at:

6 http://scisoft.fnal.gov/scisoft/bundles/toyExperiment/v0_00_29/toyExperiment-v0_00_29.html

7

1 C *art* Completion Codes

tion:codes

2 When *art* completes it prints a message to the log file:

```
3 Art has completed and will exit with status 0.
```

4 The number at the end of this line is also the exit status that *art* returns to the shell that
5 issued the *art* command.

6 Inside the *art* program, the status code is represented by a 4-byte integer. However only
7 the least significant byte is returned to the caller; the other bytes are simply discarded. It
8 is done this way because unix specifies that return codes are only one byte in size. For
9 example, if the last line of the *art* printout is an error code of 8001, the status code seen by
10 the calling shell will be 65.

11 If you ran *art* interactively you can see the return code by typing the following as the first
12 command after running *art*:

```
13 echo $?
```

14 Here the symbol `?` is a bash variable that holds the status code returned by the previously
15 executed command. If you write a shell script that runs *art*, then you can test *art*'s return
16 code using `?`.

completion:codes

Listing C.1: Error Codes when an `art::Exception` is caught in main

```
1  enum ErrorCodes {
2      OtherArt = 1,
3      StdException,
4      Unknown,
5      BadAlloc,
6      BadExceptionType,           // =5
7      ProductNotFound,
8      DictionaryNotFound,
9      InsertFailure,
10     Configuration,
11     LogicError,                 // =10
12     UnimplementedFeature,
13     InvalidReference,
14     TypeConversion,
15     NullPointerError,
16     EventTimeout,              // =15
17     DataCorruption,
18     ScheduleExecutionFailure,
19     EventProcessorFailure,
20     EndJobFailure,
21     FileOpenError,             // =20
22     FileReadError,
23     FatalRootError,
24     MismatchedInputFiles,
25     CatalogServiceError,
26     ProductDoesNotSupportViews, // =25
27     ProductDoesNotSupportPtr,
28     SQLExecutionError,
29     InvalidNumber,
30     NotFound                    // =29
31 };
```

Table C.1: *art* completion status codes. The return code is the least significant byte of the status code.

Status	Return Value	Description
7000	88	Exception from command line processing
7001	89	Check command line options failed
7002	90	Process command line options failed
7003	91	Failed to create a parameter set from parsed configuration with exception
8001	65	a <code>cet::exception</code> was caught and processed in main
8002	66	an <code>std::exception</code> was caught and processed in main
8003	67	an unknown exception was caught and processed in main (âĀĖ)
8004	68	an <code>std::bad_alloc</code> was caught in main

etion:codes

1 D Viewing and Printing Figure Files

viewing:printing

2 In many of the exercises in the Workbook you will produce files that contain figures. These
3 files will be in formats such as .pdf, .png, .jpg, and so on.



4 For those of you who are working on the Fermilab General Purpose Computing Facility
5 (GPCF), this chapter has instructions on how to view these files interactively and how to
6 print them.

7 If you are logged in to a Fermilab machine from offsite, viewing figures interactively may
8 be too slow to be practical. It depends on the quality of the network connection between
9 Fermilab and your site; and it depends on the details of the file you are viewing. In such
10 cases, the best alternative is to copy the file to a local machine and view it using the tools
11 available there.

12 **FIXME:** *Any caveat to add about making sure the tools work on the GPCF nodes?*

13 D.1 Viewing Figure Files Interactively

14 On a GPCF machine there are two interactive commands that will allow you to view a
15 figure file:

```
16 kpdf file.type  
17 display file.type
```

18 Despite its name, `kpdf` works on more than just .pdf files; it works on most types of
19 graphics formats, as does `display`.

20 `kpdf` has a more intuitive interface for navigating multipage files: it has a left-hand side-
21 bar with page thumbnails. To navigate a multipage file using `display`:

- 1 1. Click anywhere in the image; this will pop up a menu.
- 2 2. On this menu click on “File”; this will pop up another menu.
- 3 3. On this menu click on “Next” or “Former” to move forward and backward through
- 4 the pages.

5 Another interactive command that is present on some Unix machines is,

6

```
7 acroread file.pdf
```

8 Some people prefer `acroread`'s rendering of postscript and PDF files to that of the other
9 tools.

10 **FIXME:** *Add links to documentation for `kpdf`, `display` and `acroread`.*

11 D.2 Printing Figure Files

12 *A future version this document will explain how to print from the GPCF machines.*

13 **FIXME:** *Add the missing material*

14 On most Fermilab machines PDF files can be printed from the command line or from
15 a browser. Other graphics formats need to be viewed in a browser and printed from the
16 browser or saved as PDF and then printed from the command line.

17 General information about printing at Fermilab is available at fermiprint.fnal.gov.



1 E CLHEP

ch:ug:CLHEP

2 E.1 Introduction

3 The wikipedia entry for CLHEP, <http://en.wikipedia.org/wiki/CLHEP>, describes it as:

4 CLHEP (short for A Class Library for High Energy Physics) is a C++ li-
5 brary that provides utility classes for general numerical programming, vector
6 arithmetic, geometry, pseudorandom number generation, and linear algebra,
7 specifically targeted for high energy physics simulation and analysis software.
8 The project is hosted by CERN and currently managed by a collaboration of
9 researchers from CERN and other physics research laboratories and academic
10 institutions. According to the project's website, CLHEP is in maintenance
11 mode (accepting bug fixes but no further development is expected).

12 The *art* Workbook uses CLHEP, as do many of the experiments that use *art*. In both the
13 *art* run-time environment and the *art* development environment CLHEP is made avail-
14 able via UPS and is rooted at `$CLHEP_DIR`. The CLHEP header files can be found
15 at `$CLHEP_INC` and the libraries can be found at `$CLHEP_LIB_DIR`. These enviro-
16 nment variables will also be defined in the corresponding environments for your experi-
17 ment.

18 This appendix will discuss those parts of CLHEP that are important for the *art* Workbook
19 and will fill in some background information that is assumed by the CLHEP documenta-
20 tion but is not explicitly stated anywhere else.

21 CLHEP is divided in packages and the *art* Workbook uses classes from four of these
22 packages:

23 **Matrix** Support for linear algebra.

1 **Random** Support for random engines and random distributions. The distinction between
2 these two ideas will be discussed in *a section to be written in the future*.

3 **Units** Support for a standard set of units and for transformations among different units. It
4 also provides the values of many physical constants.

5 **Vector** Support for 2-vectors, 3-vectors, 4-vectors.

6 E.2 Multiple Meanings of *Vector* in CLHEP

7 CLHEP uses the word *vector* in two different senses, both of which are different from
8 the use of the word in the standard library template `std::vector`. The Matrix package
9 supports linear algebra, by providing classes to represent matrices and vectors of arbitrary
10 dimensions; the package supports operations such as matrix multiplication and the com-
11 putation of the inverse, transpose, determinant and trace of a matrix. The Vector package,
12 on the other hand, provides classes that represent a point on a plane, a point in 3-space or
13 a point in 4-dimensional space-time; the package supports operations such dot products,
14 cross products, rotations and Lorentz transformations.



15 The Vector package does not make a distinction between positions, displacements, veloc-
16 ities and momentum. The same classes are used for all four.



17 E.3 CLHEP Documentation

umentation

18 The CLHEP home page is <http://proj-clhep.web.cern.ch/proj-clhep>.

19 The following is a direct link to the CLHEP documentation page:

20 <http://proj-clhep.web.cern.ch/proj-clhep/index.html#docu>

21 In many cases the documentation for CLHEP is simply the code or the comments in the
22 code. You can view the header files by looking under `$CLHEP_INC`. You can view the
23 source files by looking under `$CLHEP_DIR/source`. A more convenient format to view
24 the header files is to use the CLHEP Doxygen site:

25 http://proj-clhep.web.cern.ch/proj-clhep/doc/CLHEP_2_1_3_1/doxygen/html/

26 Doxygen simply presents the information found in the header file in a format that is easier
27 to view than the header file itself.

- 1 To get information about a CLHEP class, go to the Doxygen page, click on the tab named
- 2 “Classes”, and use your browser’s search function to find the name the class.

3 **E.4 CLHEP Header Files**

4 **E.4.1 Naming Conventions and Syntax**

5 The syntax to include a CLHEP header file is:

```
6 #include "CLHEP/<package-name>/<filename>.h"
```

7 where `package-name` is the name of the CLHEP package to which the header file be-
8 longs and where `filename` is filename of the header file.

9 Almost all of the class names in CLHEP begin with the prefix `Hep`, for example
10 `HepLorentzVector`, which is the CLHEP representation of a 4-vector. The handful
11 of exceptions to this rule are helper classes used internally by CLHEP.


12 In most cases, the name of the header file for a class is the name of the class, ex-
13 cluding the leading `Hep`. For example, the header file for `HepLorentzVector` is
14 `CLHEP/Vector/LorentzVector.h`. Some header files also contain the declarations
15 of helper classes that are used by the main class. A few header files contain the declarations
16 of several related classes.

17 There is one important header file that follows an unusual naming pattern. The header
18 file `CLHEP/Vector/ThreeVector.h` declares the classes `Hep2Vector` and
19 `Hep3Vector` that describe a point in a plane and a point in 3-space, respectively.

20 **E.4.2 .icc Files**

21 Following a convention in use during the mid 1990’s when the CLHEP package was de-
22 veloped, CLHEP header files contain only declarations. When inline implementations are
23 required, CLHEP puts them in a file with the same name as the header file but with `.h`
24 replaced with `.icc`, which stands for “inline `cc`”. The `.icc` file is included near the
25 end of the `.h` file. For example the inline implementation for the class `Hep3Vector` is
26 included in the header file `ThreeVector.h` as:

```
27 #include "CLHEP/Vector/ThreeVector.icc"
```

1 The convention of using `.i.cc` files to segregate inline implementations from declarations
2 is no longer recommended, but CLHEP retains it for backwards compatibility. The recom- 
3 mended convention is to simply put inline implementation in the header file.

4 E.5 The CLHEP Namespace

5 All identifiers defined by CLHEP are in the `CLHEP` namespace.

6 E.5.1 `using` Declarations and Directives

7 An example of a *using* declaration is:

```
8 using CLHEP::Hep3Vector;
```


9 This tells the compiler that, when ever it sees `Hep3Vector`, it should automatically rec-
10 ognize that it means `CLHEP::Hep3Vector`.

11 An example of a *using* directive is:

```
12 using namespace CLHEP;
```

13 This tells the compiler that it should recognize all identifiers from the `CLHEP` namespace
14 without the user having to type the prefix `CLHEP::` in front of every name.

15 Using declarations and directives are a part of C++ that you can read about in any standard
16 text. Using `using` can reduce the amount of typing you have to do but over use of `using`
17 defeats the entire purpose of namespaces.

18 One very important rule is that you must never code using directives or using declarations 
19 in header files. You should only use them in source files. *A future version of this document*
20 *will refer to a complete explanation.*

21 **FIXME:** *Ref to the explanation.*

22 For the particular case of CLHEP you should not use `using` declarations or directives
23 even in source files. The remainder of this section explains why.

24 **FIXME:** *AH stopped reviewing here 7/22*

25 Consider what happens if you code:

```
1 using namespace CLHEP;
```

2 The CLHEP Units package defines many identifiers with commonly used short names, `m`, `g`
3 and `s`; in addition there are 17 two character identifiers and 24 three character identifiers,
4 such as `mm`, `m2`, `deg`, `cm3`, `amu`, `rad` and so on. Many of these short identifiers are
5 commonly used in code, `m` for mass, `s` for an arc length, `rad` for a radius and so on. If
6 you give a `using` directive for the namespace `CLHEP` then all of these short names will be
7 defined with the scope of your code.

8 A common programming error is to forget to declare a variable before using it. Normally
9 the compiler will recognize this error and issue a diagnostic message. If, on the other hand,
10 one of your undeclared variables matches one of the CLHEP variable names, and if you
11 have used `using namespace CLHEP`, then the compiler will not recognize the error
12 and will not issue a diagnostic. You will need to find the error by tedious debugging.

13 A similar problem occurs if you code:

```
14 using CLHEP::Hep3Vector;
```

15 This introduces three one letter identifiers into the scope of your code: `X`, `Y` and `Z`, which
16 are defined as the integers 0, 1, 2.

17 The bottom line is that CLHEP does not mix with `using` directives and `using` declara-
18 tions.

19 E.6 The Vector Package

20 The *art* Workbook uses the following classes from the Vector package:

21 `Hep3Vector` A vector in 3-space.

22 `HepLorentzVector` A 4-vector in 4-dimensional space time.

23 `HepBoost` Performs Lorentz boosts from one inertial frame to another; it operates on
24 objects of type `HepLorentzVector`.

Constructors

Listing E.1: The constructors of `Hep3Vector`

```

1
2 Hep3Vector();
3 explicit Hep3Vector(double x);
4 Hep3Vector(double x, double y);
5 Hep3Vector(double x, double y, double z);
6 Hep3Vector(const Hep3Vector &);

```

Hep3Vector

1 E.6.1 CLHEP::Hep3Vector

2 The class `Hep3Vector` is used to represent any 3-vector in which each element is rep-
 3 resented as `double`. There is no similar class in which each element is represented as a
 4 `float`. The header file is found at

```
5 $SCLHEP_INC/CLHEP/Vector/ThreeVector.h
```

6 You can also browse the class header using Doxygen:

7 http://proj-clhep.web.cern.ch/proj-clhep/doc/CLHEP_2_1_3_1/doxygen/html/classCLHEP_1_1Hep3Vector.html

8 Because this is a particularly simple class, the header serves as complete documentation.
 9 If you are a beginner to C++ this might not be enough. So this section will show a few of
 10 the commonly used features. The goal of this section is to guide you through the header
 11 file so that you will learn how to understand a typical CLHEP header file.

12 `Hep3Vector` has five constructors that are shown in Listing E.1. In each of the first four
 13 constructors the missing components are set to 0. If you check the `.icc` file you will
 14 see that all five are implemented inline. The `explicit` keyword is beyond the scope of this
 15 discussion. **FIXME:** *Add a reference to `explicit` when we have one.*

16 There is a single `Hep3Vector` class that can be used for positions, momenta and veloc-
 17 ities. So don't attach special meaning to the names of the arguments; the 4th constructor
 18 would have exactly the same behaviour had it been declared as:

```
1b Hep3Vector(double p_x, double p_y, double p_z);
```

20 or

```
2l Hep3Vector(double v_x, double v_y, double v_z);
```

22 There are many synonyms for access to the individual elements. If `v` is an object of type
 23 `Hep3Vector`, then following all return the value of the x component:

```

1 double x = v.x();
2 double x = v.getX();
3 double x = v(CLHEP::Hep3Vector::X);
4 double x = v[CLHEP::Hep3Vector::X];

```

5 similarly for y and z . There is an `enum` near the top of the class declaration that defines
6 the mapping of coordinates to indices

```

7 enum { X=0, Y=1, Z=2, NUM_COORDINATES=3, SIZE=NUM_COORDINATES };

```

8 The full names `CLHEP::Hep3Vector::X` are so long that it is common to see code in
9 which the integer values are substituted by hand:

```

10 double x = v(0);

```

11 Why are there so many different ways of doing the same thing? `Hep3Vector` was de-
12 signed as a common replacement for several other 3-vector packages so it's public inter-
13 face was designed as the union of all of the public interfaces. While you will see all of
14 the above used in code written by others, you should prefer to use the first form. It is pre-
15 ferred over the last two because it executes more quickly. It is preferred over the second
16 form for a stylistic reason; the *art* team recommends that accessors not contain the string
17 "get".

18 Why does documentation talk about all of these ways of doing the same thing? Because
19 you are likely to encounter them as you read code written by others.

20 The header tells you that the above accessors functions exist because it contains:

```

21 double x() const;
22 double getX() const;
23 double & operator () (int);
24 double operator [] (int) const;

```

25 All of the accessors are implemented inline. The first two accessors are member functions,
26 similar to many you have seen before. You do not need to learn all of the details about the
27 last two accessors — you only need to recognize their use, as illustrated on the previous
28 page.



31 If you do wish to learn more about the third accessor it will be described in any stan-
dard C++ text under the topic of "functors". If you wish to learn more about the fourth
accessor, it will be described in any standard C++ text; find the section on `operator []`,

1 pronounced “operator square brackets”, which may be in a more general section on operator
2 overloading.

3 If v is an object of type `Hep3Vector`, the following all set the value of the x component
4 to 123.

5

```
1b v.setX(123.);
2 v(CLHEP::Hep3Vector::X) = 123.;
3 v[CLHEP::Hep3Vector::X] = 123.;
```

9 Similarly for setting the values of the y and z components. As before, prefer to use the first
10 version because it is faster than the other two.

11 The header tells you that the above accessors functions exist because it contains:

```
1b void setX(double);
2b double & operator [] (int);
3b double & operator () (int);
```

15 You can also set the values of all three components at once. For example the following sets
16 the value of v to a unit vector in the z direction:

```
1b v.set(0., 0., 1.);
```

18 The header tells you that this function exists because it contains:

```
1b void set(double x, double y, double z);
```

20 tab:clhep:functions Table E.1 lists some of the commonly used member functions of `Hep3Vector`; it shows
21 the declaration of each function and an example of how to use it. Some other member
22 functions are shown in the text below; they do not fit nicely in the table so they are pre-
23 sented separately. In each case you will see a statement in standard mathematical notation,
24 the name of the function, followed by a usage example and the the declaration. In the
25 usage examples, u , v are all objects of type `CLHEP::Hep3Vector` and a is of type
26 `double`.

27 $\vec{v} = \vec{u}$. The assignment operator:

```
2b v = u;
3b Hep3Vector & operator = (const Hep3Vector &);
```

30 $\vec{v} = \vec{v} + \vec{u}$. Addition:

Table E.1: Selected member functions of CLHEP::Hep3Vector. In the usage examples, u , v , w are all objects of type CLHEP::Hep3Vector and a is of type double. Polar angles are measured relative to the z axis and the azimuth is measured relative to the x axis.

Usage	Declaration	Meaning
<code>a = u.mag();</code>	<code>double mag() const;</code>	$\sqrt{u_x^2 + u_y^2 + u_z^2}$
<code>a = u.mag2();</code>	<code>double mag2() const;</code>	$u_x^2 + u_y^2 + u_z^2$
<code>a = u.perp();</code>	<code>double perp() const;</code>	$\sqrt{u_x^2 + u_y^2}$
<code>a = u.perp2();</code>	<code>double perp2() const;</code>	$u_x^2 + u_y^2$
<code>a = u.theta();</code>	<code>double theta() const;</code>	Polar angle of u , θ
<code>a = u.phi();</code>	<code>double phi() const;</code>	The azimuth of u
<code>a = u.cosTheta();</code>	<code>double cosTheta() const;</code>	$\cos \theta$
<code>a = u.cos2Theta();</code>	<code>double cos2Theta() const;</code>	$\cos^2 \theta$
<code>v = u.unit();</code>	<code>Hep3Vector unit() const;</code>	A unit vector in the direction of u

```
1 v += u;
2 Hep3Vector& operator += (const Hep3Vector&);
```

3 $\vec{u} = a\vec{u}$. Multiplication by a scalar:

```
1 u *= a;
2 Hep3Vector& operator *= (double);
```

6 $\vec{u} = \vec{u}/a$. Division by a scalar:

```
1 u /= a;
2 Hep3Vector& operator /= (double);
```

9 $\vec{v} = -\vec{u}$. Unary minus:

```
1b v = -u;
1d Hep3Vector operator - () const;
```

12 $a = \vec{u} \cdot \vec{v}$. Dot product:

```
1b a = u.dot(v);
1d double dot(const Hep3Vector&) const;
```

15 $\vec{w} = \vec{u} \times \vec{v}$. Cross product:

```
1b w = u.cross(v);
1d Hep3Vector cross(const Hep3Vector&) const;
```

18 $\vec{v} = \hat{u}$. Computation of a unit vector parallel to \vec{v} .


```
1 v = u.unit();  
2 Hep3Vector unit() const;
```

3 The C++ compiler knows how to use combinations of the above functions to allow you to
4 write:

```
1 w = u + v;  
2 w = u - v;
```

7 which have the obvious meaning.

8 `Hep3Vector` has many other member functions that are not mentioned above. They provide
9 accessors for rapidity and pseudo-rapidity; they provide setters that accept arguments
10 in spherical polar coordinates or in cylindrical coordinates; they provide support for ro-
11 tations. These functions are not used in the Workbook and will not be discussed here.
12 Consult the header file or the Doxygen site for further information.

13 **FIXME:** *We should probably add rotations for completeness.*

14 E.6.1.1 Some Fragile Member Functions

15 `Hep3Vector` has functions to ask if two `Hep3Vector`'s are equal to each other, if they
16 are near to each other, if they are parallel to each other or if they are orthogonal to each
17 other. These functions are fragile and should be used with extreme care. If you think that
18 you have a reason to use one of these functions, speak with an expert to learn if there is a
19 better alternative.

20 `Hep3Vector` supports comparison for equality. The function is declared as:

```
21 bool operator == (const Hep3Vector &) const;
```

22 and is used as follows:

```
23 if ( u == v ) {  
24     // do something  
25 }
```

26 This comparison suffers from the usual dangers about doing comparisons of floating point
27 numbers for equality.

28 **FIXME:** *Reference to the problems about comparing floating types for equality.*

1 The `isNear` member function is declared as:

```
1 bool isNear (const Hep3Vector &, double epsilon=tolerance) const;
```

3 and is used as follows:

```
1 double epsilon=1.e-14; // for example
2 if ( u.isNear(v,epsilon) {
3     // do something
4 }
```

8 The notion of nearness is expressed as,

$$|\vec{v} - \vec{u}|^2 \leq \epsilon^2 |\vec{v} \cdot \vec{u}| \quad (\text{E.1})$$

9 The fragility comes from the second argument. This argument is optional and has
 10 a default value that is set to the value of a static member datum of the class,
 11 `CLHEP::Hep3Vector::tolerance`. This member datum can be modified by any-
 12 one. Therefore always specify the second argument and never let it take its default
 13 value. If you let your code take the default value then, if someone changes the value of
 14 tolerance, the behaviour of your code will change.

15 There are similar comments about the functions:

```
1 bool isParallel (const Hep3Vector & v, double epsilon=tolerance) const;  
2 bool isOrthogonal (const Hep3Vector & v, double epsilon=tolerance) const;
```

18 which are used as:

```
1 if ( u.isParallel(v,tolerance) );  
2 if ( u.isOrthogonal(v,tolerance) );
```

21 E.6.2 CLHEP::HepLorentzVector

ep:LorentzVector

22 The class `HepLorentzVector` is used to represent any 4-vector in 4-dimensional space-
 23 time. Each element is represented as `double`; there is no similar class in which each
 24 element is represented as a `float`. The header file is found at

```
2 $CLHEP_INC/CLHEP/Vector/LorentzVector.h
```

26 You can also browse the class header using Doxygen:

27 http://proj-clhep.web.cern.ch/proj-clhep/doc/CLHEP_2_1_3_1/doxygen/html/classCLHEP_1_1HepLorentzVector.

1 As for `Hep3Vector`, the only documentation for `HepLorentzVector` is the contents
2 of the header, which can be viewed either as a file or by the Doxygen web page.

3 If you are not familiar with reading C++ header files you should read Section E.6.2 before
4 reading this section. That section has more examples of how to translate a declaration in a
5 header file into example code; this section presumes that you can do that yourself.

6 In the `HepLorentzVector` the order of components is given by the enum:

```
7 enum { X=0, Y=1, Z=2, T=3, NUM_COORDINATES=4, SIZE=NUM_COORDINATES };
```

8 `HepLorentzVector` has eight constructors that may be of interest to users of the Work-
9 book:

```
10 HepLorentzVector() ;  
11 HepLorentzVector(double x, double y, double z, double t);  
12 HepLorentzVector(double x, double y, double z);  
13 HepLorentzVector(double t);  
14 HepLorentzVector(const Hep3Vector & p, double e);  
15 HepLorentzVector(double e, const Hep3Vector & p);  
16 HepLorentzVector(const HepLorentzVector & v);  
17 HepLorentzVector(const Hep3Vector & v);
```

18 In each constructor, the missing components are set to zero. As for `Hep3Vector`, there
19 is a single `HepLorentzVector` class that can be used for 4-positions, 4-momenta and
20 4-velocities. So don't attach special meaning to the names of the arguments.

21 In the following examples the objects `p`, `q`, `s` are of type `HepLorentzVector`, the
22 objects named `u`, `v`, `w` are of `Hep3Vector` and the objects `a`, `b`, `x`, `y`, `z`, `t`
23 are of type `double`.

24 `HepLorentzVector` has a set of accessor methods that are similar to those of `Hep3Vector`.
25 The following code fragments return the value of the x component:

```
26 double x = p.x();  
27 double x = p.px();  
28 double x = p.getX();  
29 double x = p(CLHEP::LorentzVector::X);  
30 double x = p[CLHEP::LorentzVector::X];
```

31 Similarly for the y , z and t components. Prefer to use one of the first two member functions
32 for same the reasons discussed for `Hep3Vector`.

1 HepLorentzVector has several accessors that return the space part of the 4-vector as
2 a Hep3Vector:

```
1 Hep3Vector u = p.vect();
2 Hep3Vector u = p.v();
3 Hep3Vector u = p;
4 Hep3Vector u = p.getV();
```

7 Prefer to use one of the first two forms. The third form is perfectly correct but it may be
8 difficult for readers of your code to understand its intended function (only real CLHEP
9 experts will understand it). The fourth form executes more slowly than the others.

10 The following code fragments set the value of the x component to 123.:

```
11 p.setX(123.);
12 p.setPx(123.);
13 p(CLHEP::HepLorentzVector::X) = 123.;
14 p[CLHEP::HepLorentzVector::X] = 123.;
```

15 Similarly for the y , z and t components. Prefer to use the first two member functions as
16 they execute more quickly than do the last two.

17 HepLorentzVector has several functions that allow you assign new values to several
18 elements at the same time:

```
19 p = u;
20 p.setVect(u);
21 p.set(x, y, z, t);
22 p.set(u);
23 p.set(u,t);
24 p.setV(u);
```

25 Two of the above lines in three previous three blocks may be difficult to understand. The
26 two lines:

```
27 Hep3Vector const& u = p; // Return the space part by const reference.
28 p = u; // Assignment of a 3-vector to the space part
29 // of a 4-vector
```

heporentzvectpr

30 are implemented, respectively, by the following operators of the class HepLorentzVector:

```
31 operator const Hep3Vector & () const;
32 operator Hep3Vector & ();
```


1 You do not need to understand how these work, just that the above two assignments work.
 2 If you do wish to, you can look in any standard C++ reference in a section that might be
 3 called type-cast operators or implicit conversions.

4 `HepLorentzVector` provides member functions that return the invariant mass:

```
1 double mass = p.mag();
2 double mass = p.m();
3 double mass = p.restMass();
4 double mass = p.invariantMass();
```

9 A common beginner's mistake involves confusion about the function `mag`. For a `Hep3Vector`,
 10 `mag`, returns the magnitude of the 3-vector. For a `HepLorentzVector`, it would make
 11 sense that there is a member function to return the magnitude of the space part but there is
 12 no such function. Instead you should do the following:

```
15 double momentumMagnitude = p.v().mag();
```

14 The common mistake is to use `mag` on a `HepLorentzVector` expecting it to return the
 15 magnitude of the space part, when it actually returns the invariant mass. 

16 `HepLorentzVector` has member functions that return the invariant mass squared:

```
17 double msq = p.mag2();
18 double msq = p.m2();
19 double msq = p.restMass2();
20 double msq = p.invariantMass2();
```

21 `HepLorentzVector` has a member function and an operator overload that return the
 22 4-dimensional dot product

$$p \cdot q = p_t q_t - p_x q_x - p_y q_y - p_z q_z \quad (\text{E.2})$$

23 This function is also provided as an overload of the multiply operator:

```
24 double dot = p.dot(q);
25 double dot = p*q;
```

26 The `HepLorentzVector` forwards many of the functions of `Hep3Vector`. With one
 27 very important exception the member functions described in Table E.1 are also available
 28 to a `HepLorentzVector` and perform the same functions. For example the member
 29 function `theta` returns the polar angle of the space-part of the 4-vector. The exception is
 30 the member function `mag`, which was discussed above.

1 `HepLorentzVector` supports arithmetic operations such as:

```
1 q = p;  
2 q = -p;  
3 p += q;  
4 p -= q;  
5 p *= a;  
6 p /= a;  
7 s = p + q;  
8 s = p - q;  
9 s = a*p + b*q;  
10 s = a*p - b*q;
```

12 `HepLorentzVector` has member functions and operators that compare for equality and
13 to ask if two vectors are almost the same, almost parallel or almost orthogonal. These have
14 signatures similar to those for `Hep3Vector` and they have the same fragile behaviour;
15 use them with care.

16 If you consult the header file or Doxygen documentation you will see that `HepLorentzVector`
17 has many member functions that are not mentioned above. These are not needed for the
18 Workbook and will not be discussed here.

19 **E.6.2.1** `HepBoost`

20 *This material will be available in a future release.*

21 **E.7 The Matrix Package**

22 The *art* Workbook uses the following classes from the Matrix package:

23 `HepMatrix` A general $n \times m$ matrix class.

24 `HepSymMatrix` A class that represents symmetric matrices.

25 `HepVector` A column-vector ($n \times 1$ matrix) class.

26 *This material will be available in a future release.*

1 E.8 The Random Package

2 The *art* Workbook uses the following classes from the Random package:

3 `HepRandomEngine` The base class from which all CLHEP random engines must in-
4 `herit`.

5 `HepJamesRandom` A random engine that implements an algorithm described by F. James
6 of CERN.

7 `RandFlat` A distribution that returns a random variate that is flat on a specified domain.

8 `RandGaussQ` A distribution that returns a random variate that is distributed as a Gaus-
9 `sian` distribution.

10 `RandPoissonQ` A distribution that returns a random variate that is distributed as a Pois-
11 `son` distribution.

12 The two classes with names ending in `Q` have no internal state except for the state of the
13 underlying engine. **FIXME:** *Ref to the full explanation.*

14 *This material will be available in a future release.*

1 F Include Guards

g:include:guards

2 All of the header files presented in the Workbook begin and end with a three line structure
3 called an include guard. An example is shown this listing:

```
4 #ifndef package_path_to_file_h  
5 #define package_path_to_file_h  
6  
7 // The C++ content of the header file  
8  
9 #endif /* package_path_to_file_h */
```

10 The three lines beginning with # are macros that will be processed by the C preprocessor
11 at the start of compilation. These lines are called *include guards* and they address the
12 following issue.

13 Suppose that you have a main program that includes two header files `A.h` and `B.h`; fur-
14 ther suppose that both of `A.h` and `B.h` include a third header file `C.h`. When you compile
15 the main program, the C preprocessor will expand all of the include directives to create a
16 temporary `.cc` file on which the compiler will do its work. This temporary file must con-
17 tain exactly one copy of the header file `C.h`; if it contains either zero copies or more than
18 one copy (as it would in this case), the compiler will issue an error. The C preprocessor,
19 by itself, is not smart enough to skip the second inclusion of `C.h` but it does provide the
20 tools for us to help it do so.

21 In the first two lines, the text `package_path_to_file_h` is the name of a C prepro-
22 cessor variable; the choice of the variable name will be described later but the important
23 feature is that it must be unique within the compilation unit (the file being compiled).
24 When the C preprocessor encounters the included file `C.h`, the line `#ifndef` tells the

1 preprocessor to check to see if the C preprocessor variable with this name is defined. If
2 the variable is not defined then the lines between the `#ifndef` line and the `#endif` line
3 will be included in the output of the C preprocessor. If it has already been defined, these
4 lines will be excluded from the output.

5 The first time that the preprocessor encounters `C.h` within a compilation unit, the variable
6 will not have been defined and the contents of the header will be included in the output of
7 the preprocessor. At the same time the second line of the above fragment will be executed;
8 it is a preprocessor directive that tells the preprocessor to define the variable.^{c0} In either
9 case, when the preprocessor encounters the second inclusion of `C.h`, the `#ifndef` test
10 will fail and the body of the header will not be copied into the output of the preprocessor.
11 And so on for subsequent inclusions of `C.h`.

12 If every header file in a code base correctly uses include guards, then every header file can
13 safely include all other header files on which it depends and one need not worry about this
14 causing compiler errors due to multiple declarations of a class or function.

15 For include guards to work, each header file must choose a C preprocessor variable name
16 that is unique within every compilation unit in which it might be included, either directly
17 included or indirectly included. The convention that is used by *art*, by other libraries man-
18 aged by the *art* team, by the toyExperiment UPS product and by the Workbook is that the
19 name of the variable is the name of the path to the header file, starting from the root of
20 the code base and with the slash and dot characters changed to underscores; the reason for
21 this change is that slash and dot characters are not legal in the name of a C preprocessor
22 variable. This works because all of these products also adopt the convention that the path
23 to their header file starts with the product name. While this is not perfect security it is a
24 very high level of security.

^{c0} The full syntax of the `#define` directive allows one to specify a value for the variable but that is not important here; the `#ifndef` test only cares that the variable is defined, not what its value is.

1

Part V

2

Index

Index

- 1 3-vector representation, 325
2 4-vector representation, 325
- 3 analyzer module, **186**
4 characteristics, 186
5 configure via parameter set, 244
6 events passed by reference, 189
7 flow of execution, 215
8 order in path, 270
9 signature, 190
10 use of override, 187
- 11 API, 13
12 art, **8**
13 API, 13
14 applicability, 8
15 artmod, 207
16 as an external product, 24
17 build systems, 204
18 C++, 8
19 command, 15
20 command line options, 138
21 long form, 138
22 short form, 138
23 configuration file, *see* configuration file
24 data persistency, 295
25 data product, *see* data product
- 26 development environment, 164, 177, 456
27 documentation suite, 10
28 dynamic library usage, 194
29 error conditions, 156
30 error status, 133
31 event, *see* event
32 event ID, *see* event ID
33 event loop, *see* event loop
34 event sharing, 190
35 executable, 127
36 execution syntax, 132
37 flow of execution, 215
38 getting help, 10
39 identifiers, 470
40 input file
41 specify in FHiCL, 139
42 specify on command line, 139
43 job status, 133
44 log file, 133
45 module, *see* module
46 module types, 19
47 operate on multiple modules, 209
48 output file, 133
49 output module, 149
50 paths used in, 41
51 post-initialization steps, 17

- 1 processing order, 270
- 2 rerun same module, 143, 267
- 3 ROOT classes, 295
- 4 ROOT support, 26
- 5 run-time configuration, 127
- 6 command line options, 134
- 7 FHiCL file, 134
- 8 run-time environment, 127, 151, 452
- 9 sample output, 132
- 10 services, *see* services
- 11 specify events to process, 138, 141
- 12 specify modules to process, 142
- 13 TFileService, *see* TFileService
- 14 Unix environment, 39
- 15 use as external package, 9
- 16 use of ROOT, 295
- 17 users, 9
- 18 art module, *see* module
- 19 art-users email list, 10
- 20 art::EDAnalyzer, 185
- 21 art::Event, 189, 270
- 22 art::EventID, 274
- 23 art::ServiceHandle, 296
- 24 artdaq, **9**
- 25 artmod, 207
- 26 options, 209
- 27 auto, 288
- 28 beginJob, 304
- 29 boost, **24**
- 30 browse command, 310
- 31 build system, 23, 164
- 32 instructions for, 171
- 33 build tools, 456
- 34 building code
- 35 clean rebuild, 200
- 36 complete, 199
- 37 finding dynamic libraries, 202
- 38 incremental, 199
- 39 linking, 202
- 40 saving output files, 200
- 41 buildtool, 175, 179
- 42 algorithm, 198
- 43 CMakeLists.txt file, 197
- 44 error, 200
- 45 functions, 197
- 46 verbose mode for experts, 204
- 47 C++, 10
- 48 -Werror, 57
- 49 .cc files, 50
- 50 .dylib files, 50
- 51 .h files, 50
- 52 .o files, 50
- 53 .so files, 50
- 54 accessors, 93, 329
- 55 base class, 17
- 56 build, 49, 67
- 57 output option, 56
- 58 build commands, 64
- 59 c++ command, 64, 66, 68
- 60 compile, 49, 54
- 61 declaration, 101
- 62 definition within declaration, 101
- 63 dereferencing operator, 332
- 64 enum, 325
- 65 exception, 254
- 66 executable program, 50
- 67 float, 59
- 68 free functions, 105, 330

- 1 function
- 2 argument list, 62
- 3 declaration, 62
- 4 definition, 63, 69
- 5 implementation, 63
- 6 return type, 62
- 7 function 'main', 55, 62, 65
- 8 header files, 50, 62
- 9 implementation within declaration, 101
- 10 implicit type conversion, 69
- 11 include directive, 68
- 12 include guards, 62
- 13 include ROOT header syntax, 298
- 14 inheritance, 17, 166
- 15 libraries, 50, 61
- 16 link, 49, 54
- 17 link list, 51
- 18 linker, 66
- 19 linker symbols, 66
- 20 looping
- 21 breaking up code, 335
- 22 const reference, 340
- 23 do-while, 338
- 24 for writing templates, 338
- 25 range-based, 331
- 26 using at function, 336
- 27 using comma operator, 338
- 28 using iterators, 337
- 29 using iterators with auto, 337
- 30 using std::vector as array, 336
- 31 while, 338
- 32 main program, 55
- 33 main program, 54, 62
- 34 module, *see* module
- 35 object files, 50
- 36 pointer, 59
- 37 prerequisites, 54
- 38 rebuild subset, 61
- 39 signature, 69
- 40 Singleton Design Pattern, 22
- 41 source code files, 50
- 42 std::vector<T>, 54
- 43 stream insertion operator, 105, 330
- 44 syntax flexibility, 194
- 45 temporary object recommendations, 195
- 46 uninitialized variable, 56
- 47 unresolved references, 66
- 48 variable addresses, 56
- 49 variable type, 59
- 50 C++11, 10
- 51 conditional exclusion, 330
- 52 calibration constants, *see* conditions infor-
- 53 mation
- 54 cetbuildtools, 23, **25**, 164, 456
- 55 CETLIB, **24**
- 56 CINT, 296
- 57 file naming convention, 312
- 58 multiple png files, 339
- 59 relation to C++, 312
- 60 script, 297
- 61 subtract two histograms, 339
- 62 , 297
- 63 class templates, *see* templates
- 64 CLHEP, **24**, 325
- 65 cmake, 23, 305
- 66 variable, 286, 307
- 67 definition, 203
- 68 CMakeLists.txt file
- 69 buildtool, 197
- 70 cmsrun, 9

- 1 coding
- 2 best practices, 34
- 3 conditional exclusion, 330
- 4 conventions, 34
- 5 rules, 34
- 6 style, 34
- 7 coding standards, **10**
- 8 C++, 10
- 9 C++ 11, 10
- 10 collection, 21
- 11 colon initializer syntax, 242, 264
- 12 conditions information, **21**, 127
- 13 configuration file, 15
- 14 constructor
- 15 explicit argument, 187

- 16 data file, 27
- 17 data members
- 18 store parameter values, 256
- 19 data product, **20**, 137, 274
- 20 collection, *see* collection
- 21 contents, 21
- 22 data type, *see* data type
- 23 distinguish from products, 25
- 24 four-part identifier, 487
- 25 full name, 487
- 26 InstanceName, *see* InstanceName
- 27 ModuleLabel, *see* ModuleLabel
- 28 name, 275, 277
- 29 operations, 26
- 30 persistency, 26
- 31 persistent representation, 26
- 32 ProcessName, *see* ProcessName
- 33 specification, 279, 284
- 34 transient representation, 26

- 35 underscore, 487
- 36 data type, 487
- 37 friendly name, 278
- 38 in data product name, 278
- 39 debugging
- 40 using optional parameters, 259
- 41 defaultExceptions, 290
- 42 development environment, 456
- 43 Doxygen, 13
- 44 dynamic libraries, 23
- 45 .dylib files, 23
- 46 .so files, 23
- 47 build system, 23
- 48 dependency lists, 203
- 49 directories, 156
- 50 file extensions, 127
- 51 naming rules, 198
- 52 paths to, 127
- 53 run-time loading, 186
- 54 use in instance creation, 194
- 55 use with buildtool, 198
- 56 dynamic library, 16
- 57 building, 182
- 58 matching module name, 182
- 59 dynamic load libraries, *see* dynamic libraries

- 60 EDAnalyzer, 143
- 61 EDM, *see* Event-Data Model
- 62 endJob, 304
- 63 error handling, 254
- 64 default vs alternate behavior, 255
- 65 event, **14**, 14
- 66 contents, 189
- 67 representation, 192
- 68 representation in art, 189

- 1 unique identifier, 14
- 2 event ID, **14**, 274
- 3 event number, 14, 191
- 4 EventID return type, 191
- 5 individual parts of, 211
- 6 run number, 14, 191
- 7 subRun number, 14, 191
- 8 event loop, 15, **19**, 19
- 9 event-data files, 27
- 10 event-data files for Workbook, 128
- 11 Event-Data Model, **26**
- 12 ROOT support, 26
- 13 exceptions, 243, 265
- 14 experiment code, *see* user code
- 15 external products, *see* products

- 16 FermiGrid, 25
- 17 Fermilab Hierarchical Configuration Language, *see* FHiCL
- 18 *see* FHiCL
- 19 FHiCL, 24, 126, 134, **466**
- 20 fhiicl::ParameterSet, 188
- 21 definition
- 22 form, 135
- 23 definitions, 241
- 24 file extension, 134
- 25 identifier
- 26 analyzers, 137
- 27 physics, 136
- 28 source, 135
- 29 numerical forms
- 30 formats, 261
- 31 numerical types, 259
- 32 output module, 149
- 33 optional parameters, 150
- 34 parameter name, 247

- 35 parameter set, **241**
- 36 error conditions, 253
- 37 print, 242, 264
- 38 parameter value
- 39 store as data member, 256
- 40 store as local variable, 257
- 41 parameter value substitution, 351
- 42 parameters, 241
- 43 canonical form, 243, 259, 265
- 44 default value, 257
- 45 default values, 243, 265
- 46 internal representation, 245
- 47 optional, 257
- 48 policies, 259
- 49 properties, 244
- 50 return type, 246
- 51 paths, 144
- 52 process name, 137
- 53 prolog, 351
- 54 scope, 135
- 55 source file, 135
- 56 special characters, 136
- 57 specify input files, 139
- 58 syntax, 135
- 59 table, 135, 241
- 60 value, 136
- 61 fhiicl::ParameterSet, 241
- 62 get, 247
- 63 file
- 64 header for read/write, 356
- 65 write to, 358
- 66 file catalog, 27
- 67 file of Monte Carlo events, *see* event-data
- 68 files
- 69 file of simulated events, *see* event-data files

- 1 filter module, 19
- 2 forward declarations, 326
- 3 four-vector conversion, 332
- 4 framework, **8**
 - 5 boundary with user code, 9
 - 6 infrastructure, 9
- 7 friendly name, 278, 284
- 8 textbf, 278
- 9 gcc, **24**
- 10 GenParticle, **276**, 281, 324, 326
- 11 GenParticleCollection, **276**, 324
- 12 geometry information, 127
- 13 geometry specification, 21
- 14 getting help, 10
- 15 git, 24, 167
- 16 about, 168
- 17 handle
 - 18 invalid, 283
- 19 handles, **282**
 - 20 default construct, 283
 - 21 service, 302
 - 22 valid, 284, 287
- 23 header files
 - 24 absence of, 212
 - 25 conflicts, 202
 - 26 finding, 200
 - 27 from UPS product, 201
 - 28 Geant4, 201
 - 29 ROOT, 201
- 30 help with art, 10
- 31 Hep3Vector, 325
- 32 HepLorentzVector, 325
- 33 histogram
 - 34 change module label, 317
 - 35 create using TFileService, 302
 - 36 filling, 304
 - 37 formats, 317
 - 38 pointer naming convention, 298
 - 39 save, 317
 - 40 structure, 303
 - 41 subtracting two, 339
 - 42 view interactively, 310
 - 43 view with ROOT, 296
 - 44 view with TBrowser, 308, 317
 - 45 write to PDF, 296
- 46 histogram file, 303, 308
 - 47 change name, 317
 - 48 overwriting, 316
- 49 ifdh_sam, 25
- 50 inheritance, 185
- 51 input tag
 - 52 initialization, 286
- 53 InstanceName, 487
- 54 jobsub_tools, 25
- 55 keyword
 - 56 auto, 288
- 57 member function
 - 58 argument names, 188
 - 59 for analyzer module, 186
 - 60 optional, 186
 - 61 override identifier, 187, 189
 - 62 template, 252
- 63 message service, 22
 - 64 purpose of, 138
- 65 MF, **24**
- 66 module, **15**

- 1 C++ class, 16
2 analyzer, *see* analyzer
3 member function requirement, 249
4 class, 241
5 communication between, 270
6 create with artmod, 207
7 dependencies, 204
8 filter, *see* filter
9 finding, 155
10 header files, 184
11 identify from art output, 180
12 instance, 267, 303
13 in data product name, 278
14 name, 278
15 label, 142, 143, 269, **271**, 303
16 identify parameter set, 271
17 in data product name, 278
18 name instance, 271
19 naming rules, 144
20 parameter set, 241
21 uniqueness, 144
22 naming rules, 198
23 optional member function, 186
24 output, *see* output
25 producer, *see* producer
26 requirements, 16
27 run simultaneously, 209
28 source, *see* source
29 source code, 182
30 three-file style, 213
31 type, **135**, 143, 186
32 types, 19
33 module types, **19**
34 ModuleLabel, 487
35 namespaces
36 fhicl, 188
37 ROOT, 298
38 tex, 185
39 toy experiment, 185
40 naming variables, 247
41 NTuple, 20
42 null pointer, 300
43 nullptr, 300
44 output directory, 180
45 output module, 19
46 overload set, 69
47 packages, *see* products
48 parameter set, 135, 269
49 for module configuration, 143
50 for service configuration, 354
51 module label, 143
52 parameters
53 default value recommendations, 286
54 Particle Data Group, *see* PDG
55 particle data table, *see* PT347
56 particles
57 generated, 324
58 parent-child relationships, 326
59 PDG identifier codes, 324
60 primary, 326
61 secondary, 326
62 paths
63 art, 145
64 different types of, 147
65 FHiCL, 144
66 module order, 270
67 PDG, 324
68 particle identifier codes, 324, 348

- 1 PDT, 347
- 2 file, 351
- 3 toy experiment, 351
- 4 PDT Service
- 5 parameters, 355
- 6 PDT service, 347, **352**
- 7 configuring, 354
- 8 plugins, *see* dynamic libraries
- 9 png files, viewing, 333
- 10 pointer, 287
- 11 bare, 287
- 12 naming convention for histogram, 298
- 13 pointers, 283
- 14 safe, 283
- 15 smart, 283
- 16 process name, 137
- 17 in data product name, 278
- 18 processing loop, *see* event loop
- 19 ProcessName, 488
- 20 producer module, 19
- 21 processing order, 270
- 22 products, **24**, 111
- 23 access to, 109
- 24 distinguish from data product, 25
- 25 distribution via UPS/UPD, 25, 109
- 26 external, 109
- 27 product directories, 109
- 28 PRODUCTS, 109
- 29 reconstruction on demand, 16
- 30 replica manager, 27
- 31 ROOT, 20, 24, 295
- 32 booking a histogram, 304
- 33 create histogram, 302
- 34 cycle numbers, 310
- 35 deleting a histogram, 304
- 36 dictionaries, 330
- 37 documentation links, 316
- 38 exit from root program, 309
- 39 file naming conventions, 304
- 40 filling a histogram, 304
- 41 genreflex tool, 330
- 42 global namespace, 298
- 43 histogram file, *see* histogram file
- 44 histogram vs event-data filenames, 304
- 45 include header syntax, 298
- 46 input vs output filenames, 304
- 47 output file structure, 303
- 48 TBrowserWindow, 309
- 49 TFileService, *see* TFileService
- 50 TNtuples, 304
- 51 treatment of bin edges, 303
- 52 TTrees, 304
- 53 ROOT files
- 54 event-data, 296
- 55 histograms, 296
- 56 RootInput module, 135
- 57 run, 14
- 58 run-time configuration
- 59 value types, 467
- 60 run-time configuration file, *see* configuration
- 61 file
- 62 run-time environment, 452
- 63 SAM, 25, 27
- 64 service handle, 302
- 65 services, **21**, 300
- 66 access via handle, 302
- 67 configuring, 354
- 68 creation of, 353

- 1 message service, *see* message service
- 2 requesting information from, 22
- 3 source code filename, 352
- 4 TFileService, *see* TFileService
- 5 set up to run Workbook, 129
 - 6 build window, 171
 - 7 source window, 167
- 8 setup to run Workbook
 - 9 login back in, 221
- 10 shareable libraries, *see* dynamic libraries
- 11 site-specific setup, 46
 - 12 procedures, 47
 - 13 Unix environment, 46
- 14 site-specific setup procedure, 152
 - 15 setup git, 168
- 16 smart pointer, 22
- 17 source code
 - 18 compile and link, 182
 - 19 filenames
 - 20 modules, 186
 - 21 services, 352
 - 22 source directory, 167
 - 23 source module, 19
 - 24 std::vector, 258, 274
 - 25 dynamic sizing, 275
 - 26 subRun, 14
- 27 TBrowser, 308
 - 28 print histogram, 317
- 29 templates, 243, 246, 265
 - 30 argument, 246, 275, 279
 - 31 dummy, 252
 - 32 type, 247
 - 33 class, 274, 277, 287
 - 34 member function, 252, 274, 285
 - 35 templates: collections of objects, use with, 277
- 36 testing
 - 37 using optional parameters, 259
- 38 TFileService, 22, **295**, 296, 300
 - 39 arguments, 302
 - 40 configure, 297, 305
 - 41 create histogram, 302
- 42 toy experiment, 12, **28**
 - 43 namespace, 185
 - 44 setup, 153
- 45 Tree, 20
- 46 typedef, 276
 - 47 recommendation for use of, 192
- 48 underscore
 - 49 as field delimiter, 488
 - 50 where forbidden, 488
- 51 Unix
 - 52 *art* Workbook environment, 39
 - 53 bash alias, 43
 - 54 bash function, 43
 - 55 bash script, 42
 - 56 bash shell, 37
 - 57 commands, 35
 - 58 computing environment, 38
 - 59 environment, 38
 - 60 environment variables, 38, 40, 46
 - 61 examine environment, 40
 - 62 execute vs source, 42
 - 63 help for commands, 35
 - 64 important concepts, 36
 - 65 login scripts, 44
 - 66 login shell, 37
 - 67 non-standard commands, 35
 - 68 path vs PATH, 40

- 1 scripts, 37
- 2 shell variables, 40
- 3 shells, 37
- 4 suggested references, 44
- 5 working environment, 38, 46
- 6 **UPS**
- 7 product conflicts, 202
- 8 product header files, 201
- 9 products area, 127
- 10 **UPS/UPD, 25, 25**
- 11 databases, 109
- 12 features, 109
- 13 **UPS/UPD:initialization, 152**
- 14 **UPS/UPD:qualifiers, 153**
- 15 **UPS:product dependency lists, 203**
- 16 **user code, 8**
- 17 viewing png files, 333
- 18 **Workbook, 12**
- 19 build window
- 20 contents, 173
- 21 setup, 171
- 22 disk space, 152
- 23 event-data files, 128
- 24 **FHiCL files**
- 25 machine-independent, 177
- 26 multi-site usage, 202
- 27 obtain code, 167
- 28 setup to run exercises, 129
- 29 initial, 129, 131
- 30 self-managed, 131
- 31 standard, 129
- 32 subsequent logins, 132
- 33 source directory contents, 170
- 34 toy experiment, *see* toy experiment
- 35 **Unix environment, 39**