

Intensity Frontier

Common Offline Documentation: *art* Workbook and Users Guide

Alpha Release 0.80

August 4, 2014

This version of the documentation is written for version v0_00_25 of the art-workbook code.

Scientific Computing Division
Future Programs and Experiments Department
Scientific Software Infrastructure Group

Principal Author: Rob Kutschke

Editor: Anne Heavey

art Developers: L. Garren, C. Green,
J. Kowalkowski, M. Paterno and P. Russo

List of Chapters

Detailed Table of Contents	iv
List of Figures	xx
List of Tables	xxii
List of Code and Output Listings	xxii
<i>art</i> Glossary	xxv
I Introduction	1
1 How to Read this Documentation	2
2 Conventions Used in this Documentation	4
3 Introduction to the <i>art</i> Event Processing Framework	7
4 Unix Prerequisites	34
5 Site-Specific Setup Procedure	45

6	Get your C++ up to Speed	47
7	Using External Products in UPS	105
II	Workbook	117
8	Preparation for Running the Workbook Exercises	118
9	Exercise 1: Running Pre-built <i>art</i> Modules	122
10	Exercise 2: Building and Running Your First Module	161
11	General Setup for Login Sessions	217
12	Keeping Up to Date with Workbook Code and Documentation	219
13	Exercise 3: Some other Member Functions of Modules	225
14	Exercise 4: A First Look at Parameter Sets	237
15	Exercise 5: Making Multiple Instances of a Module	262
16	Exercise 6: Accessing Data Products	269
17	Exercise 7: Making a Histogram	288
18	Exercise 8: Looping Over Collections	313
19	Troubleshooting	337

III	User's Guide	339
20	Obtaining Credentials to Access Fermilab Computing Resources	340
21	git	342
22	<i>art</i> Run-time and Development Environments	351
23	<i>art</i> Framework Parameters	359
24	Job Configuration in <i>art</i> : FHiCL	365
IV	Appendices	384
A	<i>art</i> Completion Codes	385
B	Viewing and Printing Figure Files	388
C	CLHEP	390
V	Index	406
	Index	407

Detailed Table of Contents

Detailed Table of Contents	iv
List of Figures	xx
List of Tables	xxii
List of Code and Output Listings	xxii
<i>art</i> Glossary	xxv
I Introduction	1
1 How to Read this Documentation	2
1.1 If you are new to HEP Software...	2
1.2 If you are an HEP Software expert...	2
1.3 If you are somewhere in between...	3
2 Conventions Used in this Documentation	4
2.1 Terms in Glossary	4
2.2 Typing Commands	4
2.3 Procedures to Follow	5
2.4 Important Items to Call Out	5
2.5 Site-specific Information	6
3 Introduction to the <i>art</i> Event Processing Framework	7
3.1 What is <i>art</i> and Who Uses it?	7

3.2	Why <i>art</i> ?	8
3.3	C++ and C++11	9
3.4	Getting Help	9
3.5	Overview of the Documentation Suite	9
3.5.1	The Introduction	11
3.5.2	The Workbook	11
3.5.3	Users Guide	12
3.5.4	Reference Manual	12
3.5.5	Technical Reference	12
3.5.6	Glossary	12
3.6	Some Background Material	12
3.6.1	Events and Event IDs	13
3.6.2	<i>art</i> Modules and the Event Loop	14
3.6.3	Module Types	18
3.6.4	<i>art</i> Data Products	19
3.6.5	<i>art</i> Services	20
3.6.6	Dynamic Libraries and <i>art</i>	21
3.6.7	Build Systems and <i>art</i>	22
3.6.8	External Products	23
3.6.9	The Event-Data Model and Persistency	24
3.6.10	Event-Data Files	25
3.6.11	Files on Tape	26
3.7	The Toy Experiment	26
3.7.1	Toy Detector Description	27
3.7.2	Workflow for Running the Toy Experiment Code	29
3.8	Rules, Best Practices, Conventions and Style	33
4	Unix Prerequisites	34
4.1	Introduction	34
4.2	Commands	34
4.3	Shells	36
4.4	Scripts: Part 1	36
4.5	Unix Environments	37
4.5.1	Building up the Environment	37
4.5.2	Examining and Using Environment Variables	38

4.6	Paths and \$PATH	39
4.7	Scripts: Part 2	41
4.8	bash Functions and Aliases	42
4.9	Login Scripts	43
4.10	Suggested Unix and bash References	43
5	Site-Specific Setup Procedure	45
6	Get your C++ up to Speed	47
6.1	Introduction	47
6.2	File Types Used and Generated in C++ Programming	47
6.3	Establishing the Environment	49
6.3.1	Initial Setup	49
6.3.2	Subsequent Logins	50
6.4	C++ Exercise 1: Basic C++ Syntax and Building an Executable	50
6.4.1	Concepts to Understand	50
6.4.2	How to Compile, Link and Run	51
6.4.3	Discussion	53
6.4.3.1	Primitive types	54
6.4.3.2	Arrays	54
6.4.3.3	Equality testing	55
6.4.3.4	Conditionals	55
6.4.3.5	Some C++ Standard Library Types	56
6.4.3.6	Pointers	56
6.4.3.7	References	57
6.4.3.8	Loops	58
6.5	C++ Exercise 2: About Compiling and Linking	58
6.5.1	What You Will Learn	58
6.5.2	The Source Code for this Exercise	58
6.5.3	Compile, Link and Run the Exercise	60
6.5.4	Alternate Script <code>build2</code>	64
6.5.5	Suggested Homework	65
6.6	C++ Exercise 3: Libraries	67
6.6.1	What You Will Learn	67
6.6.2	Building and Running the Exercise	67

6.7	Classes	71
6.7.1	Introduction	71
6.7.2	C++ Exercise 4 v1: The Most Basic Version	73
6.7.3	C++ Exercise 4 v2: The Default Constructor	79
6.7.4	C++ Exercise 4 v3: Constructors with Arguments	81
6.7.5	C++ Exercise 4 v4: Colon Initializer Syntax	84
6.7.6	C++ Exercise 4 v5: Member functions	85
6.7.7	C++ Exercise 4 v6: Private Data and Accessor Methods	90
6.7.7.1	Setters and Getters	90
6.7.7.2	What's the deal with the underscore?	95
6.7.7.3	An example to motivate private data	95
6.7.8	C++ Exercise 4 v7: The <code>inline</code> Specifier	96
6.7.9	C++ Exercise 4 v8: Defining Member Functions within the Class Declaration	98
6.7.10	C++ Exercise 4 v9: The stream insertion operator	99
6.7.11	Review	102
6.8	Overloading functions	103
6.9	C++ References	103
7	Using External Products in UPS	105
7.1	The UPS Database List: PRODUCTS	105
7.2	UPS Handling of Variants of a Product	107
7.3	The setup Command: Syntax and Function	107
7.4	Current Versions of Products	109
7.5	Environment Variables Defined by UPS	109
7.6	Finding Header Files	110
7.6.1	Introduction	110
7.6.2	Finding <i>art</i> Header Files	111
7.6.3	Finding Headers from Other UPS Products	113
7.6.4	Exceptions: The Workbook, ROOT and Geant4	114
II	Workbook	117
8	Preparation for Running the Workbook Exercises	118

8.1	Introduction	118
8.2	Getting Computer Accounts on Workbook-enabled Machines	118
8.3	Choosing a Machine and Logging In	119
8.4	Launching new Windows: Verify X Connectivity	120
8.5	Choose an Editor	120
9	Exercise 1: Running Pre-built <i>art</i> Modules	122
9.1	Introduction	122
9.2	Prerequisites	122
9.3	What You Will Learn	122
9.4	The <i>art</i> Run-time Environment	123
9.5	The Input and Configuration Files for the Workbook Exercises	124
9.6	Setting up to Run Exercise 1	125
9.6.1	Log In and Set Up	125
9.6.1.1	Initial Setup Procedure using Standard Directory	125
9.6.1.2	Initial Setup Procedure allowing Self-managed Working Directory	127
9.6.1.3	Setup for Subsequent Exercise 1 Login Sessions	128
9.7	Execute <i>art</i> and Examine Output	128
9.8	Understanding the Configuration	130
9.8.1	Some Bookkeeping Syntax	130
9.8.2	Some Physics Processing Syntax	132
9.8.3	<i>art</i> Command line Options	134
9.8.4	Maximum Number of Events to Process	135
9.8.5	Changing the Input Files	136
9.8.6	Skipping Events	138
9.8.7	Identifying the User Code to Execute	139
9.8.8	Paths and the <i>art</i> Workflow	141
9.8.8.1	Paths and the <i>art</i> Workflow: Details	143
9.8.8.2	Order of Module Execution	145
9.8.9	Writing an Output File	146
9.9	Understanding the Process for Exercise 1	147
9.9.1	Follow the Site-Specific Setup Procedure (Details)	148
9.9.2	Make a Working Directory (Details)	149
9.9.3	Setup the toyExperiment UPS Product (Details)	149

9.9.4	Copy Files to your Current Working Directory (Details)	150
9.9.5	Source <code>makeLinks.sh</code> (Details)	151
9.9.6	Run <i>art</i> (Details)	151
9.10	How does <i>art</i> find Modules?	152
9.11	How does <i>art</i> find FHiCL Files?	153
9.11.1	The <code>-c</code> command line argument	153
9.11.2	<code>#include</code> Files	154
9.12	Review	155
9.13	Test your Understanding	155
9.13.1	Tests	155
9.13.2	Answers	159
10	Exercise 2: Building and Running Your First Module	161
10.1	Introduction	161
10.2	Prerequisites	162
10.3	What You Will Learn	163
10.4	Initial Setup to Run Exercises	164
10.4.1	“Source Window” Setup	164
10.4.2	Examine Source Window Setup	165
10.4.2.1	About <code>git</code> and What it Did	165
10.4.2.2	Contents of the Source Directory	167
10.4.3	“Build Window” Setup	168
10.4.3.1	Standard Procedure	168
10.4.3.2	Using Self-managed Working Directory	169
10.4.4	Examine Build Window Setup	170
10.5	The <i>art</i> Development Environment	174
10.6	Running the Exercise	176
10.6.1	Run <i>art</i> on <code>first.fcl</code>	176
10.6.2	The FHiCL File <code>first.fcl</code>	178
10.6.3	The Source Code File <code>First_module.cc</code>	178
10.6.3.1	The <code>#include</code> Statements	179
10.6.3.2	The Declaration of the Class <code>First</code> , an Analyzer Module	181
10.6.3.3	An Introduction to Analyzer Modules	182
10.6.3.4	The Constructor for the Class <code>First</code>	184

10.6.3.5	Aside: Omitting Argument Names in Function Declarations	184
10.6.3.6	The Member Function <code>analyze</code> and the Representation of an Event	185
10.6.3.7	Representing an Event Identifier with <code>art::EventID</code>	187
10.6.3.8	<code>DEFINE_ART_MACRO</code> : The Module Maker Macros	189
10.6.3.9	Some Alternate Styles	190
10.7	What does the Build System Do?	193
10.7.1	The Basic Operation	193
10.7.2	Incremental Builds and Complete Rebuilds	195
10.7.3	Finding Header Files at Compile Time	196
10.7.4	Finding Dynamic Library Files at Link Time	198
10.7.5	Build System Details	200
10.8	Suggested Activities	201
10.8.1	Create Your Second Module	201
10.8.2	Use <i>artmod</i> to Create Your Third Module	202
10.8.3	Running Many Modules at Once	205
10.8.4	Access Parts of the EventID	206
10.9	Final Remarks	208
10.9.1	Why is there no <code>First_module.h</code> File?	208
10.9.2	The Three-File Module Style	209
10.10	Flow of Execution from Source to FHiCL File	211
10.11	Review	212
10.12	Test Your Understanding	212
10.12.1	Tests	212
10.12.2	Answers	215
10.12.2.1	FirstBug01	215
10.12.2.2	FirstBug02	215
11	General Setup for Login Sessions	217
11.1	Source Window	217
11.2	Build Window	218
12	Keeping Up to Date with Workbook Code and Documentation	219
12.1	Introduction	219

12.2	Special Instructions for Summer 2014	219
12.3	How to Update	220
12.3.1	Get Updated Documentation	221
12.3.2	Get Updated Code and Build It	221
12.3.3	See which Files you have Modified or Added	223
13	Exercise 3: Some other Member Functions of Modules	225
13.1	Introduction	225
13.2	Prerequisites	226
13.3	What You Will Learn	226
13.4	Setting up to Run this Exercise	227
13.5	The Source File <code>Optional_module.cc</code>	227
13.5.1	About the <code>begin*</code> Member Functions	228
13.5.2	About the <code>art::*ID</code> Classes	228
13.5.3	Use of the <code>override</code> Identifier	229
13.5.4	Use of <code>const</code> References	229
13.5.5	The <code>analyze</code> Member Function	230
13.6	Running this Exercise	231
13.7	The Member Function <code>beginJob</code> versus the Constructor	232
13.8	Suggested Activities	233
13.8.1	Add the Matching <code>end</code> Member functions	233
13.8.2	Run on Multiple Input Files	233
13.8.3	The Option <code>--trace</code>	234
13.9	Review	234
13.10	Test Your Understanding	235
13.10.1	Tests	235
13.10.2	Answers	236
14	Exercise 4: A First Look at Parameter Sets	237
14.1	Introduction	237
14.2	Prerequisites	238
14.3	What You Will Learn	238
14.4	Setting up to Run this Exercise	239
14.5	The Configuration File <code>pset01.fcl</code>	240
14.6	The Source code file <code>PSet01_module.cc</code>	242

14.7	Running the Exercise	245
14.8	Member Function Templates and their Arguments	248
14.8.1	Types Known to <code>ParameterSet::get<T></code>	248
14.8.2	User Defined Types	249
14.9	Exceptions	249
14.9.1	Error Conditions	249
14.9.2	Error Handling	250
14.9.3	Suggested Exercises	251
14.10	Parameters and Data Members	252
14.11	Optional Parameters with Default Values	253
14.11.1	Policies About Optional Parameters	255
14.12	Numerical Types: Precision and Canonical Forms	255
14.12.1	Why Have Canonical Forms?	257
14.13	Suggested Exercises	258
14.13.1	Formats	258
14.13.2	Fractional versus Integral Types	259
14.14	Review	259
14.15	Test Your Understanding	260
14.15.1	Tests	260
14.15.2	Answers	261
15	Exercise 5: Making Multiple Instances of a Module	262
15.1	Introduction	262
15.2	Prerequisites	262
15.3	What You Will Learn	262
15.4	Setting up to Run this Exercise	263
15.5	The Source File	
	file <code>Magic_module.cc</code>	263
15.6	The FHiCL File <code>magic.fcl</code>	264
15.7	Running the Exercise	264
15.8	Discussion	265
15.8.1	Order of Analyzer Modules is not Important	265
15.8.2	Two Meanings of <i>Module Label</i>	266
15.9	Review	266
15.10	Test Your Understanding	267

15.10.1 Tests	267
15.10.2 Answers	268
16 Exercise 6: Accessing Data Products	269
16.1 Introduction	269
16.2 Prerequisites	269
16.3 What You Will Learn	270
16.4 Background Information for this Exercise	270
16.4.1 The Data Type <code>GenParticleCollection</code>	271
16.4.2 Data Product Names	272
16.4.3 Specifying a Data Product	274
16.4.4 The Data Product used in this Exercise	275
16.5 Setting up to Run this Exercise	275
16.6 Running the Exercise	276
16.7 Understanding the First Version, <code>ReadGens1</code>	276
16.7.1 The Source File <code>ReadGens1_module.cc</code>	276
16.7.2 Adding a Link Library to <code>CMakeLists.txt</code>	280
16.7.3 The FHiCL File <code>readGens1.fcl</code>	280
16.8 The Second Version, <code>ReadGens2</code>	281
16.9 The Third Version, <code>ReadGens3</code>	282
16.10 Suggested Activities	283
16.11 Review	285
16.12 Test Your Understanding	286
16.12.1 Tests	286
16.12.2 Answers	286
17 Exercise 7: Making a Histogram	288
17.1 Introduction	288
17.2 Prerequisites	289
17.3 What You Will Learn	289
17.4 Setting up to Run this Exercise	290
17.5 The Source File <code>FirstHist1_module.cc</code>	291
17.5.1 Introducing <code>art::ServiceHandle</code>	294
17.5.2 Creating a Histogram	295
17.5.3 Filling a Histogram	296

17.5.4 A Few Last Comments	297
17.6 The Configuration File C++ firstHist1.fcl	297
17.7 The file CMakeLists.txt	298
17.8 Running the Exercise	299
17.9 Inspecting the Histogram File	301
17.9.1 A Short Cut: the browse command	304
17.9.2 Using CINT Scripts	304
17.10 Finding ROOT Documentation	308
17.10.1 Overwriting Histogram Files	308
17.10.2 Changing the Name of the Histogram File	309
17.10.3 Changing the Module Label	310
17.10.4 Printing From the TBrowser	310
17.11 Review	310
17.12 Test Your Understanding	311
17.12.1 Tests	311
17.12.2 Answers	312

18 Exercise 8: Looping Over Collections 313

18.1 Introduction	313
18.2 Prerequisites	313
18.3 What You Will Learn	313
18.4 Setting Up to Run Exercise	314
18.5 The Class GenParticle	315
18.5.1 The Included Header Files	315
18.5.2 Particle Parent-Child Relationships	317
18.5.3 The Public Interface for the Class GenParticle	317
18.5.4 Conditionally Excluded Sections of Header File	320
18.6 The Module LoopGens1	321
18.7 CMakeLists.txt	323
18.8 Running the Exercise	323
18.9 Variations on the Exercise	325
18.9.1 LoopGens2_module.cc	325
18.9.2 LoopGens3_module.cc	326
18.9.3 LoopGens3a_module.cc	329
18.10 Review	330

18.11	Test Your Understanding	330
18.11.1	Test 1	330
18.11.2	Test 2	333
18.11.3	Test 3	333
18.11.4	Answers	335
19	Troubleshooting	337
19.1	Updating Workbook Code	337
19.2	XWindows (xterm and Other XWindows Products)	337
19.2.1	Mac OSX 10.9	337
19.3	Trouble Building	337
19.4	<i>art</i> Won't Run	338
III	User's Guide	339
20	Obtaining Credentials to Access Fermilab Computing Resources	340
20.1	Kerberos Authentication	340
20.2	Fermilab Services Account	341
21	git	342
21.1	Aside: More Details about git	343
21.1.1	Central Repository, Local Repository and Working Directory . . .	343
21.1.1.1	Files that you have Added	344
21.1.1.2	Files that you have Modified	344
21.1.1.3	Files with Resolvable Conflicts	345
21.1.1.4	Files with Unresolvable Conflicts	345
21.1.2	git Branches	345
21.1.3	Seeing which Files you have Modified or Added	348
22	<i>art</i> Run-time and Development Environments	351
22.1	The <i>art</i> Run-time Environment	351
22.2	The <i>art</i> Development Environment	355
23	<i>art</i> Framework Parameters	359
23.1	Parameter Types	359

23.2	Structure of <i>art</i> Configuration Files	360
23.3	Services	362
23.3.1	System Services	362
23.3.2	FloatingPointControl	362
23.3.3	Message Parameters	364
23.3.4	Optional Services	364
23.3.5	Sources	364
23.3.6	Modules	364
24	Job Configuration in <i>art</i>: FHiCL	365
24.1	Basics of FHiCL Syntax	365
24.1.1	Specifying Names and Values	365
24.1.2	FHiCL-reserved Characters and Identifiers	368
24.2	FHiCL Identifiers Reserved to <i>art</i>	369
24.3	Structure of a FHiCL Run-time Configuration File for <i>art</i>	370
24.4	Order of Elements in a FHiCL Run-time Configuration File for <i>art</i>	374
24.5	The <i>physics</i> Portion of the FHiCL Configuration	376
24.6	Choosing and Using Module Labels and Path Names	377
24.7	Scheduling Strategy in <i>art</i>	378
24.8	Scheduled Reconstruction using Trigger Paths	381
24.9	Reconstruction On-Demand	383
24.10	Bits and Pieces	383
IV	Appendices	384
A	<i>art</i> Completion Codes	385
B	Viewing and Printing Figure Files	388
B.1	Viewing Figure Files Interactively	388
B.2	Printing Figure Files	389
C	CLHEP	390
C.1	Introduction	390
C.2	Multiple Meanings of <i>Vector</i> in CLHEP	391
C.3	CLHEP Documentation	391

C.4	CLHEP Header Files	392
C.4.1	Naming Conventions and Syntax	392
C.4.2	.icc Files	392
C.5	The CLHEP Namespace	393
C.5.1	using Declarations and Directives	393
C.6	The Vector Package	394
C.6.1	CLHEP::Hep3Vector	394
C.6.1.1	Some Fragile Member Functions	399
C.6.2	CLHEP::HepLorentzVector	400
C.6.2.1	HepBoost	404
C.7	The Matrix Package	404
C.8	The Random Package	404

V Index 406

Index 407

List of Figures

3.1	Principal components of the <i>art</i> documentation suite	10
3.2	Flowchart describing the <i>art</i> event loop	17
3.3	Geometry of the toy experiment's detector	27
3.4	Event display of a simulated event in the toy detector.	30
3.5	Event display of another simulated event in the toy detector	31
3.6	Invariant mass of reconstructed pairs of oppositely charged tracks	32
4.1	Computing environment hierarchies	38
6.1	Memory diagram at the end of a run of <code>Classes/v1/ptest.cc</code>	78
6.2	Memory diagram at the end of a run of <code>Classes/v6/ptest.cc</code>	94
9.1	Elements of the <i>art</i> run-time environment for the first exercise	124
10.1	Representation of reader's source directory structure	166
10.2	Representation of reader's build directory structure	171
10.3	Elements of the <i>art</i> development environment and information flow	175
10.4	Reader's directory structure once development environment is established	177
14.1	Parameter set <code>psetTester</code> from <code>pset01.fcl</code>	241
14.2	First part of constructor in <code>PSet01_module.cc</code>	243
14.3	Remainder of the constructor in <code>PSet01_module.cc</code>	246
17.1	TBrowser window after opening <code>output/firstHist1.root</code>	303
17.2	TBrowser window after displaying the histogram <code>hNGens;1</code>	303
17.3	Figure made by running the CINT script <code>drawHist1.C</code>	307
18.1	Histograms made by <code>loopGens1.fcl</code>	324

21.1 Illustration of git branches, simple	346
21.2 Illustration of git branches	349
22.1 <i>art</i> run-time environment (same as Figure 9.1)	352
22.2 <i>art</i> run-time environment (everything pre-built)	353
22.3 <i>art</i> run-time environment (with officially tracked inputs)	354
22.4 <i>art</i> development environment for Workbook (same as Figure 10.3)	356
22.5 <i>art</i> development environment (for building full code base)	357
22.6 <i>art</i> development environment (for building against prebuilt base)	358

List of Tables

3.1	Compiler flags for the optimization levels defined by cetbuildtools	22
3.2	Units used in the Workbook	28
5.1	Site-specific setup procedures for experiments that run <i>art</i>	46
7.1	Namespaces for selected UPS products	114
8.1	Experiment-specific Information for new users	119
8.2	Login machines for running the Workbook exercises	120
9.1	Input files provided for the Workbook exercises	125
10.1	Compiler and linker flags for a profile build	201
14.1	Canonical forms of numerical values in FHiCL files	256
23.1	<i>art</i> Floating Point Parameters	363
23.2	<i>art</i> Message Parameters	364
A.1	<i>art</i> completion status codes. The return code is the least significant byte of the status code.	387
C.1	Selected member functions of <code>CLHEP::Hep3Vector</code>	398

List of Code and Output Listings

5.1	NOvA <code>setup_art_workbook</code> script	46
6.1	Layout of a class.	72
6.2	File <code>Point.h</code> with the simplest version of the class <code>Point</code>	75
6.3	The contents of <code>v1/ptest.cc</code>	76
9.1	Sample output from running <code>hello.fcl</code>	128
9.2	Example of the value of <code>LD_LIBRARY_PATH</code>	152
10.1	Example of output created by <code>setup_for_development</code>	170
10.2	The contents of <code>First_module.cc</code>	180
10.3	An alternate layout for <code>First_module.cc</code>	192
10.4	The file <code>art-workbook/FirstModule/CMakeLists.txt</code>	200
10.5	The <code>physics</code> parameter set for <code>all.fcl</code>	206
10.6	The contents of <code>First.h</code> in the three-file model	209
10.7	The contents of <code>First.cc</code> in the three-file model	210
10.8	The contents of <code>First_module.cc</code> in the three-file model	210
12.1	Example of the output produced by <code>git pull</code>	223
13.1	Output from <code>Optional_module.cc</code> with <code>optional.fcl</code>	231
14.1	Output from <code>PSet01</code> with <code>pset01.fcl</code> (<i>art</i> -standard output not shown)	247
14.2	Output from <code>PSet01</code> with <code>pset01.fcl</code> (parameter <code>b</code> removed)	252
14.3	Output from <code>PSet01</code> with <code>pset01.fcl</code> (parameter <code>c</code> misdefined)	252
14.4	Output from <code>PSet02</code> with <code>pset02.fcl</code>	253
14.5	Parameter-related portion of output from <code>PSet03</code> with <code>pset03.fcl</code>	255
14.6	Output from <code>PSet04</code> with <code>pset04.fcl</code>	258
14.7	Output from <code>PSet04</code> with modified <code>pset04.fcl</code> (intentional error)	259
15.1	Output using <code>magic.fcl</code>	265
16.1	Contents of <code>GenParticleCollection.h</code>	271
16.2	Output using <code>readGens1.fcl</code>	276

16.3	Include statements in <code>ReadGens1_module.cc</code>	276
16.4	Configuring the module label read in <code>readGens1.fcl</code>	280
16.5	Warning message for misspelled module label of data product	284
16.6	Exception message for <code>ProductNotFound</code> , default Exceptions disabled	285
17.1	Declaration of the class <code>FirstHist1</code> from <code>FirstHist1_module.cc</code>	292
17.2	Implementation of the class <code>FirstHist1</code>	293
17.3	Creating histogram object of type <code>TH1D</code>	295
17.4	Filling the histogram	296
17.5	<code>TFileService</code> in <code>firstHist1.fcl</code>	297
17.6	<code>firstHist1.fcl</code>	298
17.7	<code>CMakeLists.tex</code> in the directory <code>FirstHistogram</code>	300
17.8	Sample CINT file <code>DrawHist1.C</code>	306
18.1	The const accessors of class <code>GenParticle</code> . 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.	318
18.2	The loop over the generated particles, from <code>LoopGens1_module.cc</code>	321
f	323
A.1	Error Codes when an <code>art::Exception</code> is caught in main	386
C.1	The constructors of <code>Hep3Vector</code>	395

art Glossary

abstraction	the process by which data and programs are defined with a representation similar in form to its meaning (semantics), while hiding away the implementation details. A system can have several abstraction layers whereby different meanings and amounts of detail are exposed to the programmer (adapted from Wikipedia’s entry for “Abstraction (computer science)”).
analyzer module	an <i>art</i> module that may read information from the current event but that may not add information to it; e.g., a module to fill histograms or make printed output
API	Application Programming Interface
art	The <i>art</i> framework (<i>art</i> is not an acronym) is the software framework developed for common use by the Intensity Frontier experiments to develop their offline code and non-real-time online code
art module	see <i>module</i>
art path	a FHiCL sequence of <i>art</i> moduleLabels that specifies the work the job will do
artdaq	a toolkit that lives on top of <i>art</i> for building high-performance event-building and event-filtering systems; this toolkit is designed to support efficient use of multi-core computers and GPUs. A technical paper on <i>artdaq</i> can be found at http://inspirehep.net/record/1229212?ln=en ; the <i>artdaq</i> home page is at https://cdcv.sfnal.gov/redmine/projects/artdaq/wiki .
bash	a UNIX shell scripting language that is used by some of the support

	scripts in the workbook exercises
boost	a class library with new functionality that is being prototyped for inclusion in future C++ standards
build system	turns source code into object files, puts them into a dynamic library, links them with other libraries, and may also run tests, deploy code to production systems and create some documentation.
buildtool	a Fermilab-developed tool (part of cetbuildtools) to compile, link and run tests on the source code of the Workbook
catch	See <i>exception</i> in a C++ reference
cetbuildtools	the build system that is used by the <i>art</i> Workbook (and by <i>art</i> itself).
CETLIB	a utility library used by <i>art</i> (developed and maintained by the <i>art</i> team) to hold information that does not fit naturally into other libraries
class	The C++ programming language allows programmers to define program-specific data types through the use of <i>classes</i> . Classes define types of data structures and the functions that operate on those data structures. Instances of these data types are known as <i>objects</i> . Other object oriented languages have similar concepts.
CLHEP	a set of utility classes; the name is an acronym for a Class Library for HEP
collection	
configuration	see <i>run-time configuration</i>
const member function	a member function of a class that does not change the value of non-mutable data members; see <i>mutable data member</i>
constructor	a function that (a) shares an identifier with its associated class, and (b) initializes the members of an object instantiated from this class
DAQ	data aquisition system
data handling	Broadly, data handling includes “whatever it takes to store data or to get at data.” This includes collection, aggregation, access control, sharing,

	transfer, bookkeeping, caching, storage, and archiving.
Data Model	see <i>Event Data Model</i>
data product	Experiment-defined class that can represent detector signals, reconstructed data, simulated events, etc. In <i>art</i> , a <i>data product</i> is the smallest unit of information that can be added to or retrieved from an event.
data type	See <i>type</i>
declaration (of a class)	the portion of a class that specifies its type, its name, and any data members and/or member functions it has
definition (of a class)	see <i>instantiation</i>
destructor	a function that (a) has the same identifier as its associated class but prefaced with a tilde (~), and (b) is used to deallocate memory and do other cleanup for a class object and its class members when the object is destroyed
Doxygen	a system of producing reference documentation based on comments in source code
ED	a prefix used in <i>art</i> (e.g., for module types) meaning <i>event-data</i>
EDAnalyzer	see <i>analyzer module</i>
EDFilter	see <i>filter module</i>
EDOutput	see <i>output module</i>
EDProducer	see <i>producer module</i>
EDSource	see <i>source module</i>
Event	In HEP there are two notions of the word <i>event</i> that are in common use; see <i>event (unit of information)</i> or <i>event (interaction)</i> . In this documentation suite, unless otherwise indicated, we mean the former.
Event (interaction)	An <i>event (unit of data)</i> may contain more than one fundamental interaction; the science goal is always to identify individual fundamental interactions and determine their properties. It is common to use the word <i>event</i> to refer to one of the individual fundamental interactions.

In the near detector of a high-intensity neutrino experiment, for example, there may be multiple neutrino interactions within the unit of time that defines a single *event (unit of information)*. Similarly, in a colliding-beam experiment, an *event (unit of information)* corresponds to the information from one beam crossing, during which time there may be multiple collisions between beam particles.

Event (unit of information) In the general HEP sense, an *event* is a set of raw data associated in time, plus any information computed from the raw data; *event* may also refer to a simulated version of same. Within *art*, the representation of an *event (unit of information)* is the class `art::Event`, which is the smallest unit of information that *art* can process. An `art::Event` contains an event identifier plus an arbitrary number of *data-products*; the information within the *data-products* is intrinsically experiment dependent and is defined by each experiment. For bookkeeping convenience, *art* groups events into a hierarchy: a *run* contains zero or more *subRuns* and a *subRun* contains zero or more events.

Event Data Model (EDM) Representation of the data that an experiment collects, all the derived information, and historical records necessary for reproduction of result

event loop within an *art* job, the set of steps to perform in order to execute the per-event functions for each event that is read in, including steps for begin/end-job, begin/end-run and begin/end-subRun

event-data all of the data products in an experiment's files plus the metadata that accompanies them. The HEP software community has adopted the word *event-data* to refer to the software details of dealing with the information found in *events*, whether the events come from experimental data or simulations.

event-data file a file of event-data, containing either experimental data or simulated events

exception, to throw a mechanism in C++ (and other programming languages) to stop the current execution of a program and transfer control up the call chain; also called *catch*

experiment code	see <i>user code</i>
external product	for a given experiment, this is a software product that the experiment's software (within the <i>art</i> framework) does not build, but that it uses; e.g., ROOT, Geant4, etc. At Fermilab external products are managed by the in-house UPS/UPD system, and are often called <i>UPS products</i> or simply <i>products</i> .
FermiGrid	a batch system for submitting jobs that require large amounts of CPU time
FHiCL	Fermilab Hierarchical Configuration Language (pronounced “fickle”), a language developed and maintained by the <i>art</i> team at Fermilab to support run-time configuration for several projects, including <i>art</i>
FHiCL-CPP	the C++ toolkit used to read FHiCL documents within <i>art</i>
filter module	an <i>art</i> module that may alter the flow of processing modules within an event; it may add information to the event
framework (<i>art</i>)	The <i>art</i> framework is an application 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. <i>art</i> also provides infrastructure for common tasks, such as reading input, writing output, provenance tracking, database access and run-time configuration.
framework (generic)	an abstraction in which software providing generic functionality can be selectively changed by additional user-written code, thus providing application-specific software (significantly abbreviated from Wikipedia's entry for “software framework”); note that the actual functionality provided by any given framework, e.g., <i>art</i> , will be tailored to the given needs.
free function	a function without data members; it knows only about arguments passed to it at run time; see <i>function</i> and <i>member function</i>
Geant4	a toolkit for the simulation of the passage of particles through matter, developed at CERN. http://geant4.cern.ch/

git	a source code management system used to manage files in the <i>art</i> Workbook; similar in concept to the older CVS and SVN, but with enhanced functionality
handle	a type of smart pointer that permits the viewing of information inside a data product but does not allow modification of that information; see <i>pointer,data product</i>
IF	Intensity Frontier
ifdh_sam	a UPS product that allows <i>art</i> to use <i>SAM</i> as an external run-time agent that can deliver remote files to local disk space and can copy output files to tape. The first part of the name is an acronym for Intensity Frontier Data Handling.
implementation	the portion of C++ code that specifies the functionality of a declared data type; where as a struct or class declaration (of a data type) usually resides in a <i>header</i> file (.h or .hh), the implementation usually resides in a separate source code file (.cc) that “#includes” the header file
instance	see <i>instantiation</i>
instantiation	the creation of an object instance of a class in an OOP language; an instantiated object is given a name and created in memory or on disk using the structure described within its class declaration; referred to as <i>definition</i> in the C++ standard.
jobsub-tools	a UPS product that supplies tools for submitting jobs to the Fermigrid batch system and monitoring them.
Kerberos	a single sign-on, strong authentication system required by Fermilab for access to its computing resources
kinit	a command for obtaining Kerberos credentials that allow access to Fermilab computing resources; see <i>Kerberos</i>
member function	(also called <i>method</i>) a function that is defined within (is a <i>member</i> of) a class; they define the behavior to be exhibited by instances of the associated class at program run time. At run time, member functions have access to data stored in the instance of the class with they are

	associated, and are thereby able to control or provide access to the state of the instance.
message facility	a UPS product used by <i>art</i> and experiments' code that provides facilities for merging messages with a variety of severity levels, e.g., informational, error, and so on; see also <i>mf</i>
message service	
method	see <i>member function</i>
mf	a namespace that holds classes and functions that make up the message facility used by <i>art</i> and by experiments that use <i>art</i> ; see <i>message facility</i>
module	a C++ class that obeys certain rules established by <i>art</i> and whose source code file gets compiled into a dynamic library that can be loaded at run-time by <i>art</i> . An <i>art</i> module “plugs into” a processing stream and performs a specific task on units of data obtained using the Event Data Model, independent of other running modules. See also <i>moduleLabel</i>
module_type	
moduleLabel	a user-defined identifier whose value is a parameter set that <i>art</i> will use to configure a module; see <i>module</i> and <i>parameter set</i>
Monte Carlo method	a class of computational algorithms that rely on repeated random sampling to obtain numerical results; i.e., by running simulations many times over in order to calculate those same probabilities heuristically just like actually playing and recording your results in a real casino situation: hence the name (Wikipedia)
mutable data member	The
namespace	a container within a file system for a set of identifiers (names); usually grouped by functionality, they are used to keep different subsets of code distinguishable from one another; identical names defined within different namespaces are disambiguated via their namespace prefix
ntuple	an ordered list of <i>n</i> elements used to describe objects such as vectors or tables

object	an instantiation of any data type, built-in types (e.g., int, double, float) or class types; i.e., a location range in memory containing an instantiation
object-oriented language	a programming language that supports OOP; this usually means support for classes, including public and private data and functions
object-oriented programming (OOP)	a programming language model organized around <i>objects</i> rather than procedures, where <i>objects</i> are quantities of interest that can be manipulated. (In contrast, programs have been viewed historically as logical procedures that read in data, process the data and produce output.) Objects are defined by <i>classes</i> that contain attributes (data fields that describe the objects) and associated procedures. See <i>C++ class</i> ; <i>object</i> .
OOP	see <i>object oriented programming</i>
output module	an <i>art</i> module that writes data products to output file(s); it may select a subset of data products in a subset of events; an <i>art</i> job contains zero or more output modules
overload resolution	the process by which a specific member of an <i>overload set</i> is chosen as the function to be used at a given function call site.
overload set	a set of functions with the same name, whether in the same namespace or in different namespaces, with differing number of types of arguments.
parameter set	a C++ class, defined by FHiCL-CPP, that is used to hold run-time configuration for <i>art</i> itself or for modules and services instantiated by <i>art</i> . In a FHiCL file, a parameter set is represented by a FHiCL <i>table</i> ; see <i>table</i>
path	a generic word based on the UNIX concept of PATH that refers to a colon-separated list of directories used by <i>art</i> when searching for various files (e.g., data input, configuration, and so on)
physics	in <i>art</i> , <i>physics</i> is the label for a portion of the run-time configuration of a job; this portion contains up to five sections, each labeled with a

	reserved
pointer	a variable whose value is the address of (i.e., that points to) a piece of information in memory. A native C++ pointer is often referred to as a <i>bare pointer</i> . <i>art</i> defines different sorts of <i>smart pointers</i> (or <i>safe pointers</i>) for use in different circumstances. One commonly used type of smart pointer is called a <i>handle</i> .
process_name	a parameter to which the user assigns a mnemonic value identifying the physics content of the associated FHiCL parameter set (i.e., the parameters used in the same FHiCL file). The process_name value is embedded into every data product created via the FHiCL file.
producer module	an <i>art</i> module that may read information from the current event and may add information to it
product	See either <i>external</i> product or <i>data</i> product
redmine	an open source, web-based project management and bug-tracking tool used as a repository for <i>art</i> code and related code and documentation
ROOT	an HEP data management and data presentation package supported by CERN; <i>art</i> , itself, uses ROOT for persistency of event-data — see ROOT files; user code often also uses ROOT to create histograms, ntuples, trees etc
ROOT files	a file written by ROOT — see ROOT. There are two types of ROOT files managed by <i>art</i> : (1) event-data files, and (2) the file managed by TFileService that holds user-defined histograms, ntuples, trees, etc.
run	a period of data collection, defined by the experiment (usually a period of time during which certain running conditions remain unchanged); a run contains zero or more subRuns
run-time configuration	(processing-related) structured documents describing all processing aspects of a single job including the specification of parameters and workflow; in <i>art</i> it is supplied by a FHiCL file; see <i>FHiCL</i>
safe pointer	see <i>pointer</i>

SAM	(Sequential data Access via Metadata) a Fermilab-supplied product that provides the functions of a file catalog, a replica manager and some functions of a batch-oriented workflow manager
scope	
sequence (in FHiCL)	one or more comma-separated FHiCL values delimited by square brackets (<div style="text-align: center; margin: 5px 0;">...</div>) in a FHiCL file is called a <i>sequence</i> (as distinct from a <i>table</i>)
service	in <i>art</i> , a singleton-like object (type) whose lifetime and configuration are managed by <i>art</i> , and which can be accessed by module code and by other services by requesting a <i>service handle</i> to that particular service. The service <i>type</i> is used to provide geometrical information, conditions and management of the random number state; it is also used to implement some internal functionality. See also <i>T File Service</i>
signature (of a function)	the unique identifier of a C++ function, which includes: (a) its name, including any class name or namespace components, (b) the number and type of its arguments, (c) whether it is a member function, (d) whether it is a const function (Note that the signature of a function does not include its return type.)
site	As used in the <i>art</i> documentation, a <i>site</i> is a unique combination of experiment and institution; used to refer to a set of computing resources configured for use by a particular experiment at a particular institution. This means that, for example, the Workbook environment on a Mu2e-owned computer at Fermilab will be different than that on an Mu2e-owned computer at LBL. Also, the Workbook environment on a Mu2e-owned computer at Fermilab will be different from that on an LBNE-owned computer at Fermilab.
smart pointer	see <i>pointer</i>
source	(refers to a <i>data</i> source) the name of the parameter set inside an FHiCL file describing the first step in the workflow for processing an event;

	it reads in each event sequentially from a data file or creates an empty event; see also <i>source code</i> ; see also <i>EDsource</i>
source code	code written in C++ (the programming language used with <i>art</i>) that requires compilation and linking to create an executable program
source module	an <i>art</i> module that can initiate an <i>art</i> path by reading in event(s) from a data file or by creating an empty event; it is the first step of the processing chain
standard library, C++	the C++ standard library, which includes, headers, templates and dynamic libraries
std	identifier for the namespace used by the C++ standard library
STL	the C++ Standard Template Library; an archaic name for the C++ Standard Library; see standard library, C++
struct	identical to a C++ class except all members are <i>public</i> (instead of <i>private</i>) by default
subRun	a period of data collection within a run, defined by the experiment (it may delineate a period of time during which certain run parameters remain unchanged); a SubRun is contained within a <i>run</i> ; a subRun contains zero or more <i>events</i>
table (in FHiCL)	a group of FHiCL definitions delimited by braces ({ ... }) is called a <i>table</i> ; within <i>art</i> , a FHiCL table gets turned into an object called a <i>parameter set</i> . Consequently, a FHiCL table is typically called a <i>parameter set</i> . See <i>parameter set</i> .
template (C++)	Templates are a feature of C++ that allows for meta-programming. In practical terms, the coder can write an algorithm that is independent of type, as long as the type supports the features required by the algorithm. For example, there is a standard library “sort” algorithm that will work for any type that provides a way to determine if one object of the type is “less than” another object of the type.
TFileService	an <i>art</i> service used by all experiments to give each module a ROOT subdirectory in which to place its own histograms, TTrees, and so on;

see *TTrees* and *ROOT*

truth information	One use of simulated events is to develop, debug and characterize the algorithms used in reconstruction and analysis. To assist in these tasks, the simulation code often creates data products that contain detailed information about the right answers at intermediate stages of reconstruction and analysis; they also write data products that allow the physicist to ask “is this a case in which there is an irreducible background or should I be able to do better?” This information is called the <i>truth information</i> , the <i>Monte Carlo truth</i> or the <i>God’s block</i> .
TTrees	a ROOT implementation of a tree; see <i>tree</i> and <i>ROOT</i>
type	Variables and objects in C++ must be classified into <i>types</i> , e.g., built-in types (integer, boolean, float, character, etc.), more complex user-defined classes/structures and typedefs; see <i>class</i> , <i>struct</i> , and <i>typedef</i> . The word <i>type</i> in the context of C++ and <i>art</i> is the same as <i>data type</i> unless otherwise stated.
typedef	A typedef is a different name, or an alias, by which a type can be identified. Type aliases can be used to reduce the length of long or confusing type names, but they are most useful as tools to abstract programs from the underlying types they use (cplusplus.com).
UPS/UPD	a Fermliab-developed system for distributing software products
user code	experiment-specific and/or analysis-specific C++ code that uses the <i>art</i> framework; this includes any personal code you write that uses <i>art</i> .
variable	a storage location and an associated symbolic name (an identifier) which contains some known or unknown quantity or information, a value. The variable name is the usual way to reference the stored value; this separation of name and content allows the name to be used independently of the exact information it represents.

Part I

Introduction

1 How to Read this Documentation

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.

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

1.3 If you are somewhere in between...

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

2 Conventions Used in this Documentation

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)!

2.1 Terms in Glossary

γ 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
```

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

```
mkdir \  
-p mydir
```

Computer output from a command is shown as:

```
command output
```

2.3 Procedures to Follow

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

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

```
mkdir -p mydir
```

2.4 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.




2.5 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 institution, and is used to refer to a set of computing resources configured for use by a particular experiment at a particular institution.



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

3.1 What is *art* and Who Uses it?

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

- analysis
- simulation

art is not designed for use in real-time environments, such as the direct interface with data-collection hardware.

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

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

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

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

3.2 Why *art*?

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

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

constraints that this separation imposes, it has improved the overall quality of the framework and reduced the duplicated effort.

3.3 C++ and C++11

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

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

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

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

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



3.4 Getting Help

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

3.5 Overview of the Documentation Suite

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

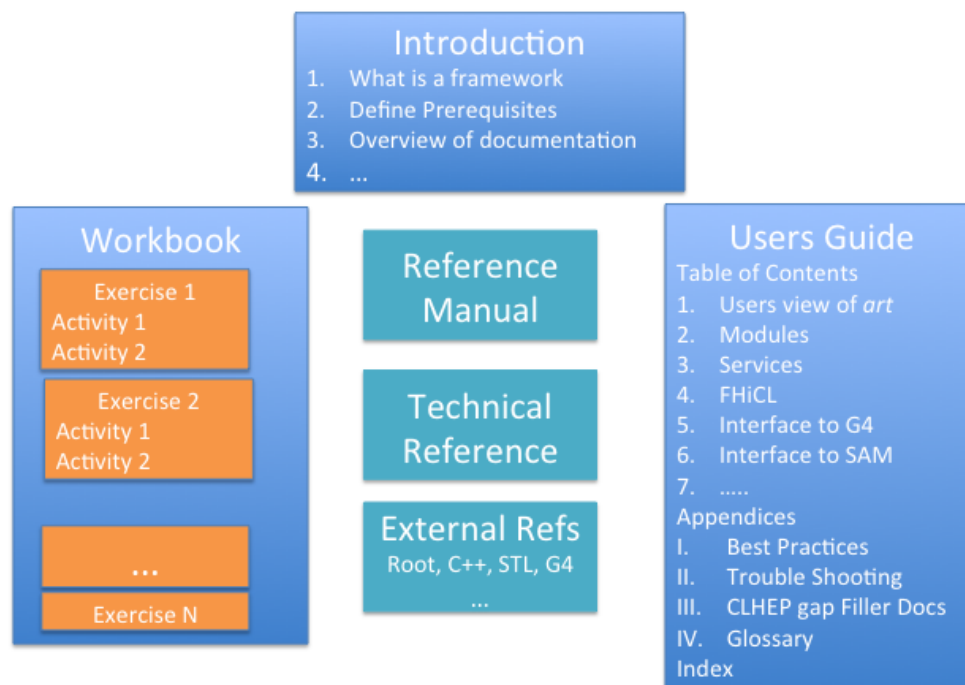


Figure 3.1: Principal components of the *art* documentation suite

3.5.1 The Introduction

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

3.5.2 The Workbook

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

One of the Workbook's primary functions is training readers how and where to find more extensive documentation on both *art* and external software tools; they will need this information as they move on to develop and use the scientific software for their experiment.

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

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

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

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, the subRun number is reset to 0 (0 not 1 for historical reasons) and the run number is incremented.
2. The second method is the same as the first except that the event number monotonically increases throughout a run and does not reset to 1 on subRun boundaries. The event number does reset to 1 at the start of each run.

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

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

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

3.6.2 *art* Modules and the Event Loop

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

The simplest command to run *art* looks like:

```
art -c filename.fcl
```

The argument to `-c` is the *run-time configuration file*(γ), a text file that tells one run of *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.

Hierarchical Configuration Language *FHiCL*(γ) (pronounced “fickle”) and the filenames end in `.fcl`. As you progress through the Workbook, this language and the conventions used in the run-time configuration file will be explained; the full details are available in Chapter 24 of the Users Guide. (The run-time configuration file is often referred to as simply the *configuration file* or even more simply as just the *configuration*(γ).)

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



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

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

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

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

- at the start of the *art* job
- at the end of the *art* job
- 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.

- at the end of each run
- at the start of each SubRun
- at the end of each SubRun



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

After *art* completes its initialization phase (intentionally not detailed here), it executes the event loop. This is illustrated in Figure 3.2, which is described in the text below:

1. calls the *constructor*(γ) of every module in the configuration
2. calls the *beginJob member function*(γ) of every module that provides one
3. reads one event from the input source, and for that event
 - (a) determines if it is from a run different from that of the previous event (true for first event in loop)
 - (b) if so, calls the *beginRun* member function of each module that provides one
 - (c) determines if the event is from a subRun different from that of the previous event (true for first event in loop)
 - (d) if so, calls the *beginSubRun* member function of each module that provides one
 - (e) calls each module’s (required) per-event member function
4. reads the next event and repeats the above per-event steps until it encounters a new subRun
5. closes out the current subRun by calling the *endSubRun* member function of each module that provides one
6. repeats steps 4 and 5 until it encounters a new run
7. closes out the current run by calling the *endRun* member function of each module that provides one




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.

8. repeats steps 3 through 7 until it reaches the end of the input source
9. calls the `endJob` member function of each module that provides one
10. calls the *destructor*(γ) of each module

This entire set of steps comprises the event loop. One of *art*'s most visible jobs is controlling the event loop.

3.6.3 Module Types

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

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

***producer module*(γ)** May inspect information found in the event and may add new information to the event.

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

***source module*(γ)** Reads events, one at a time, from some source; *art* requires that every *art* job contain exactly one source module. A source is often a disk file but other options exist and will be described in the Workbook and Users Guide.

***output module*(γ)** Reads selected data products from memory and writes them to an output destination; an *art* job may contain zero or more output modules. An output destination is often a disk file but other options exist and will be described in the Users' Guide. .



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

What does an analyzer do if it may neither alter information in an event nor add to it? Typically it creates printout and it creates ROOT files containing histograms, *trees*(γ) and *nuples*(γ) that can be used for downstream analysis. (If you have not yet encountered these terms, the Workbook will provide explanations as they are introduced.)

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

3.6.4 *art* Data Products

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

The term *data product*(γ) is used in *art* to mean the unit of information that user code may add to an event or retrieve from an event. A typical experiment will have the following sorts of data products:

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

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

A data product is just a C++ *type*(γ) (a class, *struct*(γ) or typedef) that obeys a set of rules defined by *art*; these rules are very different than the rules that must be followed for a class



to be a module; when the sections that describe these rules in detail have been prepared, we will add references here. A data product can be a single integer, a large complex class hierarchy, or anything in between.

Add the missing references alluded to in the previous para.

Very often, a data product is a *collection*(γ) of some experiment-defined type. The C++ standard libraries define many sorts of collection types; *art* supports many of these and also provides a custom collection type named `cet : :map_vector`. Workbook exercises will clarify the *data product* and *collection type* concepts.

3.6.5 *art* Services

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

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

To provide managed access to the second sort of information, *art* supports an idea named *art services* (again, shortened to *services*). Services may also be used to provide certain types of utility functions. Again, a service in *art* is just a C++ class that obeys a set of rules defined by *art*. The rules for services are different than those for modules or data products.

art implements a number of services that it uses for internal functions, a few of which you will encounter in the first couple of Workbook exercises. The *message service*(γ) is used by both *art* and experiment-specific code to limit printout of messages with a low severity level and to route messages to appropriate destinations. It can be configured

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.

to provide summary information at the end of the *art* job. The *TFileService*(γ) and the *RandomNumberGenerator* service are not used internally by *art*, but are used by most experiments. Experiments may also create and implement their own services.

After *art* completes its initialization phase and before it constructs any modules (see Section 3.6.2), it

1. reads the configuration to learn what services are requested
2. calls the constructor of each requested service

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

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

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.

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
debug	-O0 -g
prof	-O3 -g -fno-omit-frame-pointer -DNDEBUG
opt	-O3 -DNDEBUG

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 9.4. 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.

3.6.8 External Products

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

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

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

art the event processing framework

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

CETLIB a utility library used by *art*

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

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

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

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

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

git(γ) a source code management system that is used for the Workbook and by some experiments; similar in concept to the older CVS and SVN, but with enhanced functionality

cetbuildtools(γ) the build system that is used by the *art* Workbook (and by *art* itself).

UPS(γ) a Fermilab-developed system for accessing software products; it is an acronym for *Unix Product Support*.

UPD(γ) a Fermilab-developed system for distributing software products; it is an acronym for *Unix Product Distribution*.

jobsub_tools(γ) tools for submitting jobs to the Fermigrid batch system and monitoring them.

ifdh_sam(γ) allows *art* to use *SAM*(γ) as an external run-time agent that can deliver remote files to local disk space and can copy output files to tape. SAM is a Fermilab-supplied resource that provides the functions of a file catalog, a replica manager and some functions of a batch-oriented workflow manager

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

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



Finally, it is important to recognize an overloaded word, *products*. When a line of documentation simply says *products*, it may be referring either to data products or to external products. If it is not clear from the context which is meant, please let us know (see Section 3.4).

3.6.9 The Event-Data Model and Persistency

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 are many issues that are common to all data products in all experiments:

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: which module created it; what was the run-time configuration of that module; what data products were read by that module; what was 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

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

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

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

3.6.11 Files on Tape

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

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

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

<https://cdcvns.fnal.gov/redmine/projects/sam-main/wiki>.

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 sec-

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*.

tion 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 centered on the z axis; the left-hand part of Figure 3.3 shows an xy view of these shells and the right shows the radius vs z view. The inner five shells are closely spaced radially and are short in z ; the ten outer shells are more widely spaced radially and are longer in z . The detector sits in a uniform magnetic field of 1.5 T oriented in the $+z$ direction. The

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

origin of the coordinate system is at the center of the detector. The detector is placed in a vacuum.

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

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



The first six units listed in Table 3.2 are the base units defined by the CLHEP `SystemOfUnits` 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 .

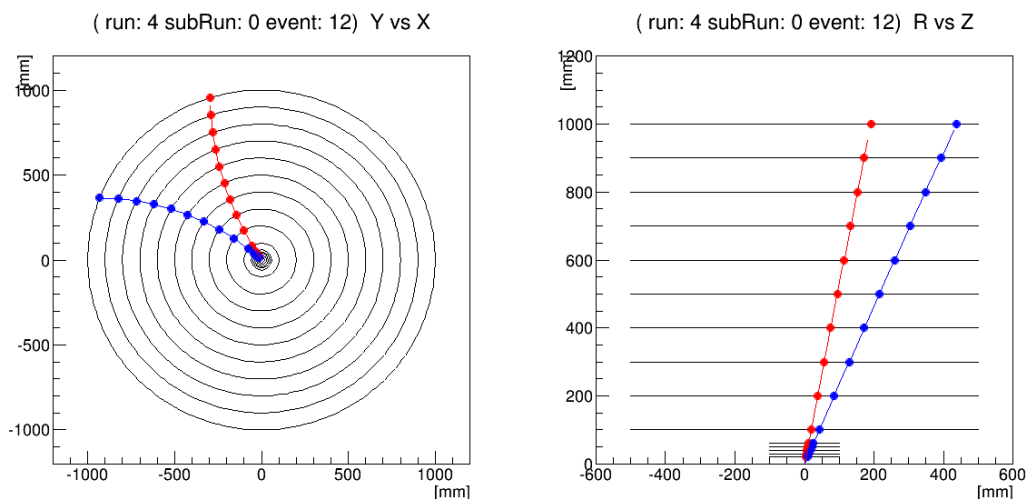


Figure 3.4: Event display of a simulated event in the toy detector.

- Its direction is chosen randomly on the unit sphere.

The above algorithm generates events with a total charge of zero but there is no concept of momentum or energy balance. About 47% of these events will not have any background tracks.

In the detector simulation step, particles neither scatter nor lose energy when they pass through the detector cylinders; nor do they decay. Therefore, the charged particles follow a perfectly helical trajectory. The simulation follows each charged particle until it either exits the detector or until it completes the outward-going arc of the helix. When the simulated trajectory crosses one of the detector shells, the simulation records the true point of intersection. All intersections are recorded; at this stage in the simulation, there is no notion of inefficiency or resolution. The simulation does not follow the trajectory of the ϕ meson because it was decayed in the generator.

Figure 3.4 shows an event display of a simulated event that has no background tracks. In this event the ϕ meson was travelling close to 90° to the z axis and it decayed nearly symmetrically; both tracks intersect all 15 detector cylinders. The left-hand figure shows an xy view of the event; the solid lines show the trajectory of the kaons, red for K^+ and blue for K^- ; the solid dots mark the intersections of the trajectories with the detector shells. The right-hand figure shows the same event but in an rz view.

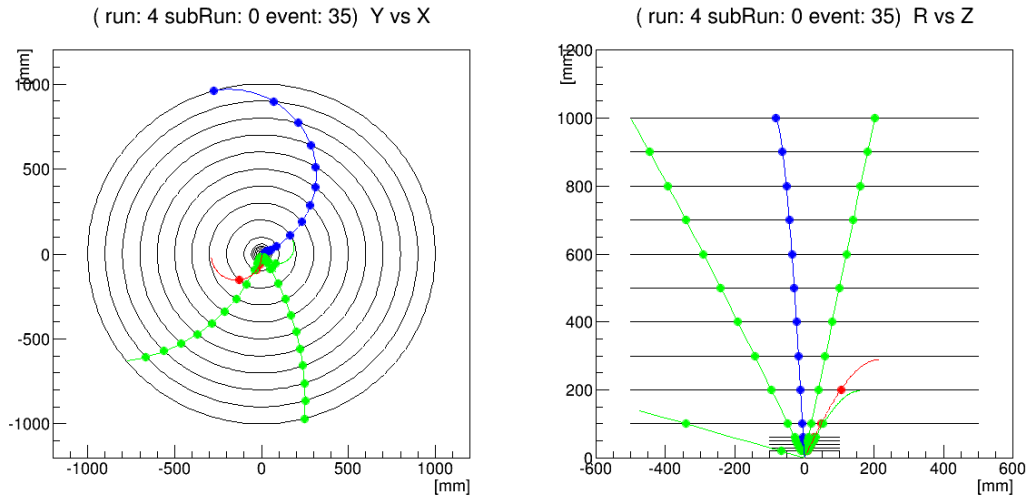


Figure 3.5: Event display of another simulated event in the toy detector; a K^- (blue) is produced with a very shallow trajectory and it does not intersect any detector shells while the K^+ (red) makes five hits in the inner detector and seven in the outer detector

Figure 3.5 shows an event display of another simulated event, one that has four background tracks, all drawn in green. In the xy view it is difficult to see the two π^- tracks, which have very low transverse momentum, but they are clear in the rz view. Look at the K^+ track, draw in red; its trajectory just stops in the middle of the detector. Why does this happen? In order to keep the exercises focused on *art* details, not geometric corner cases, the simulation stops a particle when it completes the outward-going arc of the helix and starts to curl back towards the z axis; it does this even if the the particle is still inside the detector.

The third step in the simulation chain (hit-making) is to inspect the intersections produced by the detector simulation and turn them into data-like hits. In this step, a simple model of inefficiency is applied and some intersections will not produce hits. Each hit represents a 2D measurement (φ, z) ; each component is smeared with a gaussian distribution.

The three simulation steps use tools provided by *art* to record the *truth information*(γ) about each hit. Therefore it is possible to navigate from any hit back to the intersection from which it is derived, and from there back to the particle that made the intersection.

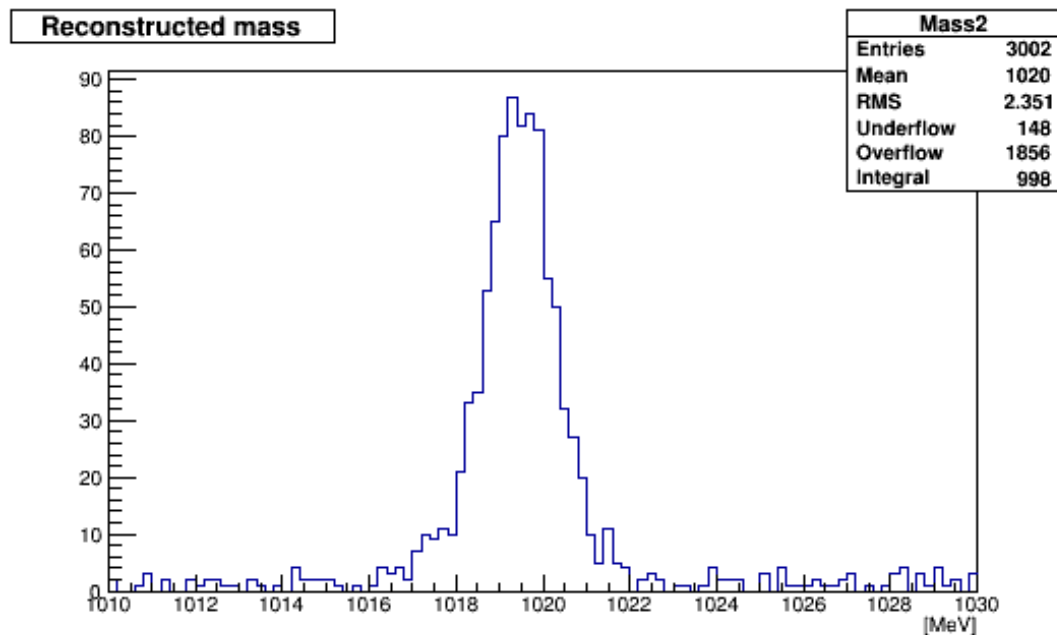


Figure 3.6: The invariant mass of all reconstructed pairs of oppositely charged tracks; for this all reconstructed tracks are assumed to be kaons.

The fourth step is the reconstruction step. The *toyExperiment* does not yet have properly working reconstruction code; instead it mocks up credible looking results. The output of this code is a data product that represents a fitted helix; it contains the fitted track parameters of the helix, their covariance matrix and collection of smart pointers that point to the hits that are on the reconstructed track. When we write proper tracking finding and track fitting code for the *toyExperiment*, the classes that describe the fitted helix will not change. Because the main point of the Workbook exercises is to illustrate the bookkeeping features in *art*, this is good enough for the task at hand. The mocked-up reconstruction code will only create a fitted helix object if the number of hits on a track is greater than some minimum value. Therefore there may be some events in which the output data product is empty.

The fifth step in the workflow does a simulated analysis using the fitted helices from the reconstruction step. It forms all distinct pairs of tracks and requires that they be oppositely charged. It then computes the invariant mass of the pair, under the assumption that both

fitted helices are kaons.⁷ This module is an analyzer module and does not make any output data product. But it does make some histograms, one of which is a histogram of the reconstructed invariant mass of all pairs of oppositely charged tracks; this histogram is shown in Figure 3.6. When you run the Workbook exercises, you will make this plot and can compare it to Figure 3.6. In the figure you can see a clear peak that is created when the two reconstructed tracks are the two true daughters of the generated φ meson. You can also see an almost flat contribution that occurs when at least one of the reconstructed tracks comes from one of the generated background particles.

3.8 Rules, Best Practices, Conventions and Style

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

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

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

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

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


```
command-name --help
```

or

```
command-name -?
```

Before starting the Workbook, make sure that you understand the basic usage of the following Unix commands:

cat, cd, cp, echo, export, gzip, head, less, ln -s, ls,
mkdir, more, mv, printenv, pwd, rm, rmdir, tail, tar

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

- filename vs pathname
- absolute path vs relative path
- directories and subdirectories (equivalent to folders in the Windows and Mac worlds)
- current working directory
- home directory (aka login directory)
- `.. /` notation for viewing the directory above your current working directory
- environment variables (discussed briefly in Section 4.5)
- *paths*(γ) (in multiple senses; see Section 4.6)
- file protections (read-write-execute, owner-group-other)
- symbolic links
- stdin, stdout and stderr
- redirecting stdin, stdout and stderr
- putting a command *in the background* via the `&` character
- pipes

4.3 Shells

When you type a command at the prompt, a Unix agent called a *Unix shell*, or simply a *shell*, reads your command and figures out what to do. Some commands are executed internally by the shell but other commands are dispatched to an appropriate program or script. A shell lives between you and the underlying operating system; most versions of Unix support several shells. The *art* Workbook code expects to be run in the *bash shell*. You can see which shell you're running by entering:

```
echo $SHELL
```



For those of you with accounts on a Fermilab machine, your login shell was initially set to the *bash* shell¹.



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

4.4 Scripts: Part 1

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



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

If you would like to learn more about *bash*, some references are listed in Section 4.10.

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

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.

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, the *art*-enabled computing environment includes some aliases and *bash* functions that have been defined; these are discussed in Section 4.8.

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



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

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

This concept of the environment hierarchy is illustrated in Figure 4.1.

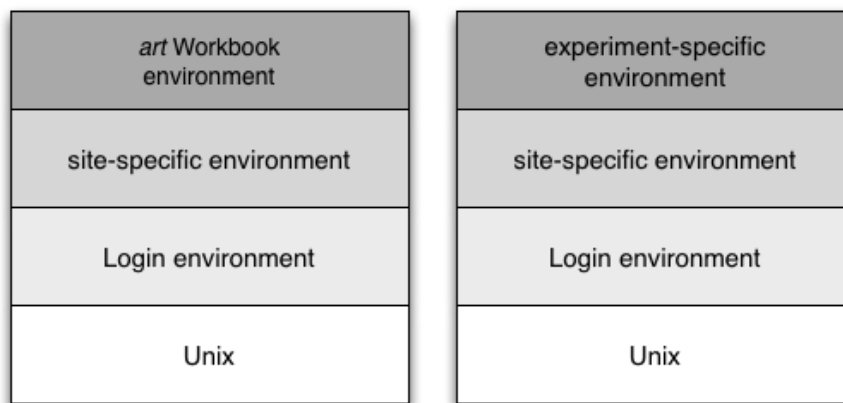


Figure 4.1: Components of the *art* Workbook (left) and experiment-specific (right) computing environments, shown in the order in which they are constructed, starting with the Unix environment

4.5.2 Examining and Using Environment Variables

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

```
printenv HOME
```

At any point in an interactive command or in a shell script, you can tell the shell that

you want the value of the environment variable by prefixing its name with the `$` character:

```
echo $HOME
```

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

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

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

```
export ENVNAME=value
```

If you read *bash* scripts written by others, you may see the following variant, which accomplishes the same thing:

```
ENVNAME=value
```

```
export ENVNAME
```

4.6 Paths and \$PATH

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

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

```
/absolute/path/to/mydir/myfile or  
relative/path/to/mydir/myfile or  
../another/relative/path/to/mydir/myfile
```

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

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

²Another type of variable, *shell variables*, are local to the current shell and are not copied to sub shells. By convention, these 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.

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

“path” generically, refers to any environment variable whose value is a colon-separated list of directory names e.g.,
`/abs/path/a:/abs/path/b:rel/path/c`

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

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

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

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

`PRODUCTS` used by UPS to resolve external products

`FHICL_FILE_PATH` use by FHiCL to resolve `#include` directives.

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

```
printenv PATH
```

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

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

In the above line, the vertical bar is referred to as a *pipe* and `tr` is a standard Unix command. A pipe takes the output of the command to its left and makes that the input of the command to its right. The `tr` command replaces patterns of characters with other patterns of characters; in this case it replaces every occurrence of the colon character with the new-line character. To learn why a double back slash is needed, read bash documentation to learn about escaping special characters.

4.7 Scripts: Part 2

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

```
file.sh
```

```
source file.sh
```

```
. file.sh
```

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

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

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

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



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

If you would like to learn more about bash, some references are listed in Section 4.10.

4.8 bash Functions and Aliases

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

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

```
declare -F
```

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

```
alias
```

You can read more about bash shell functions and aliases in any standard bash reference.

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

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

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

```
type command-name
```


4.9 Login Scripts

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

When your account on a Fermilab computer was first created, you were given standard versions of the files `.profile` and `.bashrc`; these files are used by `bash`³. You can read about login scripts in any standard bash reference. You may add to these files but you should not remove anything that is present.



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



It can be useful to inspect the login scripts of your colleagues to find useful customizations.

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

4.10 Suggested Unix and bash References

The following cheat sheet provides some of the basics:

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

A more comprehensive summary is available from:

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

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

- BASH Programming - Introduction HOW-TO
<http://tldp.org/HOWTO/Bash-Prog-Intro-HOWTO.html>

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

- Bash Guide for Beginners
<http://www.tldp.org/LDP/Bash-Beginners-Guide/html/Bash-Beginners-Guide.html>
- Advanced Bash Scripting Guide
<http://www.tldp.org/LDP/abs/html/abs-guide.html>

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

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

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

5 Site-Specific Setup Procedure

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 variables are set correctly to run the Workbook exercises or my experiment's code using *art*?

Very simply, 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 (NOvA 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 procedure when you get to Section 9.6.

For NOvA, until further notice, you need to add the bash function definition from Listing 5.1 to your login scripts. To run the site-specific setup procedure, type the following at the command line:

```
setup_art_workbook
```

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.

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 the next release of the workbook
LBNE	<code>source /grid/fermiapp/lbne/software/setup_lbne.sh</code>
MicroBoone	<code>source /grid/fermiapp/products/uboone/setup_uboone.sh</code> or <code>source /cvmfs/oasis.opensciencegrid.org/microboone/products/setup_uboone.sh</code>
Muon g-2	<code>source source /grid/fermiapp/gm2/setup</code>
Mu2e	<code>setup mu2e</code>
NO ν A	See the text

Listing 5.1: NO ν A `setup_art_workbook` script

```

1 function setup_art_workbook
2 {
3     echo "Setting_up_art_workbook"
4     source /grid/fermiapp/nova/novaart/novasvn/srt/srt.sh
5     export EXTERNALS=/nusoft/app/externals
6     export ART_WORKBOOK_OUTPUT_BASE=/nova/data/users
7     export ART_WORKBOOK_WORKING_BASE=/nova/app/users
8     export ART_WORKBOOK_QUAL="nu:e5"
9 }

```

6 Get your C++ up to Speed

6.1 Introduction

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`.

After the linker has finished, you can run your executable program typing the filename of the program at the bash command prompt. At this time, the loader does the final steps of linking the program, to allow the use of instructions and data in the dynamic libraries to which your program is linked.

A typical program links both to libraries that were built from the program’s source code

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

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

³Other conventions include `.os`.

and to libraries from other sources. Some of these other libraries might have been developed by the same programmer as general purpose tools to be used by his or her future programs; other libraries are provided by third parties, such as *art* or your experiment. Many C++ language features are made available to your program by telling the linker to use libraries provided by the C++ compiler vendor. Other libraries are provided by the operating system.

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

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

6.3 Establishing the Environment

6.3.1 Initial Setup

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.
3. Create an empty working directory and `cd` to it.
4. Run these commands (the first one is shown on two lines) to copy a gzipped tar file from the web, unpack it, and get a directory listing:

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

tar xzf C++UpToSpeed.tar.gz

rm C++UpToSpeed.tar.gz
```

```
ls
```

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

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

```
source setup.sh
```

After these steps, you are ready to begin the exercise in Section 6.4.

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

The above list explicitly does not include classes, objects and inheritance, which will be discussed in Sections 6.7 and a future section on inheritance.

6.4.2 How to Compile, Link and Run

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

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.

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

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.

```
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
```

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

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

This turned the source file `t1.cc` into an executable program, named `t1` (the argument to the `-o` (for “output”) option). We will discuss compiling and linking in Section 6.5.

6.4.3 Discussion

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

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

```
int main() {  
    ...executable code...  
}
```

Compare your output with the standard example:

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

There will almost certainly be a handful of differences, which we will discuss in Section 6.4.3.1.

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

6.4.3.1 Primitive types

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

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



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

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

6.4.3.2 Arrays

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



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

6.4.3.3 Equality testing

An important concept in C++ is the difference between two variables being *identical* and two variables being *equivalent*. Two variables are identical if they refer to the same object, and thus have the same memory address. Two variables which refer to different objects that contain the same value are equivalent. Equivalence is tested using the equality testing operator, `==`.

It is important to distinguish between the assignment operator (`=`) and the equality testing operator. Mistakenly using `=` where `==` is intended is a common mistake.

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



Another distinction to be made is that of two variables being *identical* versus *equivalent*. In contrast to equivalence, two variables are identical if they refer to the same object, and thus have the same memory address.

6.4.3.4 Conditionals

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

```
type variable-to-initialize (expression-to-evaluate) ?
                             value-if-true :
                             value-if-false;
```

An example is shown in the code.

6.4.3.5 Some C++ Standard Library Types

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

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

The `std::array` is new in C++, and should be used in preference to the older C-style array discussed in Section 6.4.3.2 due to its greater capabilities.

6.4.3.6 Pointers

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

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

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

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

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

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

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

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

The next line,

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

shows how to access the value to which a pointer points. This is called *dereferencing* the pointer, or *following* the pointer, and is done with the dereferencing operator, `*`. The expression `*pa` dereferences the pointer `pa`, yielding the value of the object to which `pa` points, in this case, the value of `a`.

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

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

In the example code, the lines

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

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

Note that a pointer to an array points to the first element of the array; again the pointer type must correspond to the type of the array elements. And for an array `l[]`, the name by itself, i.e., `l` refers to the location of its first element.

6.4.3.7 References

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

```
float a;
```

```
float& ra = a;  
float* p = &a;  
float* q = &ra;
```

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

6.4.3.8 Loops

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

6.5 C++ Exercise 2: About Compiling and Linking

6.5.1 What You Will Learn

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

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

6.5.2 The Source Code for this Exercise

The source code for this exercise is found in `Build/v1`, relative to your working directory. The relevant files are `t1.cc`, `times2.cc` and `times2.h`. Open these files and

read along with the discussion below.

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



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



To be recognized as a main program, there is one more requirement: `main` must be declared in the global namespace.

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

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

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

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

The other three lines in `times2.h` make up a *code guard*, described in Section ???. In

brief, they deal with the following scenario: suppose that we have two header files, `A.h` and `B.h`, and that `A.h` includes `B.h`; there are many scenarios in which it makes good sense for a third file, either `.h` or `.cc`, to include both `A.h` and `B.h`. The code guards ensure that, when all of the includes have been expanded, the compiler sees exactly one copy of `B.h`.

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

```
double times2(double i) {  
    return 2 * i;  
}
```

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

We now have a rich enough example to discuss another case in which the same word is frequently used to mean two different things. Sometimes people use the phrase “the source code of the function named `times2`” to refer collectively to both `times2.h` and `times2.cc`; sometimes they use it to refer exclusively to `times2.cc`. Unfortunately the only way to distinguish the two uses is from context.

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

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

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

6.5.3 Compile, Link and Run the Exercise

To perform this exercise, first log in and `cd` to your working directory if you haven’t already, 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
```

This matches the expected printout.

Look at the file `build` that you just ran. It has three steps:

1. It compiles the main program, `t1.cc`, into the object file (with the default name) `t1.o` (which will now be the thing that the term *main program* refers to):
`c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c t1.cc`
2. It (separately) compiles `times2.cc` into the object file `times2.o`:

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

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

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

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

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

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

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

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

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

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

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

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

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

information to resolve these references at the time when the program is run.

6. Once all of the machine code has been copied into the executable, the compiler knows the memory address of every function, or where to find them at run-time. The compiler can then go into the machine code, find all of the placeholders and update them with the correct memory addresses.

Sometimes resolving one unresolved reference will generate new ones. The linker iterates until (a) all references are resolved and no new unresolved references appear (success) or (b) the same unresolved references continue to appear (error). In the former case, the linker writes the output to the file specified by the `-o` option; if no `-o` option is specified the linker will write its output to a file named `a.out`. In the latter case, the linker issues an error message and does not write the output file.

After the link completes, the files `t1.o` and `times2.o` are no longer needed because everything that was useful from them was copied into the executable `t1`. You may delete the object files, and the executable will still run.

6.5.4 Alternate Script `build2`

The script `build2` shows an equivalent way of building `t1` that is commonly used for 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
```

```
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
```

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

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

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`.

6.5.5 Suggested Homework

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

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

1. In `Build/v1/t1.cc`, comment out the include directive for `times2.h`; rebuild and observe the error message.
2. In `Build/v1/times2.cc`, change the return type to `float`; rebuild and observe

the error message.

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

The first homework exercise will issue the diagnostic:

```
1 t1.cc: In function 'int main()':
2 t1.cc:9:40: error: 'times2' was not declared in this scope
```

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

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

```
1 times2.cc: In function 'float times2(double)':
2 times2.cc:3:22: error: new declaration 'float times2(double)'
3   float times2(double i) {
4       ^
5 In file included from times2.cc:1:0:
6 times2.h:4:8: error: ambiguates old declaration 'double times2(double)'
7   double times2(double);
8       ^
```

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

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

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

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

Consult the standard C++ documentation to understand when automatic type conversions may occur; see Section 6.9.

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
```

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

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

```
c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c times2.cc
```
2. Creates the library named `libpackage1.so` (on OS X, the standard suffix for a dynamic library is different, and the the library is called `libpackage1.dylib`) from `times2.o`:

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


or

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

Note that the name of the library must come before the name of the object file. The flag `-shared` directs the linker to create a dynamic library rather than an executable image; without this flag, this command would produce an error complaining about the lack of a `main` function.
3. Compiles `t1.cc`; the same as the previous exercise:

```
c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c t1.cc
```
4. Links the main program against the dynamic library (either `libpackage1.so` or `libpackage1.dylib`) and, implicitly, the necessary system libraries:


```
c++ -o t1 t1.o libpackage1.so  
or  
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:

```
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
```

3. Run `./t1` as before.

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

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

while that from `build2` is:

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

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

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

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

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

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

makes a much more readable link list.

To perform the exercise using `build3`, stay in the same directory and clean up and rebuild 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

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

6.7 Classes

6.7.1 Introduction

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

The language features that allow users to define types are the *class*(γ) and the *struct*(γ). As you work through the Workbook exercises, you will see classes that are defined by the Workbook itself; you will also see classes defined by the toyExperiment UPS product; you will see classes defined by *art* and you will see classes defined by the many UPS products

```
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.

that support *art*. You will also write some classes of your own. When you work with the software for your experiment you will work with classes defined within your experiment's software.

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

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

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

```
class MyClassName;
```

Class declarations are rarely used because a class definition also acts as a class declaration.

The remainder of Section 6.7 will give many examples of *members* of a class.



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

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

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

This documentation will use technically correct language so that you will find it easier to read the standard reference materials. We will point out colloquial usage as necessary.

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

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

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

To build and run this example:

1. Log in and follow the steps in Section 6.3.
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 the listing on page ??.

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

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
```

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

Look at the header file `Point.h`, reproduced in Listing 6.2, which shows the basic version of the class `Point`.

The three lines starting with `#` make up a code guard, described in Section ??.

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

The body of the class definition begins on line 4, with the opening brace; the body of the class definition ends on line 8, with the closing brace. The definition of the class is followed by a semicolon. Line 5 states that the following members of the class are `public`, which


```
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`.

means they are accessible from code outside the class (that is, they are accessible not only by member functions of this class, but by member functions of other classes and by free functions). Members can also be `private` or `protected`. Section 6.7.7 addresses the meaning of `private`. The use of `protected` is beyond the scope of this introduction. Lines 6 and 7 declare the data members `x` and `y`, both of which are of type `double`.

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

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

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

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

```
Point p0;
```

is executed, the program will ensure that memory has been allocated⁵ to hold the data members of `p0`. If the class `Point` contained code to initialize data members then the

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.

```
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/ptest.cc`.

program would also run that, but `Point` does not have any such code. Therefore the data members take on whatever values happened to preexist in the memory that was allocated for them.

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

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

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

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

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

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

The last section of the main program (and of `pctest.cc` itself) prints the address of each of the three objects, `p0`, `p1` and `p2`. The addresses are represented in hexadecimal (base 16) format. On almost all computers, the size of a `double` is eight bytes. Therefore

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




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

an object of type `Point` will have a size of 16 bytes. If you look at the printout made by `ptest` you will see that the addresses of `p0`, `p01` and `p2` are separated by 16 bytes; therefore the three objects are contiguous in memory.

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

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

1. a *state*, given by the values of its data members;
2. the ability to have operations performed on it: e.g., setting/reading in value of a data member, assigning value of object of a given type to another of the same type;
3. an *identity*: a unique address in memory.

6.7.3 C++ Exercise 4 v2: The Default Constructor

This exercise expands the class `Point` by adding a user-written default *constructor*(γ).

To build and run this example:

1. Log in and follow the steps in Section 6.3.

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)
```

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

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

```
Point();
```

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



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

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

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

```
Point::Point() {  
    x = 0.;  
    y = 0.;  
}
```

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

```
Point p0;
```

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

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

In the previous example, `Classes/v1/pctest.cc`, a few things happened behind the 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. Since the source code for class `Point` did not contain any user-defined constructor, the compiler (as is required by the C++ Standard) generated a default constructor for you. This will be done for any class that has no user-written constructor.
2. The compiler puts the generated constructor code directly into the object file; it does not affect the source file.
3. The generated default constructor will default construct each data member of the class.
4. Default construction of an object of a primitive type leaves that object uninitialized; this is why the data members `x` and `y` of version 1 of `Point` were uninitialized.

6.7.4 C++ Exercise 4 v3: Constructors with Arguments

This exercise introduces four new ideas:

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

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

```
pctest
```

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

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

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

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

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

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

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

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

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

The next line of code

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

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

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

If your class has data members that are pointers, or data members that manage some external resource, such as a file that you are writing to, these pointers should be *smart*

pointers, such as `std::shared_ptr<T>` or `std::weak_ptr<T>`. This will allow the compiler-generated copy constructor to give the correct behavior. For a description of smart pointers, consult the standard C++ references; see Section 6.9. There are rare cases in which you will need to write your own copy constructor, but discussing them here is beyond the scope of this document. When you need to write your own copy constructor, you can learn how to do so from any standard C++ reference.

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

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

```
Point p0;  
p0.x = 3.1;  
p0.y = 2.7;
```

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

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

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



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

4. Single-phase construction is the only way to initialize variables that are declared `const`. It is good practice to declare `const` any variable that is not intended to be changed.

6.7.5 C++ Exercise 4 v4: Colon Initializer Syntax

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

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

```
pctest
```

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

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

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

1. A colon must immediately follow the closing parenthesis of the argument list.
2. There must be a comma-separated list of data members, each one initialized by calling one of its constructors.
3. Data members will be initialized in the order in which they appear in the class declaration. Therefore it is good practice to use the same order for the initialization list.
4. The body of the constructor, enclosed in braces, must follow the initializer list. The body of the constructor will most often be empty.
5. If a data member is missing from the initializer list, that member will be default-constructed. Thus data members that are of primitive types will not be initialized.
6. If no initializer list is present, the compiler will call the default constructor of every data member, and it will do so in the order in which data members were specified in

the class declaration.

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

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

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

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

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

Now look at the second constructor in `Point.cc`; it also uses colon-initializer syntax but it is laid out differently. The difference in layout has no meaning to the compiler — whitespace is whitespace. Choose which ever seems natural to you.

Look at `ptest.cc`. It is the same as the version `v3` and it makes the same printout.

6.7.6 C++ Exercise 4 v5: Member functions

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

To build and run this exercise, `cd` to the directory `Classes/v5` and follow the same

instructions as in Section 6.7.3. When you run the `pctest` program you should see the following output:

`pctest`

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

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

```
double mag() const;  
double phi() const;  
void scale(double factor);
```

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

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

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

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



If a member function does not modify any data members, you should always declare it `const` simply as a matter of course. Any negative consequences of not doing so might

only become apparent later, at which point a lot of tedious editing will be required to make everything right.

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

Later on in `Point.cc` you will find the definition of `Point::mag`, which computes the magnitude of the 2-vector from the origin $(0, 0)$ to (x, y) . To do so, it uses `std::sqrt`, a function declared in the `<cmath>` header. This function takes the square root of its argument. The qualifier `const` that was present in the declaration of `Point::mag` must also be present in its definition.

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

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

The file `pctest.cc` contains a `main` function that illustrates these new features. The first line of this function declares and initializes an object, `p0`, of type `Point`. It then prints out the value of its data members, the value returned from calling the function `Point::mag` and the value returned from calling `Point::phi`. This shows how to invoke a member function: you write the name of the variable, followed by a dot (the *member selection operator*), followed by the *unqualified name* of the member function, followed by the argument list in the function call parentheses. The unqualified name of a member function is the part of the name that follows the double-colon scope resolution operator `::`. Thus the unqualified name of `Point::phi` is just `phi`.

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

One final comment is in order. Many other modern computer languages have ideas very similar to C++ classes and C++ member functions; in some of those languages, the name *method* is the technical term corresponding to *member function* in C++. The name *method* is not part of the formal definition of C++, but is commonly used nonetheless. In this

documentation, the two terms can be considered synonymous.

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

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

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

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

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

2. In `pctest.cc`, change the first line to

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

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

```
pctest.cc: In function 'int main()':
pctest.cc:13:14: error: no matching function for call to
'Point::scale(double) const'
pctest.cc:13:14: note: candidate is:
In file included from pctest.cc:1:0:
Point.h:13:8: note: void Point::scale(double) <near match>
Point.h:13:8: note:   no known conversion for implicit
             'this' parameter from 'const Point*' to 'Point*'
```

These first two homework exercises illustrate how the compiler enforces the contract defined by the qualifier `const` that is present at the end of the declaration of `Point::mag` and that is absent in the definition of the member function `Point::scale`. The contract says that the definition of `Point::mag` may not modify the values of any data members of the `Point` object on which it is called; users of the class `Point` may count on this behaviour. The contract also says that the definition of the member function `Point::scale` may modify the values of data members of the class `Point`; users

of the class `Point` must assume that `Point::scale` will indeed modify member data and act accordingly.⁸

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

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

3. In `Point.h`, remove the `const` qualifier from the declaration of the member function `Point::mag`:

```
double mag();
```

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

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

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

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

The third and fourth homework exercises illustrate that the compiler considers two member functions that are identical except for the presence of the `const` identifier to be different functions⁹. In homework exercise 3, when the compiler tried to compile the

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.

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

`const`-qualified version of `Point::mag` in `Point.cc`, it looked at the class definition in `Point.h` and could not find a matching member function declaration; this was a close, but not exact match. Therefore it issued a diagnostic message, telling us about the close match, and then stopped. Similarly, in homework exercise 4, it also could not find a match.

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

6.7.7.1 Setters and Getters

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

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

A 2D point class, with member data in Cartesian coordinates, is not a good example of *why* it is often a good idea to have `private` data. But it does have enough richness to illustrate the mechanics, which is the purpose of this section. Section 6.7.7.3 discusses an example in which having `private` data makes obvious sense.

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

`pctest`

```
Before p0: (1, 2)  Magnitude: 2.23607  Phi: 1.10715
After p0: (3, 6)  Magnitude: 6.7082  Phi: 1.10715
p1: (0, 1)  Magnitude: 1  Phi: 1.5708
p1: (1, 0)  Magnitude: 1  Phi: 0
p1: (3, 6)  Magnitude: 6.7082  Phi: 1.10715
```

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


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

Relative to version v5 the following changes were made:

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

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

1. the name of the class it is in;
2. the unqualified name of the member function;
3. the number, type and order of arguments in the argument list;
4. whether or not the function is qualified as `const`;
5. other qualifications (reference-qualification and `volatile`-qualification, both of which are beyond the scope of this introduction).

The two member functions named `Point::set` are completely different member functions with different signatures. A set of different functions with the same name but with different signatures is called an *overload set*. As you work through the Workbook you will encounter a many of these, and you should develop the habit of looking at the full function signature (i.e., all the parts), not just the function name. In order to distinguish between members of an overload set, C++ compilers typically rely on *name mangling*. Name mangling is the process of decorating a function name with information that encodes the signature of the function. The mangled name associated with each function is the symbol emitted by the compiler, and used by the linker to identify which member of

an overload set is associated with a specific function call. Each C++ compiler does this a little differently.



If you want to see what mangled names are created for the class `Point`, you can do the following `c++ -Wall -Wextra -pedantic -Werror -std=c++11 -c Point.cc nm Point.o`

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

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

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

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

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

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

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

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

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.

More generally, any member function that modifies the value of any member data is called a *modifier*.

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

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

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

```
diff -wb ptest.cc ../v5/
```

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

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

Presumably these are clear.

Figure 6.2 shows a diagram of the computer memory at the end of running this version of `pctest`. The two boxes with the blue outlines represent sections of the computer memory; the part on the left represents that part that is reserved for storing data (such as objects) and the part on the right represents the part of the computer memory that holds the executable code. This is a big oversimplification because, in a real running program, there are many parts of the memory reserved for different sorts of data and many parts reserved for executable code.

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

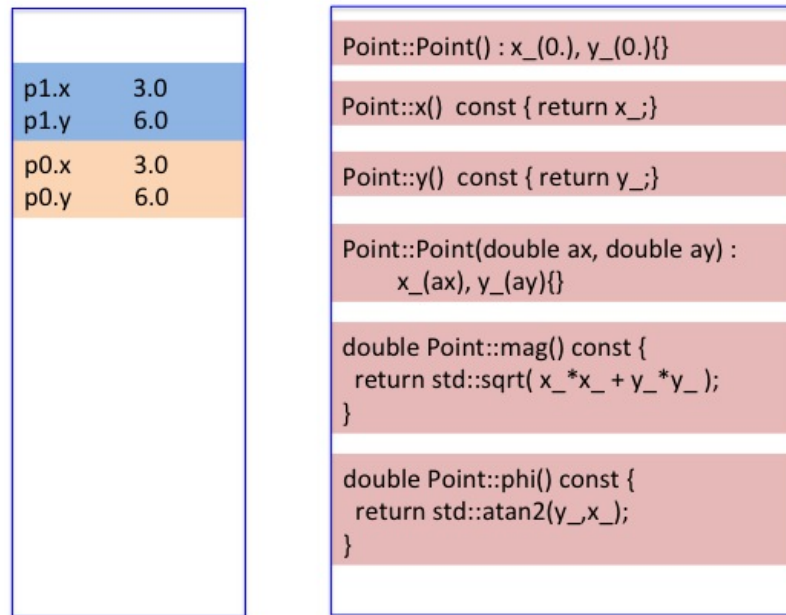


Figure 6.2: Memory diagram at the end of a run of Classes/v6/ptest.cc

declared as `const`, the expression acts as a pointer that is `const`-qualified.

Initially this sounds a little weird: the previous paragraph talks about passing an argument to the function `Point::mag` but, according to the source code, `Point::mag` does not take any arguments! The answer is that all member functions have an implied argument that always must be present — the address of the object that the member function will do work on. Because it must always be there, and because the compiler knows that it must always be there, there is no point in actually writing it in the source code! It is by using this so called *hidden argument* that the code for `Point::mag` knew that `x_` means one thing for `p0` but that it means something else for `p1`.

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

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

This version of the syntax makes it much clearer how there can be one copy of the code even though there are many objects in memory; but it also makes the code harder to read once you have understood how the magic works. There are not many places in which you need to explicitly use the keyword `this`, but there will be some. For further information,

consult standard C++ documentation (listed in Section 6.9).

6.7.7.2 What's the deal with the underscore?

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

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

```
_name;  
m_name;  
mName;  
theName;
```

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



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

6.7.7.3 An example to motivate private data

This section describes a class for which it makes sense to have private data: a 2D point class that has data members `r` and `phi` instead of `x` and `y`. The author of such a class might wish to define a standard representation in which it is guaranteed that `r` be non-negative

and that `phi` be on the domain $0 \leq \phi < 2\pi$. If the data are public, the class cannot make these guarantees; any code can modify the data members and break the guarantee.

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

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

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

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

This section introduces the `inline` specifier.

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

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

Look at `Point.h` and compare it to the version in `v6`. The new material added to this version is the implementation for the two accessors `Point::x` and `Point::y`. These accessors are defined outside of the class definition.

Look at `Point.cc` and compare it to the version in `v6`. You will see that the implementations of the accessors `Point::x` and `Point::y` have been removed.

`Point.h` now contains an almost exact copy of the the implementation of the accessor `Point::x` that was previously found in the file `Point.cc`; the difference is that it is now preceded by the specifier `inline`. This specifier tells the compiler that, at each site

where the function `Point::x` is called, it has two options that it may choose from at its discretion.

The first option is that the compiler may decline to write a call to the member function `Point::x`; instead, the compiler can insert the body of `Point::x` right into the machine code at that spot. This is called *inlining* the function. For something simple like an accessor, relative to explicitly calling a function, the inlined code is very likely to

1. have a smaller memory footprint
2. execute more quickly
3. allow additional compiler optimizations to be performed

All of these are good things.

On the other hand, if you declare a bigger or more complex function `inline`, some negative effects of inlining may appear. If the inlined function is used in many places and if the memory footprint of the inlined code is large compared to the memory footprint of a function call, then the total size of the program can increase. There are various ways in which a large program might run more slowly than a logically equivalent but smaller program. So, if you inline large functions, your program may actually run more slowly!

At a call site for a function declared `inline`, the compiler also has a second option: it can choose to ignore the `inline` qualification. When the compiler chooses this option it will write many copies of the code for the member function — one copy for each *compilation unit*¹¹ in which the function is called. Each compilation unit only knows about its own copy of the function and the compiler calls that copy as needed. The net result is completely negative: the function call is not actually elided so there is no time savings from that; moreover the code has become bigger because there are multiple copies of the function in memory; the larger memory footprint can further slow down execution; and compilation takes longer because multiple copies of the function must be compiled.

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

The bottom line is that you should always declare simple accessors and simple setters

¹¹A compilation unit is the unit of code that the compiler considers at one time. For most purposes, each source file is its own compilation unit.

`inline`. Here the adjective *simple* means that they do not do any significant computation and that they do not contain any `if` statements or loops. The decision to inline anything else should only follow careful analysis of information produced by a profiling tool.

Look at the definition of the member function `Point::y` in `Point.h`. Compared to the definition of the member function `Point::x` there is a small change in whitespace. This difference is not meaningful to the compiler. You will see several other variations on whitespace when you look at code in the Workbook and its underlying packages.

6.7.9 C++ Exercise 4 v8: Defining Member Functions within the Class Declaration

The version of `Point` in this section introduces the idea that you may provide the definition (implementation) of a member function at the point that it is declared inside the class declaration. This topic is introduced now because you will see this syntax as you work through the Workbook.

To build and run this exercise, `cd` to the directory `Classes/v8` and follow the same instructions as in Section 6.7.3. When you run the `pctest` program you should see the following output:

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

This is the same output made by `v7`.

Look at `Point.h`. The only change relative to `v7` is that the definition of the accessor methods `Point::x` and `Point::y` has been moved into the class definition.

The files `Point.cc` and `pctest.cc` are unchanged with respect to `v7`.

This version of `Point.h` shows that you may define any member function inside the class definition. When you do this, the function is implicitly declared `inline`. Section 6.7.8 discussed some cautions about inappropriate use of inlining; those same cautions apply when a member function is defined inside the class definition.



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


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

would produce a compiler diagnostic.

In summary, there are two ways to write inlined definitions of member functions. In most cases, the two are entirely equivalent and the choice is simply a matter of style. The one exception occurs when you are writing a class that will become part of an *art* data product, in which case it is recommended that you write the definitions of member functions *outside* of the class declaration.



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



6.7.10 C++ Exercise 4 v9: The stream insertion operator

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

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

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

```
pctest
```

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

This is the same output made by v7 and v8.

Look at `Point.h`. The changes relative to v7 are the following two additions:

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

Look at `Point.cc`. The changes relative to v7 are the following two additions:

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

Look at `pctest.cc`. The only change relative to v7 is that the `printout` now uses the stream insertion operator for `p0` instead of inserting each data member of `p0` by hand.

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

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



If the class whose type is used as second argument is declared in a namespace, then the stream insertion operator must be declared in the same namespace.

When the compiler sees the use of a `<<` operator that has an object of type `std::ostream` on its left hand side and an object of type `Point` on its right hand side, then the compiler will look for a function named `operator<<` whose first argument is of type `std::ostream&` and whose second argument is of type `Point const&`. If it finds such a function it will call that function to do the work; if it cannot find such a function it will issue a compiler diagnostic.

The reason that the function returns a `std::ostream&` is that this is the C++ convention that permits us to chain together multiple instances of the `<<` operator:

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

The C++ compiler parses this left to right. First it recognizes the expression `std::cout << p0`. Because `std::cout` is of type `std::ostream`, and because `p0` is a `Point`, the compiler calls our stream insertion operator to do this work. The return type of the function call is `std::ostream&`, and so the next expression is also recognized as a call to our stream insertion operator for class `Point`, this time using the object `p1`. This also returns a `std::ostream&`, allowing the last part of the expression to be recognized as a call

to the stream insertion operator for `std::endl`, which writes a newline and flushes the output stream.

Look at the implementation of the stream insertion operator in `Point.cc`. The first argument, `ost`, is a reference to an object of type `std::ostream`; the name `ost` has no special meaning to C++. When writing this operator we don't know and don't care what the output stream is connected to; perhaps it is a file; perhaps it is standard output. In any case, you send output to `ost` just as you do to `std::cout`, which is just another variable of type `std::ostream`. In this example we chose to enclose the values of `x_` and `y_` in parentheses and to separate them with a comma; this is simply our choice, not something required by C++ or by *art*.

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

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

If you want to understand more about why the operator is written the way that it is, consult the standard C++ references; see Section 6.9.

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

The choice of whether or not to put the declaration of the stream insertion operator into its own header file is a tradeoff between the following two criteria:



1. it is convenient to have it there; otherwise you would have to remember to include an additional header file when you want to use this operator
2. one can imagine many simple free functions that take an object of type `Point` as an argument. If we put them all inside `Point.h`, and if they are only infrequently

used, then the compiler will waste time processing those declarations every time `Point.h` is included somewhere.

The definition of the operator is typically put into the implementation file, rather than being inlined. Such functions are generally poor candidates for inlining.

Ultimately this is a judgement call and the code in this example follows the recommendations made by the *art* development team. Their recommendation is that the following sorts of free functions, and only these sorts, should be included in header files containing a class declaration:

1. the stream insertion operator for that class
2. out of class arithmetic and comparison operators

With one exception, if including a function declaration in `Point.h` requires the inclusion of an additional header in `Point.h`, declare that function in a different header file. The exception is that it is okay to include `<iosfwd>`.

6.7.11 Review

The class `Point` is an example of a class that is primarily concerned with providing convenient access to the data it contains. Not all classes are like this; when you work through the Workbook, you will write some classes that are primarily concerned with packaging convenient access to a set of related functions.

1. class
2. object
3. identifier
4. free function
5. member function

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.

- Herb Sutter and Andrei Alexandrescu, “C++ Coding Standards: 101 Rules, Guidelines, and Best Practices.”, Addison-Wesley, 2005. ISBN 0-321-11358-6.

7 Using External Products in UPS

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

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*,

meaning simply that each directory functions as a repository of information about one or more products. When UPS looks for a product, it checks each directory in `PRODUCTS`, in the order listed, and takes the first match.



If you are on a Fermilab machine, you can look at the value of `PRODUCTS` just after logging in, before sourcing your site-specific setup script. Run `printenv`:

```
printenv PRODUCTS
```

It should have a value of

```
/grid/fermiapp/products/common/db
```

This generic Fermilab UPS database contains a handful of software products commonly used at Fermilab; most of these products are used by all of the Intensity Frontier Experiments. This database does not contain any of the experiment-specific software nor does it contain products such as *ROOT*(γ), *Geant4*(γ), CLHEP or *art*. While these last few products are indeed used by multiple experiments, they are often custom-built for each experiment and as such are distributed via the experiment-specific (i.e., separate) UPS databases.

After you source your site-specific setup script, look at `PRODUCTS` again. It will probably contain multiple directories, thus making many more products available in your “site” environment. For example, on the DS50+Fermilab site, after running the DS50 setup script, `PRODUCTS` contains:

```
/ds50/app/products/:grid/fermiapp/products/common/db
```

You can see which products `PRODUCTS` contains by running `ls` on its directories, one-by-one, e.g.,

```
ls /grid/fermiapp/products/common/db
```

```
afs      git      ifdhc      mu2e      python      ...
cpn      gitflow  jobsub_tools  oracle_tnsnames  ...
encp     gits     login      perl      setpath      ...
```

```
ls /ds50/app/products
```

```
art      cetpkgssupport  g4neutronxs  libxml2      ...
artdaq   clhep          g4nucleonxs  messagefacility  ...
```


art_suite	cmake	g4photon	mpich	...
art_workbook_base	cpp0x	g4pii	mvapich2	...
boost	cppunit	g4radiative	python	...
caencomm	ds50daq	g4surface	root	...
...				

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.

¹`ups` appears in both listings; as always, the first match wins!

In most cases the correct flavor can be automatically detected by setup and need not be specified. However, if needed, flavor, in addition to various qualifiers and options can be specified. These are listed in the UPS documentation referenced later in this section. The version, if specified, must directly follow the product name in the command line, e.g.,:

```
setup options product-name product-version -f flavor -q qualifiers
```

Putting in real-looking values, it would look something like:

```
setup -R myproduct v3_2 -f SLF5 -q BUILD_A
```

What does the setup command actually do? It may do any or all of the following:

- define some environment variables
- define some bash functions
- define some aliases
- add elements to your PATH
- setup additional products on which it depends

Setting up dependent products works recursively. In this way, a single setup command may trigger the setup of, say, 15 or 20 products.

When you follow a given site-specific setup procedure, the PRODUCTS environment variable will be extended to include your experiment-specific UPS repository.

setup is a bash function (defined by the UPS product when it was initialized) that shadows a Unix system-configuration command also named setup, usually found in /usr/bin/setup or /usr/sbin/setup. Running the right 'setup' should work automatically as long as UPS is properly initialized. If it's not, setup returns the error message:



```
You are attempting to run ``setup`` which requires administrative
privileges, but more information is needed in order to do so.
```

If this happens, the simplest solution is to log out and log in again. Make sure that you carefully follow the instructions for doing the site specific setup procedure.

Few people will need to know more than the above about the UPS system. Those who do can consult the full UPS documentation at:

<http://www.fnal.gov/docs/products/ups/ReferenceManual/index.html>

7.4 Current Versions of Products

For some UPS products, but not all, the site administrator may define a particular fully-qualified version of the product as the default version. In the language of UPS this notion of default is called the *current* version. If a current version has been defined for a product, you can set up that product with the command:

```
setup product-name
```

When you run this, the UPS system will automatically insert the version and qualifiers of the version that has been declared current.

Having a current version is a handy feature for products that add convenience features to your interactive environment; as improvements are added, you automatically get them.

However the notion of a current version is very dangerous if you want to ensure that software built at one site will build in exactly the same way on all other sites. For this reason, the Workbook fully specifies the version number and qualifiers of all products that it requires; and in turn, the products used by the Workbook make fully qualified requests for the products on which they depend.



7.5 Environment Variables Defined by UPS

When your login script or site-specific setup script initializes UPS, it defines many environment variables in addition to PRODUCTS (Section 7.1), one of which is UPS_DIR, the root directory of the currently selected version of UPS. The script also adds \$UPS_DIR/bin to your PATH, which makes some UPS-related commands visible to your shell. Finally, it defines the bash function setup (see Sections 4.8 and 7.3). When you use the setup command, as illustrated below, it is this bash function that does the work.

In discussing the other important variables, the toyExperiment product will be used as an example product. For a different product, you would replace “toyExperiment” or “TOY-EXPERIMENT” in the following text by the product’s name. Once you have followed your appropriate setup procedure (Table 5.1) you can issue the following command this

is informational for the purposes of this section; you don't need to do it until you start running the first Workbook exercise):

```
setup toyExperiment v0_00_20 -qe2:prof
```

The version and qualifiers shown here are the ones to use for the Workbook exercises. When the setup command returns, the following environment variables will be defined:

`TOYEXPERIMENT_DIR` defines the root DIRectory of the chosen UPS product

`TOYEXPERIMENT_INC` defines the path to the root directory of the C++ header files that are provided by this product (so called because the header files are INCluded)

`TOYEXPERIMENT_LIB` defines the directory that contains all of the dynamic object LIBraries (ending in `.so`) that are provided by this product

Almost all UPS products that you will use in the Workbook define these three environment variables. Several, including `toyExperiment`, define many more. Once you're running the exercises, you will be able to see all of the environment variables defined by the `toyExperiment` product by issuing the following command:

```
printenv | grep TOYEXPERIMENT
```



Many software products have version numbers that contain dot characters. UPS requires that version numbers not contain any dot characters; by convention, version dots are replaced with underscores. Therefore `v0.00.14` becomes `v0_00_14`. Also by convention, the environment variables are all upper case, regardless of the case used in the product names.

7.6 Finding Header Files

7.6.1 Introduction

Header files were introduced in Section 6.4.2. Recall that a header file typically contains the “parts list” for its associated `.cc` source file and is “included” in the `.cc` file.

The software for the Workbook depends on a large number of external products; the same is true, on an even larger scale, for the software in your experiment. The preceeding sec-

tions in this chapter discussed how to establish a working environment in which all of these software products are available for use.

When you are working with the code in the Workbook, and when you are working on your experiment, you will frequently encounter C++ classes and functions that come from these external products. An important skill is to be able to identify them when you see them and to be able to follow the clues back to their source and documentation. This section will describe how to do that.

An important aid to finding documentation is the use of *namespaces*; if you are not familiar with namespaces, consult the standard C++ documentation.

7.6.2 Finding *art* Header Files

This subsection will use the example of the class `art::Event` to illustrate how to find header files from the *art* UPS product; this will serve as a model for finding header files from most other UPS products.

The class that holds the *art* abstraction of an HEP event is named, `art::Event`; that is, the class `Event` is in the namespace `art`. In fact, all classes and functions defined by *art* are in the namespace `art`. The primary reason for this is to minimize the chances of accidental name collisions between *art* and other codes; but it also serves a very useful documentation role and is one of the clues you can use to find header files.

If you look at code that uses `art::Event` you will almost always find that the file includes the following header file:

```
#include "art/Framework/Principal/Event.h"
```

The *art* UPS product has been designed so that the relative path used to include any *art* header file starts with the directory `art`; this is another clue that the class or function of interest is part of *art*.

When you setup the *art* UPS product, it defines the environment variable `ART_INC`, which points to the root of the header file tree for *art*. You now have enough information to discover where to find the header file for `art::Event`; it is at

```
$ART_INC/art/Framework/Principal/Event.h
```

You can follow this same pattern for any class or function that is part of *art*. This will only work if you are in an environment in which `ART_INC` has been defined, which will be described in Chapters 9 and 10.

If you are new to C++, you will likely find this header file difficult to understand; you do not need to understand it when you first encounter it but, for future reference, you do need to know where to find it.

Earlier in this section, you read that if a C++ file uses `art::Event`, it would *almost always* include the appropriate header file. Why *almost* always? Because the header file `Event.h` might already be included within one of the other headers that are included in your file. If `Event.h` is indirectly included in this way, it does not hurt also to include it explicitly, but it is not required that you do so.²

We can summarize this discussion as follows: if a C++ source file uses `art::Event` it must always include the appropriate header file, either directly or indirectly.

art does not rigorously follow the pattern that the name of file is the same as the name of the class or function that it defines. The reason is that some files define multiple classes or functions; in most such cases the file is named after the most important class that it defines.

Finally, from time to time, you will need to dig through several layers of header files to find the information you need.

There are two code browsing tools that you can use to help navigate the layering of header files and to help find class declarations that are not in a file named for the class:

1. use the *art redmine*(γ) repository browser:
<https://cdcv.sfnal.gov/redmine/projects/art/repository/revisions/master/show/art>
2. use the LXR code browser: <http://cdcv.sfnal.gov/lxr/art/>

(In the above, both URLs are live links.)

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.

7.6.3 Finding Headers from Other UPS Products

Section 3.7 introduced the idea that the Workbook is built around a UPS product named `toyExperiment`, which describes a made-up experiment. All classes and functions defined in this UPS product are defined in the namespace `tex`, which is an acronym-like shorthand for `toyExperiment` (ToyEXperiment). (This shorthand makes it (a) easier to focus on the name of each class or function rather than the namespace and (b) quicker to type.)

One of the classes from the `toyExperiment` UPS product is `tex::GenParticle`, which describes particles created by the event generator, the first part of the simulation chain (see Section 3.7.2). The include directive for this class looks like

```
#include "toyExperiment/MCDataProducts/GenParticle.h"
```

As for headers included from *art*, the first element in the relative path to the included file is the name of the UPS product in which it is found. Similarly to *art*, the header file can be found using the environment variable `TOYEXPERIMENT_INC`:

```
$TOYEXPERIMENT_INC/toyExperiment/MCDataProducts/GenParticle.h
```

With a few exceptions, discussed in Section 7.6.4, if a class or function from a UPS product is used in the Workbook code, it will obey the following pattern:

1. The class will be in a namespace that is unique to the UPS product; the name of the namespace may be the full product name or a shortened version of it.
2. The lead element of the path specified in the include directive will be the name of the UPS product.
3. The UPS product setup command will define an environment variable named `PRODUCT-NAME_INC`, where *PRODUCT-NAME* is in all capital letters.

Using this information, the name of the header file will always be

```
$PRODUCT-NAME_INC/path-specified-in-the-include-directive
```

This pattern holds for all of the UPS products listed in Table 7.1.

A table listing git- and LXR-based code browsers for many of these UPS products can be found near the top of the web page:

<https://cdcvns.fnal.gov/redmine/projects/art/wiki>

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>

7.6.4 Exceptions: The Workbook, ROOT and Geant4

There are three exceptions to the pattern described in Section 7.6.3:

- the Workbook itself
- ROOT
- Geant4

The Workbook is so tightly coupled to the `toyExperiment` UPS product that all classes in the Workbook are also in its namespace, `tex`. Note, however, that classes from the Workbook and the `toyExperiment` UPS product can still be distinguished by the leading element of the relative path found in the include directives for their header files:

- `art-workbook` for the Workbook
- `toyExperiment` for the `toyExperiment`

The ROOT package is a CERN-supplied software package that is used by *art* to write data to disk files and to read it from disk files. It also provides many data analysis and data presentation tools that are widely used by the HEP community. Major design decisions for ROOT were frozen before namespaces were a stable part of the C++ language, therefore ROOT does not use namespaces. Instead ROOT adopts the following conventions:

1. All class names by defined by ROOT start with the capital letter `T` followed by another upper case letter; for example, `TFile`, `TH1F`, and `TCanvas`.

2. With very few exceptions, all header files defined by ROOT also start with the same pattern; for example, `TFile.h`, `TH1F.h`, and `TCanvas.h`.
3. The names of all global objects defined by ROOT start with a lower case letter `g` followed by an upper case letter; for example `gDirectory`, `gPad` and `gFile`.

The rule for writing an include directive for a header file from ROOT is to write its name without any leading path elements:

```
#include "TFile.h"
```

All of the ROOT header files are found in the directory that is pointed to by the environment variable `$ROOT_INC`. For example, to see the contents of this file you could enter:

```
less $ROOT_INC/TFile.h
```

Or you can learn about this class using the reference manual at the CERN web site: <http://root.cern.ch/root/html534/ClassIndex.html>

You will not see the `Geant4` package in the Workbook but it will be used by the software for your experiment, so it is described here for completeness. `Geant4` is a toolkit for modeling the propagation particles in electromagnetic fields and for modeling the interactions of particles with matter; it is the core of all detector simulation codes in HEP and is also widely used in both the Medical Imaging community and the Particle Astrophysics community.

As with ROOT, `Geant4` was designed before namespaces were a stable part of the C++ language. Therefore `Geant4` adopted the following conventions.

1. The names of all identifiers begin with `G4`; for example, `G4Step` and `G4Track`.
2. All header file names defined by `Geant4` begin with `G4`; for example, `G4Step.h` and `G4Track.h`.

Most of the header files defined by `Geant4` are found in a single directory, which is pointed to by the environment variable `G4INCLUDE`.

The rule for writing an include directive for a header file from `Geant4` is to write its name without any leading path elements:

```
#include "G4Step.h"
```

The workbook does not set up a version of Geant4; therefore G4INCLUDE is not defined. If it were, you would look at this file by:

```
less $G4INCLUDE/G4Step.h
```



Both ROOT and Geant4 define many thousands of classes, functions and global variables. In order to avoid collisions with these identifiers, do not define any identifiers that begin with any of (case-sensitive):

- T, followed by an upper case letter
- g, followed by an upper case letter
- G4

Part II

Workbook

8 Preparation for Running the Workbook Exercises

8.1 Introduction

You will run the Workbook exercises on a computer that is maintained by your experiment, either at Fermilab or at another institution. 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 work 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.



Note that it will be possible to install the Workbook software on your local (Unix-like) machine; but this option is not yet ready. When instructions are available, a link to them will be provided here. The instructions in this document will work whether the Workbook code is installed locally or on a remote machine.

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 instructions for getting computer accounts on their machines (and various other information for

¹Remember, a *site* refers to a unique combination of experiment and institution.

Table 8.1: Experiment-specific information for new users (pages are under <https://cdcvs.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
LBNE	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

new users) on web pages that they maintain, as listed in Table 8.1. The URLs in the table are live hyperlinks.

Currently, each of the experiments using *art* has installed the Workbook code on one of its experiment machines in the Fermilab General Purpose Computing Farm (GPCF).



At time of writing, the new-user instructions for all LArSoft-based experiments are at the LArSoft site; there are no separate instructions for each experiment.

If you would like a computer account on a Fermilab computer in order to evaluate *art*, contact the *art* team (see Section 3.4).

8.3 Choosing a Machine and Logging In

The experiment-specific machines confirmed to host the Workbook code are listed in Table 8.2 In most cases the name given is not the name of an actual computer, but rather a round-robin alias for a cluster. For example, if you log into `mu2evm`, you will actually be connected to one of the five computers `mu2egpvm01` through `mu2egpvm05`. These Mu2e machines share all disks that are relevant to the Workbook exercises, so if you need to log in multiple times, it is perfectly OK if you are logged into two different machines; you will still see all of the same files.

Each experiment's web page has instructions on how to log in to its computers from your local machine.

Table 8.2: Login machines for running the Workbook exercises

Experiment	Name of Login Node
ArgoNeut	argoneutvm.fnal.gov
Darkside	ds50.fnal.gov
LBNE	lbnevm.fnal.gov
MicroBoone	uboonevm.fnal.gov
Muon g-2	gm2gpvm.fnal.gov
Mu2e	mu2evm.fnal.gov
NOvA	nova-offline.fnal.gov

8.4 Launching new Windows: Verify X Connectivity

Some of the Workbook exercises will launch an X window from the remote machine that opens in your local machine. To test that this works, type `xterm &`:

```
xterm &
```



This should, without any messages, give you a new command prompt. After a few seconds, a new shell window should appear on your laptop screen; if you are logging into a Fermilab computer from a remote site, this may take up to 10 seconds. If the window does not appear, or if the command issues an error message, contact a computing expert on your experiment.

To close the new window, type `exit` at the command prompt in the new window:

```
exit
```



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

As you work through the Workbook exercises you will need to edit files. Familiarize yourself with one of the editors available on the computer that is hosting the Workbook. Most Fermilab computers offer four reasonable choices: `emacs`, `vi`, `vim` and `nedit`. Of these, `nedit` is probably the most intuitive and user-friendly. All are very powerful once you have learned to use them. Most other sites offer at least the first three choices. You can always

contact your local system administrator to suggest that other editors be installed.

A future version of this documentation suite will include recommended configurations for each editor and will provide links to documentation for each editor.

9 Exercise 1: Running Pre-built *art* Modules

9.1 Introduction

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

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 Chapter 8 to be able to find information as needed.

If you are following the instructions below on a Mac computer, 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

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)

- how to set up the toyExperiment UPS product (Section 9.6.1)
- how to run an *art* job (Section 9.6.1)
- how to control the number of events to process (Section 9.8.4)
- how to select different input files (Section 9.8.5)
- how to start at a run, subRun or event that is not the first one in the file (Section 9.8.6)
- how to concatenate input files (Section 9.8.5)
- how to write an output file (Section 9.8.9)
- some basics about the grammar and structure of a FHiCL file (Section 9.8)
- how *art* finds modules and configuration (FHiCL) files. (Sections 9.10 and 9.11)

9.4 The *art* Run-time Environment

This discussion is aimed to help you understand the process described in this chapter as a whole and how the pieces fit together in the *art* run-time environment. This environment is summarized in Figure 9.1. In this figure the boxes refer either to locations in memory or to files on a disk.

At the center of the figure is a box labelled “*art* executable;” this represents the *art* main program resident in memory after being loaded. When the *art* executable starts up, it reads its run-time configuration (FHiCL) file, represented by the box to its left. Following instructions from the configuration file, *art* will load dynamic libraries from toyExperiment, from *art*, from ROOT, from CLHEP and from other UPS products. All of these dynamic libraries (*.so* or *.dylib* files) will be found in the appropriate UPS products in LD_LIBRARY_PATH (DYLD_LIBRARY_PATH for OS X), which points to directories in the UPS products area (box at upper right). Also following instructions from the FHiCL file, *art* will look for input files (box labeled “Event-data input files” at right). The FHiCL file will tell *art* to write its event-data and histogram output files to a particular directory (box at lower right).

One remaining box in the figure (at right, second from bottom) is not encountered in the first Workbook exercise but has been provided for completeness. In most *art* jobs it is necessary to access experiment-related geometry and conditions information; in a mature

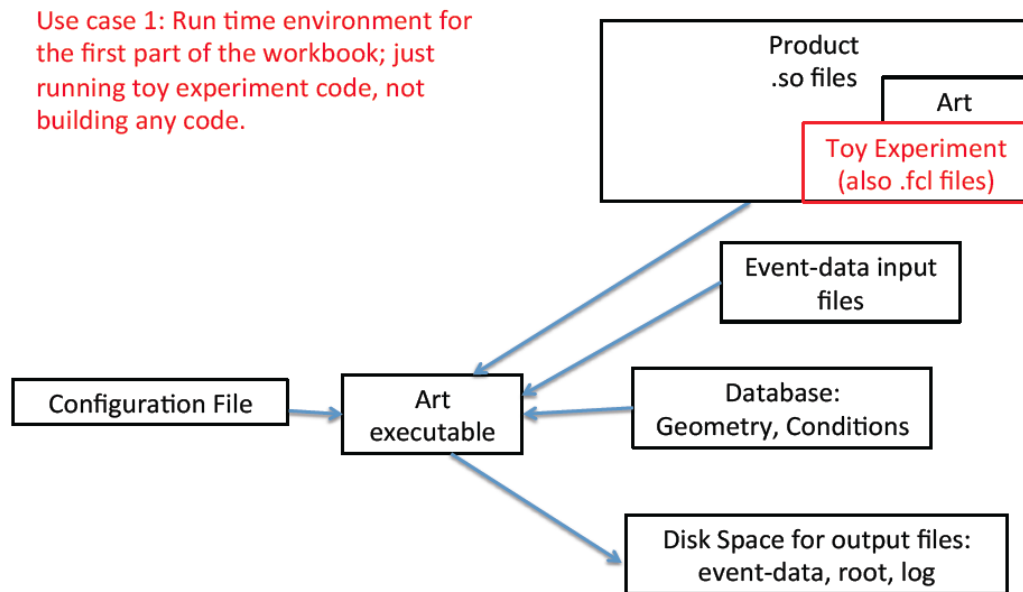


Figure 9.1: Elements of the *art* run-time environment for the first Workbook exercise

experiment, these are usually stored in a database that stands apart from the other elements in the picture.

The arrows in Figure 9.1 show the direction in which information flows. Everything but the output flows into the *art* executable.

9.5 The Input and Configuration Files for the Workbook Exercises

Several event-data input files have been provided for use by the Workbook exercises. These input files are packaged as part of the toyExperiment UPS product. Table 9.1 lists the range of event IDs found in each file. You will need to refer back to this table as you proceed.

A run-time configuration (FHiCL) file has been provided for each exercise. For Exercise 1 it is `hello.fcl`.

Table 9.1: Input files provided for the Workbook exercises

File Name	Run	SubRun	Range of Event Numbers
input01_data.root	1	0	1 ... 10
input02_data.root	2	0	1 ... 10
input03_data.root	3	0	1 ... 5
	3	1	1 ... 5
	3	2	1 ... 5
input04_data.root	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 11.

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 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_20 -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
```

Proceed to Section 9.7.

9.6.1.2 Initial Setup Procedure allowing Self-managed Working 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 a working directory and `cd` to it.
4. Setup the toyExperiment UPS product:

```
setup toyExperiment v0_00_20 -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
ls
bug01.fcl bug04.fcl inputFiles output
bug02.fcl helloExample.log inputs.txt
skipEvents.fcl bug03.fcl hello.fcl
makeLinks.sh writeFile.fcl
```

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_20 -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
2 %MSG-i MF_INIT_OK:  art 27-Apr-2013 21:22:13 CDT JobSetup
3 Messagelogger initialization complete.
4 %MSG
5 27-Apr-2013 21:22:14 CDT  Initiating request to open file
6 inputFiles/input01_data.root
7 27-Apr-2013 21:22:14 CDT  Successfully opened file
8 inputFiles/input01_data.root
9 Begin processing the 1st record. run: 1 subRun: 0 event: 1 at
10 27-Apr-2013 21:22:14 CDT
11 Hello World! This event has the id: run: 1 subRun: 0 event: 1
12 Begin processing the 2nd record. run: 1 subRun: 0 event: 2 at
13 27-Apr-2013 21:22:14 CDT
14 Hello World! This event has the id: run: 1 subRun: 0 event: 2
15 Hello World! This event has the id: run: 1 subRun: 0 event: 3
16 Hello World! This event has the id: run: 1 subRun: 0 event: 4
17 Hello World! This event has the id: run: 1 subRun: 0 event: 5
18 Hello World! This event has the id: run: 1 subRun: 0 event: 6
19 Hello World! This event has the id: run: 1 subRun: 0 event: 7
20 Hello World! This event has the id: run: 1 subRun: 0 event: 8
21 Hello World! This event has the id: run: 1 subRun: 0 event: 9
22 Hello World! This event has the id: run: 1 subRun: 0 event: 10
23 27-Apr-2013 21:22:14 CDT  Closed file inputFiles/input01_data.root
24
25 TrigReport ----- Event  Summary -----
26 TrigReport Events total = 10 passed = 10 failed = 0
27
28 TrigReport ----- Modules in End-Path: e1 -----
29 TrigReport  Trig Bit#      Visited      Passed      Failed      Error Name
30 TrigReport      0      0          10          10          0          0 hi
31
32 TimeReport ----- Time  Summary ---[sec]----
33 TimeReport CPU = 0.004000 Real = 0.002411
34
35 Art has completed and will exit with status 0.

```

Every time you run *art*, the first thing to check is the last line in your output or log file. It should be Art has completed and will exit with status 0. If the status is not 0, or if this line is missing, it is an error; please contact the *art* team as described in Section 3.4.

A future version of these instructions will specify how much disk space is needed, including space for all output files.

9.8 Understanding the Configuration

The file `hello.fcl` gives *art* its run-time configuration. This file is written in the Fermilab Hierarchical Configuration Language (FHiCL, pronounced “fickle”), a language that was developed at Fermilab to support run-time configuration for several projects, including *art*. By convention, files written in FHiCL end in `.fcl`. As you work through the Workbook, the features of FHiCL that are relevant for each exercise will be explained.

art accepts some command line options that can be used in place of items in the FHiCL file. You will encounter some of these in this section.



The full details of the FHiCL language, plus the details of how it is used by *art*, are given in the Users Guide, Chapter 24. Most people will find it much easier to follow the discussion in the Workbook documentation than to digest the full documentation up front.

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 : {
```

At the outermost scope of the FHiCL file, *art* will interpret the `source` parameter set as the description of the source of events for this run of *art*.

```
module_type : RootInput
```

`module_type` is an identifier that tells *art* the name of a module to load and run, `RootInput` in this case. `RootInput`, a standard source module provided by *art*, reads disk files containing event-data written in *art*-defined ROOT-based format.

```
fileNames : [
  "inputfiles/input01_data.root"
]
}
```

The string `fileNames` defined in the `RootInput` module scope gives the module a list of filenames from which to read events.

The name `source` 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. The notion of *scope* in FHiCL is discussed further in Chapter 14. When *art* sees a parameter set named `source` at the outermost scope, then *art* will interpret that parameter set to be the description of the source of events for this run of *art*.



In the `source` parameter set, `module_type` 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

The string `fileNames` is again an identifier, but this time defined in the `RootInput` module, that gives the module a list of filenames from which to read events. The list is delimited by square brackets and contains a comma-separated list of filenames. This exam-

¹ In the Workbook, the only source `module_type` that you will see will be `RootInput`. Your experiment may have a source module that reads events from the live experiment and other source modules that read files written in experiment-defined formats.

ple shows only one filename, but the square brackets are still required. The proper FHiCL name for a comma-separated list delimited by square brackets is a *sequence*(γ).

In most cases the filenames in the sequence must be enclosed in quotes. FHiCL, like many other languages has the following rule: if a string contains white space or any special characters, then quoting it is required, otherwise quotes are optional.

FHiCL has its own set of special characters; these include anything *except* all upper and lower case letters, the numbers 0 through 9 and the underscore character. *art* restricts the use of the underscore character in some circumstances; these will be discussed as they arise.

It is implied in the foregoing discussion that a FHiCL *value* need not be a simple thing, such as a number or a quoted string. For example, in `hello.fcl`, the value of `source` is a parameter set (of two parameters) and the value of `fileNames` is a (single-item) sequence.

9.8.2 Some Physics Processing Syntax

The identifier *physics*(γ), when found at the outermost scope, is an identifier reserved to *art*. The `physics` parameter set is so named because it contains most of the information needed to describe the physics workflow of an *art* job.

The fragment of `hello.fcl` below shows a rather long-winded way of telling *art* to find a module named `HelloWorld` and execute it.

```
physics : {
```

At the outermost scope of the FHiCL file, *art* will interpret the `physics` parameter set as the description of the physics workflow for this run of *art*.

```
  analyzers : {
```

```
    hi: {
```

```
      module_type : HelloWorld
```

```
    }
```

```
  }
```

As a top-level identifier within the `physics` scope, `analyzers` defines for *art* the run-time configuration for all the analyzer modules in the job, e.g., `HelloWorld`.

```
  e1 : [ hi ]
```

```
  end_paths : [e1 ]
```

```
}
```

`e1` is called a path; it is a FHiCL sequence of module labels; it is an arbitrary identifier. `end_paths` is a FHiCL sequence of path names; it is an identifier reserved to *art*. Together they specify the workflow; see Section 9.8.8.

Why so long-winded? *art* has very powerful features that enable execution of multiple complex chains of modules; the price is that specifying something simple takes a lot of keystrokes.

Within the `physics` parameter set, notice the identifier `analyzers`. When found as a top-level identifier within the `physics` scope, it is recognized as a keyword reserved to *art*. The `analyzers` parameter set defines the run-time configuration for all of the analyzer modules that are part of the job – only `HelloWorld` in this case.

For our current purposes, the module `HelloWorld` does only one thing of interest, namely for every event it prints one line (shown here as three):

```
Hello World! This event has the id: run: <RR>
                                   subRun: <SS>
                                   event: <EE>
```

where `RR`, `SS` and `EE` are substituted with the actual run, subRun and event number of each event.

If you look back at Listing 9.1, you will see that this line appears ten times, once each for events 1 through 10 of run 1, subRun 0 (as expected, according to Table 9.1). The remainder of the listing is standard output generated by *art*.

The remainder of the lines in `hello.fcl` appears below. The line starting with `process_name(γ)` tells *art* that this job has a name and that the name is “hello”; it has no real significance in these simple exercises. However the name of the process must not contain any underscore characters; the reason for this restriction will be explained in Section 16.4.2.

The `services` parameter set provides run-time configuration information for all *art* services.

```
#include "fcl/minimalMessageService.fcl"
```

```
process_name : hello
```

The identifier `process_name` tells *art* that the name of this job is `hello`.

```
services : {  
  message : @local::default_message }
```

The `services` parameter set provides the run-time configuration for all the required *art* services, in this case the message service. Its configuration is set in the file indicated in the `#include` directive.

For our present purposes, it is sufficient to know that the configuration for the message service is found inside the file that is included via the `#include` line. The message service controls the limiting and routing of debug, informational, warning and error messages generated by *art* or by user code. The message service does not control information written directly to `std::cout` or `std::cerr`.

9.8.3 *art* Command line Options

art supports some command line options. To see what they are, type the following command at the bash prompt:

```
art --help
```

Note that some options have both a short form and a long form. This is a common convention for Unix programs; the short form is convenient for interactive use and the long form makes scripts more readable. It is also a common convention that the short form of

an option begins single dash character, while the long form of an option begins with two dash characters, for example `--help` above.

9.8.4 Maximum Number of Events to Process

By default *art* will read all events from all of the specified input files. You can set a maximum number of events in two ways, one way is from the command line:

```
art -c hello.fcl -n 5 >& output/hello-n5.log
```

```
art -c hello.fcl --nevt 4 >& output/hello-nevt4.log
```

Run each of these commands and observe their output.

The second way is within the FHiCL file. Start by making a copy of `hello.fcl`:

```
cp hello.fcl hi.fcl
```

Edit `hi.fcl` and add the following line anywhere in the source parameter set:

```
maxEvents      : 3
```

By convention this is added after the `fileNames` definition but it can go anywhere inside the source parameter set because the order of parameters within a FHiCL table is not important. Run *art* again, using `hi.fcl`:

```
art -c hi.fcl >& output/hi.log
```

You should see output from the `HelloWorld` module for only the first three events.

To configure the file for *art* to process all the events, i.e., to run until *art* reaches the end of the input files, either leave off the `maxEvents` parameter or give it a value of `-1`.



If the maximum number of events is specified both on the command line and in the FHiCL file, then the command line takes precedence. Compare the outputs of the following commands:

```
art -c hi.fcl >& output/hi2.log
```

```
art -c hi.fcl -n 5 >& output/hi-n5.log
```

```
art -c hi.fcl -n -1 >& output/hi-neg1.log
```

9.8.5 Changing the Input Files

For historical reasons, there are multiple ways to specify the input event-data file (or the list of input files) to an *art* job:

- within the FHiCL file's `source` parameter set
- on the *art* command line via the `-s` option (you may specify one input file only)
- on the *art* command line via the `-S` option (you may specify a text file that lists multiple input files)
- on the *art* command line, after the last recognized option (you may specify one or more input files)



If input file names are provided both in the FHiCL file and on the command line, the command line takes precedence.

Let's run a few examples.

We'll start with the `-s` command line option (second bullet). Run *art* without it (again), for comparison (or recall its output from Listing 9.1):

```
art -c hello.fcl >& output/hello.log
```

To see what you should expect given the following input file, check Table 9.1, then run:

```
art -c hello.fcl -s inputFiles/input02_data.root >& output/hello-s.log
```

Notice that the 10 events in this output are from run 2 subRun 0, in contrast to the previous printout which showed events from run 1. Notice also that the command line specification overrode that in the FHiCL file. The `-s` (lower case) command line syntax will only permit you to specify a single filename.

This time, edit the source parameter set inside the `hi.fcl` file (first bullet); change it to:

```
source : {  
  module_type : RootInput  
  fileNames   : [ "inputFiles/input01_data.root",  
                  "inputFiles/input02_data.root" ]  
  maxEvents   : -1  
}
```

(Notice that you also added `maxEvents : -1`.) The names of the two input files could have been written on a single line but this example shows that, in FHiCL, newlines are treated simply as white space.

Check Table 9.1 to see what you should expect, then rerun *art* as follows:

```
art -c hi.fcl >& output/hi-2nd.log
```

You will see 20 lines from the `HelloWorld` module; you will also see messages from *art* at the open and close operations on each input file.

Back to the `-s` command-line option, run:

```
art -c hi.fcl -s inputFiles/input03_data.root >& output/run3.log
```

This will read only `inputFiles/input03_data.root` and will ignore the two files specified in the `hi.fcl`. The output from the `HelloWorld` module will be the 15 events from the three `subRuns` of run 3.

There are several ways to specify multiple files at the command line. One choice is to use the `-S` (upper case) `[--source-list]` command line option (third bullet) which takes as its argument the name of a text file containing the filename(s) of the input event-data file(s). An example of such a file has been provided, `inputs.txt`. Look at the contents of this file:

```
cat inputs.txt
```

```
inputFiles/input01_data.root  
inputFiles/input02_data.root  
inputFiles/input03_data.root
```

Now run *art* using `inputs.txt` to specify the input files:

```
art -c hi.fcl -S inputs.txt >& output/file010203.log
```

You should see the `HelloWorld` output from the 35 events in the three files; you should also see the messages from *art* about the opening and closing of the three files.

Finally, you can list the input files at the end of the *art* command line (fourth bullet):

```
art -c hi.fcl inputFiles/input02_data.root inputFiles/input03_data.root >& output/file0203.log
```

(Remember the Unix convention about a trailing backslash marking a command that continues on another line; see Chapter 2.) In this case you should see the `HelloWorld` output from the 25 events in the two files.



In summary, there are three ways to specify input files from the command line; all of them override any input files specified in the `FHiCL` file. Do not try to use two or more of these methods on a single *art* command line; the *art* job will run without issuing any messages but the output will likely be different than you expect.

9.8.6 Skipping Events

The `source` parameter set supports a syntax to start execution at a given event number or to skip a given number of events at the start of the job. Look, for example, at the file `skipEvents.fcl`, which differs from `hello.fcl` by the addition of two lines to the source parameter set:

```
firstEvent   : 5
maxEvents    : 3
```

art will process events 5, 6, and 7 of run 1, subRun 0. Try it:

```
art -c skipEvents.fcl >& output/skipevents1.log
```

An equivalent operation can be done from the command line in two different ways. Try the following two commands and compare the output:

```
art -c hello.fcl -e 5 -n 3 >& output/skipevents2.log
```

```
art -c hello.fcl --nskip 4 -n 3 >& output/skipevents3.log
```


You can also specify the initial event to process relative to a given event ID (which, recall, contains the run, subRun and event number). Edit `hi.fcl` and edit the `source` parameter set as follows:

```
source : {  
  module_type : RootInput  
  fileNames   : [ "inputFiles/input03_data.root" ]  
  firstRun     : 3  
  firstSubRun  : 1  
  firstEvent   : 6  
}
```

When you run this job, *art* will process events starting from run 3, subRun 2, event 1, – because there are only 5 events in subRun 1.

```
art -c hi.fcl >& output/startatrun3.log
```

9.8.7 Identifying the User Code to Execute

Recall from Section 9.8.2 that the `physics` parameter set contains the physics content for the *art* job. Within this parameter set, *art* must be able to determine which (user code) modules to process. These must be referenced via *module labels*(γ), which as you will see, represent the pairing of a module name and a run-time configuration.

Look back at the listing on page 133, which contains the `physics` parameter set from `hello.fcl`. The `analyzer` parameter set, nested inside the `physics` parameter set, contains the definition:

```
hi : {  
  module_type : HelloWorld  
}
```

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:

```
module_type : the-right-module-name
```

Here `module_type` is an identifier reserved to *art* and `the-right-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 `the-right-module-name` must be `EDAnalyzer`.)

Module labels are fully described in Section 24.5.

In this example *art* will look for a module named `HelloWorld`, which it will find as part of the `toyExperiment UPS` product. Section 9.10 describes how *art* uses `the-right-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.

```
analyzers : {  When art processes this fragment it will look for a
                module named MyAnalysis and instantiate it twice...

  loose: {
    module_type : MyAnalysis
    mass_cut : 20.
  }           ... once using the parameter set labeled loose ...

  tight: {
    module_type : MyAnalysis
    mass_cut : 15.
  }           ... and once using the parameter set labeled tight.

}
```

The two instances of the module `MyAnalysis` are distinguished by the module labels `tight` and `loose`. *art* requires that module labels be unique within a FHiCL file. Module labels may contain only upper- and lower-case letters and the numerals 0 to 9.



In the FHiCL files in this exercise, all of the modules are analyzer modules. Since analyzers do not make data products, these module labels are nothing more than identifiers inside the FHiCL file. For producer modules, however, which *do* make data products, the module label becomes part of the data product identifier and as such has a real significance. All module labels must conform to the same naming rules.

Within *art* there is no notion of reserved names or special names for module labels; however your experiment will almost certainly have established some naming conventions.

9.8.8 Paths and the *art* Workflow

In the `physics` parameter set in `hello.fcl` there are two parameter definitions that, taken together, specify the *workflow* of the *art* job:

```
physics {  
    e1          : [ hi ]  
    end_paths  : [ e1 ]  
}
```

where *workflow* means which modules *art* should run and in which order.²

In this exercise there is only one module to run so the workflow is trivial: for each event, run the module with the label `hi`. As you work through the Workbook you will encounter workflows that are more complex and they will be described as you encounter them.

The FHiCL parameter `e1` is called a *path*. A path is simply a FHiCL sequence of module labels. The name of a path can be any user-defined name that satisfies the following:

1. It must be defined as part of the `physics` parameter set.

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 sub-section will give only cover the parts needed to understand Workbook Exercise 1.

2. It must be a valid FHiCL name.
3. It must be unique within the *art* job.
4. It must NOT be one of the following 5 names that are reserved to *art*: `analyzers`, `filters`, `producers`, `end_paths` and `trigger_paths`.

Another way of saying the first item is that it must be defined “at physics scope”.

An *art* job may contain many paths, each of which is a FHiCL sequence of module labels. When many groups are working on a common project, this helps to maximize the independence of each work group.

The parameter `end_paths` is not itself a path! Rather it is a FHiCL sequence of path names. It is also a name that is reserved to *art*. When *art* wants to learn the workflow that it should execute it starts by looking at the `end_paths` parameter. Using this information it can find all of the module labels that are part of the workflow.

If a path is listed in the `end_paths` parameter, then it may contain module labels only for analyzer and output modules. A similar mechanism is used to specify the workflow of producer and filter modules; that mechanism will be discussed when you encounter it. If you need a reminder about the types of modules, see Section 3.6.3.

If the `end_paths` parameter is absent or defined as:

```
end_paths : [ ]
```

art will understand that this job has no analyzer modules and no output modules to execute. It is legal to define a path as an empty FHiCL sequence.

As is standard in FHiCL, if the definition of `end_paths` appears more than once, the last definition takes precedence.

The notion of path introduced in this section is the third thing in the *art* documentation suite that is called a path. The other two are the notion of a path in a filesystem and the notion of an environment variable that is a colon-delimited set of directory names. The use should be clear from the context; if it is not, please let the authors of the workbook know; see Section 3.4.

If you find the above description unsatisfying or incomplete. the next section, 9.8.8.1, has a more detail. You can choose to read it now or come back to it later.

9.8.8.1 Paths and the *art* Workflow: Details

This material is optional; it contains more details about the material just described in Section 9.8.8. It is not really a dangerous bend section — just a side trip.

Exercise 1 is not rich enough to illustrate how to specify an *art* workflow so let's construct a richer example.

Suppose that there are two groups of people working on a large collaborative project, the project leaders are Anne and Rob. Each group has a workflow that requires running 5 or 6 module instances; some of the module instances may be in the workflow for both groups. Recall that an instance of a module refers to the name of a module plus its parameter set; a module instance is specified by giving its module label. For this example let's have 8 module instances with the unimaginative names, a through h. The workflow for this example might look something like:

```
anne      : [ a, b, c, d, e, h]
rob       : [ a, b, f, c, g ]
end_paths : [ anne, rob ]
```

That is, Anne defines the modules that her group needs to run and Rob defines the modules that his group needs to run. Anne does not need to know anything about Rob's list and Rob does not need to know anything about Anne's list. The parameter definitions `anne` and `rob` are called *paths*; each is a list of module labels. The rules for legal path names were given in Section 9.8.8.

The parameter named `end_paths` is not itself a path, rather it is a FHiCL sequence of paths. Moreover it has a special meaning to *art*. During *art*'s initialization phase, it needs to learn the workflow for the job. The first step is to find the parameter named `end_paths`, defined within the `physics` parameter set. When *art* processes the definition of `end_paths` it will form the set of all module labels found in the contributing paths, with any duplicates removed. For this example, the list might look something like: `[a, b, c, d, e, h, f, g]`. When *art* processes an event, this is the set of module instances that it will execute. The order in which the module instances are executed is discussed in Section 9.8.8.2.

The above machinery probably seems a little heavyweight for the example given. But consider a workflow like that needed to design the trigger for the CMS experiment: it

had a about 200 paths and many hundreds of modules. Finding the set of unique modules labels is not a task that you would want to do by hand! By introducing the idea of paths, the design allows each group can focus on it own work, without needing ongoing, detailed consultation with every other group.

Actually, the above story is only half of the story: the module labels given in the paths `anne` and `rob` may only be the labels of analyzer or output modules. There is a parallel mechanism to specify the workflow for producer and filter modules.

To illustrate this parallel mechanism let's continue the above example of two work groups lead by Rob and Anne. In this case let there be filter modules with labels given by, `f0`, `f1`, `f2` ... and producer modules with labels given by `p0`, `p1`, `p2` In this example, a workflow might look something like:

```
t_anne      : [ p0, p1, p2, f0, p3, f1 ]
t_rob       : [ p0, p1, f2, p2, f0, p4 ]
trigger_paths : [ t_anne, t_rob ]

e_anne      : [ a, b, c, d, e ]
e_rob       : [ a, b, f, c, g ]
end_paths   : [ e_anne, e_rob ]
```

Here the parameters `t_anne`, `e_anne`, `t_rob`, and `e_rob` are all the names of paths. All must be legal FHiCL parameter names, be unique within an *art* job and not conflict with identifiers reserved to *art* at physics scope. In this example the path names are prefixed with `t_` for paths that will be put into the `trigger_paths` parameter and with `e_` for paths that will be put into the `end_paths` parameter. This is just to make it easier for you to follow the example; the prefixes have no intrinsic meaning.

During *art*'s initialization phase it processes `trigger_paths` in the same way that it processed `end_paths`: it forms the set of all module labels found in the contributing paths, with duplicates removed. The order in which the module instances are executed is discussed in Section 9.8.8.2.

What happens if you define a path with a mix of modules from the two groups:

```
bad_path     : [ p0, p1, p2, f0, p3, f1, a, b ]
end_paths    : [ e_anne, e_rob, bad_path ]
```

In this case *art* (not FHiCL) will recognize that there are producer and filter modules specified in a path that contributes to `end_paths`; *art* will then print a diagnostic message and stop. This will occur very early in *art*'s initialization phase so you will get reasonably prompt feedback. Similarly, if *art* discovers that there are analyzer or output modules in any of the paths contributing to `trigger_paths`, it will then print a diagnostic message and stop.

If you put a module label into either `end_paths` or `trigger_paths`, *art* will print a diagnostic message and stop. If you put a path name into the definition of another path *art* will print a diagnostic message and stop.

Now it's time to define two really badly chosen names:³ *trigger paths* and *end paths*, both without underscores! In the above fragment the paths prefixed with `t_` are called *trigger paths*, without an underscore; they are so named because they contain module labels for only producer and filter modules; therefore they are paths that satisfy the rules for inclusion in the definition of `trigger_paths` parameter. Similarly the paths prefixed with `e_` are called *end paths* because they satisfy the rules for inclusion in the definition of `end_paths` parameter.

While this documentation will try to avoid avoid confusion between *end paths* and `end_paths` you will certainly encounter careless use of the two ideas when you read other documentation or discuss your FHiCL files with your colleagues. There should always be enough context to sort out the meaning but it might take work.

9.8.8.2 Order of Module Execution

If the `trigger_paths` parameter contains a single *trigger path*, then *art* will execute the modules in that *trigger path* in the order that they are specified.

When more than one *trigger path* is present in `trigger_paths`, *art* will choose one of the *trigger paths* and execute its module instances in order. It will then choose a second *trigger path*. If any module instances in this path were already executed in the first *trigger path*, *art* will not execute them a second time; it will execute the remaining module instances in the order specified by the second *trigger path*. And so on for any remaining *trigger paths*.

³ It could have been worse. They could have been named on Opposite Day!



The rules for order of execution of module instances named in an *end path* are different. Since analyzer and output modules may neither add new information to the event nor communicate with each other except via the event, the processing order is not important. By definition, then, *art* may run analyzer and output modules in any order. In a simple *art* job with a single path, *art* will, in fact, run the modules in the order of appearance in the path, but do not write code that depends on execution order because *art* is free to change it.

9.8.9 Writing an Output File

The file `writeFile.fcl` gives an example of writing an output file. Open the file in an editor and find the parts of the file that are discussed below.

Output files are written by output modules; one module can write one file. An *art* job may run zero or more output modules.

If you wish to add an output module to an *art* job there three steps:

1. Create a parameter set named `outputs` at the outermost scope of the FHiCL file. The name `outputs` is prescribed by *art*.
2. Inside the `outputs` parameter set, add a parameter set to configure an output module. In `writeFile.fcl` this parameter set has the module label `output1`.
3. Add the module label of the output module to an *end path* (not to the `end_paths` parameter but to one of the paths that is included in `end_paths`). In `writeFile.fcl` the module label `output1` is added to the *end path* `e1`.

If you wish to write more output files, repeat steps 2 and 3 for each additional output file.

The parameter set `output1` tells *art* to make a module whose type is `RootOutput`. The class `RootOutput` is a standard module that is part of *art* and that writes events from memory to a disk file in an *art*-defined, ROOT-based format. The `fileName` parameter specifies the name of the output file; this parameter is processed by the `RootOutput` module. Files written by the module `RootOutput` can be read by the module `RootInput`. The identifier `output1` is just another module label that obeys the rules discussed in Section 9.8.7.

In the example of `writeFile.fcl` the output module takes its default behaviour: it will write all of the information about each input event to the output file. `RootOutput` can be configured to:

1. write only selected events
2. for each event write only a subset of the available data products.

How to do this will be described in section that will be written later.

Before running the exercise, look at the source parameter set of `writeFile.fcl`; note that it is configured to read only events 4, 5, 6, and 7.

To run `writeFile.fcl` and check that it worked correctly:

```
art -c writeFile.fcl
```

```
ls -s output/writeFile_data.root
```

```
art -c hello.fcl -s output/writeFile_data.root
```

The first command will write the output file; the second will check that the output file was created and will tell you its size; the last one will read back the output file and print the event IDs for all of the events in the file. You should see the `HelloWorld` printout for events 4, 5, 6 and 7.

9.9 Understanding the Process for Exercise 1

Section 9.6.1 contained a list of steps needed to run this exercise; this section will describe each of those steps in detail. When you understand what is done in these steps, you will understand the run-time environment in which *art* runs. As a reminder, the steps are listed again here. The 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 for the string *username*.

4.

```
cd $ART_WORKBOOK_WORKING_BASE/username/\  
workbook-tutorial/pre-built
```
5.

```
setup toyExperiment v0_00_20 -q$ART_WORKBOOK_QUAL:prof
```
6.

```
cp $TOYEXPERIMENT_DIR/HelloWorldScripts/* .
```
7.

```
source makeLinks.sh
```
8. Run *art*:

```
art -c hello.fcl >& output/hello.log
```

Steps 1 and 4 should be self explanatory and will not be discussed further.



When reading this section, you do not need to run any of the commands given here; this is a 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 replace every occurrence of the colon character with the newline character; this makes the output much easier to read.)

9.9.4 Copy Files to your Current Working Directory (Details)

The step:

```
cp $TOYEXPERIMENT_DIR/HelloWorldScripts/* .
```

only needs to be done only the first time that you log in to work on the Workbook.

In this step you copied the files that you will use for the exercises into your current working directory. You should see these files:

```
hello.fcl  makeLinks.sh  skipEvents.fcl  writeFile.fcl
```

9.9.5 Source *makeLinks.sh* (Details)

This step:

```
source makeLinks.sh
```

only needs to be done only the first time that you log in to work on the Workbook. It created some symbolic links that *art* will use.

The FHiCL files used in the Workbook exercises look for their input files in the subdirectory `inputFiles`. This script made a symbolic link, named `inputFiles`, that points to:

```
$TOYEXPERIMENT_DIR/inputFiles
```

in which the necessary input files are found.

This script also ensures that there is an output directory that you can write into when you run the exercises and adds a symbolic link from the current working directory to this output directory. The output directory is made under the directory `$ART_WORKBOOK_OUTPUT_BASE`; this environment variable was set by the site-specific setup procedure and it points to disk space that will have enough room to hold the output of the exercises.

9.9.6 Run *art* (Details)

Issuing the command:

```
art -c hello.fcl
```

runs the *art* main program, which is found in `$ART_FQ_DIR/bin`. This directory was added to your `PATH` when you setup `toyExperiment`. You can inspect your `PATH` to see that this directory is indeed there.

9.10 How does *art* find Modules?

When you ran `hello.fcl`, how did *art* find the module `HelloWorld`?

It looked at the environment variable `LD_LIBRARY_PATH`, which is a colon-delimited set of directory names defined when you setup the `toyExperiments` product. We saw the value of `LD_LIBRARY_PATH` in Section 9.9.3; to see it again, type the following:

```
printenv LD_LIBRARY_PATH | tr : \\n
```

The output should look similar to that shown in Listing 9.2.

Listing 9.2: Example of the value of `LD_LIBRARY_PATH`

```
1 /ds50/app/products/tbb/v4_1_2/Linux64bit+2.6-2.12-e2-prof/lib
2 /ds50/app/products/sqlite/v3_07_16_00/Linux64bit+2.6-2.12-prof/lib
3 /ds50/app/products/libsigcpp/v2_2_10/Linux64bit+2.6-2.12-e2-prof/lib
4 /ds50/app/products/cppunit/v1_12_1/Linux64bit+2.6-2.12-e2-prof/lib
5 /ds50/app/products/clhep/v2_1_3_1/Linux64bit+2.6-2.12-e2-prof/lib
6 /ds50/app/products/python/v2_7_3/Linux64bit+2.6-2.12-gcc47/lib
7 /ds50/app/products/libxml2/v2_8_0/Linux64bit+2.6-2.12-gcc47-prof/lib
8 /ds50/app/products/fftw/v3_3_2/Linux64bit+2.6-2.12-gcc47-prof/lib
9 /ds50/app/products/root/v5_34_05/Linux64bit+2.6-2.12-e2-prof/lib
10 /ds50/app/products/boost/v1_53_0/Linux64bit+2.6-2.12-e2-prof/lib
11 /ds50/app/products/cpp0x/v1_03_15/slf6.x86_64.e2.prof/lib
12 /ds50/app/products/cetlib/v1_03_15/slf6.x86_64.e2.prof/lib2
13 /ds50/app/products/fhiclcpp/v2_17_02/slf6.x86_64.e2.prof/lib
14 /ds50/app/products/messagefacility/v1_10_16/slf6.x86_64.e2.prof/lib
15 /ds50/app/products/art/v1_06_00/slf6.x86_64.e2.prof/lib
16 /ds50/app/products/toyExperiment/v0_00_14/slf6.x86_64.e2.prof/lib
17 /grid/fermiapp/products/common/prd/git/v1_8_0_1/Linux64bit-2/lib
```

Of course the leading element of each directory name, `/ds50/app` will be replaced by whatever is correct for your experiment. The last element in `LD_LIBRARY_PATH` is not relevant for running *art* and it may or may not be present on your machine, depending on details of what is done inside your site-specific setup procedure.

If you compare the names of the directories listed in `LD_LIBRARY_PATH` to the names of the directories listed in the `PRODUCTS` environment variable, you will see that all of these directories are part of the UPS products system. Moreover, for each product, the version, flavor and qualifiers are embedded in the directory name. In particular, both *art* and `toyExperiment` are found in the list.

If you ls the directories in LD_LIBRARY_PATH you will find that each directory contains many dynamic object library (.so files).

When *art* looks for a module named HelloWorld, it looks through the directories defined in

LD_LIBRARY_PATH and looks for a file whose name matches the pattern,

```
lib*HelloWorld_module.so
```

where the asterisk matches (zero or) any combination of characters. *art* finds that, in all of the directories, there is exactly one file that matches the pattern, and it is found in the directory (shown here on two lines):

```
/ds50/app/products/toyExperiment/v0_00_14/  
    slf6.x86_64.e2.prof/lib/
```

The name of the file is:

```
libtoyExperiment_Analyzers_HelloWorld_module.so
```

If *art* had found no files that matched the pattern, it would have printed a diagnostic message and stopped execution. If *art* had found more than one file that matched the pattern, it would have printed a different diagnostic message and stopped execution. If this second error occurs it is possible to tell *art* which of the matches to choose; how to do this will be covered in a future chapter.

9.11 How does *art* find FHiCL Files?

This section will describe where *art* looks for FHiCL files. There are two cases: looking for the file specified by the command line argument -c and looking for files that have been included by a #include directive within a FHiCL file.

9.11.1 The -c command line argument

When you issued the command

```
art -c hello.fcl
```

art looked for a file named `hello.fcl` in the current working directory and found it. You may specify any absolute or relative path as the argument of the `-c` option. If *art* had not found `hello.fcl` in this directory it would have looked for it relative to the path defined by the environment variable `FHICL_FILE_PATH`. This is just another path-type environment variable, like `PATH` or `LD_LIBRARY_PATH`. You can inspect the value of `FHICL_FILE_PATH` by:

```
printenv FHICL_FILE_PATH
```

```
.:<some-directory-structure>products//toyExperiment/v0_00_20
```

In this case, the output will show the translated value of the environment variable `TOY_EXPERIMENT_DIR`. The presence of the current working directory (dot) in the path is redundant when processing the command line argument but it is significant in the case discussed in the next section.



Some experiments have chosen to configure their version of the *art* main program so that it will not look for the command line argument `FHiCL` file in `FHICL_FILE_PATH`. It is also possible to configure *art* so that only relative paths, not absolute paths, are legal values of the `-c` argument. This last option can be used to help ensure that only version-controlled files are used when running production jobs. Experiments may enable or disable either of these options when their main program is built.

9.11.2 `#include` Files

Section 9.8 discussed the listing on page 134, which contains the fragments of `hello.fcl` that are related to configuring the message service. The first line in that listing is an include directive. *art* will look for the file named by the include directive relative to `FHICL_FILE_PATH` and it will find it in:

```
$TOYEXPERIMENT_DIR/fcl/minimalMessageService.fcl
```

This is part of the `toyExperiment` UPS product.

The version of *art* used in the Workbook does not consider the argument of the include directive as an absolute path or as a path relative to the current working directory; it only looks for files relative to `FHICL_FILE_PATH`. This is in contrast to the choice made when processing the `-c` command line option.

When building *art*, one may configure *art* to first consider the argument of the `include` directive as a path and to consider `FHiCL_FILE_PATH` only if that fails.



9.12 Review

- 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)
- how to set up the toyExperiment UPS product (Section 9.6.1)
- how to run an *art* job (Section 9.6.1)
- how to control the number of events to process (Section 9.8.4)
- how to select different input files (Section 9.8.5)
- how to start at a run, subRun or event that is not the first one in the file (Section 9.8.6)
- how to concatenate input files (Section 9.8.5)
- how to write an output file (Section 9.8.9)
- some basics about the grammar and structure of a FHiCL file (Section 9.8)
- how *art* finds modules and configuration (FHiCL) files. (Sections 9.10 and 9.11)

9.13 Test your Understanding

When you got set up to run Exercise 1 in Section 9.6.1, you may have noticed the four FHiCL files with *bug* in their names: `bug01.fcl`, `bug02.fcl`, `bug03.fcl` and `bug04.fcl`. These are files into which errors have expressly been inserted. You will now find and fix these errors. Answers are provided at the end.

9.13.1 Tests

For each of these exercises, your job is to figure out what's wrong, fix it and rerun the command. With practice you will learn to decipher the error messages that *art* throws your

way and solve the problems.

I am not aware of any place that lists the values of the completion codes.

Note that in the output of each command the `MSG-i` messages are just informational – they may or may not be relevant to the actual error. On the other hand, the `MSG-s` messages – that appear only when the FHiCL files have errors – are severe and cause *art* to shut down.

The `MSG-s` message is what actually caused *art* to shutdown.

1. `art -c bug01.fcl`

This will make an error message that includes the following:

```
Input file not found: input01_data.root.  
...  
Art has completed and will exit with status 24.  
...  
%MSG-s  
...
```

2. `art -c bug02.fcl`

This will make an error message that includes the following:

```
The following module label is not assigned to any path:  
'hi'  
...  
Path configuration: The following were encountered while
```

```
processing path configurations:
ERROR: Entry hello in path e1 refers to a module label hello
which is not configured.
...
Art has completed and will exit with status 9.
...
%MSG-s
...
```

3. `art -c bug03.fcl`

This will make an error message that includes the following:

```
ERROR: Configuration of module with label hi encountered
the following error:
--- Configuration BEGIN
UnknownModule
--- Configuration BEGIN
Library specificaton "Helloworld": does not correspond to
any library in LD_LIBRARY_PATH of type "module"
--- Configuration END
Module Helloworld with version v1_09_03 was not registered.
Perhaps your module type is misspelled or is not a framework
plugin.
...
Art has completed and will exit with status 9.
```

4. `art -c bug04.fcl`

This will make an error message that includes the following:

```
ERROR: Entry e1 in path end_path refers to a module label  
e1 which is not configured.
```

```
...
```

```
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. The critical part of the error message is:

```
Input file not found: input01_data.root
```

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:

```
fileNames : [ "inputFiles/input01_data.root" ]
```

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.

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`.

3. The key parts of this error message are:

```
Library specifiction "Helloworld": does not correspond to  
any library in LD_LIBRARY_PATH of type "module"
```

and

Perhaps your module type is misspelled or is not a framework

plugin.

The answer is that in the definition of the module label `hi` the value of the `module_type` parameter is misspelled. It has a lower case `w` when it should have upper case.

4. This fourth one is tricky. The `MSG-s` message is the same in `bug02.fcl`! The clue is that it thinks `end_path` should contain module labels not a sequence of paths. The answer is that `end_path` is misspelled. It is missing the final `s` and should be `end_paths`.

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**. This 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*:

- In the *source window* you will check out and edit source code.
- In the *build window* you will build and run code.

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.



Exercise 2 and all subsequent Workbook exercises will use the setup instructions found in Sections 10.4 and 11.

10.2 Prerequisites

Before running this exercise, you need to be familiar with the material in Part I (Introduction) of this documentation set and Chapter 9 from Part II (Workbook). Concepts that this chapter refers to include:

- namespace
- `#include` directives
- header file
- class
- constructor
- destructor
- the C preprocessor
- member function (aka method)
- const vs non-const member function
- argument list of a function
- signature of a function
- *declaration vs definition* of a class
- arguments passed by reference
- arguments passed by const reference
- notion of *type*: e.g., a class, a struct, a free function or a typedef

In this chapter you will also encounter the C++ idea of *inheritance*. Understanding inheritance is not a prerequisite; it will be described as you encounter it in the Workbook exercises. Inheritance includes such ideas as,

- base class
- derived class
- virtual function
- pure virtual function
- concrete class

10.3 What You Will Learn

In this exercise you will learn:

- how to establish the *art* development environment
- how to checkout the Workbook exercises from the `git` source code management system
- how to use the **cetbuildtools** build system to build the code for the Workbook exercises
- how include files are found
- what a *link list* is
- where the build system finds the link list
- what the `art::Event` is and how to access it
- what the `art::EventID` is and how to access it
- what makes a class an *art module*
- 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://cdcv.s.fnal.gov/projects/art-workbook`
- (c) `cd art-workbook`
- (d) `git checkout -b work v0_00_25`

6. Source the script that sets up the environment properly:

```
source ups/setup_deps -p $ART_WORKBOOK_QUAL
```

The git commands are discussed in Section 10.4.2.1. The final step sources a script that defines a lot of environment variables — the same set that will be defined in the build window.

10.4.2 Examine Source Window Setup

10.4.2.1 About git and What it Did

Git is a source code management system² that is used to hold the source code for the Workbook exercises. A source code management system is a tool that looks after the book-keeping of the development of a code base; among many other things it keeps a complete history of all changes and allows one to get a copy of the source code as it existed at any time in the past. Because of git’s many advanced features, many HEP experiments are moving to git. git is fully described in the git manual³.

Some experiments set up git in their site-specific setup procedure; others do not. In running setup git, you have ensured that a working copy of git is in your PATH⁴.

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>.

⁴No version needs to be supplied because the git UPS product has a current version declared; see Section 7.4.

The `git clone` and `git checkout` commands produce a working copy of the Workbook source files in your source directory. Figure 10.1 shows a map of the source directory structure created by the `git` commands. It does not show all the contents in each subdirectory. Note that the `.git` (hidden) directory under the source directory is colored differently; this is done to distinguish it from the rest of the contents of the source directory structure:

- When you ran `git clone` in Section 10.4.1, it copied the entire contents of the remote repository into this directory. The `.git` directory contains your local copy of the repository.
- When you ran `git checkout`, it created the rest of the structure under the source directory (what we call your “working area”) and copied the requested version of everything you need from `.git` into this structure.

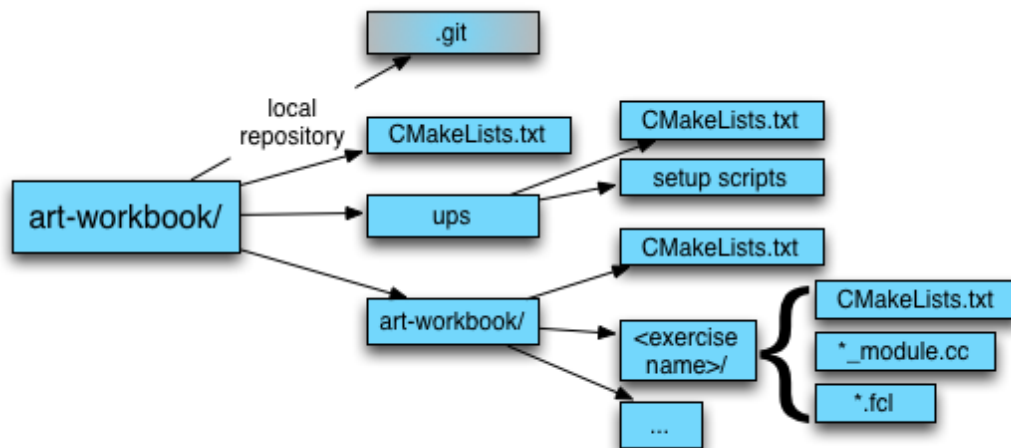


Figure 10.1: Representation of the reader’s source directory structure (an `admin` directory is not shown)

; `git clone` should produce the following output:

```
Cloning into 'art-workbook'...
```

Executing the `git checkout` command should produce the following output:

```
Switched to a new branch 'v0_00_25'
```

If you wish to learn about `git` branches, for the time being, you will need to consult a

`git` manual.

If you do not see the expected output, contact the *art* team as described in Section 3.4.

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 correspondance 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.
- The `ups` directory contains information about what UPS products this product depends on; it contains additional information used to configure the development environment.

Look inside the `art-workbook` (sub)directory (via `ls`) and see that it contains several files and subdirectories. The file `CMakeLists.txt` contains more instructions for the build system. Actually, every directory contains a `CMakeLists.txt`; each contains additional instructions for the build system. The subdirectory `FirstModule` contains the files that will be used in this exercise; the remaining subdirectories contain files that will be used in subsequent Workbook exercises.

If you look inside the `FirstModule` directory, you will see

```
CMakeLists.txt FirstAnswer01_module.cc First_module.cc  
firstAnswer01.fcl first.fcl
```

The file `CMakeLists.txt` in here contains yet more instructions for the build system and will be discussed later. The file `First_module.cc` is the first module that you will look at and `first.fcl` is the FHiCL file that runs it. This exercise will suggest that you try to write some code on your own; the answer is provided in

`FirstAnswer01_module.cc` and the file `firstAnswer01.fcl` runs it. These files will be discussed at length throughout the exercises.

10.4.3 “Build Window” Setup



Again, advanced users wanting to manage their own working directory may skip to Section 10.4.3.2.

10.4.3.1 Standard Procedure

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.
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`

This step may take a few minutes.

Skip Section 10.4.3.2 and move on to Section 10.4.4.

10.4.3.2 Using Self-managed Working Directory



The steps in this procedure that are the same as for the “standard” procedure are explained in Section 10.4.4.

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.
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`

10.4.4 Examine Build Window Setup

Logging in and sourcing the site-specific setup script should be clear by now. Notice that next you are told to `cd` to the same `workbook` directory as in Step 4 of the instructions for the source window. From there, you make a directory in which you will run builds (your *build* directory), and `cd` to it. (The name `build-prof` can be any legal directory name but it is suggested here because this example performs a profile build; this is explained in Section 3.6.7). Figure 10.2 shows roughly what the build directory contains.

Step 6 sources a script called `setup_for_development` found in the `ups` subdirectory of the source directory. This script, run exactly as indicated, defines `build-prof` to be your build directory. This command selects a profile build (via the option `-p`); it also requests that the UPS qualifiers defined in the environment variable `ART_WORKBOOK_QUAL` be used when requesting the UPS products on which it depends; this environment variable was discussed in Section 9.9.3. The expected output is shown in Listing 10.1.

Check that there are no error messages in the indicated block. The listing concludes with a request for you to run a `cmake` command; *do not run* `cmake` (this line is an artifact of layering **cetbuildtools** on top of `cmake`).



Listing 10.1: Example of output created by `setup_for_development`

```

1
2 The working build directory is /ds50/app/user/kutschke/workbook/build-prof
3 The source code directory is /ds50/app/user/kutschke/workbook/art-workbook
4 ----- check this block for errors -----
5 -----
```

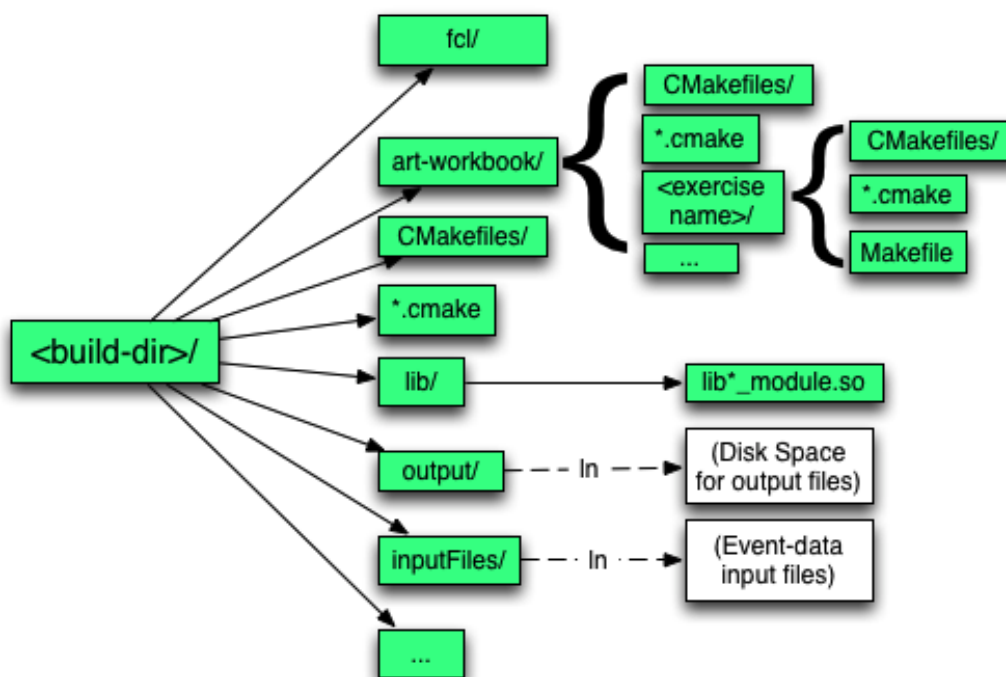



Figure 10.2: Representation of the reader's build directory structure (the `fcl/` directory is a sym-link to `art-workbook/art-workbook/` in the source area)

```

6 /ds50/app/user/kutschke/workbook/build-prof/lib has been added to LD_LIBRARY_PATH
7 /ds50/app/user/kutschke/workbook/build-prof/bin has been added to PATH
8
9 CETPKG_SOURCE=/ds50/app/user/kutschke/workbook/art-workbook
10 CETPKG_BUILD=/ds50/app/user/kutschke/workbook/build-prof
11 CETPKG_NAME=art_workbook
12 CETPKG_VERSION=v0_00_15
13 CETPKG_QUAL=e2:prof
14 CETPKG_TYPE=Prof
15
16 Please use this cmake command:
17 cmake -DCMAKE_INSTALL_PREFIX=/install/path
18           -DCMAKE_BUILD_TYPE=$CETPKG_TYPE $CETPKG_SOURCE

```

This script sets up all of the UPS products on which the Workbook depends; this is analogous to the actions taken by Step 6 in the first exercise (Section 9.6.1.1) when you were working in the *art* run-time environment. This script also creates several files and directories in your *build-prof* directory; these comprise the working space used by **cetbuildtools**.

After sourcing this script, the contents of *build-prof* will be

```

art_workbook-v0_00_25  bin lib
cetpkg_variable_report diag_report

```

At this time the two subdirectories *bin* and *lib* will be empty. The other files are used by the build system to keep track of its configuration.

Step 7 (buildtool) tells **cetbuildtools** to build everything found in the source directory; this includes all of the Workbook exercises, not just the first one. The build process will take two or three minutes on an unloaded (not undergoing heavy usage) machine. Its output should end with the lines:

```

-----
INFO: Stage build successful.
-----

```

After the build has completed do an *ls* on the directory *lib*; you will see that it contains a large number of dynamic library (*.so*) files; for *v0_00_25* there will be about 30 *.so* files (subject to variation as versions change); these are the files that *art* will load as you work through the exercises.

Also do an `ls` on the directory `bin`; these are scripts that are used by **cetbuildtools** to maintain its environment; if the Workbook contained instructions to build any executable programs, they would have been written to this directory.

After running `buildtool`, the build directory will contain:

<code>admin</code>	<code>CMakeFiles</code>	<code>fcl</code>
<code>art-workbook</code>	<code>cmake_install.cmake</code>	<code>inputFiles</code>
<code>art_workbook-v0_00_15</code>	<code>CPackConfig.cmake</code>	<code>lib</code>
<code>bin</code>	<code>CPackSourceConfig.cmake</code>	<code>Makefile</code>
<code>cetpkg_variable_report</code>	<code>CTestTestfile.cmake</code>	<code>output</code>
<code>CMakeCache.txt</code>	<code>diag_report</code>	<code>ups</code>

Most of these files are standard files that are explained in the **cetbuildtools** documentation, <https://cdcvs.fnal.gov/redmine/projects/cetbuildtools/wiki>. However, three of these items need special attention here because they are customized for the Workbook.

An `ls -l` on the files `fcl`, `inputFiles` and `output` will reveal that they are symbolic links to

```
inputFiles -> ${TOYEXPERIMENT_DIR}/inputFiles
output -> ${ART_WORKBOOK_OUTPUT_BASE}/
          <username>/art_workbook_output
fcl -> <your source directory>/art-workbook
```

These links are present so that the FHiCL files for the Workbook exercises do not need to be customized on a per-user or per-site basis.

- The link `inputFiles` points to the directory `inputFiles` present in the toyExperiment UPS product; this directory contains the input files that *art* will read when you run the first exercise. These are the same files used in the first exercise; if you need a reminder of the contents of these files, see Table 9.1. These input files will also be used in many of the subsequent exercises.
- The link `outputFiles` points to a directory that was created to hold your output files; the environment variable `ART_WORKBOOK_OUTPUT_BASE` was defined by your site-specific setup procedure.
- The symlink `fcl` points into your source directory hierarchy; it allows you to ac-

cess the FHiCL files that are found in that hierarchy with the convenience of tab completions.

10.5 The *art* Development Environment

In the preceeding sections of this chapter you established what is known as the *art development environment*; this is a superset of the *art* run-time environment, which was described in Section 9.4. This section summarizes the new elements that are part of the development environment but not part of the run-time environment.

In Section 10.4.1, step 4b (`git clone ...`) contacted the central source code repository for the *art* Workbook code and made a clone of the repository in your source area under `art-workbook`; the clone contains the complete history of the repository, including all versions of `art-workbook`. Step 4d (`git checkout ...`) examined your clone of the repository, found the requested version of the code and put a working copy of that version into your source directory. The central repository is hosted on a Fermilab server and is accessed via the network. The upper left box in Figure 10.3 denotes the central repository and the box below it denotes the clone of the repository in your disk space; the box below that denotes the checked out working copy of the Workbook code. The flow of information during the clone and checkout processes is indicated by the green arrows (at left) in the figure.

In step 7 of Section 10.4.3, you ran `buildtool` in your build area, which read the source code files from your working copy of the Workbook code and turned them into dynamic libraries. The script `buildtool` is part of the build system, which is denoted as the box in the center left section of Figure 10.3. When you ran `buildtool`, it wrote dynamic library files to the `lib` subdirectory of your build directory; this directory is denoted in the figure as the box in the top center labeled `<build-directory>/lib`. The orange arrows in the figure denote the information flow at build-time. In order to perform this task, `buildtool` also needed to read header files and dynamic libraries found in the UPS products area, hence the orange arrow leading from the UPS Products box to the build system box.

In the figure, information flow at run-time is denoted by the blue lines. When you ran the *art* executable, it looked for dynamic libraries in the directories defined by `LD_LIBRARY_PATH`. In the *art* development environment, `LD_LIBRARY_PATH` contains

1. the `lib` subdirectory of your build directory

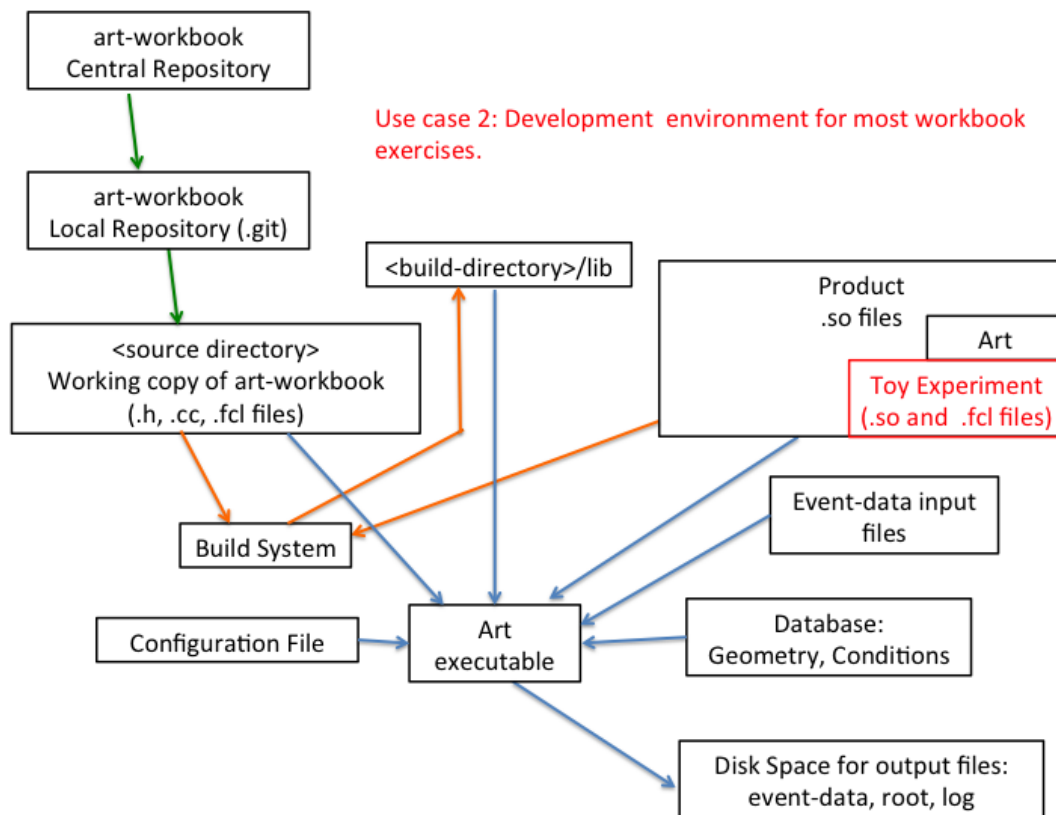


Figure 10.3: Elements of the *art* development environment as used in most of the Workbook exercises; the arrows denote information flow, as described in the text.

2. all of the directories previously described in Section 9.10

In all environments, the *art* executable looks for FHiCL files in

1. in the file specified in the `-c` command line argument
2. in the directories specified in `FHICL_FILE_PATH`

The first of these is denoted in the figure by the box labeled “Configuration File.” In the *art* development environment, `FHICL_FILE_PATH` contains

1. some directories found in your checked out copy of the source
2. all of the directories previously described in Section 9.11

The remaining elements in Figure 10.3 are the same as described for Figure 9.1.

Figure 10.4, a combination of Figures 10.1 and 10.2), illustrates the distinct source and build areas, and the relationship between them. It does not show all the contents in each subdirectory.

10.6 Running the Exercise

10.6.1 Run *art* on `first.fcl`

In your build window, make sure that your current working directory is your build directory. From here, run the first part of this exercise by typing the following:

```
art -c fcl/FirstModule/first.fcl >& output/first.log
```

(As a reminder, we suggest you get in the habit of routing your output to the `output` directory.) The output of this step will look much like that in Listing 9.1, but with two significant differences. The first difference is that the output from `first.fcl` contains an additional line

```
Hello from First::constructor.
```

The second difference is that the words printed out for each event are a little different; the printout from `first.fcl` looks like

```
Hello from First::analyze. Event id: run: 1 subRun: 0 event: 1
```




Figure 10.4: Representation of the reader's directory structure once the development environment is established.

while that from `hello.fcl` looked like

```
Hello World! This event has the id: run: 1 subRun: 0 event: 1
```

The reason for changing this printout is so that you can identify, from the printout, which module was run.

10.6.2 The FHiCL File `first.fcl`

Compare the FHiCL file used in this exercise, `fcl/FirstModule/first.fcl`, with `hello.fcl` from the first exercise (i.e., run `cat` or `diff` on them). Other than comments, the only difference is that the `module_type` has changed from `HelloWorld` to `First`:

```
diff $TOYEXPERIMENT_DIR/HelloWorldScripts/hello.fcl fcl/FirstModule/first.fcl
```

```
...
<      module_type : HelloWorld
---
>      module_type : First
```

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:

```
lib/libart-workbook_FirstModule_First_module.so
```

This dynamic library file was created when you ran `buildtool`.

10.6.3 The Source Code File `First_module.cc`

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!):

```
art-workbook/FirstModule/First_module.cc
```


When you ran `buildtool`, it compiled and linked this source file into the following dynamic library (relative to your build directory):

```
lib/libart-workbook_FirstModule_First_module.so
```

This is the dynamic library that was loaded by *art* when you ran code for this exercise, in Section 10.6.2.

Look at the file `First_module.cc`, shown in Listing 10.2. In broad strokes, it does the following:

- declares a class named `First`
- provides the implementation for the class
- contains a call to the C-Preprocessor macro named `DEFINE_ART_MODULE`, discussed in Section 10.6.3.8

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 separate files; the *art* team does not recommend this practice, but it is in use and it will be discussed in Section 10.9.2.

10.6.3.1 The `#include` Statements

```
#include "art/Framework/Core/EDAnalyzer.h"  
#include "art/Framework/Core/ModuleMacros.h"  
#include "art/Framework/Principal/Event.h"
```

The first three lines of code in the file `First_module.cc` are *include directives* that pull in header files. All three of these files are included from the *art* UPS product (determining the location of included header files is discussed in Section 7.6).

```
#include <iostream>
```

The next line, `#include <iostream>`, includes the C++ header that enables this code to write output to the screen; for details, see any standard C++ documentation.

Listing 10.2: The contents of `First_module.cc`

```
1
2 #include "art/Framework/Core/EDAnalyzer.h"
3 #include "art/Framework/Core/ModuleMacros.h"
4 #include "art/Framework/Principal/Event.h"
5
6 #include <iostream>
7
8 namespace tex {
9
10   class First : public art::EDAnalyzer {
11
12   public:
13
14     explicit First(fhicl::ParameterSet const& );
15
16     void analyze(art::Event const& event) override;
17
18   };
19
20 }
21
22 tex::First::First(fhicl::ParameterSet const& pset) : art::EDAnalyzer(pset) {
23   std::cout << "Hello_from_First::constructor." << std::endl;
24 }
25
26 void tex::First::analyze(art::Event const& event){
27   std::cout << "Hello_from_First::analyze._Event_id:_ "
28             << event.id()
29             << std::endl;
30 }
31
32 DEFINE_ART_MODULE(tex::First)
```

Those of you with some C++ experience may have noticed that there is no file named `First_module.h` in the directory `art-workbook/FirstModule`. The explanation for this will be given in Section 10.9.1.

If you are a C++ beginner you will likely find these header files difficult to understand; you do not need to understand them at this time but you do need to know where to find them for future reference.



10.6.3.2 The Declaration of the Class `First`, an Analyzer Module

Let's start with short explanations of each line and follow up with more information.

```
namespace tex {
```

Open a namespace named `tex`.

```
    class First : public art::EDAnalyzer{
```

- The first line of the declaration of the class `First`.
- Analyzer modules must inherit publicly from the base class `EDAnalyzer`.

```
    public:
```

Class members below here are public; any above would be private.

```
        explicit First(fhicl::ParameterSet const& );
```

- Declaration of a constructor.
- Its argument list is prescribed by `art`.
- `art` will call the constructor once at the start of each job.

```
        void analyze(art::Event const& event) override;
```

- Declaration of the `analyze` member function.
- Its argument list is prescribed by `art`.
- `art` will call this member function once per event.
- The `override` contextual identifier is an important safety feature. Use it!

```
    };
```

Close the declaration of the class `First`.

```
}
```

Close the namespace `tex`.

All of the code in the toyExperiment UPS product was written in the namespace `tex`; the name `tex` is an acronym-like shorthand for the toyExperiment (ToyEXperiment) UPS

product. In order to keep things simple, all of the classes in the Workbook are also declared in the namespace `tex`. For more information about this choice, see Section 7.6.4. If you are not familiar with namespaces, consult the standard C++ documentation.

In the first line of the class declaration, the fragment `: public art::EDAnalyzer` tells the compiler that the class `First` is a *derived class* that *inherits publicly* from the *base class* `art::EDAnalyzer`. At this time it is not necessary to understand inheritance, which is fortunate, because it takes a long, long time to explain. You just need to recognize and follow the pattern. You can read about C++ inheritance in the standard C++ documentation.

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"
```

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.

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 requirement that the class provide the prescribed constructor is enforced by the macro `DEFINE_ART_MODULE`, which will be described in Section 10.6.3.8.

The declaration of the constructor begins with the keyword `explicit`. This is a safety feature this relevant only for constructors that have exactly one argument. A proper explanation would take too long so just follow a simple guideline: all constructors that have exactly one argument should be declared `explicit`. There will be rare circumstances in which you need to go against this guideline but you will not encounter any in the Workbook.

The *override* contextual identifier in the `analyzer` member function definition is a feature that is new in C++ 11 so older references will not discuss it. It is a new safety feature that we recommend you use; we cannot give a proper explanation until we have had a chance to discuss inheritance further. For now, just consider it a rule that, in all analyzer modules, you should provide this identifier as part of the declaration of `analyze`.

For those who are knowledgeable about C++, the base class `art::EDAnalyzer` declares the member function `analyze` to be pure virtual; so it must be provided by the derived class. The optional member functions of the base class are declared virtual but not pure virtual; do-nothing versions of these member functions are provided by the base class.



In a future version of this documentation suite, more information will be available in the Users Guide.

10.6.3.4 The Constructor for the Class `First`

The next code in the source file (Listing 10.2) is the *definition* of the constructor for the class `First`. This constructor simply prints some information (via `std::cout`) to let the user know that it has been called.

```
tex::First::First(fhicl::ParameterSet const& pset) : art::EDAnalyzer(pset) {  
    std::cout << "Hello from First::constructor." << std::endl; }
```

The fragment `tex::First::First` says that this definition is for a constructor of the class `First` from the namespace `tex`.

The argument to the constructor is of type `fhicl::ParameterSet const&` as required by *art*. The class `ParameterSet`, found in the namespace `fhicl`, is a C++ representation of a FHiCL parameter set (aka FHiCL *table*). You will learn how to use this parameter set object in Chapter 14.

The argument to the constructor is passed by `const` reference, `const&`. This is a requirement specified by *art*; if you write a constructor that does not have exactly the correct argument type, then the compiler will issue a diagnostic and will stop compilation.

The first line of the constructor contains the fragment “`: art::EDAnalyzer(pset)`”. This is the constructor’s initializer list and it tells the compiler to call the constructor of the base class `art::EDAnalyzer`, passing it the parameter set as an argument. This is required by rule 3 in the list in Section 10.6.3.3.

The requirement that the constructor of an analyzer module pass the parameter set to the constructor of `art::EDAnalyzer` started in *art* version 1.08.09. If you are using an earlier version of *art*, constructors of analyzer modules must NOT call the constructor of `art::EDAnalyzer`.

10.6.3.5 Aside: Omitting Argument Names in Function Declarations

In the declaration of the class `First`, you may have noticed that the declaration of the member function `analyze` supplied a name for its argument (`event`) but the declaration

of the constructor did not supply a name for its argument.

In the declaration of a function, a name supplied for an argument is ignored by the compiler. So code will compile correctly with or without a name. Remember that a constructor is just a special kind of function so the rule applies to constructors too. It is very common for authors of code to provide an argument name as a form of documentation. You will code written both with and without named arguments in declarations.

The above discussion only applied to the *declarations* of functions, not to their definition (aka implementation).

10.6.3.6 The Member Function `analyze` and the Representation of an Event

The definition of the member function `analyze` comes next in the source file and is reproduced here

```
void tex::First::analyze(art::Event const& event){
    std::cout << "Hello from First::analyze. Event id: "
               << event.id()
               << std::endl;
}
```

If the type of the argument is not exactly correct, including the `const&`, the compiler will issue a diagnostic and stop compilation. The compiler is able to do this because of one of the features of *inheritance*; the details of how this works is beyond the scope of this discussion.

Note that the `override` contextual identifier that was present in the declaration of this member function is not present in its definition; this is standard C++ usage.

Section 3.6.1 discussed the HEP idea of an event and the *art* idea of a three-part event identifier. The class `art::Event` is the representation within *art* of the HEP notion of an event. For the present discussion it is safe to consider the following over-simplified view of an event: it contains an event identifier plus a collection of data products (see Section 3.6.4). The name of the argument `event` has no meaning either to *art* or to the compiler — it is just an identifier — but your code will be easier to read if you choose a meaningful name.



At any given time in a running *art* program there is only ever one `art::Event` object; in the rest of this paragraph we will call this object *the event*. It is owned and managed by *art*, but *art* lets analyzer modules see the contents of the event; it does so by passing the event by `const` reference when it calls the `analyze` member function of analyzer modules. Because the event is passed by reference (indicated by the `&`), the member function `analyze` does not get a copy of the event; instead it is told where to find the event. This makes it efficient to pass an event object even if the event contains a lot of information. Because the argument is a `const` reference, if your code tries to change the contents of the event, the compiler will issue a diagnostic and stop compilation.



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 the answer, this section will guide you through the process of discovering the answer for yourself.

Before you work through this section, you may wish to review Section 7.6 which discusses how to find header files.

In Section 10.6.3.6 you learned that the member function `art::Event::id()` returns an object that represents the event identifier. To see this for yourself, look at the header file for `art::Event`. Enter:

```
less $ART_INC/art/Framework/Principal/Event.h
```

or use one of the code browsers discussed in 7.6.2. In this file you will find the definition of the member function `id()`:⁶

```
EventID  
id() const {return aux_.id();}
```

The important thing to look at here is the return type, `EventID`; you do not need to (or want to) know anything about the data member `aux_`. If you look near the beginning of `Event.h` you will see that it has the line:

```
#include "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.

which is the header file that declares `EventID`. Look at this file, e.g.,

```
less $ART_INC/art/Persistency/Provenance/EventID.h
```

and find the declaration for `EventID`; you will see that the class `EventID` is within the namespace `art`, making its full name `art::EventID`. Near the top of the file you will also see the comments:

```
// An EventID labels an unique readout of the data
// acquisition system, which we call an ``event``.
```

Look again at `EventID.h`; you will see that it has accessor methods that permit you see the three components of the event identifier:

```
RunNumber_t    run()    const;
SubRunNumber_t subRun() const;
EventNumber_t  event()  const;
```

Earlier in `EventID.h` the C++ *type*⁷ `EventNumber_t` was defined as:

```
namespace art {
    typedef std::uint32_t EventNumber_t;
}
```

meaning that the event number is represented as a 32-bit unsigned integer. A *typedef*(γ) is a different name, or an alias, by which a type can be identified. If you are not familiar with the C++ concept of *typedef*, or if you are not familiar with the definite-length integral types defined by the `<cstdint>` header, consult any standard C++ documentation. If you dig deeper into the layers included in the `art::EventID` header, you will see that the run number and subRun number are also implemented as 32-bit unsigned integers.

The authors of *art* might have chosen an alternate definition of `EventNumber_t`

```
namespace art {
    typedef unsigned EventNumber_t;
}
```

The difference is the use of `unsigned` rather than `std::uint32_t`. This alternate


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.

version was not chosen because it runs the risk that some computers might consider this type to have a length of 32 bits while other computers might consider it to have a length of 16 or 64 bits. In the definition that is used by *art*, an event number is guaranteed to be exactly 32 bits on all computers.

Why did the authors of *art* insert the extra level of indirection and not simply define the following member function inside `art::EventID`?

```
std::unit32_t event() const;
```

The answer is that it makes it easy to change the definition of the type should that be necessary. If, for example, an experiment requires that event numbers be of length 64 bits, only one change is needed, followed by a recompilation.

It is good practice to use typedefs for every concept for which the underlying data type is not absolutely certain. 

It is a very common, but not universal, practice within the HEP C++ community that typedefs that are used to give context-specific names to the C++ built-in types (`int`, `float`, `char`, etc.) end in `_t`.

One last observation about `EventID.h`. Near the top of this file you can find the following fragment, with a few lines omitted for clarity:

```
namespace art {
    std::ostream &
    operator<<(std::ostream & os, EventID const & iID);
}
```

This tells the compiler that the class `art::EventID` has a stream insertion operator (see Section 6.7.10). Because this operator exists, the compiler knows how to use `std::cout` to print an object of type `art::EventID`. You have already used this capability — near the end of Section 10.6.3.6 see the discussion of the line

```
<< event.id()
```

10.6.3.8 DEFINE_ART_MACRO: The Module Maker Macros

The final line in `First_module.cc`,

```
DEFINE_ART_MODULE(tex::First)
```

invokes a C preprocessor macro. This macro is defined in the header file that was pulled in by

```
#include "art/Framework/Core/ModuleMacros.h"
```

If you are not familiar with the C preprocessor, don't worry; you do not need to look under the hood. But if you would like to learn about it, consult any standard C++ reference.

The `DEFINE_ART_MODULE` macro instructs the compiler to put some additional code into the dynamic library made by `buildtool`. This additional code provides the glue that allows *art* to create instances of the class `First` without ever seeing the header or the source for the class; it only gets to see the `.so` or `.dylib` file and nothing else.

The `DEFINE_ART_MODULE` macro adds two pieces of code to the `.so` file. It adds a factory function that, when called, will create an instance of `First` and return a pointer to the base classes `art::EDAnalyzer`. In this way, *art* never sees the derived type of any analyzer module; it sees all analyzer modules via pointer to base. When *art* calls the factory function, it passes as an argument the parameter set specified in the FHiCL file for this module instance. The factory function passes this parameter set through to the constructor of `First`. The second piece of code put into the `.so` file is a static object that will be instantiated at load time; when this object is constructed, it will contact the *art* module registry and register the factory function under the name `First`. When the FHiCL file says to create a module of type `First`, *art* will simply call the registered factory function, passing it the parameter set defined in the FHiCL file. This is the last step in making the connection between the source code of a module and the *art* instantiation of a module.



10.6.3.9 Some Alternate Styles

C++ allows some flexibility in syntax, which can be seen as either powerful or confusing, depending on your level of expertise. Here we introduce you to a few alternate styles that you will need to recognize and may want to use.

Look at the `std::cout` line in the `analyze` method of Listing 10.2:

```
std::cout << "Hello from First::analyze. Event id: "  
          << event.id()  
          << std::endl;  
}
```

This could have been written:

```
art::EventID id = event.id();  
std::cout << "Hello from First::analyze. Event id: "  
          << id  
          << std::endl;
```

This alternate version explicitly creates a temporary object of type `art::EventID`, whereas the original version created an *implicit* temporary object. When you are first learning C++ it is often useful to break down compound ideas by introducing *explicit temporaries*. However, the recommended best practice is to not introduce explicit temporaries unless there is a good reason to do so.

You will certainly encounter the first line of the above written in a different style, too, i.e.,

```
art::EventID id(event.id());
```

Here `id` is initialized using *constructor syntax* rather than using *assignment syntax*. For almost all classes these two syntaxes will produce exactly the same result.

You may also see the argument list of the `analyze` function written a little differently,

```
void analyze( const art::Event& );
```

instead of

```
void analyze( art::Event const& );
```

The position of the `const` has changed. These mean exactly the same thing and the compiler will permit you to use them interchangeably. In most cases, small differences in the placement of the `const` identifier have very different meanings but, in a few cases, both variants mean the same thing. When C++ allows two different syntaxes that mean the same thing, this documentation suite will point it out.



Finally, Listing 10.3 shows the same information as Listing 10.2 but using a style in which the namespace remains open after the class declaration. In this style, the leading `tex::` is no longer needed in the definitions of the constructor and of `analyze`. Both layouts of the code have the same meaning to the compiler. Many experiments use this style in their source code.

Listing 10.3: An alternate layout for `First_module.cc`

```

1
2 #include "art/Framework/Core/EDAnalyzer.h"
3 #include "art/Framework/Core/ModuleMacros.h"
4 #include "art/Framework/Principal/Event.h"
5
6 #include <iostream>
7
8 namespace tex {
9
10   class First : public art::EDAnalyzer {
11
12   public:
13
14     explicit First(fhicl::ParameterSet const& );
15
16     void analyze(art::Event const& event) override;
17
18   };
19
20   First::First(fhicl::ParameterSet const& pset ) : art::EDAnalyzer(pset) {
21     std::cout << "Hello_from_First::constructor." << std::endl;
22   }
23
24   void First::analyze(art::Event const& event){
25     std::cout << "Hello_from_First::analyze._Event_id:_ "
26               << event.id()
27               << std::endl;
28   }
29
30 }
31
32 DEFINE_ART_MODULE(tex::First)

```

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 **cetbuildtools**. 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`, it saved the name of the source directory. The instructions in `CMakeLists.txt` tell **cetbuildtools** to look for more instructions in the subdirectory `ups` and in the file `art-workbook/CMakeLists.txt`, which, in turn, tells it to look for more instructions in the `CMakeLists.txt` files in each subdirectory of `art-workbook`.

When **cetbuildtools** has digested these instructions it knows the rules to build everything that it needs to build.

The object file created by the compilation step is a temporary file and, once it has been inserted into the dynamic library, it is not used any more. Therefore the name of the object file is not important.

On the other hand, the name of the dynamic library file is very important. *art* requires that for every module source file (ending in `_module.cc`) the build system must create exactly one dynamic library file (ending in `_module.so`). It also requires that the name of each `_module.so` file conform to a pattern. Consider the example of the file `First_module.cc`; *art* requires that the dynamic library for this file match the pattern

```
lib*First_module.so
```

where the `*` wildcard matches 0 or more characters.

When naming dynamic libraries, *buildtool* uses the following algorithm, which satisfies the *art* requirements and adds some additional features; the algorithm is illustrated using the example of `First_module.cc`:

1. find the relative path to the source file, starting from the source directory
`art-workbook/FirstModule/First_module.cc`
2. replace all slashes with underscores
`art-workbook_FirstModule_First_module.cc`
3. change the trailing `.cc` to `.so`
`art-workbook_FirstModule_First_module.so`
4. add the prefix `lib`
`libart-workbook_FirstModule_First_module.so`
5. put the file into the directory `lib`, relative to the build directory
`lib/libart-workbook_FirstModule_First_module.so`

You can check that this file is there by issuing the following command from your build directory:

```
ls -l lib/libart-workbook_FirstModule_First_module.so
```

This algorithm guarantees that every module within `art-workbook` will have a unique name for its dynamic library.

The experiments using *art* have a variety of build systems. Some of these follow the minimal *art*-conforming pattern, in which the wildcard is replaced with zero characters. If the Workbook had used such a build system, the name of the dynamic library file would have been

```
lib/libFirst_module.so
```

Both names are legal.

10.7.2 Incremental Builds and Complete Rebuilds

When you edit a file in your source area you will need to rebuild that file in order for those changes to take effect. If any other files in your source area depend on the file that you edited, they too will need to be rebuilt. To do this, reissue the command:

```
buildtool
```

Remember that this command must be executed from your build directory and that, before executing it, you must have setup the environment in your build window. When you run this command, **cetbuildtools** will automatically determine which files need to be rebuilt and will rebuild them; it will not waste time rebuilding files that do not need to be rebuilt. This is called an *incremental build* and it will usually complete much faster than the initial build.

If you want to clean up everything in your build area and rebuild everything from scratch, use the following command:

```
buildtool -c
```

This command will give you five seconds to abort it before it starts removing files; to abort, type ctrl-C in your build window. It will take about the same time to execute as did your initial build of the Workbook. The name of the option -c is a mnemonic for “clean”.

When you do a clean build it will remove all files in your build directory that are not managed by **cetbuildtools**. For example, if you redirected the output of *art* as follows,

```
art -c fcl/FirstModule/first.fcl >& first.log
```



then, when you do a clean build, the file `first.log` will be deleted. This is why the instructions earlier in this chapter told you to redirect output to a log file by

```
art -c fcl/FirstModule/first.fcl >& output/first.log
```

When you ran `buildtool`, it created a directory to hold your output files and you created a symbolic link, named `output`, from your build directory to this new directory. Both the other directory and the symbolic link survive clean builds and your output files will be preserved. The Workbook exercises write all of their root and event-data output files to this directory.

If you edit certain files in the `ups` subdirectory of your source directory, rebuilding requires an extra step. If you edit one of these files, the next time that you run `buildtool`, it will issue an error message saying that you need to re-source `setup_for_development`. If you get this message, make sure that you are in your build directory, and

```
source ../art-workbook/ups/setup_for_development \  
-p $ART_WORKBOOK_QUAL  
  
buildtool
```

10.7.3 Finding Header Files at Compile Time

When `setup_for_development` establishes the working environment for the build directory, it does a UPS setup on the UPS products that it requires; this triggers a chain of additional UPS setups. As each UPS product is set up, that product defines many environment variables, two of which are `PRODUCT-NAME_INC` and `PRODUCT-NAME_LIB`. The first of these points to a directory that is the root of the header file hierarchy for that version of that UPS product. The second of these points to a single directory that holds all of the dynamic library files for that UPS product.

You can spot-check this by doing, for example,

```
ls $TOYEXPERIMENT_INC/*
```

```
ls $TOYEXPERIMENT_LIB
```

```
ls $ART_INC/*
```

```
ls $ART_LIB
```

You will see that the `_INC` directories have a subdirectory tree underneath them while the `_LIB` directories do not.

There are a few small perturbations on this pattern. The most visible is that the ROOT product puts most of its header files into a single directory, `$ROOT_INC`. The Geant4 product does a similar thing.

When the compiler compiles a `.cc` file, it needs to know where to find the files specified by the `#include` directives. The compiler looks for included files by first looking for arguments on the command line, of the form

- *Ipath-to-a-directory*

There may be many such arguments on one command line. If you compiled the code by hand, you would add the `-I` options yourself. In the Workbook, one of the jobs of `buildtool` is to figure out which `-I` options are needed when. The compiler assembles the set of all `-I` arguments and uses it as an include path; that is, it looks for the header files by trying the first directory in the path and if it does not find it there, it tries the second directory in the path, and so on. The choice of `-I` for the name of the argument is a mnemonic for Include.

When `buildtool` compiles a `.cc` file it adds many `-I` options to the command line; it adds one for each UPS product that was set up when you sourced `setup_for_development`. When building `First_module.cc`, `buildtool` added `-$ART_INC`, `-$TOYEXPERIMENT_INC` and many more.

A corollary of this discussion is that when you wish to include a header file from a UPS product, the `#include` directive must contain the relative path to the desired file, starting from the `_INC` environment variable for that UPS product.



This system illustrates how the Workbook can work the same way on many different computers at many different sites. As the author of some code, you only need to know paths of include files relative to the relevant `_INC` environment variable. This environment variable may have different values from one computer to another but the setup and build systems will ensure that the site-specific information is communicated to the compiler using environment variables and the `-I` option.

This system has the potential weakness that if two products each have a header file with exactly the same relative path name, the compiler will get confused. Should this happen,

the compiler will always choose the file from the earlier of the two `-I` arguments on the command line, even when the author of the code intended the second choice to be used. To mitigate this problem, the *art* and UPS teams have adopted the convention that, whenever possible, the first element of the relative path in an `#include` directive will be the UPS package name. It is the implementation of this convention that led to the repeated directory name `art-workbook/art-workbook` that you saw in your source directory. There are a handful of UPS products for which this pattern is not followed and they will be pointed out as they are encountered.

The convention of having the UPS product name in the relative path of `#include` directives also tells readers of the code where to look for the included file.

10.7.4 Finding Dynamic Library Files at Link Time

The module `First_module.cc` needs to call methods of the class `art::Event`. Therefore the compiler left a notation in the object file saying “to use this object file you need to tell it where to find `art::Event`.” The technical way to say this is that the object file contains a list of *undefined symbols* or *undefined external references*. When the linker makes the dynamic library

```
libart-workbook_FirstModule_First_module.so
```

it must resolve all of the undefined symbols from all of the object files that go into the library. To resolve a symbol, the linker must learn what dynamic library defines that symbol. When it discovers the answer, it will write the name of that dynamic library into something called the *dependency list* that is kept inside the dynamic library. **cetbuildtools** tells the linker that the dependency list should contain only the filename of each dynamic library, not the full path to it. If, after the linker has finished, there remain unresolved symbols, then the linker will issue an error message and the build will fail.

If library A depends on library B and library B depends on library C, but library A does not directly depend on library C, then the dependency list of library A should contain only library B. In other words, the dependency should contain only *direct dependencies* (also called *first order dependencies*).

To learn where to look for symbol definitions, the linker looks at its command line to find something called the *link list*. The link list can be specified in several different ways

and the way that **cetbuildtools** uses is simply to write the link list as the absolute path to every `.so` file that the linker needs to know about. The link list can be different for every dynamic library that the build system builds. However it is very frequently true that if a directory contains several modules, then all of the modules will require the same link list. The bottom line is that the author of a module needs to know the link list that is needed to build the dynamic library for that module.



For these Workbook exercises, the author of each exercise has determined the link list for each dynamic library that will be built for that exercise. In the **cetbuildtools** system, the link list for `First_module.cc` is located in the `CMakeLists.txt` file from same directory as `First_module.cc`; the contents of this file are shown in Listing 10.4. This `CMakeLists.txt` file says that all modules found in this directory should be built with the same link list and it gives the link list; the link list is the seven lines that begin with a dollar sign; these lines each contain one `cmake` variable. Recall that **cetbuildtools** is a build system that lives on top of `cmake`, which is another build system. A `cmake` variable is much like an environment variable except that is only defined within the environment of the running build system; you cannot look at it with `printenv`.

The five `cmake` variables beginning with `ART_` were defined when `buildtool` set up the UPS *art* product. Each of these variables defines an absolute path to a dynamic library in `$ART_LIB`. For example `${ART_FRAMEWORK_CORE}` resolves to

```
$ART_LIB/libart_Framework_Core.so
```

Almost all *art* modules will depend on these five libraries. Similarly the other two variables resolve to dynamic libraries in the **fhiclepp** and **cetlib** UPS products.

When **cetbuildtools** constructs the command line to run the linker, it copies the link list from the `CMakeLists.txt` file to the command linker line.

The experiments that use *art* use a variety of build systems. Some of these build systems do not require that all external symbols be resolved at link time; they allow some external symbols to be resolved at run-time. This is legal but it can lead to certain difficulties. *A future version of this documentation suite will contain a chapter in the Users Guide that discusses linkage loops and how use of closed links can prevent them. This section will then just reference it.*

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 )
```

Consult the `cmake` and **cetbuildtools** documentation to understand the remaining details of this file.

10.7.5 Build System Details



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.

If you want to see what buildtool is actually doing, you can enable verbose mode by issuing the command:

```
buildtool VERBOSE=TRUE
```

For example, if you really want to know the name of the object file, you can find it in the verbose output. For this exercise, the object file is

```
./art-workbook/FirstModule/CMakeFiles/  
art-workbook_FirstModule_First_module.dir/First_module.cc.o
```

where the above is really just one line.

Also, you can read the verbose listing to discover the flags given to the compiler and linker. The more instructive compiler and linker flags valid at time of writing are given in Table 10.1. The C++ 11 features are selected by the presence of the `-std=c++11` flag and a high level of error checking is specified. The linker flag,

```
-Wl,--no-undefined
```

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

tells the linker that it must resolve all external references at link time. This is sometime referred to as a *closed link*.

10.8 Suggested Activities

This section contains some suggested exercises in which you will make your own modules and learn more about how to use the class `art::EventID`.

10.8.1 Create Your Second Module

In this exercise you will create a new module by copying `First_module.cc` and making the necessary changes; you will build it using `buildtool`; you will copy `first.fcl` and make the necessary changes; and you will run the new module using the new `FHiCL` file.

Go to your source window and `cd` to your source directory. If you have logged out, out remember to re-establish your working environment; see Section 11 Type the following commands:

```
cd art-workbook/FirstModule
```

```
cp First_module.cc Second_module.cc
```

```
cp first.fcl second.fcl
```

Edit the files `Second_module.cc` and `second.fcl`. In both files, change every occurrence of the string “First” to “Second”; there are eight places in the source file and two

in the FHiCL file, one of which is in a comment.

The new module needs the same link list as did `First_module.cc` so there is no need to edit `CMakeLists.txt`; the instructions in `CMakeLists.txt` tell buildtool to build all modules that it finds in this directory and to use the same link list for all modules.

Go to your build window and `cd` to your build directory. Again, remember to re-establish your working environment as necessary. Rebuild the Workbook code:

```
buildtool
```

This should complete with the message:

```
-----
INFO: Stage build successful.
-----
```

If you get an error message, consult a local expert or the *art* team as described in Section 3.4.

When you run buildtool it will perform an incremental build (see Section 10.7.2) during which it will detect `Second_module.cc` and build it.

You can verify that buildtool created the expected dynamic library:

```
ls lib/*Second*.so
```

```
lib/libart-workbook_FirstModule_Second_module.so
```

Stay in your build directory and run the new module:

```
art -c fcl/FirstModule/second.fcl >& output/second.log
```

Compare `output/second.log` with `output/first.log`. You should see that “First” has been replaced by “Second” everywhere and the date/time lines are different.

10.8.2 Use *artmod* to Create Your Third Module

This exercise is much like the previous one; the difference is that you will use a tool named `artmod` to create the source file for the module.

Go to your source window and `cd` to your source directory. If you have logged out, remember to re-establish your working environment; see Section 11

The command `artmod` creates a file containing the skeleton of a module. It is supplied by the UPS product **cetpkgssupport**, which was set up when you performed the last step of establishing the environment in the source window, sourcing `setup_deps`. You can verify that the command is in your path by using the `bash` built-in command `type` (output shown on two lines):

```
type artmod
```

```
artmod is hashed (/ds50/app/products/cetpkgssupport/  
v1_02_00/bin/artmod)
```

The leading elements of the directory name will reflect your UPS products area, and may be different from what is shown here. The version number, `v1_02_00`, may also change with time.

From your source directory, type the following commands:

```
cd art-workbook/FirstModule
```

```
artmod analyzer tex::Third
```

```
cp first.fcl third.fcl
```

The second command tells `artmod` to create a source file named `Third_module.cc` that contains the skeleton for an ‘analyzer’ module, to be named `Third` in the namespace `tex`.

If you compare `Third_module.cc` to `First_module.cc` you will see a few differences:

1. `Third_module.cc` is longer and has more comments
2. the layout of the class is a little different but the two layouts are equivalent
3. there are some extra `#include` directives
4. the include for `<iostream>` is missing

5. in the `analyze` member function, the name of the argument is different (`event` vs `e`)
6. `artmod` supplies the skeleton of a destructor (`~Third`)

The `#include` directives provided by `artmod` are a best guess, made by the author of `artmod`, about which ones will be needed in a “typical” module. Other than slowing down the compiler by an amount you won’t notice, the extra `#include` directives do no harm; keep them or leave them as you see fit.

Edit `Third_module.cc`

1. add the `#include` directive for `<iostream>`
2. copy the bodies of the constructor and the `analyze` member function from `First_module.cc`; change the string “First” to “Third”
3. in the definition of the member function `analyze`, change the name of the argument to `event`.

When you built `First_module.cc`, the compiler wrote a destructor for you that is identical to the destructor written by `artmod`; so you can leave the destructor as `artmod` wrote it, i.e., with an empty body. Or you can delete it; if you decide to do so, you must delete both the declaration and the implementation.

Edit `third.fcl` Change every occurrence of the string “First” to “Third”; there are two places, one of which is in a comment.

Go to your build window and `cd` to your build directory. If you have logged, out remember to re-establish your working environment; see Section 11. Rebuild the Workbook code:

`buildtool`

Refer to the previous section to learn how to identify a successful build and how to verify that the expected library was created.

Stay in your build directory and run the third module:

```
art -c fcl/FirstModule/third.fcl >& output/third.log
```

Compare `output/third.log` with `output/first.log`. You should see that the `printout` from `First_module.cc` has been replaced by that from `Third_module.cc`.

`artmod` has many options that you can explore by typing:

```
artmod --help
```

10.8.3 Running Many Modules at Once

In this exercise you will run four modules at once, the three made in this exercise plus the `HelloWorld` module from Chapter 9.

Go to your source window and `cd` to your source directory. Type the following commands:

```
cd art-workbook/FirstModule
```

```
cp first.fcl all.fcl
```

Edit the file `all.fcl` and replace the `physics` parameter set with the contents of Listing 10.5. This parameter set:

1. defines four module labels and
2. puts all four module labels into the `end_paths` sequence.

When you run `art` on this FHiCL file, `art` will first look at the definition of `end_paths` and learn that you want it to run four module labels. Then it will look in the `analyzers` parameter set to find the definition of each module label; in each definition `art` will find the class name of the module that it should run. Given the class name and the environment variable `LD_LIBRARY_PATH`, `art` can find the right dynamic library to load. If you need a refresher on module labels and `end_paths`, refer to Sections 9.8.7 and 9.8.8.

Go to your build window and `cd` to your build directory. If you have logged out, remember to re-establish your working environment; see Section 11. You do not need to build any code for this exercise.

Run the exercise:

```
art -c fcl/FirstModule/all.fcl >& output/all.log
```

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 }

```

Compare `output/all.log` with the log files from the previous exercises. The new log file should contain printout from each of the four modules. Once, near the start of the file, you should see the printout from the three constructors; remember that the `HelloWorld` module does not make any printout in its constructor. For each event you should see the printout from the four `analyze` member functions.

Remember that *art* is free to run analyzer modules in any order; this was discussed in Section 9.8.8.

10.8.4 Access Parts of the EventID

In this exercise, you will access the individual parts of the event identifier.

Before proceeding with this section, review the material in Section 10.6.3.7 which discusses the class `art::EventID`. The header file for this class is:

```
$ART_INC/art/Persistency/Provenance/EventID.h
```

In this exercise, you are asked to rewrite the file `Second_module.cc` so that the print-

out made by the `analyze` method looks like the following (lines split here due to space restrictions):

```
Hello from FirstAnswer01::analyze.  run number: 1
    sub run number: 0 event number: 1
Hello from FirstAnswer01::analyze.  run number: 1
    sub run number: 0 event number: 2
```

and so on for each event.

To do this, you will need to reformat the text in the `std::cout` statement and you will need to separately extract the `run`, `subRun` and event numbers from the `art::EventID` object.

You will do the editing in your source window, in the subdirectory `art-workbook/FirstModule`.

When you think that you have successfully rewritten the module, you can test it by going to your build window and `cd`'ing to your build directory. Then:

```
buildtool
```

```
art -c fcl/FirstModule/second.fcl >& output/eventid.log
```

If you have not figured out how to do this exercise after about 15 minutes, you can find one possible answer in the file `FirstAnswer01_module.cc`, in the same directory as `First_module.cc`.

To run the answer module and verify that it makes the requested output, run:

```
art -c fcl/FirstModule/firstAnswer01.fcl >& output/firstAnswer01.log
```

(The command can be typed on a single line.) You did not need to build this module because it was already built the first time that you ran `buildtool`; that run of `buildtool` built all of the modules in the Workbook.

There is a second correct answer to this exercise. If you look at the header file for `art::Event`, you will see that this class also has member functions

```
EventNumber_t  event()  const {return aux_.event();}
SubRunNumber_t subRun() const {return aux_.subRun();}
RunNumber_t    run()    const {return id().run();}
```

So you could have called these directly,

```
std::cout << "Hello from FirstAnswer01::analyze. "
          << " run number: "      << event.run()
          << " sub run number: "  << event.subRun()
          << " event number: "    << event.event()
          << std::endl;
```

instead of

```
std::cout << "Hello from FirstAnswer01::analyze. "
          << " run number: "      << event.id().run()
          << " sub run number: "  << event.id().subRun()
          << " event number: "    << event.id().event()
          << std::endl;
```

But the point of this exercise was to learn a little about how to dig down into nested header files to find the information you need.

10.9 Final Remarks

10.9.1 Why is there no `First_module.h` File?

When you performed the exercises in this chapter, you saw, for example, the file `First_module.cc` but there was no corresponding `First_module.h` file. This section will explain why.

In a typical C++ programming environment there is a header file (`.h`) for each source file (`.cc`). As an example, consider the files `Point.h` and `Point.cc` that you saw in Section 6.7.10.

The reason for having `Point.h` is that the implementation of the class, `Point.cc`, and the users of the class need to agree on what the class `Point` is. In the Section 6.7.10 example, the only user of the class is the main program, `ptest.cc`. The file `Point.h` serves as the unique, authoritative declaration of what the class is; both `Point.cc` and `ptest.cc` rely on on this declaration.

If you think carefully, you are already aware of a very common exception to the pattern of one `.h` file for each `.cc` file: there is never a header file for a main program. For example,


in the examples that exercised the class `Point`, `pctest.cc` had no header file. Why not? No other piece of user-written code needs to know about any classes or functions declared or defined inside `pctest.cc`.

The `First_module.h` file is omitted simply because every entity that needs to see the declaration of the class `First` is already inside the file `First_module.cc`. There is no reason to have a separate header file. Recall the “dangerous bend” paragraph at the end of Section 10.6.3.8 that described how *art* is able to use modules without needing to know about the declaration of the module class.

art is designed such that only *art* may construct instances of module classes and only *art* may call member functions of module classes. In particular, modules may not construct other modules and may not call member functions of other modules. The absence of a `First_module.h`, provides a physical barrier that enforces this design.

10.9.2 The Three-File Module Style

In this chapter, the source for the module `First` was written in a single file. You may also write it using three files, `First.h`, `First.cc` and `First_module.cc`.

Some experiments use this three-file style. The authors of *art* do not recommend it, however, because it exposes the declaration of `First` in a way that permits it to be misused (as was discussed in Section 10.9.1). The build system distributed with the Workbook has not been configured to build modules written in this style. 

In this style, `First.h` contains the class declaration plus any necessary `#include` directives; it now also requires code guards; this is shown in Listing 10.6.

Listing 10.6: The contents of `First.h` in the three-file model

```
1
2 #ifndef art-workbook_FirstModule_First_h
3 #define art-workbook_FirstModule_First_h
4
5 #include "art/Framework/Core/EDAnalyzer.h"
6 #include "art/Framework/Principal/Event.h"
7
8 namespace tex {
9
10     class First : public art::EDAnalyzer {
11
```

```

12     public:
13
14         explicit First(fhicl::ParameterSet const& );
15
16         void analyze(art::Event const& event) override;
17
18     };
19
20 }
21 #endif

```

The file `First.cc` contains the definitions of the constructor and the `analyze` member function, plus the necessary `#include` directives; this is shown in Listing 10.7.

Listing 10.7: The contents of `First.cc` in the three-file model

```

1
2 #include "art-workbook/FirstModule/First.h"
3
4 #include <iostream>
5
6 tex::First::First(fhicl::ParameterSet const& pset) : art::EDAnalyzer(pset) {
7     std::cout << "Hello_from_First::constructor." << std::endl;
8 }
9
10 void tex::First::analyze(art::Event const& event){
11     std::cout << "Hello_from_First::analyze._Event_id:_ "
12               << event.id()
13               << std::endl;
14 }

```

And `First_module.cc` is now stripped down to the invocation of the `DEFINE_ART_MODULE` macro plus the necessary `#include` directives; this is shown in Listing 10.8.

Listing 10.8: The contents of `First_module.cc` in the three-file model

```

1
2 #include "art-workbook/FirstModule/First.h"
3 #include "art/Framework/Core/ModuleMacros.h"
4
5 DEFINE_ART_MODULE(tex::First)

```


10.10 Flow of Execution from Source to FHiCL File

The properties that a class must have in order to be an analyzer module are summarized in Section 10.6.3.2 for reference. This section reviews how the source code found in an analyzer module, e.g., `First_module.cc`, is executed by *art*:

1. The script `setup_for_development` defines many environment variables that are used by *buildtool*, *art* and *toyExperiment*.
2. `LD_LIBRARY_PATH`, an important environment variable, contains the directory `lib` in your build area plus the `lib` directories from many UPS products, including *art*.
3. *buildtool* compiles `First_module.cc` to a temporary object file.
4. *buildtool* links the temporary object file to create a dynamic library in the `lib` sub-directory of your build area:
`lib/libart-workbook_FirstModule_First_module.so`
5. When you run *art* using file `first.fcl`, this file tells *art* to find and load a module with the “`module_type`” `First`.
6. In response to this request, *art* will search the directories in `LD_LIBRARY_PATH` to find a dynamic library file whose name matches the pattern:
`lib*First_module.so`
7. If *art* finds either zero or more than one match to this pattern, it will issue an error message and stop.
8. If *art* finds exactly one match to this pattern, it will load the dynamic library.
9. After *art* has loaded the dynamic library, it has access to a function that can, on demand, create instances of the class `First`.

The last bullet really means that the dynamic library contains a factory function that can construct instances of `First` and return a pointer to the base class, `art::EDANalyzer`. The dynamic library also contains a static object that, at load-time, will contact the *art* module registry and register the factory function under the `module_type` `First`.



10.11 Review

- how to establish the *art* development environment
- how to checkout the Workbook exercises from the `git` source code management system
- how to use the **cetbuildtools** build system to build the code for the Workbook exercises
- how include files are found
- what a *link list* is
- where the build system finds the link list
- what the `art::Event` is and how to access it
- what the `art::EventID` is and how to access it
- what makes a class an *art module*
- where the build system puts the `.so` files that it makes

10.12 Test Your Understanding

10.12.1 Tests

Two modules are provided that intentionally contain bugs and fail to build:

```
FirstModule/FirstBug01_module.cc.nobuild  
FirstModule/FirstBug02_module.cc.nobuild
```

Your job in each case is to figure out what's wrong with the module and fix it. The build system will ignore these files until the `.nobuild` is removed. It's a good idea to work on these two files one at a time. In the procedure below, start with the one labeled 01. Repeat the procedure for the second file, changing 01 to 02 as needed. Answers are provided at the end of the chapter, in Section 10.12.2.



Pay attention to which window you do each operation in; there's a fair bit of back-and-forth between the source and build windows!

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 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
```

Repeat the procedure for the second file. Because this is the first time you are doing this, we will give you a hint: you only need to fix one line in the first file but you need to fix two in the second.

The answers are intentionally placed on a new page (remember to try before you read 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 file is no longer present. However, the compiler does detect that the base class has a virtual

function with the same name but a different argument list; this is what the error text means by “hidden”.

While there are legitimate reasons to hide a virtual function, most cases are, in fact, errors. Therefore the build system for the Workbook is configured as follows:

1. When the compiler detects a hidden virtual function it will issue a warning (not an error).
2. The compiler will promote all warnings to errors, which will stop the build. (hence the `[-Werror=overloaded-virtual]` message at the end).

This forces you to think about each case and decide if it is really what you intended. To learn how to allow hiding of virtual functions consult the `cetbuildtools` documentation.

11 General Setup for Login Sessions

After you've done the initial setup described in Section 10.4, there are some steps that don't need to be repeated for subsequent login sessions. To begin a general login session for Exercise 2 or any subsequent exercise, you need to follow the instructions in this chapter.

If during your initial setup you chose to manage your own directory names, then the names of your source and build directories will be different than those shown here.



11.1 Source Window

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
```

The contents of the source directory is discussed in Section 10.4.2.2.

11.2 Build Window

In your build window:

1. Log in to the computer you chose in Section 8.3.
2. Follow the site-specific setup procedure; see Chapter 5
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
```

The build window setup is discussed in Section 10.4.4 and the *art* development environment is described in Section 10.5.

12 Keeping Up to Date with Workbook Code and Documentation

12.1 Introduction

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.

12.3.1 Get Updated Documentation

First, check which documentation release you’re currently using: it’s noted on the title page of this document¹. Then go to the *art* Documentation website and compare your documentation release number to the latest available.

Download a new copy of the documentation, as needed.

12.3.2 Get Updated Code and Build It

Also noted on the title page of the documentation is the release² of the art-workbook code that the documentation is intended for. Recall from Figure 10.1 that git commands are used to clone the code in the remote repository into your local copy, then copy the requested release from that local copy into your working area. The git system is described in more detail in Chapter 21.

Chances are that you’re using the code release that goes with the documentation you have been using. You can check by looking in the file `art-workbook/ups/product_deps`. From your source directory run:

```
grep art_workbook ups/product_deps
```

```
parent art_workbook v0_00_13
```

This shows version `v0_00_13` as an example. If your version is earlier than the one listed on the cover of the latest documentation, you will need to get new code and build it.

These instructions illustrate updating the working version of the art-workbook code from version `v0_00_13` to version `v0_00_15`. There is nothing special about these two versions; the instructions serve as a model for a change between any pair of versions.

1. Start from (or `cd` to) your source directory (see Section 10.4.1):

```
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.

2. Use `git status` and make a note of the files that you have modified and/or added (see Section 21.1.3 for instructions).

```
git status [-s]
```

3. Switch from your tagged version branch back to the `develop` branch (“branches” are discussed in Chapter 21, you don’t need to understand them at this stage³).

```
git checkout develop
```

```
Switched to branch 'develop'
```

4. Update your local copy of the repository (the `.git` directory)

```
git pull
```

The output from this command is shown in Listing 12.1.

5. Switch your working code to the new branch:

```
git checkout -b v0_00_15 v0_00_15
```

```
Switched to a new branch 'v0_00_15'
```

Use the new version number twice in this command. In the messages produced in this step, watch for the names of files that you have modified. Check for conflicts that `git` did not merge correctly.

To rebuild your updated working code:

1. In your build window, `cd` to your build directory

```
cd $ART_WORKBOOK_WORKING_BASE/username/workbook/build-prof
```

2. Tell **cetbuildtools** to look for, and act on, any changes in your checked out version of the code (command shown on two lines):

```
source ../art-workbook/ups/setup_for_development \-p $ART_WORKBOOK_QUAL
```

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.

Listing 12.1: Example of the output produced by `git pull`

```

1 From http://cdcvcs.fnal.gov/projects/art-workbook
2   e79d9ef..81d2a76  develop    -> origin/develop
3   6435ecc..c0c1af5  master     -> origin/master
4 From http://cdcvcs.fnal.gov/projects/art-workbook
5 * [new tag]         v0_00_14   -> v0_00_14
6 * [new tag]         v0_00_15   -> v0_00_15
7 Updating e79d9ef..81d2a76
8 Fast-forward
9  art-workbook/ModuleInstances/magic.fcl      | 26 ++++++-----
10 art-workbook/ParameterSets/PSet01_module.cc | 36 ++++++-----
11 art-workbook/ParameterSets/PSet02_module.cc | 53 ++++++-----
12 art-workbook/ParameterSets/PSet03_module.cc | 28 ++++++-----
13 art-workbook/ParameterSets/PSet04_module.cc | 44 ++++++-----
14 art-workbook/ParameterSets/pset01.fcl       |  6 +---
15 art-workbook/ParameterSets/pset02.fcl       | 14 +++++---
16 art-workbook/ParameterSets/pset03.fcl       |  6 +---
17 art-workbook/ParameterSets/pset04.fcl       |  7 +++---
18 ups/product_deps                           |  2 +-
19 10 files changed, 109 insertions(+), 113 deletions(-)

```

3. Rebuild:

`buildtool`

If this step does not complete successfully, the first thing to try is a clean rebuild:

`buildtool -c`

12.3.3 See which Files you have Modified or Added

At any time you can check to see which files you have modified and which you have added. The code is structured in such a way that when you checkout a new version, these files will remain in your working directory and will not be modified or deleted. The `git checkout` command will generate some informational messages about them, but you do not need to take any action.

To see the new/modified files, `cd` to your source directory and issue the `git status` command. Suppose that you have checked out version `v0_00_13`, modified `first.fcl` and added `second.fcl`. The `git status` command will produce the following output:

git status

```
# On branch v0_00_13
# Changes not staged for commit:
#   (use "git add <file>..." to update what will be committed)
#   (use "git checkout -- <file>..." to discard changes in
working directory)
#
# modified:   first.fcl
#
# Untracked files:
#   (use "git add <file>..." to include in what will be committed)
#
# second.fcl
no changes added to commit (use "git add" and/or "git commit -a")
```



Do not issue the `git add` or `git commit` commands that are suggested in the command output above.

In the rare case that you have neither modified nor added any files, the output of `git status` will look like:

git status

```
# 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`.

13.2 Prerequisites

The prerequisites for this chapter include all of the material in Part I (Introduction) and all of the material up to this point in Part II (Workbook).

In particular, make sure that you understand the *event loop* (see Section 3.6.2).

13.3 What You Will Learn

This chapter will show you *how* to provide the optional member functions in your *art* modules to execute special functionality at the beginning and end of jobs, runs and/or subRuns. These include

1. `beginJob()`
2. `beginRun(art::Run const&)`
3. `beginSubRun(art::SubRun const&)`
4. `endJob()`
5. `endRun(art::Run const&)`
6. `endSubRun(art::SubRun const&)`

As you gain experience, you will gain proficiency at knowing *when* to provide them.

You will also be introduced to the classes

1. `art::RunID`
2. `art::Run`
3. `art::SubRunID`
4. `art::SubRun`

that are analogous to the `art::EventID` and `art::Event` classes that you have already 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
`void beginJob () override;`

```
void beginRun ( art::Run const& run ) override;
void beginSubRun( art::SubRun const& subRun ) override;
```

4. the text printed by the constructor and analyze member functions has changed
5. the file contains the definitions of the three new member functions, each of which simply makes some identifying printout

13.5.1 About the `begin*` Member Functions

The optional member functions `beginJob`, `beginRun` and `beginSubRun`, described in the Introduction to this chapter (Section 13.1), must have exactly the argument list prescribed by *art* as shown in list item 3 above.

art knows to call the `beginJob` member function of each module, if present, once at the start of the job; it knows to call `beginRun`, if present, at the start of each run and, likewise, `beginSubRun` at the start of each subRun.

13.5.2 About the `art::*ID` Classes

In Section 10.6.3.7 you learned about the class `art::EventID`, which describes the three-part event identifier. *art* also provides two related classes:

- `art::RunID`, a one-part identifier for a run number
- `art::SubRunID`, a two-part identifier for a subRun

The header files for these classes are found at:

```
$ART_INC/art/Persistency/Provenance/RunID.h
$ART_INC/art/Persistency/Provenance/SubRunID.h
```

Similar to the `art::Event` class discussed in Section 10.6.3.6, *art* provides `art::Run` and `art::subRun`. These contain the IDs, e.g., `art::RunID`, plus the data products for the entire run or subRun, respectively. You can find their header files at:

```
$ART_INC/art/Framework/Principal/Run.h
$ART_INC/art/Framework/Principal/SubRun.h
```

In the call to `beginSubRun` the argument is of type `art::SubRun const&`. A simplified description of this object is that it contains an `art::SubRunID` plus a collection of data products that describe the subRun. All of the comments about the class `art::Run` in the preceding few paragraphs apply to `art::SubRun`. You can find the header file for `art::SubRun` at:

```
less $ART_INC/art/Framework/Principal/SubRun.h
```

13.5.3 Use of the `override` Identifier

The `override` identifier on each of these member functions instructs the compiler to check that both the name (and spelling) of the member function and its argument list are correct; if not, the compiler will issue an error message and stop. This is a very handy feature. Without it, a misspelled function name or incorrect argument list would cause the compiler to assume that you intended to define a *new* member function unrelated to one of these optional *art*-defined member functions. This would result in a difficult-to-diagnose run-time error: *art* would simply not recognize your member function and would never call it.

Always provide the `override` identifier when using any of the optional *art*-defined member functions.



For those with some C++ background, the three member functions `beginJob`, `beginRun` and `beginSubRun` are declared as `virtual` in the base class, `art::EDAnalyzer`. The `override` identifier is new in C++-11 and will not be described in older text books. It instructs the compiler that this member function is intended to override a virtual function from the base class; if the compiler cannot find such a function in the base class, it will issue an error.



13.5.4 Use of `const` References

In `Optional_module.cc` the argument to the `beginRun` member function is a `const` reference to an object of type `art::Run` that holds the current run ID and the collection of data products that together describe the run. If you take a snapshot of a running *art* job you will see that, at any time, there is exactly one object of type `art::Run`. This object is

owned by *art*. *art* gives modules access to it when it (*art*) calls the modules' `beginRun` and `endRun` member functions.

Because the object is passed by reference, the `beginRun` member function does not get a copy of the object; instead it is given access to it. Because it is passed by `const` reference in this example, your analyzer module may look at information in the object but it may not add or change information to the `art::Run` object.



There is a very important habit that you need to develop as a user of *art*. Many member functions in *art*, in the Workbook code and very likely in your experiment's code, will return information by `&` or by `const&`. If you receive these by value, not by reference, then you will make copies that waste both CPU and memory; in some cases these can be significant wastes. Unfortunately there is no way to tell the compiler to catch this mistake. The only solution is your own vigilance.

To access the `art::Run` and `art::SubRun` objects through, for example, an `art::Event` named `event`, you can use

```
art::SubRun const& subRun = event.getSubRun();
```

for the `subRun` and

```
art::Run const& run = subRun.getRun();
```

for the `run`.

13.5.5 The `analyze` Member Function

In your `analyze` member function, if you have an `art::Event`, named `event`, you can access the associated run information by:

```
art::Run const& run = event.getRun();
```

You may sometimes see this written as:

```
auto const& run = event.getRun();
```



Both versions mean exactly the same thing. When a type is long and awkward to write, the `auto` identifier is very useful; however it is likely to be very confusing to beginners. When you encounter it, check the header files for the classes on the right hand side of the

assignment; from there you can learn the return type of the member function that returned the information.

13.6 Running this Exercise

Look at the file `optional.fcl`. This FHiCL file runs the module `Optional` on the input file `inputFiles/input03_data.root`. Consult Table 9.1 and you will see that this file contains 15 events, all from run 3. It contains events 1 through 5 from each of subRuns 0, 1 and 2. With this knowledge, and the knowledge of the source file `Optional_module.cc`, you should have a clear idea of what this module will print out.

In your build directory, run the following command

```
art -c fcl/OptionalMethods/optional.fcl >& output/optional.log
```

The part of the printed output that comes from the module `Optional` is given in Listing 13.1. Is this what you expected to see? If not, understand why this module made the printout that it did. If you did not get this printout, double check that you followed the 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`

```

1
2 Hello from Optional::constructor.
3 Hello from Optional::beginJob.
4 Hello from Optional::beginRun: run: 3
5 Hello from Optional::beginSubRun: run: 3 subRun: 0
6 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 1
7 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 2
8 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 3
9 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 4
10 Hello from Optional::analyze. Event id: run: 3 subRun: 0 event: 5
11 Hello from Optional::beginSubRun: run: 3 subRun: 1
12 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 1
13 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 2
14 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 3
15 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 4
16 Hello from Optional::analyze. Event id: run: 3 subRun: 1 event: 5
17 Hello from Optional::beginSubRun: run: 3 subRun: 2
18 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 1
19 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 2

```

```

20 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 3
21 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 4
22 Hello from Optional::analyze. Event id: run: 3 subRun: 2 event: 5

```

13.7 The Member Function `beginJob` versus the Constructor

The member function `beginJob` gets called once at the start of the job. The constructor of the each module is also called once at the start of the job. This brings up the question: What code belongs in the constructor and what code belongs in the `beginJob` member function?

A small number of things must be done in the constructor — see below. Other tasks can be done in either place but most experiments have found it useful to follow the rough guideline that you should put initializers and code related to *art* bookkeeping in the constructor and that you should put physics-related code in `beginJob`. Hopefully the meaning of this advice will become clear as you work through the Workbook. Your experiment may have additional, more specific, guidelines.

The correct place to initialize data members is in the constructor and, whenever possible, you should use the initializer list syntax. Never defer initialization of a data member to the `beginJob` member function or later. When you encounter producer modules, you will learn about some more tasks that must be performed in the constructor. *This chapter has not yet been written.*



For those of you familiar with ROOT, we can provide an example of something physics-related. You should create histograms, ntuples and trees in one of the `begin` member functions, not in the constructor. In many cases you can create them in `beginJob` but there are cases in which you will need to defer creation until `beginRun` or `beginSubRun`. 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 indentify 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_data.root  
inputFiles/input02_data.root  
inputFiles/input03_data.root
```

If you need a reminder about how to tell *art* to run on three input files in one job, consult Section 9.8.5.

Make sure that the printout from this job matches the description of the event loop found in Section 3.6.2.

13.8.3 The Option `--trace`

The *art* command supports a command line option named `--trace`. This creates additional printout that identifies every step in the event loop. Use this option to trace what *art* is doing when you run this exercise. For example

```
art -c fcl/OptionalMethods/optional.fcl --trace >& output/trace.log
```

You should be able to identify your printout among the printout from *art* and see that your printout appears in the expected place.

When you are getting an error from *art* and you don't understand which module is causing the problem, you can use `--trace` to narrow your search.

13.9 Review

1. `beginJob()`
2. `beginRun(art::Run const&)`
3. `beginSubRun(art::SubRun const&)`
4. `endJob()`
5. `endRun(art::Run const&)`
6. `endSubRun(art::SubRun const&)`

As you gain experience, you will gain proficiency at knowing *when* to provide them.

You will also be introduced to the classes

1. `art::RunID`
2. `art::Run`

3. `art::SubRunID`

4. `art::SubRun`

13.10 Test Your Understanding

13.10.1 Tests

Two files are provided, one of which intentionally contain bugs. The module should compile and run, but some printout will be missing. The files are:

```
OptionalMethods/OptionalBug01_module.cc  
OptionalMethods/bug01.fcl
```

Your job is to figure out what's missing and why, then fix it. The answer is provided at the end.

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.

The answer is intentionally placed on a new page (remember to try before you read further!).

13.10.2 Answers

The code compiles and runs but the printout from the `beginRun` member function is missing. Why? Because `beginRun` was misspelled as `beginrun` (lower case R) in both the declaration and the definition. They matched, but they weren't recognized by the `std::cout` statement that goes with the definition.

The warning that we tripped over in the Exercise 2 test only comes up if the name of the 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 `fhiCL::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       : {  
                        a : 1  
                        b : 2  
                    }  
}
```

This functionality allows you to write modules whose behaviour is run-time configurable. For example, if you have a reconstruction algorithm that depends on some cuts, the values of those cuts can be provided in this way.

14.2 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 10. You can read this chapter without necessarily having read Chapter 12 or 13.

14.3 What You Will Learn

In Section 10.6.3.4 you saw that the constructor of a module is required to take an argument of type `fhiCL::ParameterSet const&`.

In this chapter you will learn how to use this argument to read additional parameters in a parameter set. In particular, you will learn about the class `fhiCL::ParameterSet` and after working through the exercises in this section, you should know how to:

1. read parameter values from a FHiCL file into a module

2. require that a particular parameter be present in a parameter set
3. use data members to communicate information from the constructor to other member functions of a module
4. print a parameter set
5. use the colon initializer syntax
6. provide a default value for a parameter (if the parameter is absent from a parameter set)
7. modify the precision of the printout of floating point types
8. recognize the error messages for a missing parameter or for a value that cannot be converted to the requested type

You will also learn:

1. that you should find out your experiment’s policy about what sorts of parameters are allowed to have default values
2. an extra parameter automatically added by *art*, but only in parameter sets that are used to configure modules
3. the canonical forms of parameters

Finally, you will learn a small amount about C++ templates and C++ exceptions, just enough to understand the exercise.

14.4 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 `ParameterSets`:

```
ls art-workbook/ParameterSets
```

```
CMakeLists.txt      pset02.fcl          PSet03_module.cc  
pset01.fcl          PSet02_module.cc    pset04.fcl  
PSet01_module.cc    pset03.fcl          PSet04_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 `PSet01_module.cc` and the FHiCL file to run it is `pset01.fcl`. The file `CMakeLists.txt` is identical to that used by the previous two exercises. The remaining files are the source and FHiCL files for additional steps in this exercise.

14.5 The Configuration File `pset01.fcl`

The FHiCL file that you will run in this exercise is `pset01.fcl`. Look at this file in your source window. You will see that `pset01.fcl` defines a parameter set `psetTester`, shown below, that configures an analyzer module named `PSet01`.

```
analyzers: {  
  psetTester : {  
    module_type : PSet01  
    a : "this is quoted string"  
    b : 42  
    c : 3.14159  
    d : true  
    e : [ 1, 2, 3 ]  
    f : {  
      a : 4  
      b : 5  
    }  
  }  
}
```

Figure 14.1: The FHiCL definition of the parameter set `psetTester` from `pset01.fcl`.

The parameter `module_type` is processed by *art*. All of the other parameters are processed by code in the module class `PSet01`. Additional definitions like these in a FHiCL file have the following properties:

1. The module specified by the `module_type` parameter defines which parameters must be present in this list, and which parameters are optional.
2. Each definition must be a legal FHiCL definition.
3. These definitions have no meaning, per se, to *art* or to FHiCL; they only have meaning to the C++ code in `PSet01_module.cc`.
4. Each definition may use the full power of FHiCL and may contain nested parameter sets to arbitrary depth.

Looking at the parameter set, it appears that the parameter `a` has a value that is a string of text, parameter `b`'s value is an integer number, parameter `c`'s is a floating point number, parameter `d`'s is one of the two possible boolean values, parameter `e`'s is an array of integers and that parameter `f`'s is a nested parameter set. You will learn in Section 14.6 that, from the point of view of the code in `PSet01_module.cc`, this intuition is cor-

rect. But there is one subtlety: FHiCL itself has no notion of type and, inside FHiCL, all parameter values are just strings. The interpretation of a parameter value as a particular type is done by code inside `PSet01_module.cc`. The computer-science-speak for this is that FHiCL is a *type-free* language; this is in contrast to C++ which is a *strongly-typed* language.

14.6 The Source code file `PSet01_module.cc`

The source code for this exercise is found in the file `PSet01_module.cc`. The new features seen in this exercise are all in the definition of the constructor.

When *art* starts up, it reads the file `pset01.fcl` and, among many other things, copies the FHiCL table `psetTester` into an object of type `fhiCL::ParameterSet`. When *art* calls the constructor of `PSet01`, it passes this `fhiCL::ParameterSet` as the argument of the constructor, named `pset`. That is, the table named `psetTester` in the FHiCL file appears in the module as a parameter set named `pset`.

Let's examine the first part of the constructor; see Figure 14.2.

Recall from Section 14.5 that the object `pset` internally represents the value of each parameter as a string. If you ask that the value of a parameter be returned as a string, `pset` will simply return a copy of its internal representation of that parameter. On the other hand, if you ask that the value of a parameter be returned as any other type, then `pset` needs to do some additional work. For example, if you ask that a parameter be returned as an `int`, then `pset` must first find its internal string representation of that parameter; it must then convert that string into a temporary variable of the requested type and return the temporary variable. Therefore, when your code asks `pset` to return the value of a parameter, it must tell `pset` two things:

1. the name of the parameter
2. the type to which the string representation should be converted

The angle bracket syntax `<>` is the signature of a feature of C++ called *templates*(γ). *art* and FHiCL use templates in several prominent places. You do not need to fully understand templates — just how to use them when you encounter them. The following pages describe how to use templates when getting the value of a parameter from `pset`.

The type of the argument is `fhicl::ParameterSet const&`. The class `ParameterSet` is in the namespace `fhicl`; all identifiers in the namespace `fhicl` are found in the UPS package named **fhiclcpp**. You can find the header file for `ParameterSet` at `$FHICLCPP_INC/fhiclcpp/ParameterSet.h`.

```
tex::PSet01::PSet01(fhicl::ParameterSet const& pset )
    : art::EDAnalyzer(pset) {
```

The argument named `pset` is passed by `const` reference; i.e., the module is not allowed to modify it (due to `const`), and the module is given access to it, but not given a copy of it (due to the `&`, meaning “reference”).

The fragment below illustrates how to extract values from `pset` and copy those values into local variables. The angle brackets, `<>`, indicate the use of a feature of C++ called *templates*(γ). This section of code is further discussed within the chapter text.

```
std::string a = pset.get<std::string>("a");
int b = pset.get<int> ("b");
double c = pset.get<double>("c");
bool d = pset.get<bool> ("d");
std::vector<int> e = pset.get< std::vector<int> >("e");
fhicl::ParameterSet f = pset.get<fhicl::ParameterSet>("f");

int fa = f.get<int>("a");
int fb = f.get<int>("b");
```

In most cases the `fhicl::ParameterSet` representation of a FHiCL table does NOT contain the name of the FHiCL table (`psetTester` in this case); it contains only the parameters defined between the braces. The exception is for FHiCL tables used to configure *art* modules; in this case *art* adds an extra parameter named `module_label` whose value is the the module label.

```
std::string module_type =
    pset.get<std::string>("module_type");
std::string module_label =
    pset.get<std::string>("module_label");
```

Figure 14.2: First part of constructor in `PSet01_module.cc`

When you ask for the value of a parameter, the name of the parameter is specified as a familiar function argument while the return type is specified between the angle brackets. The name between the angle brackets is called a *template argument*. If you do not supply a template argument, then your code will not compile.

For example, the line that reads the parameter `a` that reads

```
std::string a = pset.get<std::string>("a");
```

It first declares a local variable named `a` that is of type `std::string` and then asks `pset` to do the following:

1. Check if it has a parameter named `a`.
2. If it has this parameter, return it as a string.

The returned value is used to initialize the local variable, `a`. Section 14.9 will describe what happens if `pset` does not have a parameter named `a`.

It is not required that the local variable, `a`, have the same name as the FHiCL parameter `a`. But, with rare exceptions, it is a good practice to make them either exactly the same or close to the same.

The following line, that sets the parameter `b` and that reads

```
int b = pset.get<int>("b");
```

is similar to the previous line; the main difference is that `pset` will convert the string to an `int` before returning it. `pset` knows that it must perform the conversion to `int` because the template argument tells it to. Section 14.9 will describe what happens if the string cannot be converted to an `int`.

It is beyond the scope of this chapter to discuss how the template mechanism is used to trigger automatic type conversions. It is sufficient to always remember the following: when you use the `get` member function of the class `fhiCL::ParameterSet`, the template argument must always match the type of the variable on the left hand side. Templates will be discussed in Section 14.8.



The authors of FHiCL could have designed a different interface, such as

```
std::string a = pset.get_as_string("a");  
std::string b = pset.get_as_int    ("b");
```

Instead they chose to write it using templates. The reason for this choice is that it allows one to add new types to FHiCL without needing to recompile FHiCL. How you do this is beyond the scope of this chapter. You now know everything that you need to know about templates in order to use `fhiCL::ParameterSet` effectively.

The rest of the lines in that section of code extract the remaining parameters from `pset` and make copies of them in local variables. The remainder of the constructor, shown in Figure 14.3, prints the values of these parameters; these make the printout in lines 1 through 11 of Listing 14.1.

Your code may ask for the values of parameters from a `ParameterSet` in any order, and any number of times, including zero.

Two final comments on `PSet01_module.cc`. First, the `analyze` member function is empty. Nevertheless, it must be present because *art* requires all analyzer modules to provide a member function named `analyze`. If we removed this member function from the class `PSet01`, then the module would not compile. Second, the argument of the `analyze` member function is not used; therefore it is not given a name. Were it given a name, the compiler would complain that the argument was never used. When no name is given the compiler understands that it is your intention not to use the argument. Even though the code does not use the argument, its type must be present because the number, type and order of the arguments is part of the signature of a function.

14.7 Running the Exercise

Now let's see what happens when you run the job. In your build directory, run the following command

```
art -c fcl/ParameterSets/pset01.fcl >& output/pset01.log
```

The expected output from this command is shown in Listing 14.1.

The module reads in the parameter set and then prints out each of the values in several different ways. Check that the printout matches the definitions of the parameters from `pset01.fcl`. Understand the relationship between the printout and the lines in the source file `PSet01_module.cc`.

```

std::cout << "\n—————\nPart 1:\n";
std::cout << "a : " << a << std::endl;
std::cout << "b : " << b << std::endl;
std::cout << "c : " << c << std::endl;
std::cout << "d : " << d << std::endl;

std::cout << "e :";
for ( int i: e ){
    std::cout << " " << i; }
std::cout << std::endl;

```

The following lines show one way to print the two values, `a` and `b` from the parameter set `f`.

```

std::cout << "f.a : " << fa << std::endl;
std::cout << "f.b : " << fb << std::endl;

std::cout << "module_type: " << module_type << std::endl;
std::cout << "module_label: " << module_label << std::endl;

```

The following three lines show two other ways, using the `to_string()` and `to_indented_string()` member functions of the class `fhicl::ParameterSet`. These lines make the printout found in lines 13 through 18 of Listing 14.1.

```

std::cout << "\n—————\nPart 2:\n";
std::cout << "f as string: " << f.to_string() << std::endl;
std::cout << "f as indented-string:\n" << f.to_indented_string() << std::endl;

std::cout << "\n—————\nPart 3:\n";
std::cout << "pset:\n" << pset.to_indented_string() << std::endl;

```

The last two lines use the `to_indented_string()` member function to print everything found in the parameter set `psetTester`. These lines make the printout found in lines 21 through 36 of Listing 14.1.

Figure 14.3: Remainder of the constructor in `PSet01_module.cc`

Listing 14.1: Output from PSet01 with `pset01.fcl` (*art*-standard output not shown)

```
1 -----
2 Part 1:
3 a : this is quoted string
4 b : 42
5 c : 3.14159
6 d : 1
7 e : 1 2 3
8 f.a : 4
9 f.b : 5
10 module_type: PSet01
11 module_label: psetTester
12
13 -----
14 Part 2:
15 f as string:          a:4 b:5
16 f as indented-string:
17 a: 4
18 b: 5
19
20
21 -----
22 Part 3:
23 pset:
24 a: "this is quoted string"
25 b: 42
26 c: 3.14159
27 d: true
28 e: [ 1
29     , 2
30     , 3
31     ]
32 f: { a: 4
33     b: 5
34     }
35 module_label: "psetTester"
36 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 by its proper name:

```
ParameterSet::get<T>
```

or, sometimes, just `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 `fhcil::ParameterSet` in the files:

```
$FHICL_CPP_INC/fhiclcpp/ParameterSet.h
$FHICL_CPP_DIR/source/fhiclcpp/ParameterSet.cc
```

In particular, 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.

- For a parameter whose value is another parameter set, `T` must be `fhcil::ParameterSet`.
- For a parameter with a value that is a sequence of items, all items in the sequence must be of the same type and `get<T>` allows `T` to be `std::vector<S>`, where the template argument `S` is any of the types given in the previous two bullets.

14.8.2 User Defined Types

This sub-section on reading user defined types from a parameter set is for experts only.



You can write helper functions that will allow the type `T` to be almost any type that you might want. How to do this is beyond the scope of this chapter. For an example see the files:

```
$TOYEXPERIMENT_DIR/source/toyExperiment/Utilities/ParameterSetHelpers.h
$TOYEXPERIMENT_DIR/source/toyExperiment/Utilities/ParameterSetHelpers.cc
```

These files allow you to define a FHiCL parameter like:

```
zaxis : [ 0., 0., 1.]
```

and to read it as

```
auto zaxis = pset.get<CLHEP::Hep3Vector>("zaxis");
```

14.9 Exceptions

14.9.1 Error Conditions

There are two sorts of error conditions that may occur when reading parameters from a parameter set:

1. The requested parameter is not present in the parameter set.
2. The requested parameter is present but cannot be converted into the requested type.

To give an example of the second sort, suppose that on line 6 of Listing 14.1 you change the FHiCL definition of the parameter `c` from `3.14159` to `"test"`: Now consider what

happens when you try to read this parameter as a double, as is done on the line

```
double c = pset.get<double>("c");
```

from the listing in Figure 14.2. The code will correctly find that parameter *c* exists but it will produce an error when it tries to convert the string “test” to a double.

In both of these cases, the code inside *pset* will tell *art* to stop processing events and to perform an orderly shutdown, which will be described in the next sub-section.

14.9.2 Error Handling

From time to time code within *art* will discover that, because of some error condition, it cannot continue to process events. When this happens *art* can be configured to stop processing events and then to do one of several different things:

1. It can attempt an orderly shutdown, described below.
2. It can write the offending event to a separate output file and continue normal with the next event.
3. It can skip the module in which the problem occurred and continue normal processing with the next module.
4. There are several other options that cannot be described here because the necessary background information has not yet been established.

When *art* attempts an orderly shutdown, it will:

1. Write a message to the log file that describes what happened.
2. Record the error condition that stopped processing; this information will be written to all output event-data files.
3. Call the `endSubRun` member function of every module.
4. Call the `endRun` member function of every module.
5. Call the `endJob` member function of every module.
6. Properly flush and close all output and log files.

7. Perform a few other clean up and shutdown actions for parts of *art* that have not yet been discussed.
8. Return a non-zero status code to the parent process; the status code is the number that appears on the last line of your *art* output, the line that begins "Art has completed ...".

For most sorts of errors, the orderly shutdown will be successful and your work up to the error will be preserved.

But there are circumstances for which the orderly shutdown will fail. One example of this is if you have reached your disk quota and there is no disk space to hold more output.

For all cases but one, *art*'s default behaviour is to attempt an orderly shutdown. The one non-standard case is when *art* is unable to find a requested data product; in that case the default behaviour is to continue with the next module.

These default behaviours can be changed by adding lines to the FHiCL file. A discussion of these features is beyond the scope of this chapter. When the section that describes how to do this is written, a link to that section will be added here.

.

The technology that *art* uses to interrupt event processing and to take one of the possible follow-on actions is a feature of C++ called exceptions. When *art* stops event processing it is said to *throw an exception*; this phrase will be used throughout the Workbook. The topic of exceptions is much too complex to even sketch in this chapter; at this time you only need to understand that the phrase *throw an exception* means to stop event processing and to take one of the allowed follow-on actions, usually an orderly shutdown. A chapter yet to be written will describe how to use exceptions in your own code to tell *art* to interrupt processing.

.

14.9.3 Suggested Exercises

In `pset01.fcl`, remove the definition of the parameter `b`. Rerun *art*. You should see an error message like that shown in Listing 14.2. Read the error message and understand

what it is telling you so that you will recognize the error message if you make this mistake in the future.

Listing 14.2: Output from PSet01 with `pset01.fcl` (parameter `b` removed)

```

1 %MSG-s ArtException:  PSet01:psetTester@Construction 14-Jul-2013 19:38:19 CDT
  ModuleConstruction
2 cet::exception caught in art
3 ---- Can't find key BEGIN
4   b
5 ---- Can't find key END
6 %MSG
7 Art has completed and will exit with status 8001.
```

Note too that the completion status is non-zero.

In `pset01.fcl`, restore the definition of `b` and change the definition of `c` to `"test"` . Rerun `art`. You should see an error message like that shown in Listing 14.3. Again, read the error message and understand it.

Listing 14.3: Output from PSet01 with `pset01.fcl` (parameter `c` misdefined)

```

1 %MSG-s ArtException:  PSet01:psetTester@Construction 14-Jul-2013 19:42:54 CDT
  ModuleConstruction
2 cet::exception caught in art
3 ---- Type mismatch BEGIN
4   c
5   ---- Type mismatch BEGIN
6     error in float string:
7     test
8     at or before:
9   ---- Type mismatch END
10 ---- Type mismatch END
11 %MSG
12 Art has completed and will exit with status 8001.
```

14.10 Parameters and Data Members

Very often information from the parameter set is needed in a member function of the module class. The way to propagate this information from the parameter set to the member function is to store the values of these parameters as data members of the module class. This is illustrated in the two files `PSet02_module.cc` and `pset02.fcl`. Open these files with an editor and follow along with the description below.

If you need to refamiliarize yourself with the concept of data members of a class, refer to Section 6.7.2.

There are three things to notice in `PSet02_module.cc`.

1. The class declares three data members named `b_`, `c_`, and `f_`. These are declared in the private section so that only the module itself can see them.
2. In the constructor, these three data members are initialized to values extracted from the module’s parameter set.
3. In the `analyze` member function all three data members are printed out.

If you need to refamiliarize yourself with the colon initializer syntax, refer to Section 6.7.5.

If you need to refamiliarize yourself with the conventions about underscore characters in the names of data members refer to Section 6.7.7.2.

To run this example, enter

```
art -c fcl/ParameterSets/pset02.fcl >& output/pset02.log
```

The expected output from this is given in Listing 14.4.

Listing 14.4: Output from `PSet02` with `pset02.fcl`

```
1 Event number: run: 1 subRun: 0 event: 1  b: 42  c: 3.14159  f: a:4 b:5
2 Event number: run: 1 subRun: 0 event: 2  b: 42  c: 3.14159  f: a:4 b:5
3 Event number: run: 1 subRun: 0 event: 3  b: 42  c: 3.14159  f: a:4 b:5
```

This example is only relevant when parameters are actually used in member functions. If a parameter is used only inside the constructor, do not store it as a data member; instead you should store it as a local variable of the constructor. This brings up a “best practice:” always declare a variable in the narrowest scope that works.



14.11 Optional Parameters with Default Values

It is sometimes convenient to provide a default value for a parameter. Default values may be provided in the source code that reads the parameter set. This mechanism is illustrated by the files `PSet03_module.cc` and `pset03.fcl`. Open these files with an editor and follow along with the description below.

You have already seen that the member function template `ParameterSet::get<T>` takes one function argument, the name of the parameter. For example,

```
int b = pset.get<int>("b");
```

It also takes an optional second function argument, a default value for the parameter. For example,

```
int b = pset.get<int>("b", 0);
```

If the second argument is present, there two cases:

1. If the parameter is not defined in the FHiCL file, then the second argument is returned as the value of the call to `get`.
2. If the parameter is defined in the FHiCL file, then the second argument is ignored and the value read from the FHiCL file is returned as the value of the call to `get`.

When reading the code in this example you will encounter the expression:

```
std::vector<double>(5, 1.0);
```

This tells the compiler to instantiate an object of type `std::vector<double>`, set its size to 5 and initialize elements 0 through 4 to have the value 1.0. If you are not familiar with this syntax, you can read about it in the documentation for the C++ Standard Library (see Section 6.9).

This expression appears as the second argument of the second call to the member function `pset.get<T>`. Therefore the compiler will create an unnamed temporary object (the vector of doubles) and pass that object to the member function `get<std::vector<double>>` as its the second argument; the compiler ensures that, once function call has completed, the temporary object is deleted.

With the above explanations, the source code for this example should be reasonably self-explanatory; it looks for two parameters named `debugLevel` and `g` and supplies default values for each of them. Look at the file `pset03.fcl`; you will see that the parameters `debugLevel` and `g` are not present in the `testPSet` parameter set; therefore `printout` will show the default values.

To run this example,

```
art -c fcl/ParameterSets/pset03.fcl >& output/pset03.log
```

Listing 14.5: Parameter-related portion of output from `PSet03` with `pset03.fcl`

```
1 debug level: 0
2 g:  1 1 1 1 1
```

The expected output from this is given in Listing 14.5.

As a suggested exercise, edit `pset03.fcl` and, in the parameter set `testPSet`, provide definitions for the parameters `debugLevel` and `g`. Make their values different from the default values. Rerun *art* and verify that the module has correctly read in and printed out the values you defined.

14.11.1 Policies About Optional Parameters

Allowing *optional* parameters is important for developing, debugging and testing; if all parameters were required all of the time, the complete list of parameters could become unwieldy². On the other hand, the use of optional parameters can make it difficult to audit the physics content of a job. Therefore experiments typically have policies for what sorts of parameters may have defaults and what sorts may not. For example, your experiment may prohibit default values for parameters that define the physics behavior, but allow them for parameters that control printout and other diagnostics.



Consult your experiment to learn what policies you should follow.

14.12 Numerical Types: Precision and Canonical Forms

FHiCL recognizes numbers in both fixed point and exponential notation, for example `123.4` and `1.234e2`; the letter `e` that separates the exponent can be written in either upper or lower case.

In the preceding exercises you defined some numerical values in a FHiCL file, read them into your code and printed them out; the printed values exactly matched the input values. The values used in those exercises were carefully chosen to avoid a few surprises: there

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	Canonoical 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

are cases in which the printed value will be an equivalent, but not identical, form. This section discusses those cases and provides some examples.

When FHiCL recognizes that a parameter value is a number it converts the number into a *canonical form* and stores the canonical form as a string. The transformation to the canonical form preserves the full precision of the number and involves the following steps:

1. The canonical form has no insignificant characters:
 - (a) no insignificant trailing zeros
 - (b) no insignificant trailing decimal point
 - (c) no insignificant leading plus sign
 - (d) no insignificant leading plus sign in the exponent
2. If a number is specified in exponential notation and if the number can be represented as a integer without loss of precision, and if the resulting integer has 6 or fewer digits, then the canonical form is the integer. For example, the canonical form of 1.23456E5 is 123456 but the canonical form of 1.23456E6 is 1.23456e6.
3. The canonical form of all other floating point numbers is exponential notation with a single, non-zero digit to the left of the decimal point.
4. The canonical form of all strings includes beginning and ending quotes; this is true even if the string contains no embedded whitespace or other special characters.

Some examples of numbers and their canonical forms are given in Table 14.1.

If a numerical value, when expressed as a fixed point number, has no fractional part, your code may ask for the parameter to be returned as either a floating point type (such as

double or float) or as an integral type (such as `int`, `short unsigned` or `std::size_t`). For example, fourth non-blank line in the listing in Figure 14.2 was written

```
int b = pset.get<int>("b");
```

It might also have been written

```
double b = pset.get<double>("b");
```

which would do the expected thing: given the input from `pset01.fcl`, it would read the value 42 into a variable of type `double`.

On the other hand, if a numerical value, when expressed as a fixed point number, does have a fractional part, you may only ask for the parameter to be returned as a floating point type. If you ask for such a value as an `int`, the `ParameterSet::get<int>` member function will throw an exception; similarly for all other integral types. This behavior may not be intuitive: the authors of *art* could have decided, instead, to discard the fractional part and return the integer part. They chose not to do this because when this situation occurs, it is almost always an error.

14.12.1 Why Have Canonical Forms?

What is the point of having canonical forms for numbers?

There are times when it is necessary for *art*, or user of *art*, to ask if two parameter sets are the same. An example is when looking at the processing history of a data product, which includes the parameter set used to configure the module that produced the data product. In this example you might wish to ask if a particular data product was created with the official version of the parameter set or with an unofficial version.

When comparing two parameter sets a trivial, but pernicious, complication is when two parameter sets differ only by meaningless differences in the representation of numbers, such as “1” vs “1.0”. If one compares the canonical forms of the two parameter sets, this complication is removed.

Listing 14.6: 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
```

14.13 Suggested Exercises

14.13.1 Formats

The above ideas are illustrated by the files `PSet04_module.cc` and `pset04.fcl`. To run this example,

```
art -c fcl/ParameterSets/pset04.fcl >& output/pset04.log
```

The expected output from this is given in Listing 14.6. Read the source code and the FHiCL file; then examine the output. The first three lines show three different printed formats of the parameter `a`, with the first being the canonical form. While all forms are equal to the number found in the FHiCL file, they all have different formats. Understand why each line has the format it does.

Line 4 shows the canonical form of the parameter `b`. Line 5 shows what is printed using the default C++ settings; the two least significant characters were dropped. The code that produces line 6 shows how to use the `precision` function from the C++ Standard Library to tell C++ to print more significant figures.

If you modify the precision of `cout`, it will change the format of the printout for the rest of the job; usually this is a bad thing. To avoid this, `PSet04_module.cc` illustrates how to save and restore the precision of `cout`.

Line 7 shows the canonical form of the parameter `c` and line 8 shows the default C++ printed form of the integer.

Listing 14.7: Output from PSet04 with modified `pset04.fcl` (intentional error)

```
1 %MSG-s ArtException: PSet04:pset@Construction 28-Jul-2013 23:39:31 CDT
  ModuleConstruction
2 cet::exception caught in art
3 ---- Type mismatch BEGIN
4     c
5     narrowing conversion
6 ---- Type mismatch END
7 %MSG
8 Art has completed and will exit with status 8001.
```

14.13.2 Fractional versus Integral Types

For the next exercise, edit `pset04.fcl` and change the value of `c` to something with a fractional part. Rerun *art*; you should see that it throws an exception because it is illegal to read a numeric value with a fractional part into a variable of integral type. The error message from *art* is shown in Listing 14.7. Read the error message and understand what it is telling you so that you will recognize the error message if you make this mistake yourself.

14.14 Review

1. read parameter values from a FHiCL file into a module
2. require that a particular parameter be present in a parameter set
3. use data members to communicate information from the constructor to other member functions of a module
4. print a parameter set
5. use the colon initializer syntax
6. provide a default value for a parameter (if the parameter is absent from a parameter set)
7. modify the precision of the printout of floating point types

8. recognize the error messages for a missing parameter or for a value that cannot be converted to the requested type

You will also learn:

1. that you should find out your experiment's policy about what sorts of parameters are allowed to have default values
2. an extra parameter automatically added by *art*, but only in parameter sets that are used to configure modules
3. the canonical forms of parameters

Finally, you will learn a small amount about C++ templates and C++ exceptions, just enough to understand the exercise.

14.15 Test Your Understanding

14.15.1 Tests

The answers are intentionally placed on a new page (remember to try before you read further!).

14.15.2 Answers

15 Exercise 5: Making Multiple Instances of a Module

15.1 Introduction

In a typical HEP experiment 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 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).
5. The parameter `magicNumber` is a required parameter.
6. Both the constructor and the `analyze` method print an informational message that includes the value of `magicNumber_`.

15.6 The FHiCL File `magic.fcl`

The FHiCL file used to run this exercise is `magic.fcl`. Look at this file and you should see the following features:

1. Compared to previous exercises, The FHiCL names `process_name`, `source` and `services` have no important differences.
2. In the `analyzers` parameter set, inside the `physics` parameter set, you will see the definition of four module labels, `boomboom`, `rocket`, `flower` and `bigbird`¹. The value of each definition is a parameter set.
3. The first three of these parameter sets tell *art* to run the module `MagicNumber` and each provides a value for the required `magicNumber` parameter².
4. The last parameter set tells *art* to run the module `First`, the source for which was discussed in Chapter 10 and is listed in Listing 10.2; this module does not need any additional parameters.
5. The path `e1` contains the names of all of the module labels from the `analyzers` parameter set.

15.7 Running the Exercise

In your build directory, run the following command

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.

Listing 15.1: Output using `magic.fcl`

```

1 MagicNumber::constructor: magic number: 9
2 MagicNumber::constructor: magic number: 5
3 Hello from First::constructor.
4 MagicNumber::constructor: magic number: 10
5 MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 9
6 MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 5
7 Hello from First::analyze. Event id: run: 1 subRun: 0 event: 1
8 MagicNumber::analyze: event: run: 1 subRun: 0 event: 1 magic number: 10
9 MagicNumber::analyze: event: run: 1 subRun: 0 event: 2 magic number: 9
10 MagicNumber::analyze: event: run: 1 subRun: 0 event: 2 magic number: 5
11 Hello from First::analyze. Event id: run: 1 subRun: 0 event: 2
12 MagicNumber::analyze: event: run: 1 subRun: 0 event: 2 magic number: 10
13 MagicNumber::analyze: event: run: 1 subRun: 0 event: 3 magic number: 9
14 MagicNumber::analyze: event: run: 1 subRun: 0 event: 3 magic number: 5
15 Hello from First::analyze. Event id: run: 1 subRun: 0 event: 3

```

```
art -c fcl/ModuleInstances/magic.fcl >& output/magic.log
```

The expected output from this command is shown in Listing 15.1; for clarity, the printout made by *art* has been elided. Compare this printout to the printout from your run; it should be exactly the same. Inspect the printout and the files `MagicNumber_module.cc` and `../FirstModule/First_module.cc`; understand why the printout is what it is.

15.8 Discussion

15.8.1 Order of Analyzer Modules is not Important

As it happens, *art* runs the four analyzer modules in the order specified in the path definition `e1`. But you must not count on this behaviour! Two of the design rules of *art* are:



1. Modules may only communicate with each other by putting information into, and reading information from, the `art::Event`.
2. Analyzer modules may not put information into the `art::Event`.

Therefore *art* is free to run analyzer modules in any order.

For producer modules, which may add information to the event, the order of execution is often very important. When you reach the exercises that run producer modules, you will be told how to specify the order of execution.

You may wish to review some of the other ideas about *art* paths that are described in Section 9.8.8.

15.8.2 Two Meanings of *Module Label*

In the preceding discussion, the name *module label* was used in two subtly different ways, as is illustrated by the module label `rocket`:

1. `rocket` identifies a parameter set that is used to configure an instance of the module `MagicNumber`.
2. `rocket` is also used as the name of the module instance that is configured using this parameter set; the elements in the path `e1` are all the names of module instances.

Clearly these two meanings are very closely related, which is why the same name, *module label*, is used for both ideas. Throughout the remainder of this document suite the name *module label* will be used for both meanings; the authors believe it will be clear from the context which meaning is intended. This is standard usage within the *art* community.

15.9 Review

After working through this exercise, you should:

1. Know how to run multiple instances of the same module within one *art* job.
2. Understand that *art* does not guarantee the order in which analyzer modules will be run.
3. Understand the two senses in which the name *module label* is used: as the name of a parameter set and as the name of the corresponding instance of a module.

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`.

15.10 Test Your Understanding

15.10.1 Tests

Remember to switch from your source window to your build window as you edit files and run *art*! (Need a refresher? Review the procedure under Section 10.12.) Answers are provided in Section 15.10.2.

Edit `magic.fcl` and do the following:

1. Add a new analyzer module label that configures an instance of the module `Optional` from Chapter 13.
2. Add the new module label to `e1`.

Then re-run

```
art -c fcl/ModuleInstances/magic.fcl
```

Do you see the expected additional printout?

Now run the configuration with the intentional errors.

```
art -c fcl/ModuleInstances/bug01.fcl
```

This will run but produce an error. There are two bugs. When you fix the first one and rerun, a new one will pop up. Your job is to figure out each problem, in turn, and fix it.

The answers are intentionally placed on a new page (remember to try before you read 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.

2. The identifier inside the angle brackets is called a *template argument* and it is usually the name of a C++ type.¹
3. The dynamic sizing occurs in the middle of a running program; not at compile time.
4. This expression sets `nEntries` to the number of entries in `t`:

```
std::size_t nEntries = t.size();
```

16.3 What You Will Learn

In this exercise you will learn about:

1. the data type `tex::GenParticleCollection`
2. the four-part name of an *art* data product
3. the class `art::InputTag`
4. the class template `art::Handle`
5. the class template `art::ValidHandle`
6. the member function templates of `art::Event`:
 - `getByLabel(art::InputTag, art::Handle<T>) const;`
 - `getValidHandle<T>(art::InputTag) const;`

16.4 Background Information for this Exercise

The input files used for the *art* workbook contain data products created by a workflow that simulates the response of the toy detector to a generated event, described in Section 3.7.2. The first step in this workflow is to use an event generator to create a collection of generated 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.

a single data product that holds a collection of generated particles; there is not one data product per generated particle.

In this exercise you will retrieve this data product and print the number of generated particles in each event. *A future chapter will look at the properties of individual generated particles.*

16.4.1 The Data Type `GenParticleCollection`

Each generated particle in the simulated event is described by an object of type `tex::GenParticle`. All of the generated particles in a given event are stored in an object of type `tex::GenParticleCollection`. This object is written to the `art::Event` as a data product.

The header files that describe these two classes, `GenParticle.h` and `GenParticleCollection.h`, are found under:

```
$TOYEXPERIMENT_INC/toyExperiment/MCDataProducts/
```

The content of `GenParticleCollection.h` is shown in Listing 16.1; the code guards and comments have been omitted. This header uses a typedef to declare that the name `tex::GenParticleCollection`; is a synonym for `std::vector<tex::GenParticle>`.

Listing 16.1: Contents of `GenParticleCollection.h`

```
1
2 #include "toyExperiment/MCDataProducts/GenParticle.h"
3
4 #include <vector>
5
6 namespace tex {
7
8     typedef std::vector<GenParticle> GenParticleCollection;
9 }
```

Why did the authors of the workbook decide to use a typedef and not simply ask you to code `std::vector<GenParticle>` when needed? The reason is future-proofing. Suppose that down the road the authors find that they need to change the definition of `tex::GenParticleCollection`; if you used the typedef, it is much more likely

that your code will continue to compile and work correctly as is. If, on the other hand, you used `std::vector<GenParticle>`, then you would need to identify and edit every instance.



Please use the typedef `GenParticleCollection` in your own code and do not hand-substitute its definition.

Why did the authors of the workbook decide to call this typedef `GenParticleCollection` and not, for example, `GenParticleVector`? The answer is a different sort of future-proofing. The C++ standard library provides class templates other than vectors that are collections of objects, and one can imagine a scenario in which it would make sense to change `GenParticleCollection` to use a collection type, such as `std::deque` for example. In such a scenario, the following definition would make perfect sense to the C++ compiler but would be misleading to human readers:

```
typedef std::deque<GenParticle> GenParticleVector
```

The generic name *Collection* avoids this problem.

16.4.2 Data Product Names

Each *art* data product has a name that is a text string with four fields, delimited by underscore characters (`_`) that represent, in order, the data type, module label, instance name and process name, e.g.,:

```
MyDataType_MyModuleLabel_MyInstanceName_MyProcessName
```

Each data product name must be unique within an *art* event-data file. The fields in the data product name may only contain the following characters²:

- `a...z`
- `A...Z`
- `0...9`
- `::` (double colon)

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.

In particular, periods, dashes, commas, underscores, semicolons, white space and single colons are not allowed; underscores are only allowed as the field separator, not within a field.

About each field:

1. The data type field is the so-called *friendly name* name of the data type for the data product; friendly names are discussed below.
2. The module label field is the label of the module that created the data product. Note that it is the module *label* as specified in the FHiCL file, not the `module_type`.
3. A given module instance in a given *art* process may make many data products of the same type. These are distinguished by giving each a unique instance name. An empty string is a valid instance name and in fact is the default. The other three fields must be non-empty strings.
4. The process name field holds the value of the `process_name` parameter from the FHiCL file for the *art* job that created the data product.

The friendly name of a data type is a concept that *art* inherited from the CMS software suite. You will never need to write friendly names but you will need to recognize them. Knowing the following rules will be sufficient in most cases:

1. If a type is not a collection type, then its friendly name is the fully qualified name of the class.
2. If a type is `std::vector<T>`, its friendly name is `Ts`; the mnemonic is that adding the letter "s" makes it plural.
3. If a type is `std::vector<std::vector<T> >`, its friendly name is `Tss`. And so on.
4. If a type is `cet::map_vector<T>`, its friendly name is `Tmv`.

The full set of rules is given in the Users' Guide.

Corollaries of the above discussion include:

- None of the four fields in a product name may contain an underscore character: otherwise the parsing of the name into its four fields is ambiguous.

- If an *art* event-data file is populated by running several *art* jobs, each of which adds some data products, then each *art* job in the sequence must have a unique `process_name`.

16.4.3 Specifying a Data Product

To identify a data product, *art* requires that you specify the data type, module label and instance name fields (an empty string is a valid instance name). If the event contains exactly one data product that matches this specification, then *art* allows a wild card match on the process name field. If the event contains more than one data product that matches this specification, then *art* requires that you also specify the process name. I.e., *art* allows a wild card match only on the process name field, not on the others.

To tell *art* which data type you want, you use a template argument. To specify the other three fields, module label, instance name and process name, you use an object of type `art::InputTag`. Why is the data-type field treated differently than the others? This method allows *art* to look after the translation of the data type to its friendly name. Users of *art* never need to learn how to do this translation.

The header for `art::InputTag` is found in the file

`$ART_INC/art/Utilities/InputTag.h`.

You can construct an input tag by passing it a string with the three fields separated by colons, e.g.,:

```
art::InputTag tag("MyModuleLabel:MyInstanceName:MyProcessName");
```

For this exercise, the full specification of the input tag includes only the module label and the process name:

```
art::InputTag tag("evtgen::exampleInput");
```

The double colon indicates that the instance name (which would come between the colons) is an empty string. The process name rarely needs to be specified, and in fact it is not needed in this exercise. It will be sufficient to specify the input tag as

```
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`.

In this exercise you will run three modules that differ in only a few lines. The three source files use different syntax to accomplish the same thing. Most of the subsequent exercises in the workbook will use the syntax shown in the third version, `ReadGens3_module.cc`, and we recommend using this syntax in most cases. A description of the first two here serves as a pedagogical progression. You will likely see all three types of syntax in your experiment's code.

16.6 Running the Exercise

You will run the exercise from your build directory in your build window. To run this exercise, `cd` to your build directory and type the command:

```
art -c fcl/ReadGenParticles/readGens1.fcl >& output/readGens1.log
```

This will make the usual *art* output, interspersed with the output made by `readGens1.fcl`. The output from this module is shown in Listing 16.2. For each event it prints the event number and the number of `GenParticles` in that event.

Listing 16.2: Output using `readGens1.fcl`

```
ReadGens1::analyze event: 1 GenParticles: 7
ReadGens1::analyze event: 2 GenParticles: 3
ReadGens1::analyze event: 3 GenParticles: 3
ReadGens1::analyze event: 4 GenParticles: 3
ReadGens1::analyze event: 5 GenParticles: 5
```

16.7 Understanding the First Version, `ReadGens1`

16.7.1 The Source File `ReadGens1_module.cc`

The module `ReadGens1_module.cc` contains a new `include` statement for the `GenParticleCollection.h` header file. Here is the set of `include` statements at the top of the file:

Listing 16.3: Include statements in `ReadGens1_module.cc`

```
1 #include "toyExperiment/MCDataProducts/GenParticleCollection.h"
2
3 #include "art/Framework/Core/EDAnalyzer.h"
```

```

4 #include "art/Framework/Core/ModuleMacros.h"
5 #include "art/Framework/Principal/Event.h"
6
7 #include <iostream>
8 #include <string>

```

In the next portion of the file, notice the new data member `gensTag_` on line 16, which is initialized in the constructor using a string value that is taken from the parameter set:

```

9  class ReadGens1 : public art::EDAnalyzer {
10
11  public:
12  explicit ReadGens1(fhicl::ParameterSet const& );
13  void analyze(art::Event const& event) override;
14
15  private:
16  art::InputTag gensTag_;
17  };
18 }

```

Notice two things in the remainder of the file, below: Lines 25-26 introduce the concept of a *handle*(γ), setting `gens` as a handle to the requested `GenParticleCollection`. Lines 28-30 print out the number of entries in the data product — the same number as the number of generated particles in the event.

```

19 tex::ReadGens1::ReadGens1(fhicl::ParameterSet const& pset ):
20   art::EDAnalyzer(pset),
21   gensTag_(pset.get<std::string>("genParticlesInputTag")) {
22 }
23 void tex::ReadGens1::analyze(art::Event const& event ) {
24
25   art::Handle<GenParticleCollection> gens;
26   event.getByLabel(gensTag_, gens);
27
28   std::cout << "ReadGens1::analyze_event:_ " << event.id().event()
29             << "_GenParticles:_ " << gens->size()
30             << std::endl;
31 }
32
33 DEFINE_ART_MODULE(tex::ReadGens1)

```

As you work through the *art* workbook you will encounter several types of handles. All of the handle types behave like pointers with additional features:

1. They have safety features that make it impossible for your code to look at a pointee that is either not valid or not available.
2. They may also have an interface that lets you access metadata that describes the pointee.

The handle is an example of a broader idea sometimes called a *safe pointer*(γ) and sometimes called a *smart pointer*(γ).

The header for the class template `art::Handle` is found in the file `$ART_INC/art/Framework/Principal/Handle.h`. This file is automatically included by the include for `Event.h`.

The `art::Handle` line tells the compiler to default construct an object of type: `art::Handle<GenParticleCollection>`. The name of the default-constructed object is `gens`. A default-constructed handle does not point at anything and, if you try to use it as a pointer, it will throw an exception. A handle in this state is said to be invalid.

The following line calls `getByLabel`, which uses its first argument (`gensTag_`) to learn three of the four elements of the name of the requested data product. It can deduce the fourth element, the data type, from the *type* of its second argument (`gens`): that is, it knows that it must look for a data product of type `tex::GenParticleCollection`. *art* has tools to compute the friendly name from the full class name, which is why you will never need to write a friendly name.

When this line is executed, the event object looks to see if it contains a data product that matches the request. There are three possible outcomes:

1. the event contains exactly one product that matches
2. the event contains no product that matches
3. the event contains more than one product that matches

In the first case, the event object will give the handle a pointer to the requested `tex::GenParticleCollection`; the handle `gens` can then be used as a pointer, as is done in the second line of the `std::cout` section. When the handle has received the pointer, it is said to be in a valid state. In the second and third cases, the event object will leave the handle in its default-constructed state and, if you try to use it as a pointer, it will

throw an exception.

If the event object finds exactly one match, it will also add two pieces of metadata to the handle. One is a pointer to an object of type `art::Provenance`, which contains information about the processing history of the data product. The second is an object of type `art::ProductID`; this is essentially a synonym for the four-field string form of the product name. Both of these will be illustrated in future exercises. .

The third case bears one more comment: the developers of *art* made a careful decision that, except for the process name field, `getByLabel` will not have a notion of “best match”. When you use `getByLabel` you must unambiguously specify the data product you want or *art* will leave the handle in its default-constructed state.

If the `getByLabel` member function does not find the requested data product, e.g., if you run it on a different input file or if you misspell any of the fields in the input tag, the handle will be left in its default-constructed state. In this case, the `gens->size()` call will know that the handle is invalid and will throw an exception.

In all cases but one, *art*’s response to an exception is to attempt a graceful shutdown. The one unusual case is *ProductNotFound*, which is the exception thrown by an invalid handle when you try to use it as a pointer. In this case *art* will print a warning message, skip this module and attempt to run the remaining modules in the trigger paths and end paths.

It is possible to test the state of `gens` by using the member function `gens.isValid()`, which returns a `bool`. This not illusrated in the example because in most cases we recommend that you let *art* deal with this for you.

In the preceding discussion we did not mention that `getByLabel` is actually a *member function template*. There is no explicit template argument in the `event.getByLabel` line because the C++ template mechanism is able to deduce the template argument from the type of the second argument.



The `art::Event` object supports several other ways to request data products from the event, including a way to get handles to all data products that match a partial specification. This material is beyond the scope of this exercise. .

16.7.2 Adding a Link Library to `CMakeLists.txt`

`ReadGens1_module.so` requires linking to a dynamic library that was not needed by previous exercises, namely

`$TOYEXPERIMENT_LIB/libtoyExperiment_MCDataProducts.so`.

This library contains the object code for the classes and functions defined in the `MCDataProducts` subdirectory of the `toyExperiment` UPS product. In particular it contains object code needed by the data product `tex::GenParticleCollection`.

Adding this library to the link list required a one-line modification to `CMakeLists.txt`. If you compare this file to the corresponding file for the previous exercise, you will see that `CMakeLists.txt` for this exercise contains one additional line:

```
    ${TOYEXPERIMENT_MCDATAPRODUCTS}
```

The string `TOYEXPERIMENT_MCDATAPRODUCTS` is a `cmake` variable that was defined when you first ran the `buildtool` command. The translated value of this variable is the name of the required link library.

16.7.3 The FHiCL File `readGens1.fcl`

There is only one fragment of `readGens1.fcl` that contains any new ideas. It is the fragment that configures the module label `read`, reproduced in Listing 16.4

Listing 16.4: Configuring the module label `read` in `readGens1.fcl`

```

1
2   read : {
3       module_type           : ReadGens1
4       genParticlesInputTag   : "evtgen"
5   }
```

On line 4 of this fragment, the parameter `genParticlesInputTag` specifies the input tag that identifies the data product to be read by this exercise.



We recommend that you always initialize input tags using parameters from the parameter set and that you never initialize them using strings defined within the code. This will allow you run the same module on data products with different input tags; this is a widely used feature.

We further recommend that you not provide a default value in the call to get the parameter value from the parameter set. This derives from a general recommendation that parameters affecting physics output should never have default values; the only parameters with default values should be those that control debugging and diagnostics.



16.8 The Second Version, ReadGens2

Version 2 of this exercise consists of the files `ReadGen2_module.cc` and `readGen2.fcl`. To run this version, `cd` to your build directory and type the command:

```
art -c fcl/ReadGenParticles/readGens2.fcl >& output/readGens2.log
```

It will produce the same output as the previous two versions.

The only significant change from version 1 to version 2 is that lines

```
art::Handle<GenParticleCollection> gens;
event.getByLabel(gensTag_, gens);
```

have been replaced by the single (long) line:

```
art::ValidHandle<GenParticleCollection> gens =
    event.getValidHandle<GenParticleCollection>(gensTag_);
```

This version is a little verbose but that aspect will be addressed in version 3. Note that the class template `art::Handle` has been replaced by a new class template `art::ValidHandle`. Both class templates are defined in the same header file, `$ART_INC/art/Framework/Principal/Handle.h`.

The above line has functionality very similar to that of the two lines from version 1: the net result is that `gens` can be used as a pointer to the requested data product. It also has an interface to access the `art::Provenance` and the `art::ProductID`.

However, there are several significant differences between `art::Handle<T>` and `art::ValidHandle<T>`:

1. Unlike an `art::Handle<T>`, which may be either valid or invalid, an `art::ValidHandle<T>` is guaranteed to be valid. It cannot be default-constructed.

2. A call to `getValidHandle<T>` will either return a properly constructed `art::ValidHandle<T>` or it will throw a `ProductNotFound` exception.
3. `art::ValidHandle` does not have an `isValid()` method.
4. Everytime that you use an `art::Handle<T>` as a pointer, it first checks that the pointer is valid. On the other hand, when you use an `art::ValidHandle<T>` as a pointer, no check is necessary; using an `art::ValidHandle<T>` as a pointer 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):

```
auto gens =  
    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 requested module label, for example

```
genParticlesInputTag : "genevent"
```

Run *art* and observe the warning messages, which should look like the message in Listing 16.5. This tells you that:

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.
2. The error occurred while calling the member function `getByLabel`.
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.
4. *art* found zero data product matching this request.

You would also get an error if *art* had found more than one data product matching this request.

Observe that for each event *art* prints the warning message and continues with the next event.

Look at the last line of the *art* output and observe that *art* completed with status 0. This is because *art* treats `ProductNotFound` as a warning, not as an error that will initiate a shutdown.

You can reconfigure *art* so that a `ProductNotFound` exception will cause *art* to shutdown gracefully. To do this, edit your modified `readGens3.fcl` and add the following line

Listing 16.5: Warning message for misspelled module label of data product

```
1 %MSG
2 %MSG-w FailModule:  ReadGens3:read 15-Jun-2014 10:00:24 CDT
3                               run: 1 subRun: 0 event: 5
4 Module failed due to an exception
5 ---- ProductNotFound BEGIN
6   getByLabel: Found zero products matching all criteria
7   Looking for type: std::vector<tex::GenParticle>
8   Looking for module label: genevent
9   Looking for productInstanceName:
10
11 ---- ProductNotFound END
```

inside the `services` parameter set:

```
scheduler : { defaultExceptions : false }
```

This line tells *art* that its response to all exceptions should be to attempt a graceful shutdown. When you rerun *art* you should see output like that shown in Listing 16.6.

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
7. how to tell *art* to treat the `ProductNotFound` exception as a hard error that will

initiate a graceful shutdown.

16.12 Test Your Understanding

16.12.1 Tests

Two source files and one FHiCL file are provided that contain intentional bugs. Your job, of course, is to find the bugs and fix them. The answers are provided in Section 16.12.2.

The source files listed below fail to build. Follow essentially the same procedure as outlined in Section 10.12 on each of these files. To run them in *art*, modify an existing FHiCL file, as needed, or create a new one.

```
ReadGenParticles/ReadGensBug01_module.cc.nobuild
ReadGenParticles/ReadGensBug02_module.cc.nobuild
```

The FHiCL file `ReadGenParticles/bug03.fcl` runs to completion, printing lots of warning messages. However, it does not make any of the expected printout from the `analyze` member function. Run it, find the problem, and fix it.

16.12.2 Answers

For the file `ReadGensBug01_module.cc`, when you build you should find the following error message:

```
/home/aheavey/workbook/art-workbook/art-
workbook/ReadGenParticles/ReadGensBug01_module.cc:43:57: error: 'class
art::ValidHandle<std::vector<tex::GenParticle> >' has no member named 'size'
    << " GenParticles: "          << gens.size()
```

As the error message indicates, the error is indeed in the line:

```
<< "_GenParticles:_"          << gens.size()
```

Remember that `gens` is a handle to `GenParticleCollection`, and does not itself have a member function named `size`. The entity that does have a member function named `size` is the collection itself. Therefore the solution is that `gens.size()` should be `gens->size()`.

More generally, the syntax `gens->foo()` says “go find the thing to which `points`, find its function `foo()` and call it.” The (incorrect) `gens.foo()` says to call the function `foo` of the handle itself, which doesn’t exist.

For the file `ReadGensBug02_module.cc`, you will find the error:

```
/home/ahavey/workbook/art-workbook/art-
workbook/ReadGenParticles/ReadGensBug02_module.cc:34:57:
error: no matching function for call to
'art::EDAnalyzer::EDAnalyzer()'
    gensTag_(pset.get<std::string>("genParticlesInputTag")){
```

The problem here is that the constructor of an analyzer module must pass the parameter set to the constructor of `art::EDAnalyzer`, and it’s missing in this file. This was described in Sections 10.6.3.3 and ??

To make the build work, fix the constructor to look like

```
tex::ReadGensBug01::ReadGensBug01(fhicl::ParameterSet const& pset ):
    art::EDAnalyzer(pset),
    gensTag_(pset.get<std::string>("genParticlesInputTag")){
}
```

Now it should work fine.

Finally, run *art* on `bug03.fcl`. It will produce the error:

```
---- ProductNotFound BEGIN
getByLabel: Found zero products matching all criteria
Looking for type: std::vector<tex::GenParticle>
Looking for module label: Evtgen
Looking for productInstanceName:
---- ProductNotFound END
```

The parameter value for `genParticlesInputTag` has an upper-case `E` in `Evtgen`; it should be `evtgen`.

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-data ¹. 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.

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 and the file type `.root` for the names of input and output event-data files.

This chapter will introduce you to *art* services in general and to the `TFileService` in particular.

All user interactions with ROOT should happen via this service.



Note that it is possible to use ROOT as an event-processing framework, e.g., the AliRoot 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 documentation is at <http://aliweb.cern.ch/Offline/AliRoot/Manual.html>.



17.2 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 16.

17.3 What You Will Learn

In this exercise you will learn:

1. What the `art::TFileService` is and what it does for you.
2. How to configure the `art::TFileService`.
3. What an `art::ServiceHandle` is and what it does for you.
4. How to access ROOT via the `art::TFileService`.
5. How to create and fill a ROOT TH1D histogram.
6. How to use the interactive ROOT browser to view the histogram.
7. How to run a CINT script to view the histogram and to write the histogram to a PDF file.
8. The naming convention used by the Workbook to distinguish event-data ROOT files from ROOT files containing histograms, ntuples, and so on. This convention is 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`

The C++ source code for this exercise is found in the file `FirstHist1_module.cc`. Open the file in your source window to see it as a whole. Listing 17.1 contains a fragment of this file, showing the included headers and the declaration of the module class `FirstHist1`. Compared to the file `ReadGens3_module.cc` from the previous exercise, four new lines have been added; they appear in the listing as:

1. line 7, which includes the header for the *art* `TFileService`
2. line 9, which includes the header for the ROOT class `TH1D`
3. line 22, which declares the member function `beginJob`
4. line 29, which declares a new member datum, named `hNGens_`, of type *pointer* to an object of type `TH1D`.

The name `hNGens_` was chosen because this pointer will eventually point at a histogram object that contains a histogram of the number of generated particles per event.

The *art* workbook has adopted the style that all names for pointers to histograms begin with the lower case letter “h”.



The two new headers can be found at:

```
$ART_INC/art/Framework/Services/Optional/TFileService.h  
$ROOT_INC/TH1D.h
```

The conventions for including header files from ROOT differ from those for including header files from *art* and from *toyExperiment*. To remind you, the conventions for *art* and the *toyExperiment* UPS product are:

1. The names of all classes and functions are inside a namespace, *art* or *tex*, respectively.
2. In the header file `#include` lines, the name of the package to which the header belongs is always the first element of the path.

When ROOT was developed, namespaces were not supported robustly by many C++ compilers. Therefore a different set of conventions were adopted – and remain – for ROOT:

Listing 17.1: Declaration of the class `FirstHist1` from `FirstHist1_module.cc`

```
1
2 #include "toyExperiment/MCDataProducts/GenParticleCollection.h"
3
4 #include "art/Framework/Core/EDAnalyzer.h"
5 #include "art/Framework/Core/ModuleMacros.h"
6 #include "art/Framework/Principal/Event.h"
7 #include "art/Framework/Services/Optional/TFileService.h"
8
9 #include "TH1D.h"
10
11 #include <iostream>
12 #include <string>
13
14 namespace tex {
15
16     class FirstHist1 : public art::EDAnalyzer {
17
18     public:
19
20         explicit FirstHist1(fhicl::ParameterSet const& );
21
22         void beginJob() override;
23         void analyze(art::Event const& event) override;
24
25     private:
26
27         art::InputTag gensTag_;
28
29         TH1D* hNGens_;
30
31     };
32
33 }
```

Listing 17.2: Implementation of the class `FirstHist1`

```

1
2 tex::FirstHist1::FirstHist1(fhicl::ParameterSet const& pset ):
3     art::EDAnalyzer(pset),
4     gensTag_(pset.get<std::string>("genParticlesInputTag")),
5     hNGens_(nullptr) {
6 }
7
8 void tex::FirstHist1::beginJob() {
9
10     art::ServiceHandle<art::TFileService> tfs;
11     hNGens_ = tfs->make<TH1D>( "hNGens",
12         "Number_of_generated_particles_per_event", 20, 0., 20.);
13
14 }
15
16 void tex::FirstHist1::analyze(art::Event const& event ) {
17
18     auto gens = event.isValidHandle<GenParticleCollection>(gensTag_);
19
20     hNGens_->Fill(gens->size());
21
22 }

```

1. The names of all ROOT classes and functions are in the *global* namespace, i.e., they are not part of a namespace defined by ROOT.
2. The names of all ROOT classes begin with a capital letter T followed by an upper case letter (this serves as a weak substitute for using a namespace).
3. The syntax to include a file from ROOT is to give the filename without any leading path elements. The clue that the file is a ROOT header file comes from the leading capital T.

Listing 17.2 shows the implementation section of the file `FirstHist1_module.cc`.

The new features in this listing are:

1. line 5, which initializes `hNGens_` to have the value of a null pointer
2. lines 8 through 14, which create an empty histogram
3. line 20, which fills the histogram with the number of generated particles in the cur-

rent event

The identifier `nullptr`, used in line 5, was added to the C++ core language in the 2011 Standard. It is the value of a pointer that points to nothing; in practice it has a value of zero. You will very likely encounter code written prior to the 2011 Standard. In such code you will see the equivalent of line 5 written in one of the following two ways: `hNGens_(0)` or `hNGens_(NULL)`. In the second form, the value `NULL` is a C-Preprocessor **MACRO** variable that is defined to have a value of 0.



We strongly recommend, first, that you use `nullptr` for this purpose, and second that you never use the C-Preprocessor `NULL`.

17.5.1 Introducing `art::ServiceHandle`

Section 3.6.5 discussed the idea of *art* services. These are classes that provide some functionality (i.e., a service) that can be used by any module or by other *art* services. In this exercise you will see your first example of an *art* service, the `art::TFileService`, which provides a bookkeeping layer to ensure that your use of ROOT does not interfere with other uses of ROOT within the same *art* job.

In a similar way that access to data products is provided by the class templates `art::Handle` and `art::ValidHandle`, access to services is provided by the class template `art::ServiceHandle`. Line 10 in Listing 17.2 tells the compiler to default construct an object, named `tfs`, of type `art::ServiceHandle<art::TFileService>`. The constructor of `tfs` will contact the internals of *art* and ask *art* to find a service of type `art::TFileService`. If *art* can find such a service, it will give the service handle a pointer to the service. If not, it will throw an exception and attempt a graceful shutdown.

Once a service handle has been constructed, the downstream code can use the service handle as a pointer to the pointee, i.e., to `art::TFileService`.

The header file for `art::ServiceHandle` is found at:

```
$ART_INC/art/Framework/Services/Registry/ServiceHandle.h
```

It is automatically included by one of the files that are already included in `FirstHist1_module.cc`.

Listing 17.3: Creating histogram object of type TH1D

```
10  art::ServiceHandle<art::TFileService> tfs;  
11  hNGens_ = tfs->make<TH1D>( "hNGens",  
12  "Number_of_generated_particles_per_event", 20, 0., 20.);
```

17.5.2 Creating a Histogram

Lines 10 through 12 of Listing 17.2, use `art::TFileService` to create a new histogram object of type `TH1D`. In the call to the member function template `tfs->make` (lines 11 and 12), the type of object to be created is specified using a template argument. The function arguments, listed below, are the arguments needed by a constructor of that type of object. You do not need to understand why things are done this way or how it all works. You just need to follow the pattern. The return value of the call to `tfs->make` is a pointer to the newly created histogram object and this value is assigned the member datum `hNGens_`.

In the case of creating a `TH1D`, the five function arguments are:

1. the name by which ROOT will know this histogram; the *art* workbook has adopted the convention that this name will always be the name of the corresponding member datum, excluding the underscore (in this case `hNGens`)
2. the title that will be displayed when the histogram is drawn (given on line 12 of Listing 17.2)
3. the number of bins in the histogram (20)
4. the lower edge of the lowest bin of the histogram (0.)
5. the upper edge of the uppermost bin of the histogram (20.)

ROOT defines that the low edge of a bin is within that bin, while the upper edge of a bin is part of the next bin up. Therefore the lower edge of the lowest bin is inside the histogram but the upper edge of the uppermost bin is outside of the histogram.



If you would like to learn more about the `TH1D` class you can look at its header file or you can read about it on the ROOT web site: <http://root.cern.ch/root/html534/TH1D.html>.

Where is the histogram created? The histogram is created in memory that is owned and

Listing 17.4: Filling the histogram

```
20 hNGens_ -> Fill(gens -> size());
```

managed by ROOT. ROOT also knows that when the job is finished, it should write the histogram to a ROOT output file that you can inspect at a later time. The name of the output file is specified in the FHiCL file for the *art* job; more on that later. We will call this file the *histogram output file* or *histogram file*. Although histogram files often contain much more than just histograms, the name is in fairly common usage among the experiments that use *art*. Histogram files do not contain *art* data products.

Just as file systems have the notion of directories and subdirectories (or folders and subfolders if you prefer), a ROOT file has the notion of directories and subdirectories that are internal to the ROOT file. If a module makes at least one histogram, then the `TFileService` will first create a new top-level directory in the histogram file. The name of this top-level directory is the name of the module label of the module that created the histogram. All ROOT objects that are created by that module will be created within this top level directory. When the contents of ROOT-managed memory are written to the histogram file, this directory structure is preserved.

Recall that within a given *art* job each module label must be unique. This ensures that, for every module instance that uses the `TFileService`, a uniquely named top-level directory will be created in the histogram file. It is this strategy that ensures that the histogram names of my module will never collide with the histogram names of your module.

17.5.3 Filling a Histogram

Line 20 of Listing 17.2, 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.

By convention both the histogram files (output) and the *art* event-data files (input) end in `.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.

The *art* Workbook has adopted the convention that *art* event-data files always end in `_data.root`. All other files ending in `.root` are histogram files.



Some experiments have adopted a similar convention while others have adopted precisely the opposite convention: files ending in `_hist.root` are histogram files and all other files ending in `.root` are *art* event-data files.

17.6 The Configuration File C++ `firstHist1.fcl`

The file `firstHist1.fcl`, shown in Listing 17.6, is very much like the file `readGens3.fcl` from the previous exercise.

The most important new feature is at line 13,

Listing 17.5: `TFileService` in `firstHist1.fcl`

```
13 TFileService : { fileName : "output/firstHist1.root" }
```

which configures the `TFileService`. This service has one required parameter, which is the name of the histogram file that contains the histograms, trees, and so on that are created by the *art* job.

If this parameter is missing, or if the configuration for the `TFileService` is missing entirely, then the first attempt to get a service handle to the `TFileService` will throw an exception, and *art* will attempt a graceful shutdown.

Unlike in the previous exercises, the FHiCL file runs on the large input event-data file, `inputFiles/input04_data.root`, which contains 1000 events.

Listing 17.6: `firstHist1.fcl`

```

1 #include "fcl/minimalMessageService.fcl"
2
3 process_name : firstHist1
4
5 source : {
6   module_type : RootInput
7   fileNames   : [ "inputFiles/input04_data.root" ]
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 }
```

17.7 The file `CMakeLists.txt`

The `CMakeLists.txt` file used for this exercise is shown in Listing 17.7. Compared to the corresponding file for the previous exercise, there are two new features.

1. Two link libraries have been added.

2. There is a directive indicating that `cmake` should not build files ending in `.C`.

The two new link libraries are specified by

`${ART_FRAMEWORK_SERVICES_OPTIONAL_TFILESERVICE_SERVICE}`
and `${ROOT_HIST}` (lines 12 and 15). These items are both `cmake` variables that were defined for you when established the development environment. The first variable translates to

`$ART_LIB/libart_Framework_Services_Optional_TFileService_service.so`
and the second translates to `$ROOTSYS/lib/libHist.so`, which contains the object code for the class `TH1D`, among others.

Many projects use the convention that files ending in `.C` contain code written in the **C** programming language. By default `cmake` will assume that files ending in `.C` follow this convention and, therefore, it will try to compile and link them. We have already encountered a `CINT` script that ends in `.C`. The `cmake` needs to ignore `CINT` files. `CMakeLists.txt` includes code to effect this.

You do not need to understand the details of how the `CMakeLists.txt` excludes `drawHist1.C` from the build. For those who wish to look up the details, the high level explanation follows. Lines 1 and 2 in the listing of `CMakeLists.txt` tell `cmake` to define a new `cmake` variable named `ROOT_MACROS_DO_NOT_BUILD`. This variable is the set of all filenames that end in `.C` from the same directory as the `CMakeLists.txt` file. Line 4 in `CMakeLists.txt` tells `cmake` that it should do nothing for all files that appear in the translation of this variable.



17.8 Running the Exercise

Listing 17.7: 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 )
```

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
```

This module does not make any of its own printout. You should see the standard printout from *art*, including the final line saying that *art* will exit with status 0. Remember to add `>& output/<filename>.log` to the end of the command to send the printout to a file. The Workbook will not always show this in subsequent exercises, but it is always recommended.

You should see that the *art* job created the file `output/firstHist1.root`. This is the histogram file.

17.9 Inspecting the Histogram File

In this section you will inspect the file `output/firstHist1.root`.

First, look again at `fcl/FirstHistogram/firstHist1.fcl`. Note that the module label of the `FirstHist1` module is `hist1`.

To inspect the histogram you will remain in your build directory and you will run the interactive ROOT program, using the command `root`. This command was put into your 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”.

In step 4 you should have recognized the name of the folder, `hist1;1`. Ignoring the trailing `;1`, it is the name of the module label used in `firstHist1.fcl`. In step 5 you should have recognized the name of the histogram, `hNGens`; ignoring `;1`, it is the name that you gave the histogram when you created it.

About the `;1` that ROOT has stitched onto `hist1` and `hNGens`: ROOT calls these *cycle numbers*. They are part of a checkpointing mechanism that is beyond the scope of this exercise; if you ever see more than one cycle number for a ROOT object, the highest number is the one that you want. Consult the ROOT documentation for more details.

Now look at the histogram in the right hand panel of the TBrowser window. In the statistics box on the upper right you should see that it has 1000 entries, one for each event in the input file. You should also notice that only the odd bins are populated: this is because the

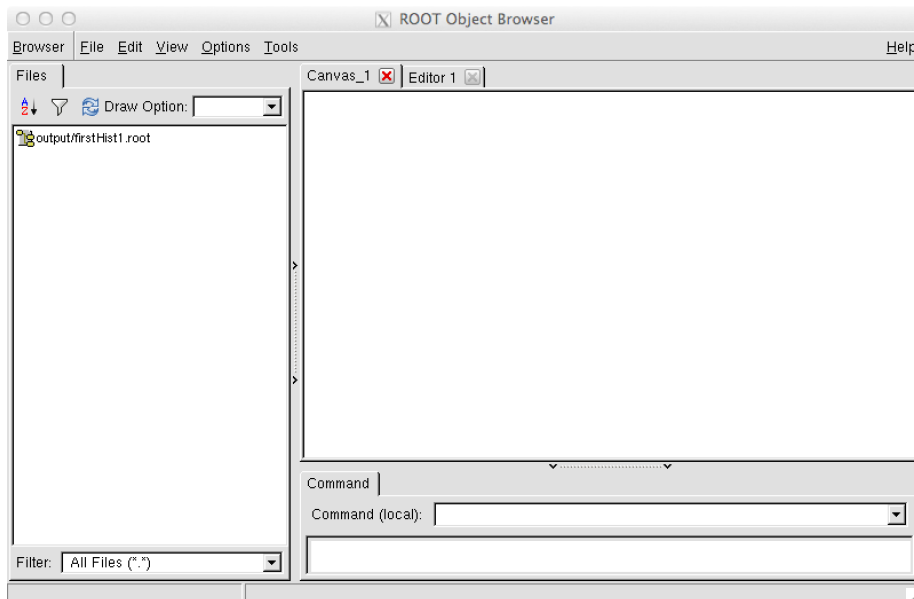


Figure 17.1: The TBrowser window immediately after opening `output/firstHist1.root`.

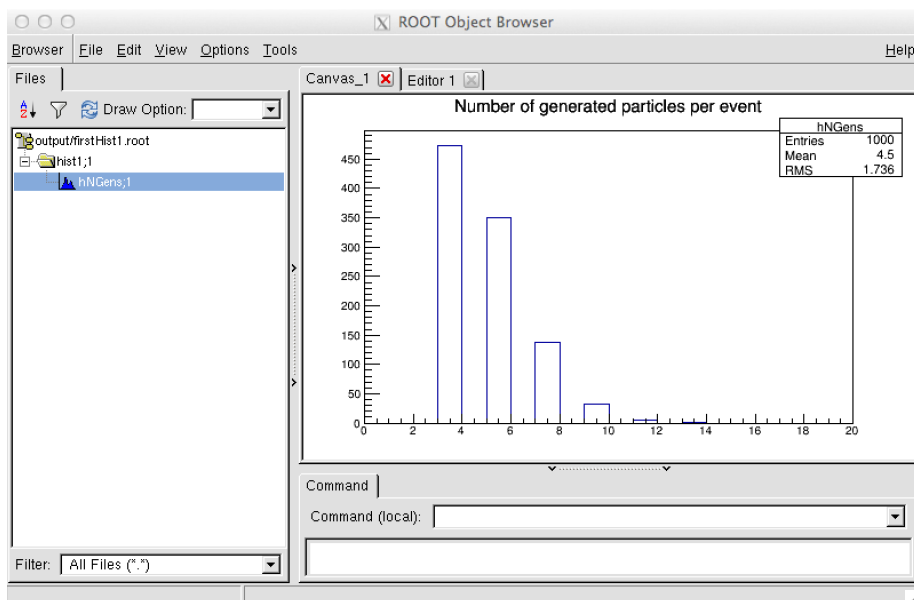


Figure 17.2: TBrowser window after displaying the histogram `hNGens;1`.

generated events always contains three signal particles, plus a random number of pairs of background particles ($3 + 2n$). The three signal particles are φ meson and the two kaons into which it decays. You should also recognize the title and the name that you set when you created the histogram. Finally you should recognize that the binning matches the binning you requested when you created the histogram.

17.9.1 A Short Cut: the `browse` command

The above description for viewing a ROOT file interactively requires a lot of tedious typing at step 2. The toyExperiment UPS product provides a command named `browse` that does the typing for you. To use this command:

```
browse output/firstHist1.root
```

Then follow the instructions from the previous section, starting at step 3.

When you created your *art* build environment, the toyExperiment UPS product put the command `browse` into your path. This command is implemented as a bash script and you can find its definition using the following bash command:

```
type browse
browse is <...>/scripts/browse
```

```
normal command myfilename
```

where `<...>` will change from one site to the next; it will have the value of `$TOYEXPERIMENT_DIR` as defined at that site.

```
[aheavey@cluck build-prof]$ type browse
browse is hashed (/home/kutschke/products//toyExperiment/v0_00_15/scripts)
[aheavey@cluck build-prof]$ echo $TOYEXPERIMENT_DIR
/home/kutschke/products//toyExperiment/v0_00_15
[aheavey@cluck build-prof]$
```

17.9.2 Using CINT Scripts

When you type a command at the root prompt, you are typing commands in a ROOT-defined language called CINT. Scripts in this language are sometimes called ROOT scripts

and other times CINT scripts.² As with many interpreted languages, you may write CINT commands in a file and execute that file as script.

It is a common convention that files that contain CINT scripts have a file type of `.C`. This convention is followed throughout the *art* workbook.

This exercise provides an example of a CINT script, `drawHist1.C`, in Listing 17.8.

To use this script, run:

```
root -l drawHist1.C
```

This will open a window on your display, draw the histogram in that window and save the window to the PDF file `output/NumberGenerated.pdf`, which is shown as Figure 17.3. When the script is complete, it returns control to the root prompt in your build window. At this prompt you can issue more CINT commands. To exit ROOT, type `.q` at the root prompt.

Appendix B has instructions on how to view the PDF file interactively. For those of you at Fermilab, it also has instructions on how to print the PDF file.

If you compare this figure to the histogram in Figure 17.2 you will see that there is difference in the statistics box in the upper right. In this figure, the name of the histogram is not shown but three new fields are: the number of entries below the lower limit (Underflow), the number of entries above the upper limit (Overflow) and the number of entries between the limits (Integral). The field named “Entries” is the sum of Integral plus Underflow plus Overflow.

It is beyond the scope of this writeup to describe all of the features in Listing 17.8, but we will describe some of the code.

Regarding lines 9 and 12, we will let comments in the code be a guide, and give two additional hints. The object `gROOT` is a pointer to an instance of the class `TROOT` and the object `gStyle` is a pointer to an instance of the class `TStyle`. To find the documentation for these classes, see Section 17.10.

²*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++.

Listing 17.8: Sample CINT file DrawHist1.C

```
1 //
2 // Root script to draw the histogram made by FirstHist1_module.cc
3 //
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 }
```

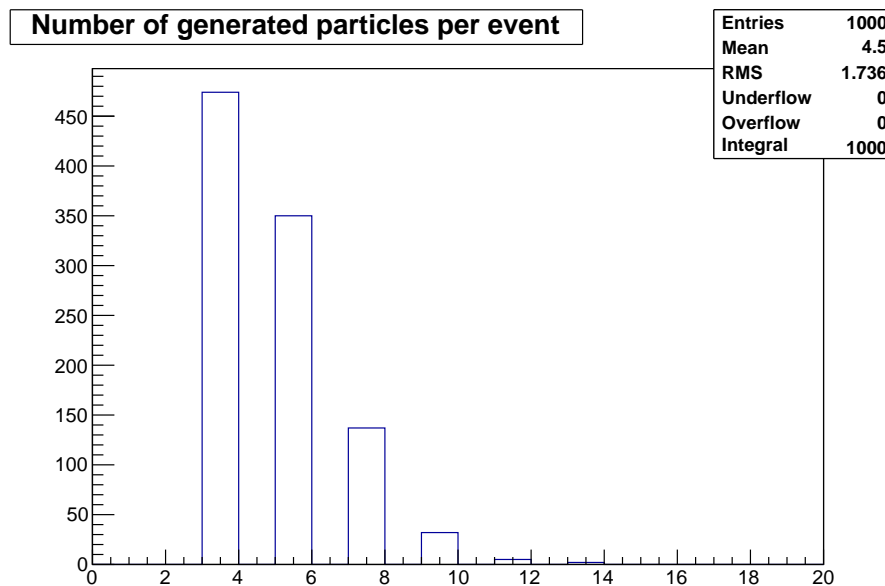



Figure 17.3: Figure made by running the CINT script `drawHist1.C`.

Line 17 tells ROOT what to draw in the statistics box in the upper right of every histogram. The comments in the code describe the mnemonics of the letter codes. The full set of letter codes is described on the ROOT web site: <http://root.cern.ch/root/html534/TStyle.html> as part of the documentation for the member function `SetOptStat`.

Line 20 opens the input file. The ROOT type `TFile` is ROOT's interface to ROOT information that lives in a disk file; it allows a ROOT program to write ROOT objects to the file and read ROOT objects from the file. You can learn more about the class `TFile` from the ROOT web site <http://root.cern.ch/root/html534/TFile.html>.

Line 23 has two statements on it. The first declares `hNGens` as an object of type pointer to `TH1D` and initializes it to 0 (ROOT does not support `nullptr`). The second asks the `TFile` to find a ROOT object named "hist1/hNGens", copy it from the file into memory and to set the pointer `hNGens` to point to that object. If ROOT cannot find the requested object, or if the type of the requested object does not match the type of the pointer, ROOT will leave `hNGens` with a value of 0. This is reminiscent of asking an `art::Event` to fill an `art::Handle`.

Line 26 tells ROOT to open a new window on your display. The first argument is an

arbitrary name that must be unique within the job; ROOT uses it internally to differentiate multiple canvases. The second argument is the title that will be drawn on the title bar of the window.

Line 29 tells ROOT to draw the histogram on the canvas. If, at line 21, ROOT was unable to properly set the pointer, then this line will produce an error message and return control to the root prompt in the build window.

Line 31 tells ROOT to flush its internal buffers and make sure that everything that is in the queue to be drawn on the canvas is actually drawn.

Line 33 tells ROOT to save the canvas by writing it to the file specified as the function argument. The format in which the file will be written is governed by the file type field in the filename, `.pdf` in this case. Many other formats are supported and a full list is 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 v5.34/18.

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>
```

To restore the file to a writeable state the unix command is:

```
chmod o+w <filename>
```

17.10.2 Changing the Name of the Histogram File

You can change the name of the histogram file by editing the FHiCL file but you can also do so from the *art* command line by using the `--TFileName` option; the short form of this option is `-T`.

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
```

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.

17.10.3 Changing the Module Label

In `firstHist1.fcl`, change the name of the module label, `hist1`. Rerun the job and browse the histogram file. You should see that the name of the directory in the histogram file has changed to match the new module label.

17.10.4 Printing From the TBrowser

You can use the `browse` command to open the histogram file and view the histogram.

In the TBrowser window, click on the “File” button. This will open a pull-down menu.

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`.

17.11 Review

In this exercise you have learned:

1. How to configure the `TFileService`.
2. How to use an `art::ServiceHandle` to access the `TFileService`.
3. How to use the `TFileService` to create a histogram that will automatically be written to the histogram file.
4. Three different ways to view the contents of the histogram file: by launching a TBrowser by hand, by using `browse` command and by running a CINT script.

5. The convention used by the *art* Workbook to differentiate histogram files from *art* event data files.

17.12 Test Your Understanding

17.12.1 Tests

Two source files and one FHiCL file are provided that contain intentional bugs. Your job, of course, is to find the bugs and fix them. The answers are provided in Section 17.12.2.

The source files listed below fail to build. Follow essentially the same procedure as outlined in Section 10.12 on each of these files. To run them in *art*, use the FHiCL file that matches the number 1 or 2, respectively.

```
FirstHistogram/FirstHistBug01_module.cc.nobuild  
FirstHistogram/FirstHistBug02_module.cc.nobuild
```

Next, run *art* on this FHiCL file that has a problem; then diagnose and fix it:

```
FirstHistogram/bug03.fcl
```

The answers are intentionally placed on a new page (remember to try before you read further!).

17.12.2 Answers

When you run `buildtool` after you have made the file `FirstHistBug01_module.cc`, the output includes the following error message:

```
/home/ahavey/workbook/art-workbook/art-workbook/FirstHistogram/FirstHistBug01_module.cc:56:11:
/home/ahavey/workbook/art-workbook/art-workbook/FirstHistogram/FirstHistBug01_module.cc:56:11:
  hNGens_>Fill(gens->size());
```

The cause of the problem is that `hNGens_` was declared as type `TH1D`, not as a *pointer* to type `TH1D`. The solution is to add an asterisk to the type declaration so that it reads

```
private:
...
TH1D* hNGens_;
```

When you run `buildtool` after you have made the file `FirstHistBug02_module.cc`, the output includes the following error message:

```
/home/ahavey/workbook/art-workbook/art-workbook/FirstHistogram/FirstHistBug02_module.cc:56:11:
In member function 'virtual void tex::FirstHistBug02::analyze(const art::Event&)' at:
/home/ahavey/workbook/art-workbook/art-workbook/
FirstHistogram/FirstHistBug02_module.cc:56:11:
error: request 'for' member 'Fill' in '(tex::FirstHistBug02*)this' ->tex::FirstHistBug02
which is of pointer type 'TH1D*' (maybe you meant to use '->' ?)
```

The solution is to replace the dot in the `Fill` line with a `->` in order to dereference the pointer `hNGens_`:

```
hNGens_>Fill(gens->size());
```

`hNGens_` is not the histogram class (`TH1D`) that has a `Fill` member function; it is a pointer to that class.

When you try to run:

```
art -c fcl/FirstHistogram/bug03.fcl
```

you will obtain the following error message:

```
RootOutput cannot ascertain a unique temporary filename for output based on stem
"outputs/TFileService": No such file or directory.
```

To fix `bug03.fcl`, correct the spelling of `outputs` to `output`.

18 Exercise 8: Looping Over Collections

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 ??.

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

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
2. the recommended way to write a loop over a `GenParticle` collection

3. seven other ways to write the same loop; you will encounter many of these in code written by others
4. about the classes `CLHEP::Hep3Vector` and `CLHEP::HepLorentzVector` and how to use them

18.4 Setting Up to Run 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 `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.

The main body of the exercise consists of the files `LoopGens1_module.cc`, `loopGens1.fcl` and `loopGens1.C`. The files with the numerals 2 and 3 in their names show alternate ways to write the loop over the generated particles. The files with the numeral 4 in their names collectively provide the solution to an exercise that will be assigned towards the end of this chapter. The files with the numerals 5, 6, 7, 8 in their

names contain deliberate errors that you will be asked to find and fix.

18.5 The Class `GenParticle`

This chapter discusses the data product `GenParticleCollection`, which you learned in Section 16.4.1 is a `std::vector<GenParticle>`. This section will describe the class `GenParticle`, which contains information about one generated particle.

The header file and the source file can be found at:

```
$TOYEXPERIMENT_INC/toyExperiment/MCDataProducts/GenParticle.h  
$TOYEXPERIMENT_DIR/source/toyExperiment/MCDataProducts/GenParticle.cc
```

Open the files to follow along with the discussion.

18.5.1 The Included Header Files

The header file `GenParticle.h` includes six header files that have not yet been used in the Workbook:

```
1 #include "toyExperiment/DataProducts/PDGCode.h"  
2 #include "art/Persistency/Common/Ptr.h"  
3 #include "art/Persistency/Common/PtrVector.h"  
4 #include "CLHEP/Vector/LorentzVector.h"  
5 #include "CLHEP/Vector/ThreeVector.h"  
6 #include <iosfwd>
```

The PDG in the name of first header file refers to the *Particle Data Group*, which has defined a standard set of about 400 integer codes to identify particles in simulations of HEP events. This scheme provides identifiers for individual quarks and leptons, for their anti-partners, for the gauge bosons, for the Higgs, for the graviton, for hadrons and for many proposed particles that are beyond the Standard Model. This scheme has been broadly adopted by the HEP community and is described on the PDG web site:

<http://pdg.lbl.gov/2007/reviews/montecarlo/rpp.pdf>

The Workbook includes a class that declares and initializes a small subset of the PDG codes, only those that are of interest to the toy experiment. It does so by way of an *enum*, which is a standard C++ type that sets a restricted range of possible values to a variable. This class is defined in:

```
$TOYEXPERIMENT_INC/toyExperiment/DataProducts/PDGCode.h
```

The only content of this class is the enum that defines mnemonic names for each code. By using the mnemonic names instead of the integer codes, your code will be understandable by people who do not have all of the codes committed to memory. To use one of the codes in the Workbook, you need to give its full name; for example, use `PDGCode::pi_minus` to specify the code for a π^- . If you are not familiar with enums, consult any standard C++ text.

The class templates from the next header files, `Ptr.h` and `PtrVector.h`, will be discussed in Chapter ?? . For this exercise you need to know only that these classes provide a type of smart pointer that can be written to disk and read back. In the `GenParticle` class they are used to record the parent-child relationships among the generated particles.

The header files `ThreeVector.h` and `LorentzVector.h` are part of a package named CLHEP, which is a library of classes widely used in HEP. If you are not familiar with CLHEP, there is general discussion in Appendix C. The first file contains the declaration of the class `CLHEP::Hep3Vector` that represents 3-vectors in 3-dimensional space. The second file contains the declaration of the class `CLHEP::HepLorentzVector` that represents a 4-vector in 4-dimensional space-time. These classes are used to describe, respectively, the position and the 4-momentum of the `GenParticle` at the time that it was created. The header files can be found at:

```
$CLHEP_INC/CLHEP/Vector/ThreeVector.h  
$CLHEP_INC/CLHEP/Vector/LorentzVector.h
```

You can find specific discussions about `Hep3Vector` in Section C.6.1 and about `LorentzVector` in Section C.6.2.

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.

18.5.2 Particle Parent-Child Relationships

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.

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.

18.5.3 The Public Interface for the Class `GenParticle`

The first element of the class declaration for `GenParticle` in `GenParticle.h` is a typedef:

```
typedef art::PtrVector<GenParticle> children_type;
```

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 ??.

The `GenParticle` class has two constructors:

```
GenParticle();
```

which is the default, and another that declares a few data members:

```
GenParticle( PDGCode::type          pdgId,
             art::Ptr<GenParticle> const& parent,
             CLHEP::Hep3Vector const& position,
             CLHEP::HepLorentzVector const& momentum );
```

It also has some private member functions with access to them provided by (public) accessor functions; accessors are discussed in Section 6.7.7. The accessor functions are given in Table 18.1.

The public interface has one non-const member function that is used to add a child particle to a `GenParticle`:

```
void addChild( art::Ptr<GenParticle> const& child );
```

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 }

```

This member function will not be discussed here but it will be discussed *in a future chapter that discusses the EventGenerator module in the toy experiment UPS package*. It is not needed when you are using `GenParticles` that already exist, as in this exercise, only when making new ones. For this exercise, you only need to concern yourself with the eight accessors. the remainder of this section will discuss each of them.

Finally there is a do-nothing definition of operator less-than; this is required to work around a bug in a tool named `genreflex` that is part of ROOT's persistency system. This operator should never be called because it does not do anything meaningful. Once the bug is fixed, we will remove this operator.

The member function `position` returns the position at which the particle was generated. The return type is `CLHEP::Hep3Vector` which was described in Section 18.5.3; additional details can be found in Section C.6.1.

The member function `momentum` returns the 4-momentum of the particle at the place of its creation. The return type is `CLHEP::HepLorentzVector` which was described in

Section 18.5.3; additional details can be found in in Section C.6.2.

As was described in Table 3.2, the position is given in mm and the 4-momentum is given in MeV.

The member function `pdgId` returns the value of the `PDGCode`; the next exercise, in Chapter ??, will describe the type `PDGCode : type` in more detail.

The member function `children` returns a `const` reference to the collection of smart pointers to the children of this `GenParticle`. For this exercise you only need to know one thing: `art::PtrVector` has a member function:

```
std::size_t size() const;
```

that returns the size of the collection, which corresponds to the number of children of the `GenParticle`.

The member function `hasParent` returns `true` if the `GenParticle` has a parent. In the `toyExperiment`, this will only be true for the daughters of the ϕ meson; all others are primary particles.

The member function `hasChildren` returns `true` if the `GenParticle` has children. In the `toyExperiment`, this will only be true for the ϕ meson.

You will notice that the implementations of the member functions `hasParent` and `hasChildren` are simple, almost trivial, logic on data members. So why bother to provide them? The reason is that coding

```
if ( gen.hasParent() ) { ... }
```

makes the programmer's intent manifest, while coding

```
if ( gen.parent().isNonnull() ) { ... }
```

requires that the reader draw some context dependent inferences to understand the programmer's intent.

The remaining two member functions have to do with parent-child navigation and will be described in Chapters ?? and ??.

In `GenParticle`, all accessors follow the recommendation that accessors for large objects return their information by `const` reference while accessors for small objects return

their information by value. With very few exceptions, the data product classes in the toy-Experiment and in the Workbook follow this pattern.¹

.

The final part of the public interface comes after the declaration of the class.² It is the free function, called the stream insertion operator:

```
std::ostream& operator<<(std::ostream& ost, const tex::GenParticle& genp );
```

If `gen` is an object of type `GenParticle`, this allows you to print it using the syntax:

```
std::cout << gen << std::endl;
```

18.5.4 Conditionally Excluded Sections of Header File

ROOT has a tool called `genreflex` that is used to generate ROOT dictionaries. These are needed for all classes that are part of a data product but *art* does not need a complete dictionary: it only needs to know about the data members, the default constructor and, if present, the destructor. To reduce both the time needed to create the dictionaries and the memory footprint of the dictionaries, it makes sense to hide code from `genreflex` that it doesn't need. In addition, `genreflex` does not understand C++11 syntax; therefore C++11 features in a data product class must be hidden from `genreflex`. To enable this exclusion, `genreflex` sets a C-preprocessor macro `__GCCXML__` that can be used to hide portions of code from `genreflex`. Two sections of `GenParticle.h` are conditionally marked for exclusion from `genreflex` in this way:

```
#ifndef __GCCXML__
    // Code to be excluded
#endif
```

A future chapter will describe a use case for which it is useful to generate complete dictionaries and will describe how to do so.

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.

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.

18.6 The Module `LoopGens1`

The new features in this module are:

1. The class declaration declares two new histogram pointers, `hP_` and `hcz_`.
 - (a) `hP_` will point to a histogram of the magnitude of the 3-momentum of each generated particle.
 - (b) `hcz_` will point to a histogram of the cosine of the polar angle of the 3-momentum of each generated particle.
2. The constructor initializes these pointers to `nullptr`.
3. The `beginJob` member function creates the two histograms and sets `hP_` and `hcz_` to point to them.
4. The `analyze` member function loops over the generated particles and fills these histograms with one entry per particle.

Listing 18.2 shows the loop over the generated particles, found in the `analyze` member function of `LoopGens1_module.cc`. This loop uses a feature of C++ that was introduced in C++-11 called the *ranged for loop*. You can read about ranged for loops in a standard C++ text that is up-to-date with C++-11 features.

The simple picture of this ranged for loop is that the body of the loop will be executed once for each element in the collection specified by the expression to the right of the colon. On the first pass, `gen` will be a const reference to the first element in the collection. On the

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  }
```

second pass it will be a const reference to the second element in the collection and so on.

Had the first line of the loop been written,

```
1  for ( GenParticle gen : *gens ){
```

The code would have compiled and it would have produced the correct results. However, on each pass of the loop, `gen` would be a copy of the corresponding element of the collection. This copy is unnecessary and will cause the code to be slower.

Why is there a `*` in front of `gens`? In the opening line of a ranged for loop, the expression on the right hand side of the colon must be a collection type. Recall that the object named `gens` is a handle to a `GenParticleCollection`. Therefore it is not itself a collection type. On the other hand, the expression `*gens` dereferences the handle and returns a const reference to the pointee of the handle, which in this case is a `GenParticleCollection`. This *is* a collection type and therefore it can be put to the right of the colon. If you are unsure of how this works, the key is to know that, used in this context, the character “`*`” is known as the *dereferencing operator*. You can look this up in any standard C++ text. Then you can consult the header for `art::ValidHandle` to see that it obeys the standard conventions for the dereferencing operator.

What sort of collection types may be on the right hand side of the colon? Any standard library collection type is allowed. User-written collection types may also be used.

The first two lines of the body of the loop are

```
3  CLHEP::HepLorentzVector const& p4 = gen.momentum();
4  CLHEP::Hep3Vector const& p      = p4;
```

Recall that a reference behaves like a compile-time alias. So there is no run-time cost for these lines to be present. The code is written this way because the author thinks that it makes the loop body easier to read.

The last two lines of the loop body fill the histograms with the requested information. If you are not familiar with the accessor member functions of `Hep3Vector` consult Section C.6.1.

As a last comment, the line that defines the variable `p` uses a somewhat obscure feature of CLHEP: a `HepLorentzVector` vector has an implicit conversion operator

to a `Hep3Vector`. This operator returns a `const` reference to the space part of the `HepLorentzVector`. One might have considered writing this line differently, for example

```
1 CLHEP::Hep3Vector p = p4.vect();
```

While this version makes the intent more explicit, and perhaps makes it a little easier to understand, it also forces an unnecessary copy. So the authors of the Workbook elected to choose the more efficient, albeit more obscure, variant.

18.7 CMakeLists.txt

The file `CMakeLists.txt` is unchanged from the previous exercise.

18.8 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/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
```

The last step will make a figure like that shown in Figure 18.1. It will also write the figure to the file `output/loopGens1.png`. Appendix B has instructions on how to view the `.png` file interactively.

For those of you at Fermilab, it also has instructions on how to print the `.png` file.



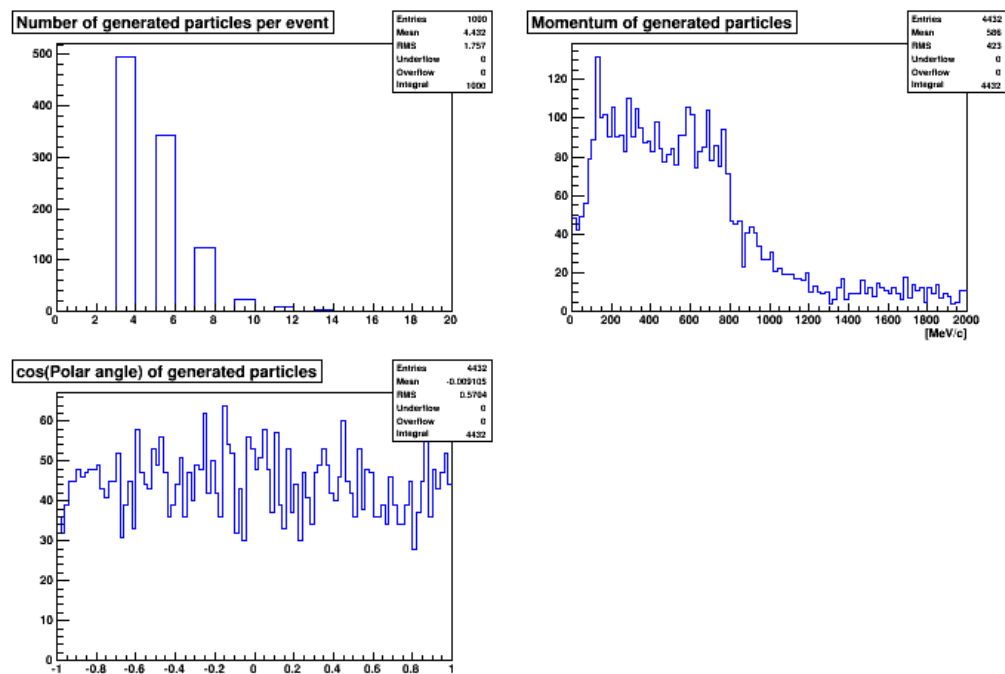


Figure 18.1: Histograms made by `loopGens1.fcl`

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 ranged 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
```

This will produce a page of histograms that should be identical to those made by `loopGens1.C`. It will also write the figure to the file `output/loopGens2.png`.

About the second item on the list. When a type is as simple as `GenParticle` there is little benefit to be gained by using `auto`; it may even obfuscate things a little. The benefit of `auto` arises when the correct type is a long complicated name that is hard to get right.

Regarding the third item on the list, the Workbook authors felt it would be confusing to use this form before showing you the first form. From now on we will use this form to get a `const` reference to its 3-momentum from a `GenParticle`.

18.9.2 `LoopGens3_module.cc`

The file `LoopGens3_module.cc` shows six other ways to write the loop over generated particles. These are shown because you will likely see them in code written by others. There is no need, at this time, to become proficient in their use; just know where to look them up if you encounter them.

This module makes seven histograms of the momentum spectrum of the generated particles. The first is made using the familiar `range` for loop and is provided for reference.

The second method uses operator square brackets:

```
for ( size_t i=0; i!=gens->size(); ++i ){
    CLHEP::Hep3Vector const& p = (*gens)[i].momentum();
    hP2_->Fill( p.mag() );
}
```

Recall that the type of `*gens` is a `GenParticleCollection`, which is a typedef for `std::vector<GenParticle>`. This example shows that a `std::vector` can be used as an array; as for an array, `(*gens)[i]` returns a reference to the i^{th} element in the collection.

The third method uses the `at` member function:

```
for ( size_t i=0; i!=gens->size(); ++i ){
    CLHEP::Hep3Vector const& p = gens->at(i).momentum();
    hP3_->Fill( p.mag() );
}
```

The member function `at(i)` is a run-time bounds-checked version of `[i]`. That is, `at(i)` first checks that the value of `i` is within the bounds of the `std::vector`. If it is, then `at(i)` returns a reference to the i^{th} element in the collection. If it is not, then `at(i)` will throw an exception; *art* will catch the exception and will attempt a graceful shutdown. The check for a valid value of `i` is performed at run-time during each call to `at(i)`; therefore the check will safely follow changes to the size of the collection while the program is running.

If you use the operator `[i]` with a value of `i` that is out of bounds, then the calling code will receive a reference to undefined memory. If you are lucky your code will crash quickly. If you are unlucky you will get subtly incorrect results. If you are really unlucky your Mars lander might join Romana in E-Space.

We strongly recommend that, when possible, you use a *ranged for* loop. If you need to use one of the indexed versions, we strongly recommend `at(i)` over `[i]`; only when code is very well tested and is proven to be time-critical should you switch to `[i]`.

The fourth method uses iterators:

```
GenParticleCollection::const_iterator i=gens->begin();
GenParticleCollection::const_iterator end=gens->end();
for ( ; i!=end; ++i ){
    CLHEP::Hep3Vector const& p = i->momentum();
    hp4_->Fill( p.mag() );
}
```

An iterator behaves like a pointer to the selected element. If you wish to learn about iterators, consult any standard C++ reference.

The fifth method also uses iterators:

```
for ( auto j=gens->begin(); j!=gens->end(); ++j ){
    CLHEP::Hep3Vector const& p = j->momentum();
    hp5_->Fill( p.mag() );
}
```

Here we have used the `auto` keyword to avoid typing:

```
GenParticleCollection::const_iterator
```

This is the first time in the Workbook that using `auto` has really paid off. All that we need to know about this iterator is that `j->` will do the right thing; we don't need to spend a lot of time spelling its type correctly.

We recommend the fifth method over the fourth method. In the fifth method, the scope of the loop iterator `j` is strictly inside the loop. In the fourth method two iterators, `i` and `end` are defined outside the scope of the loop. Their scope extends to the closing brace of the member function `analyze`. After the body of the loop has ended, this pollutes the remaining scope of `analyze` with identifiers that no longer serve any purpose.

The sixth method introduces the comma operator:

```
for ( auto j=gens->begin(), jend=gens->end(); j!=jend; ++j ){
    CLHEP::Hep3Vector const& p = j->momentum();
    hP6_->Fill( p.mag() );
}
```

This version says that it will evaluate, sequentially from left to right, the two comma-separated expressions as part of the initializer phase of the `for` loop. Compared to the previous version, this version needs to evaluate `gens->end()` only once, not once per loop iteration. Therefore this version is slightly more efficient.

Finally, the seventh version uses the `std::begin` and `std::end` free functions:

```
for ( auto j=std::begin(*gens), jend=std::end(*gens); j!=jend; ++j ){
    CLHEP::Hep3Vector const& p = j->momentum();
    hP7_->Fill( p.mag() );
}
```

The benefit to using this form only appears if you are writing templates — so whereas you are unlikely to use it yourself, you may see others using it.

The C++ language also supports `while` and `do-while` loops. There are appropriate places to use these features but writing a simple loop over a `std::vector` is not one of them. Why? These loop styles require that at least one identifier be defined outside of the scope of the loop. You can learn about these types of loops in any standard C++ reference.

The bottom line is this: if a *ranged for* loop will do, use it. Because this is a new feature in C++-11, legacy code will contain many loops written in the other ways.

To run this variation of the exercise, do the following in your build directory:

```
art -c fcl/LoopOverGens/loopGens3.fcl >& output/loopGens3.log
```

```
root -l fcl/LoopOverGens/loopGens3.C
```

This will make two files of histogram output:

```
output/loopGens3_1.png outout/loopGens3_2.png
```

Each file will show a 2×2 array of histograms; four on the first page, three on the second. On the first page, the upper left histogram is the histogram made by method 1, which should be the same as that from the previous two exercises. The remaining six histograms show, for each method, the difference between the histogram made by that method and the histogram made by method 1. All of these histograms should be flat lines at zero.

You can inspect the file `fcl/LoopOverGens/loopGens3.C` to learn how to use a CINT script to subtract two histograms. It will also show you how one CINT script can produce multiple png files.

18.9.3 LoopGens3a_module.cc

`LoopGens3a_module.cc` shows a minor variation on `LoopGens3_module.cc` that many people prefer. You will certainly see code written this way and you are free to write it this way if you find it more natural than the earlier version.

In `LoopGens3_module.cc` the first line of the member function `analyze` was:

```
auto gens = event.getValidHandle<GenParticleCollection>(gensTag_);
```

In `LoopGens3a_module.cc` this has been changed too:

```
auto const& gens(*event.getValidHandle<GenParticleCollection>(gensTag_));
```

You might find it easier to understand the second version if you break it into two steps:

```
auto gensHandle = event.getValidHandle<GenParticleCollection>(gensTag_);
GenParticleCollection const& gens(*gensHandle);
```

Note that both blocks of code end by defining a variable named `gens`, which is used by the subsequent code. The difference is that, in the first version you use `gens` as a pointer to the `const` data product whereas, in the second version, you use `gens` as a `const` reference to the data product.

In the remainder of the member function `analyze` there two sets of changes: `gens->` was changed to `gens.` and `(*gens)` was changed to `gens`.

18.10 Review

In this exercise you have learned:

1. How to use most of the accessor member functions of the class `GenParticle`.
2. How to use the stream insertion operator of `GenParticle`.
3. How to use some of the accessors functions of `CLHEP::Hep3Vector` and `CLHEP::HepLorent`.
4. CLHEP's use of `.icc` files to hold inline implementation.
5. Eight different ways loop over a `GenParticleCollection`.
6. That, for most purposes, the preferred way to write the loop is to use a ranged for loop.
7. That if the body of a loop is long, or if it will be repeated elsewhere, then put it into a separate function.
8. How to write a CINT script that writes multiple figure files.
9. How to write a CINT script that writes a multi-page PDF file.
10. How to use a CINT script to subtract two histograms.

18.11 Test Your Understanding

18.11.1 Test 1

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.

For each generated particle print to `cout`

1. The event number
2. The index of the `GenParticle` within the `GenParticleCollection`
3. The full information about the `GenParticle`; see the hint below.

Add a parameter set parameter to limit this printout to a number of events specified in the parameter set. The default should be for no printout.

Hint: To print all of the information about the `GenParticle` named `gen`, you can use the stream insertion operator (operator `<<`):

```
std::cout << gen << std::endl;
```

When you are ready to build your module, go to your build directory and give the command:

```
buildtool
```

This will compile your module and link it into a shared library.

Prepare a FHiCL file to run this module on the file `inputs/input04_data.fcl`. Be sure to pick a unique name for the histogram file created by the `TFileService`. Keep the FHiCL file in the source directory for this exercise.

To run your module, go to your build directory and give the command:

```
art -c fcl/LoopOverGens/<your-file-name>.fcl >& output/<your-file-name>.log
```

When you are done, you can compare your solution to that given in the files `LoopGens4_module.cc` and `loopGens4.fcl`. To run this module, be sure that you are in your build directory and give the command:

```
art -c fcl/LoopOverGens/loopGens4.fcl >& output/loopGens4.fcl
```

Compare the printout made by `LoopGens4_module.cc` with your printout. Compare the histograms made by `LoopGens4_module.cc` with your histograms; a discussion of the histograms will follow the discussion of the printout.

The first line of printout, when split over two lines, will look like:

```
Event:      1 GenParticle:      0 : [ pdg: 211 Position: (0,0,0)
    4-momentum (229.7,536.378,-81.9304;605.521) parent: none  children: none ]
```

The printout for each `GenParticle` is enclosed in square brackets and consists of:

1. The PDG Id code
2. The position at which the particle was created: (x, y, z)
3. The 4-momentum at which the particle was created: $(p_x, p_y, p_z; E)$
4. An identifier of the parent, or “none” if the particle has no parent
5. Identifiers of the children, in parentheses, or “none” if the particle has no children

You can inspect the printout to see that only five different PDG Id codes occur: ± 211 , ± 321 and 333. The next exercise, in Chapter ?? will explain how to interpret these codes. You will learn that the code 333 is for the ϕ meson, that ± 211 are for π^\pm and that ± 321 are for K^\pm . The meaning of the parent/child identifiers will be discussed in Chapter ??.

To help compare histograms you can make a multipage PDF file containing all of the reference histograms with the command:

```
root -l fcl/LoopOverGens/loopGens4.C
```

This will make a file named `output/loopGens4.pdf` that has 3 pages of histograms.

Appendix B has instructions on how to view a multipage PDF file interactively. For those of you at Fermilab it also has instructions on how to print it using the Fermilab printers.

You can also study the file `fcl/LoopOverGens/loopGens4.C` to learn how to use ROOT to make a multipage PDF file of histograms from a ROOT file.

18.11.2 Test 2

In the source directory for this exercise there are three files that end in `.cc.nobuild`. Each of these has an error expressly inserted. Your assignment is to find and fix each of the errors. Work on one of these at a time. The answers are given in Section 18.11.4.

To start work on the first file, go to the source directory for this exercise and give the command:

```
mv LoopGens5_module.cc.nobuild LoopGens5_module.cc
```

Then go to your build directory and give the command:

```
buildtool
```

This will start to build the module but will fail with a compiler error. The error message is short.

When you have found and fixed the error in `LoopGens5_module.cc`, follow the directions in the preceding paragraph for the file `LoopGens6_module.cc.nobuild`. When you build this module it will produce a very long error message. Scroll back to first few lines of the error message because they contain the most useful information. This is generally good advice when C++ produces very long error messages: the most useful information is in the first few lines.

If you have trouble finding the error and wish to try one of the other files, rename the `.cc` file back to `.cc.nobuild` and start to work on the other file.

What's with the `.nobuild` suffix? The build system looks at each file in each subdirectory and figures out what to do with the file based on the file type of the file; the file type is that last part of the filename, the part that comes after the dot. It has no information on what to do with files that end in `.nobuild` so it simply ignores them. When you rename the file to `.cc` the next run of `buildtool` will discover the file and attempt to build it.

18.11.3 Test 3

The code in `LoopGens8_module.cc` builds and runs but it produces incorrect output. To run the code and look at its output, go to your build directory and run the com-

mands:

```
art -c fcl/LoopOverGens/loopGens8.fcl >& output/loopGens8.log
```

```
browse output/loopGens8.root
```

Navigate to the histogram `gens/hP` and view it. The title says that it is the momentum spectrum of the generated particles; that is it should look like the corresponding histogram in any of the other root files from this exercise. But it looks very different. Find out why and fix it. The answers are given in Section 18.11.4.

The answers are intentionally placed on a new page (remember to try before you read further!).

18.11.4 Answers

In `LoopGens5_module.cc` the error is in the ranged 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 originally appeared as one long line, which has been broken into two for this document.

In `LoopGens6_module.cc` the error is again in the ranged 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 322. 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. The first line of this error message originally appeared as one long line, which has been broken into two for this document.

In `LoopGens7_module.cc` the error is in the body of the ranged for loop:

```
CLHEP::Hep3Vector& p = gen.momentum();
```

The error is that the type of `p` should be a const reference, not a non-const reference. The underlying reason is the same as for `LoopGens5_module.cc`; the immediate reason is that `gen` is itself a const reference so the code may only call const functions of `gen`. This time the error message is short and on point:

```
error: invalid initialization of reference of type 'CLHEP::Hep3Vector&'
from expression of type 'const CLHEP::HepLorentzVector'
    CLHEP::Hep3Vector& p = gen.momentum();
                        ^
```

An inferior solution would be to remove the ampersand and make `p` a copy of the space part of the generated momentum, not a const reference to it. This would allow the code to compile and run correctly but it will execute more slowly than it should.

In `LoopGens8_module.cc` the error is in the line

```
hP_-> Fill( gen.momentum().mag() );
```

The issue is that both `HepLorentzVector` and `Hep3Vector` have a function named `mag()`. The former returns the invariant mass of a 4-vector and the latter returns the magnitude of a 3-vector. The line in the above listing very clearly says to call the member function `mag` from the class `HepLorentzVector`, which was not the intent stated in the title of the histogram.

The preferred solution is write the code the way that `LoopGens2_module.cc` is written:

```
CLHEP::Hep3Vector const& p = gen.momentum();
hP_-> Fill( p.mag() );
```

This makes it clear that we want to call the `mag` function of a `Hep3Vector`. Also read the section that begins “As a last comment” on page 322.

19 Troubleshooting

19.1 Updating Workbook Code

If the remote machine that you log onto to run the Workbook exercises runs into problems during the setup procedure, it's possible that the admin for that machine has not installed the most recent versions of the Workbook code, or some dependent code. Contact the administrator.

19.2 XWindows (xterm and Other XWindows Products)

19.2.1 Mac OSX 10.9

The XWindows products, xterm, xclock and so on, likely reside in the directory `/opt/x11/bin/`. You will need to add this to your PATH. When your machine is connected to a second monitor, the XWindows products may not position properly on the screen. You may need to contact a Macintosh support person to configure the X11 setup properly so that it works with the multiple screen configuration.

At Fermilab, open a service desk ticket at <https://fermi.service-now.com/navpage.doService-Now>.



19.3 Trouble Building

If the buildtool doesn't seem to find your module, check that it can see it:

ls lib

Make sure it's not some other module that's causing the build failure.

19.4 *art* Won't Run

Make sure you're in your build window

Make sure the path to the FHiCL file is correct, e.g.,

```
art -c fcl/<dir-for-exercise>/<filename>.fcl
```


Part III

User's Guide

20 Obtaining Credentials to Access Fermilab Computing Resources

To request your Fermilab computing account(s) and permissions to log into the your experiment's nodes, fill out the form Request for Fermilab Visitor ID and Computer Accounts. Typically, experimenters that are not Fermilab employees are considered *visitors*. You will be required to read the Fermilab Policy on Computing.

After you submit the form, an email from the Fermilab Service Desk should arrive within a week (usually more quickly), saying that your Visitor ID (an identifying number), Kerberos Principal and Services Account have been created. You will need to change the password for both Kerberos and Services.

20.1 Kerberos Authentication

Your Kerberos Principal is effectively a username for accessing nodes that run Kerberos in what's called the `FNAL.GOV` *realm* (all non-PC Fermilab machines).¹

To change your Kerberos password, first choose one (minimum 10 characters with mixture of upper/lower case letters and numbers and/or symbols such as `!`, `,`, `#`, `$`, `?`, `&`, `*`, `%`). From your local machine, log into the machine using `ssh` or `slogin` and run the `kpasswd` command. Respond to the prompts, as follows:

```
$ kpasswd <username>@FNAL.GOV
```

```
    Password for username@FNAL.GOV: <--- type your current password here
```

¹The `FERMI.WIN.FNAL.GOV` realm is available for PCs.

```
New password:                <--- type your new password here

New password (again):         <--- type your new password here for c

Kerberos password changed.
```

Your Kerberos password will remain valid for 400 days.

20.2 Fermilab Services Account

The Services Account enables you to access a number of important applications at Fermilab with a single username/password (distinct from your Kerberos username/password). Applications available via the Services Account include SharePoint, Redmine, Service Desk, VPN and others.

To get your initial Services Account password, a user must first contact the Service Desk to get issued a first time default password. Once a default password is issued, users can access <http://password-reset.fnal.gov/> to change it.

If you are not on-site or connected to the Fermi VPN, call the Service Desk at 630-840-2345. You will be given a one-time password and a link to change it.

21 git

The source code for the exercises in the *art* workbook is stored in a source code management system called *git* and maintained in a repository managed by Fermilab. Think of *git* as an enhanced *svn* or (a VERY enhanced) *cvs* system. The repository is located at . You will be shown how to access it with *git*.

If you want some background on *git*, we suggest the Git Reference.

You will need to know how to install *git*, download the workbook exercise files initially to your system and how to download updates. You will not be checking in any code.

To install *git* on a Mac:

```
$ http://git-scm.com/download/mac
```

This will automatically download a disk image. Open the disk image and click on the .pkg file.

In your home directory, edit the file `.bash_profile` and add the line:

```
$ export PATH=/usr/local/git/bin:${PATH}
```

```
$ git clone ssh://p-art-workbook@cdcvns.fnal.gov/cvs/projects/art-workb
```

and how to download updates as the developers make them:

```
$ git pull
```

21.1 Aside: More Details about git

To bring your working copy of the workbook code up to date, you need to use `git`. Before describing the required `git` commands, we need to explain a little more about how `git` works and how the art-workbook team have chosen to use it. If you are familiar with `git` you can skip this section.

21.1.1 Central Repository, Local Repository and Working Directory

At any given time, there are three copies of the code that you need to be aware of, the central repository, your local clone of the central repository and the working copy of the code in your source directory.

1. The central git repository that contains all of the versions of the workbook is hosted by a machine named `cdcvcs.fnal.gov`¹ The art-workbook team updates this repository as it develops and maintains the exercises.
2. In section 10.4.1, in step 5b) you used the `git clone` command to make a copy of the central repository in your source directory. This clone contains a complete history of the development of art-workbook *as it existed at the time that you made the clone*. The local clone is found in the `.git` subdirectory of your source directory.
3. In section 10.4.1, in step 5d), you used the `git checkout` command to choose one of the tagged versions of art-workbook. This command looked into your local clone of the central repository, found all of the files in the requested version and put copies of them in the correct spot in the directory tree rooted at your source directory.

There are two other source code management systems that are widely used in HEP, `cvs` 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 contact the central repository but, when you are using `git`, you need only to contact your local clone of the central repository.



¹ Originally this machine hosted only `cvs` repositories, hence its name. It now hosts `cvs`, `svn` and `git` repositories.

To bring your working code up to date you need to do two steps:

1. Update your local clone of the central repository.
2. Checkout the new version from the local clone.

The discussion of the checkout has several cases. Each is discussed in one of the following sub-sections. It is possible for all four of these cases to occur on any given checkout.

21.1.1.1 Files that you have Added

When you worked on Exercise 2, you added some files to your working directories; for example you added the files `Second_module.cc` and `second.fcl`. When you do the checkout of the new version, these files will remain in your working directory and will not be modified; however the checkout command will generate some informational messages telling you that your working directories contain files that are not part of the checked out version.

You do not need to take any action; just be aware of the situation.

21.1.1.2 Files that you have Modified

Another case occurs for files that have the following properties:

1. They were part of the old version.
2. You have modified them.
3. They have not been modified in the central repository since you cloned the repository.

For example, suppose that you modified `first.fcl`; it is very unlikely that this file would have been modified in the central repository after you cloned the repository.

In this case, the checkout command will issue a warning message to let you know that your working version contains changes that are not part of the release you checked out.

You do not need to take any action; just be aware of the situation.

21.1.1.3 Files with Resolvable Conflicts

Another case occurs for files that have the first two properties from the list in Section 21.1.1.2 but which have been modified in the central repository since you cloned the repository. This will happen from time to time when we update exercises based on suggestions from users.

When this happens there are two cases, one of which is discussed here, while the other is discussed in the next sub-section.

If the two sets of changes (yours and those in the repository) are on different lines of the file `git`, will usually successfully merge these changes; `git` will then issue a warning message telling you what it has done.

It is your responsibility to identify these cases, understand the changes made in the repository and understand if `git` did the merge correctly.

21.1.1.4 Files with Unresolvable Conflicts

The final case is a variant of the previous case; it occurs when `git` is unable to automatically merge conflicting changes. This will happen when the changes you made and the changes made in the repository affect the same line, or lines, of code. When `git` does not know how to merge the changes it will give up, add markup to the offending files to mark the conflict and issue an error message. This leaves the offending files in an unusable state and you must correct the conflicts, by hand, before continuing.

The art-workbook has been designed so that this should happen very, very rarely. Most readers should bookmark this spot for future reference and only read it when they need to.

21.1.2 `git` Branches

`git` supports a concept known as *branches*. This is a very powerful feature that simplifies the task of having many developers collaboratively working on a single code base. Moreover, different experiments can choose to use branches in different ways; therefore a full description of branches is very open ended topic.

Fortunately, to use the workbook you do not need to know very much about branches; all that you need to know is summarized in Figure 21.1, which shows a simplified view of the way that the art-workbook team uses `git` branches. In Figure 21.1 time starts at the bottom




Figure 21.1: A figure to illustrate the idea of git branches, as used in the Workbook; the figure is described in the text.

of the figure and runs upward. The art-workbook team has adopted the convention that the most up to date version of the art-workbook code will always be found by checking out a branch named *develop*. In Figure 21.1 the vertical line represents the develop branch.

At the earliest time represented in Figure 21.1, the develop branch existed in some state that the art-workbook team liked. So they tagged the develop branch with the name `v_1`, for version 1. Shortly afterwards, the art team needed to add some improvements. To do this they did:

1. Use the `git pull` command to make sure that their local copy of the repository is up to date.

2. Check out the develop branch.
3. Start a new branch; in this example the new branch has the name `f1`, for “feature number 1”.
4. Do all the development work on this branch. When they change files and commit their changes, the changes stay local to the branch.
5. Once the new code has been tested it is merged back into the develop branch.
6. Use the `git pull` command to make sure that their local copy of the repository is up to date; in this example, no one else has made.
7. Finally the developer must push their local copy of the repository to the central repository.

This is illustrated in Figure 21.1 by the blue line labelled `f1`. While the developer is working on `f1`, he can change back and forth between the develop branch and the `f1` branch.

In Figure 21.1 the red line labelled `f2` represents a second feature that is added to the code base following the same pattern as the first.

In this example, the development team decided to tag the develop branch after the `f2` branch was merged back in; the tag was given the name `v2`; this is represented by the green filled circle in the figure.

The next items on the figure are the branches named `f3` and `f4`. In this example, someone started with the develop branch and began work on the feature `f3`. A little later someone else (or maybe the same person) started with the develop branch and began work on the feature `f4`. The person starting work on `f4` did so before the changes from `f3` were merged back into the develop branch; therefore the two branches `f3` and `f4` both start from the same place, the `v2` tag of develop. In this example, the next item on the timeline is that the developer of `f3` commits their changes back to the develop branch. Sometime after that the developer of `f4` merges their changes back. At this time the developer of `f4` has the responsibility to check for conflicts that occurred during the merge and fix them; this may or may not require consultation with the author of `f3`.

After this, the develop branch is again tagged, this time with a version named `v_3`.

The next items on the timeline are the branches named `f5` and `f6`. This example was

included to show that it is legal for `f6` both start and end during the time that `f5` is alive.

Finally, the `develop` branch is tagged one more time, this time with the name `v_4`.

In Figure 21.1 consider a time when both branches `f5` and `f6` are active.

Fortunately, to use the workbook you do not need to know very much about branches. You really need to know only two things.

1. In the art-workbook repository, there is a branch named *develop*; this branch is the head of the project.
2. When the art-workbook team decides that a new stable version of the code is available, they `checkout` the `develop` branch and then start a new branch. This new branch, called a release branch has the same name as the version number of the release. For example, the code for version `v0_00_13` is found in branch `v0_00_13`, and so on. New work, towards the next release, continues on the `develop` branch.

This is a bit of simplification but it captures the big ideas. Users of art-workbook should always work with one of the release branches; and they should always consult the documentation to learn which version of the code is matched to that version of the documentation.

Users of the art-workbook should never work in the `develop` branch; at any given time that branch may contain code that is still under development.

21.1.3 Seeing which Files you have Modified or Added

At any time you can check to see which files you have modified and which you have added. To do this, `cd` to your source directory and issue the `git status` command. Suppose that you have checked out version `v0_00_13`, modified `first.fcl` and added `second.fcl`. The `git status` command will produce the following output:

```
$ git status

# On branch v0_00_13
# Changes not staged for commit:
#   (use "git add <file>..." to update what will be committed)
```




Figure 21.2: A figure to illustrate the idea of git branches, as used in the Workbook; the figure is described in the text.

```
# (use "git checkout -- <file>..." to discard changes in
working directory)
#
# modified:   first.fcl
#
# Untracked files:
# (use "git add <file>..." to include in what will be committed)
#
# second.fcl
no changes added to commit (use "git add" and/or "git commit -a")
```

You should not issue the `git add` or `git commit` commands that are suggested above.

In the rare case that you have neither modified nor added any files, the output of `git status` will be:

```
$ git status
# On branch v0_00_13
```

22 *art* Run-time and Development Environments

22.1 The *art* Run-time Environment

Your *art* run-time environment consists of:

- your current working directory
- all of the directories that you can see and that contain relevant files, including system directories, project directories, product directories, and so on
- the files in the above directories
- the environment variables in your environment (not sure how to say this nicely)
- any aliases or shell functions that are defined

Figures 22.1, 22.2 and 22.3 show the elements of the run-time environment in various scenarios, and a general direction of information flow for job execution.

When you are running *art*, there are three environment variables that are particularly important:

- PATH
- LD_LIBRARY_PATH
- FHICL_FILE_PATH

They are colon-separated lists of directory names. When you type a command at the command prompt, or in a shell script, the (bash) shell splits the line using whitespace and the

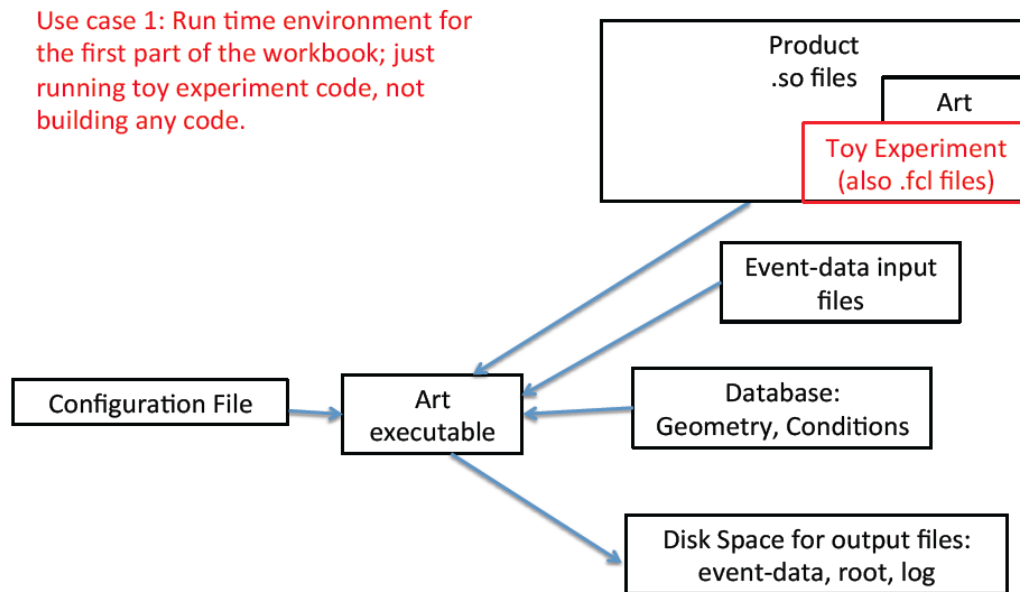


Figure 22.1: Elements of the *art* run-time environment, just for running the Toy Experiment code for the Workbook exercises

first element is taken as the name of a command. It looks in three places to figure out what you want it to do. In order of precedence:

1. it first looks at any aliases that are defined
2. secondly, it looks for shell keywords in your environment with the command name you provide
3. thirdly, it looks for shell functions in your environment with that name
4. then it looks for shell built-ins in your environment with that name
5. finally, it looks in the first directory defined in `PATH` and looks for a file with that name; if it does not find a match, it continues with the next directory, and so on, followed by the paths defined in the other two variables.

Some parts of the run-time environment will be established at login time by your login scripts. This is highly site-dependent. We will describe what happens at Fermilab - consult your site experts to find out if anything is provided for you at your remote site.

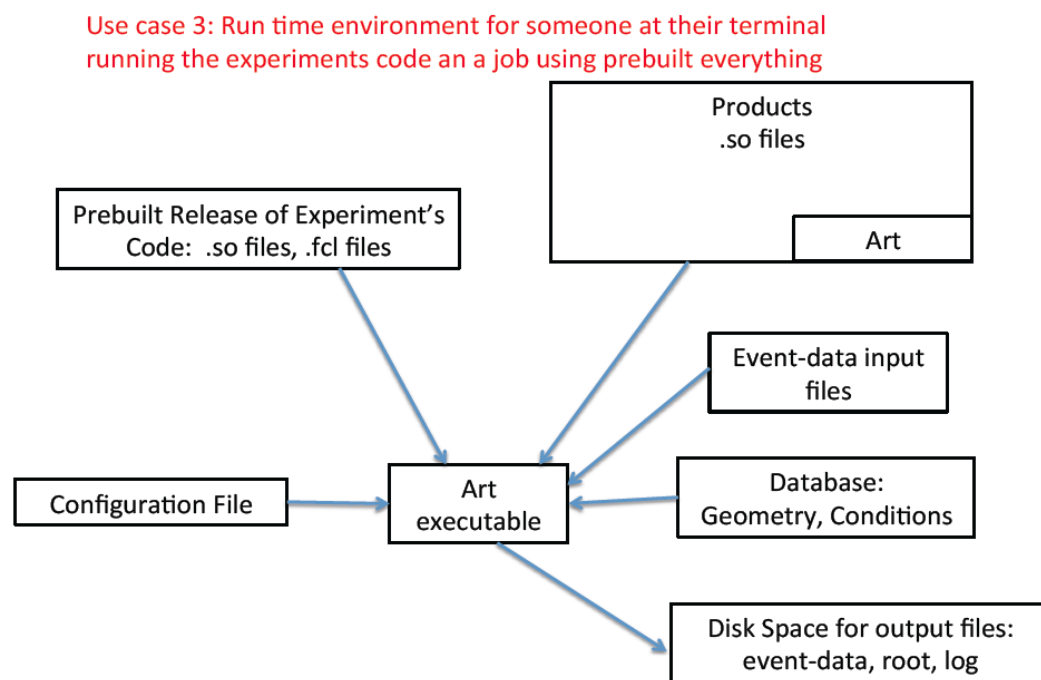


Figure 22.2: Elements of the *art* run-time environment for running an experiment's code (everything pre-built)

Use case 4: Run time environment for someone running a production job with officially tracked inputs.

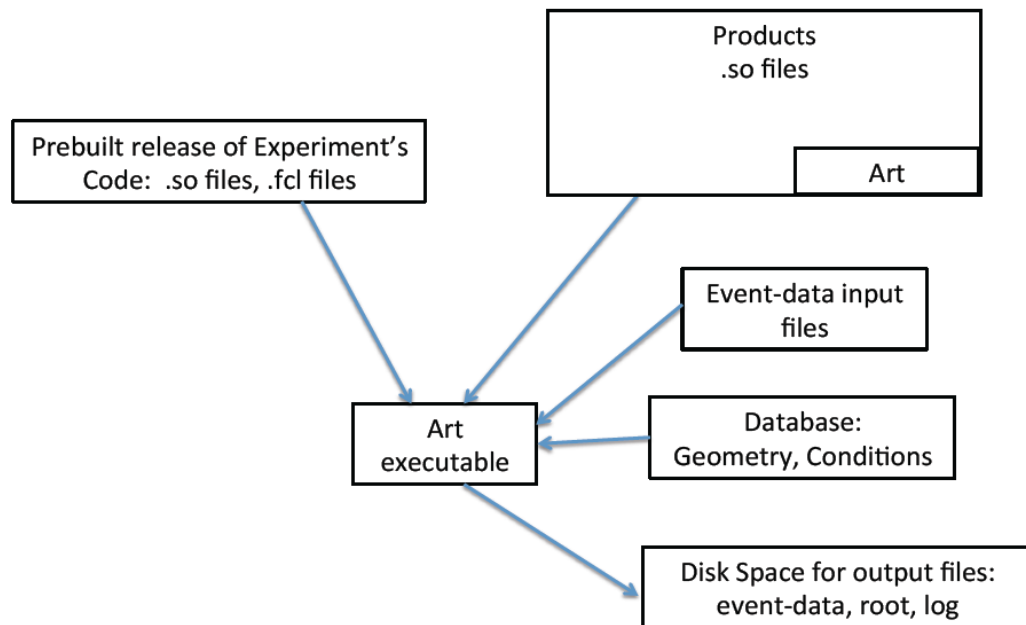


Figure 22.3: Elements of the *art* run-time environment for a production job with officially tracked inputs

When running the workbook, the interesting parts of your environment are established in two steps:

- source a site-specific setup script
- source a project-specific setup script

The Workbook, and the software suites for most IF experiments, are designed so that all site dependence is encoded in the site-specific setup script; that script adds information to your environment so that the project-specific scripts can be written to work properly on any site.

22.2 The *art* Development Environment

The development environment includes the run-time environment in Section 22.1 plus the following.

- the source code repository
- the build tools (these are the tools that know how to turn `.h` and `.cc` files in to `.so` files)
- additional environment variables and `PATH` elements that simplify the use of the above

Figures 22.4, 22.5 and 22.6 illustrate the development environment for various scenarios.

In some experiments the run-time and development environments are identical.

It turns out that there is no perfect solution for the job that build tools do. As a result, several different tools are widely used. Every tool has some pain associated with it. You never get to avoid pain entirely but you do get to pick where you will take your pain.

The workbook uses a build tool named **cetbuildtools**. Other projects have chosen `make`, `cmake`, `scons` and Software Release Tools (SRT). Here is something to watch out for: “build tools” written as two words refers generically to the above set of tools; but “build-tools” written as one word is the name of the executable that runs the build for **cetbuildtools**.

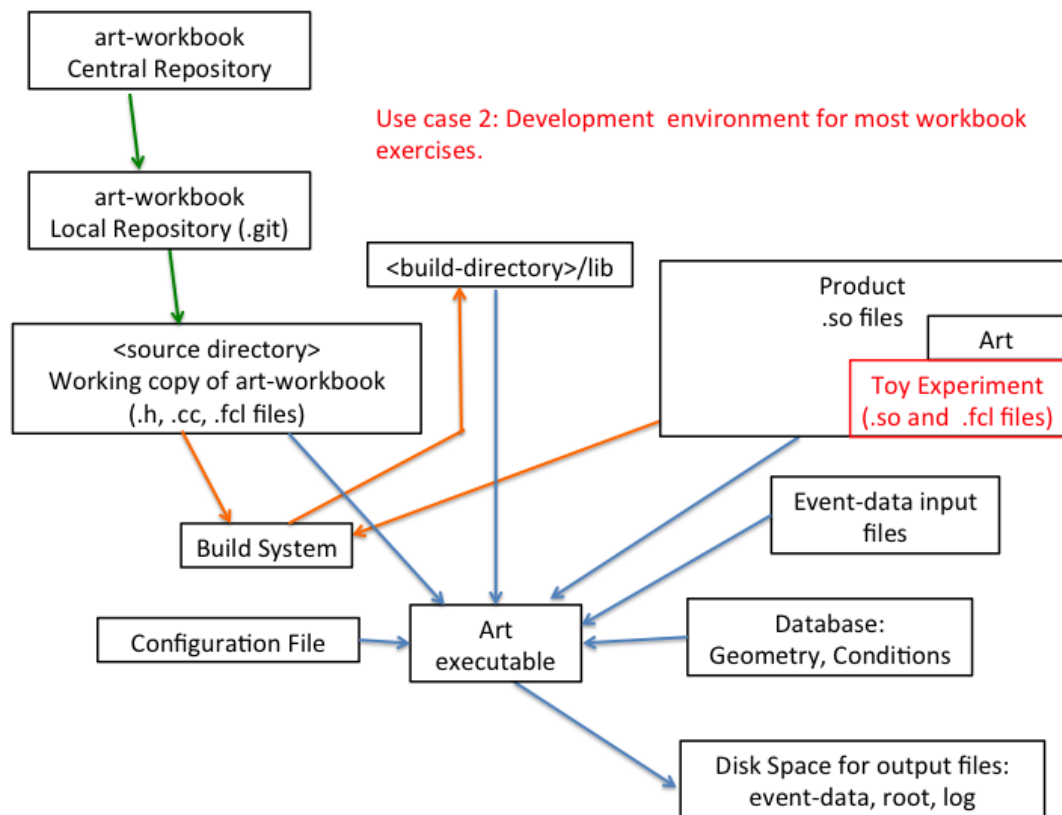


Figure 22.4: Elements of the *art* development environment as used in most of the Workbook exercises

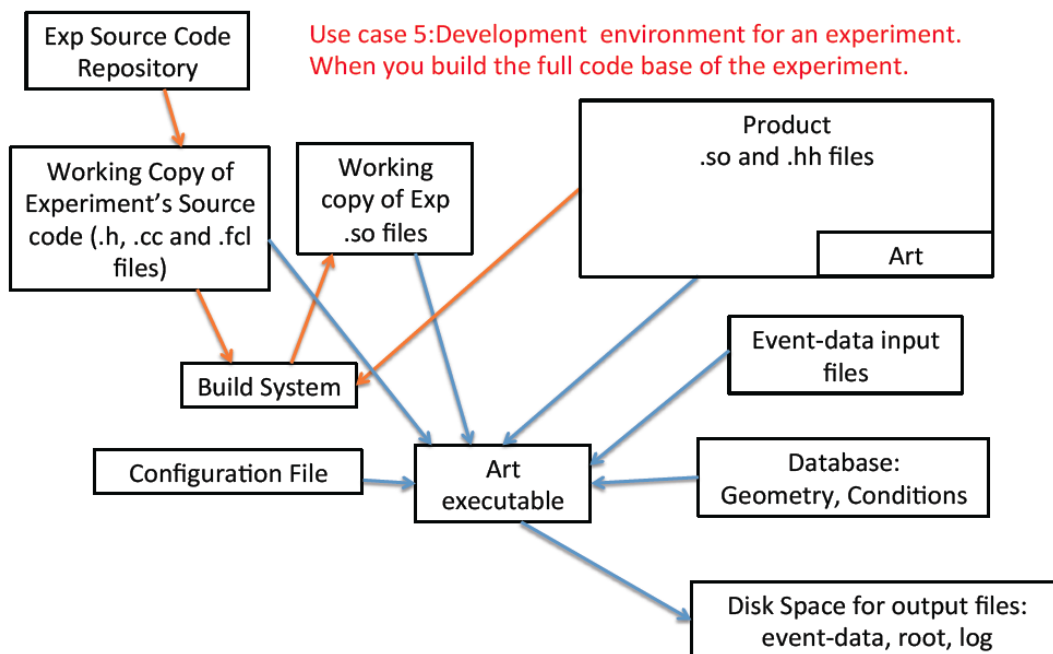


Figure 22.5: Elements of the *art* development environment for building the full code base of an experiment

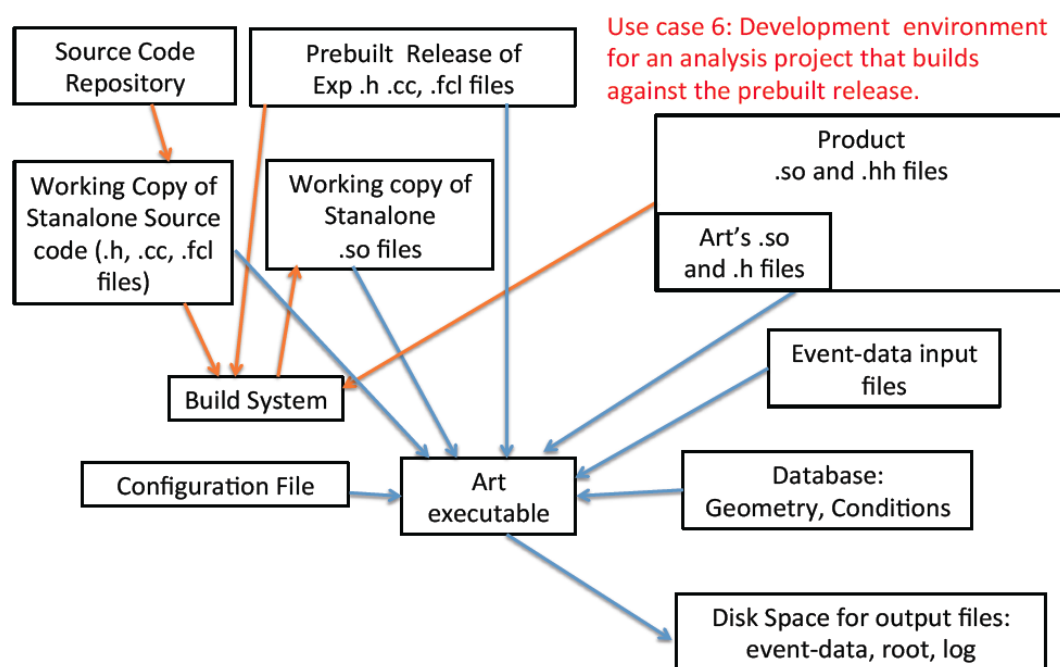


Figure 22.6: Elements of the *art* development environment for an analysis project that builds against prebuilt release

23 *art* Framework Parameters

This chapter describes all the parameters currently understood by the *art* framework, including by framework-provided services and modules. The parameters are organized by category (module, service or miscellaneous), and preceded by a general introduction to the expected overall structure of an *art* FHiCL configuration document.

23.1 Parameter Types

The parameters are described in tables for each module. The type of a defined parameter may be:

- **TABLE:** A nested parameter set, e.g., `set: { par1: 3 }`
- **SEQUENCE:** A homogeneous sequence of items, e.g., `list: [1, 1, 2, 3, 5, 8]`
- **STRING:** A string (enclosing double quotes not required when the string matches `[A-Za-z_][A-Za-z0-9_]*`). (Note: Identifiers when quoted do not function as special identifiers.) E.g.,

```
simpleString: g27
harderString: "a-1"
sneakysting1: "nil"
sneakysting2: "true"
sneakysting3: "false"
```

- **COMPLEX:** A complex number; e.g., `cnum: (3, 5)`
- **NUMBER:** A scalar (integer or floating point), e.g., `num: 2.79E-8`

- **BOOL:** A boolean, e.g.,

```
tbool: true
fbool: false
```

23.2 Structure of *art* Configuration Files

The expected structure of an *art* configuration file

Note, any parameter set is optional, although certain parameters or sets are expected to be in particular locations if defined.

```
# Prolog (as many as desired, but they must all be contiguous with only
# whitespace or comments inbetween.
```

```
BEGIN_PROLOG
```

```
pset:
```

```
{
  nested_pset:
  {
    v1: [ a, b, "c-d" ]
    b1: false
    c1: 29
  }
}
```

```
END_PROLOG
```

```
# Defaulted if missing: you should define it in most cases.
```

```
process_name: PNAME
```

```
# Descriptions of service and general configuration.
```

```
services:
```

```
{
  # Parameter sets for known, built-in services here.
  # ...
}
```

```
# Parameter sets for user-provided services here.
user:
{
}

# General configuration options here.
scheduler:
{
}
}

# Define what you actually want to do here.
physics:
{
  # Parameter sets for modules inheriting from EDProducer.
  producers:
  {
    myProducer:
    {
      module_type: MyProducer
      nested_pset: @local::pset.nested_pset
    }
  }
}

# Parameter sets for modules inheriting from EDFilter.
filters:
{
  myFilter: { module_type: SomeFilter }
}

# Parameter sets for modules inheriting from EDAnalyzer.
analyzers:
{
}
```

```
# Define parameters which are lists of names of module sets for
# inclusion in end_paths and trigger_paths.

p1: [ myProdroducer, myFilter ]
e1: [ myAnalyzer, myOutput ]

# Compulsory for now: will be computed automatically in a future
# version of ART.

trigger_paths: [ p1 ]
end_paths: [ e1 ]
}

# The primary source of data: expects one and only one input source
parameter set.
source:
{
}

# Parameter sets for output modules should go here.
outputs:
{
}
```

23.3 Services

23.3.1 System Services

These services are always loaded regardless of whether a configuration is specified.

23.3.2 FloatingPointControl

These parameters control the behavior of floating point exceptions in different modules.

Table 23.1: *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	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:
	reportSettings	BOOL	false	
	moduleNames	SEQUENCE	[]	
<module-name>	enableDivByZeroEx	BOOL	false	
	enableInvalidEx	BOOL	false	
	enableOverFlowEx	BOOL	false	
	enableUnderFlowEx	BOOL	false	

23.3.3 Message Parameters

These parameters configure the behavior of the message logger (this is a pseudo-service – not accessible via `ServiceHandle`).

Table 23.2: *art* Message Parameters

Enclosing Table Name	Parameter Name	Type	Default	Notes
services	message	TABLE		Top-level parameter set for the service
message				

23.3.4 Optional Services

These services are only loaded if a configuration is specified (although it may be empty).

23.3.5 Sources

23.3.6 Modules

Output modules

24 Job Configuration in *art*: FHiCL

Run-time configuration for *art* is written in the Fermilab Hierarchical Configuration Language (FHiCL, pronounced “fickle”), a language that was developed at Fermilab to support run-time configuration for several projects, including *art*. For this reason, this chapter will need to discuss FHiCL both as a standalone language and as used by *art*.

By convention, the names of FHiCL files end in `.fcl`. Job execution is performed by running *art* on a FHiCL configuration file, which is specified via an argument for the `-c` option:

```
$ art -c run-time-configuration-file.fcl
```

See Figure 22.1 in Section 22.1 to see how the configuration file fits into the run-time environment.

The FHiCL concept of *sequence*, as listed in brackets `[]`, maps onto the C++ concept of `std::vector`, which is a sequence container representing an array that can change in size. Similarly, the FHiCL idea of *table*, as listed in curly brackets `{ }`, maps onto the idea of `fhicl::ParameterSet`. Note that `ParameterSet` is not part of *art*; it is part of a utility library *used* by *art*, FHiCL-CPP, which is the C++ toolkit used to read FHiCL documents within *art*. FHiCL files provide the parameter sets to the C++ code, specified via module labels and paths, that is to be executed.

24.1 Basics of FHiCL Syntax

24.1.1 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'

name7  : 1 name8 : 2          # Two definitions on a line, separated by
                                # whitespace.

name9  :                      # Same as name9:3 ; newlines are just
:                                           # whitespace, which is not important.
```

3

```
namea : [ abc, def, ghi, 123 ]  # A sequence of atomic values.
                                   # FHiCL allows heterogeneous
                                   # sequences, which are not,
                                   # however, usable via the C++ API.

nameb :                          # A table of definitions; tables may nest.
{
    name0: 456
    name1: [7, 8, 9, 10 ]
    name2:
    {
        name0: 789
    }
}

namec : [ name0:{ a:1 b:2 } name1:{ a:3 c:4 } ]
                                   # A sequence of tables.

named : []                       # An empty sequence
namee : {}                       # An empty table

namef : nil                      # An atomic value that is undefined.

abc : 1                          # If a definition is repeated twice within
abc : 2                          # the same scope, the second definition
def : [ 1, 2, 3 ]                # will win (e.g., "abc" will be 2 and
def : [ 4, 5, 6 ]                # "def" will be [4,5,6])
name : {
    abc : 1
    abc : 2
}
```

```

cont1:{x: 1.0 y: 2.0 z: 3.0} # Hierarchical (compound) names denote
cont1.x : 5                  # levels of scope; here set x in cont1 t
OR
cont2:[1, 2, 3]
cont2[0] : 1                 # Here, redefine the first (atomic) valu
                             # for cont2, assign it the value 1. I.e.
                             # no action. Indices of PHiCL sequences
                             # begin with 0. \fixme{right?}

name0:{ a:1 b:2 }
x : @local::name0.a         # Using reference notation "@local," this assign
                             # to xthe value of a in table name0, in the
                             # line above, this value is 1.

```

24.1.2 FHiCL-reserved Characters and Identifiers

Several identifiers, characters and strings are *reserved to* FHiCL. What does this mean? Whenever FHiCL encounters a *reserved* string, FHiCL will interpret it according to the *reserved* meaning. Nothing prevents you from using these reserved strings in a name or value, but if you do, it is likely to confuse FHiCL. FHiCL may produce an error or warning, or it may silently do something different than what you intended. Bottom line: don't use reserved strings or symbols in the FHiCL environment for other than their intended uses.

The following characters, including the two-character sequence `::`, are reserved to FHiCL:

```
, : :: @ [ ] { } ( )
```

The following strings have special meaning to FHiCL. They can be used as parameter values to pass to classes, e.g., to initialize a variable within a program, but their uses will not be fully described here because of subtleties and variations. As you work with C++ and FHiCL, the way to use them will become clearer.

true, false These values convert to a boolean

nil This value is associated with no data type. E.g. if `a : nil`, then `a` can't be converted

to any data type, and it must be redefined before use

infinity, +infinity, -infinity These values initialize a variable to positive (the first two) or negative (the third) infinity

BEGIN_PROLOG, END_PROLOG ()

The first six strings (three lines) above function as identifiers reserved to *art* only when entered as lower case and unquoted; the last two strings (the last line) are reserved to *art* only when they are in upper case, unquoted and at the start of a line. Otherwise these are just strings. You may include any of the above reserved characters and identifiers in a “quoted” string to prevent them from being recognized as reserved to *art*.

24.2 FHiCL Identifiers Reserved to *art*

FHiCL supports run-time configuration for several projects, not only for *art*. *art* reserves certain FHiCL names as identifiers that it uses in well-defined ways. (Other projects may use FHiCL names differently.) Within FHiCL files used by *art*, these FHiCL names obey scoping rules similar to C++. These identifiers appear in the FHiCL file with a scope, i.e.,

```
identifier : {
...
}
```

if they define a list of modules or a processing block, or with square brackets

```
identifier :[
...
]
```

if they define a list of paths.

The following is a list of the identifiers reserved to *art* and their meanings. In the outermost scope within a FHiCL file, the following can appear:

process_name A user-given name to identify the configuration defined by the FHiCL file (it is recommended to make it similar to the FHiCL file name). This must appear at the top of the file. It may not contain the underscore character (_).

source Identifies the data source, e.g., a file in ROOT format containing HEP events.

services Identifies ...

physics Identifies the block of code that configures the scientific work to be done on every event (as contrasted with the “bookkeeping” portions).

outputs List of output modules.

The following may appear within the `physics` scope:

producers Specifies the configurations of producer modules

analyzers Specifies the configuration of analyzer modules

filters Specifies the configuration of filter modules

trigger_paths Sequence of pathnames; the paths named here may contain only producer and/or filter modules.

end_paths Sequence of pathnames; the paths named here may contain only analyzer and output modules.

The last two elements specify which of the modules will be executed in an *art* job. (It is legal for a module to be configured but not to be executed.) To understand order of execution see Sections 24.4 and 24.7.

The identifier `process_name` is really only reserved to *art* within the outermost scope (but it would seem to be needlessly confusing to use `process_name` as the name of a parameter within some other scope). The names `trigger_paths` and `end_paths` are artifacts of the first use of the CMS framework, to simulate the several hundred parallel paths within the CMS trigger; their meaning should be come clear after reading the remainder of this page.

24.3 Structure of a FHiCL Run-time Configuration File for *art*

Here is a sample FHiCL file called `ex01.fcl` that will do a physics analysis using the code in the *art* module `Ex01_module.so` (the object file of the C++ source file

Ex01_module.cc). In this configuration, *art* will operate sequentially on the first three events contained in the source file `inputFiles/input01_data.root`.

```
#include "fcl/minimalMessageService.fcl"

process_name : ex01

source : {
  module_type : RootInput
  fileNames   : [ "inputFiles/input01_data.root" ]
  maxEvents   : 3
}

services : {
  message : @local::default_message
}

physics :{
  analyzers: {
    hello : {
      module_type : Ex01
    }
  }

  e1      : [ hello ]
  end_paths : [ e1 ]
}
```

Let's look at it step-by-step.

```
#include "fcl/minimalMessageService.fcl"
```

Similar to C++ syntax, this effectively replaces the ‘`#include`’ line with the contents of the named file. This particular file sets up a messaging service.

```
process_name : ex01
```

The value of the parameter `process_name` (ex01, here, the same as the FHiCL file name) identifies this *art* job. It is used as part of the identifier for data products produced in this job. For this reason, the value that you assign may not contain underscore (`_`) characters. If the `process_name` is absent, *art* substitutes a default value of “DUMMY.”

```
source : {
  module_type : RootInput
  fileNames   : [ "inputFiles/input01_data.root" ]
  maxEvents   : 3
}
```

This `source` parameter describes where events come from. There may be at most one source module declared in an *art* configuration. At present there are two options for choosing a source module:

module_type : RootInput *art*::Events will be read from an input file or from a list of input files; files are specified by giving their pathname within the file system.

module_type : EmptyEvent Internally *art* will start the processing of each event by incrementing the event number and creating an empty *art*::Event. Subsequent modules then populate the *art*::Event. This is the normal procedure for generating simulated events.

Here `RootInput` is used; the data input file, in ROOT format, is assigned to the variable `fileNames`. The `maxEvents` parameter says: Look at only the first three events in this file. (A value of -1 here would mean “read them all.”)

Note that if no source parameter set is present, *art* substitutes a default parameter set of:

```
source : {
  module_type : EmptyEvent
  maxEvents   : 1
}
```

See the web page about configuring input and output modules for details about what other parameters may be supplied to these parameter sets.

```
services : {  
  message : @local::default_message  
}
```

Before starting processing, this puts the message logger in the recommended configuration.

```
physics :{  
  analyzers: {  
    hello : {  
      module_type : Ex01  
    }  
  }  
}
```

In *art*, `physics` is the label for a portion of the run-time configuration of a job. It contains the “meat” of the configuration, i.e., the scientific processing instructions, in contrast to the more administrative or bookkeeping information. The `physics` block of code may contain up to five sections, each labeled with a reserved identifier (that together form a parameter set within the FHiCL language); the strings are *analyzers*, *producers*, *filters*, *trigger_paths* and *end_paths*. In our example it’s set to *analyzers*.

The `analyzers` identifier takes values that are FHiCL tables of parameter sets (this is true also for `filters` and `producers`). Here it takes the value `hello`, which is defined as a table with one parameter, namely `module_type`, set to the value `Ex01`. The setup defined a variable called `LD_LIBRARY_PATH`; *art* knows to match the value defined by the name `module_type` to a C++ object file with the name `Ex01_module`. so somewhere in the path defined by `LD_LIBRARY_PATH`.

We will expand on the `physics` portion of the FHiCL configuration in Section 24.5.

```
e1      : [ hello ]  
end_paths : [ e1 ]
```

24.4 Order of Elements in a FHiCL Run-time Configuration File for *art*

In FHiCL files there are very, very few places in which order is important. Here are the places where it matters:



- A `#include` must come before lines that use names found inside the `#include`.
- A later definition of a name overrides an earlier definition of the same name.
- The definition of a name resolved using `@local` needs to be earlier in the file than the place(s) where it is used.
- Within a trigger path, the order of module labels is important.

Here is a list of *a few places* (of many) where order does not matter. This list is by no means exhaustive.



- Inside the physics scope, the order in which modules are defined does NOT matter for filters and analyzers blocks. These blocks define the run-time configurations of *instances* of modules.
- The five *art*-reserved words that appear in the outermost scope of a FHiCL file can be in any order. You could put outputs first and process_name last, as far as FHiCL cares. It may be more difficult for humans to follow, however.
- Within the services block, the services may appear in any order.

Regarding trigger_paths and end_paths, the following is a conceptual description of how *art* processes the FHiCL file:

1. *art* looks at the trigger_paths sequence. It expands each trigger path in the sequence, removes duplicate entries and turns the result into an ordered list of module labels. The final list has to obey the order of each contributing trigger path, but there are no other ordering constraints.
2. It does the same for the end_paths sequence but there is no constraint on order.
3. It makes one big sequence that contains everything in 1 followed by everything in 2.

4. It looks throughout the file to find parameter sets to match to each module label in the big list in 3.
5. It gives warning messages if there are left over parameter set definitions not matched to any module label in 3.
6. It then parses the rest of the physics block to make a “dictionary” that matches module labels to their configuration.

A conceptual description for the processing of services is as follows:

1. *art* first makes a list of all services, sorted alphabetically.
2. It makes a dictionary that matches service names to their parameter sets. A corollary is that service names must be unique within an *art* job. 
3. *art* has some “magic” services that it knows about internally. It loads the `.so` file for each of them and constructs the services.
4. It loads the `.so` files for all of the services and calls their constructors, passing each service its proper parameter set.
5. It works through its list of modules in 5 - it loads the `.so` and calls the constructor, passing the constructor the right parameter set.
6. It gives warning messages if there are left-over parameter set definitions not matched to any module label in 3. 

When one service relies on another, things get a bit more complicated. If service A requires that service B be constructed first, then the constructor of service A must ask *art* for a handle to service B. When this happens, *art* will start to construct service A since it is alphabetically first. When the constructor of A asks for a handle to B, *art* will interrupt the construction of service A, construct service B, and return to finish service A. Next, *art* will see that the next thing in the list is B, but it will notice that B has already been constructed and will skip to the next one.

Got that? Whew!

24.5 The *physics* Portion of the FHiCL Configuration

art looks for the experiment code in *art* modules. These must be referenced in the FHiCL file via *module labels*, which are just variable names that take particular values, as this section will describe. The structure of the FHiCL file – or a portion thereof – therefore defines the event loop for *art* to execute. The event loop, as defined in the FHiCL file, is collected into a scope labeled *physics*.

For a module label you may choose any name, as long as it is unique within a job, contains no underscore (`_`) characters and is not one of the names reserved to *art*. In the sample physics scope code below, we define `aProducer`, `bProducer`, `checkAll`, `selectMode0` and `selectModel` as module labels.

```
physics: {

  producers : {
    aProducer: { module_type: MakeA }
    bProducer: { module_type: MakeB }
  }

  analyzers : {
    checkAll: { module_type: CheckAll }
  }

  filter : {
    selectMode0: {
      module_type: Filter1
      mode: 0
    }
    selectModel: {
      module_type: Filter1
      mode: 1
    }
  }
}
```

The minimum configuration of a module is:

```
<moduleLabel> : { module_type : <ClassName> }
```

for example, in our code above:

```
aProducer: { module_type: MakeA }
```

`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`.

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:

```
moduleLabel : {
  module_type : MakeA
  pname0 : 1234.
  pname1 : [ abc, def]
  pname2 : {
    name0: {}
  }
}
```

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.

Note that *paths* are lists of module labels, while the two reserved names, `trigger_paths` and `end_paths` are lists of paths.

24.6 Choosing and Using Module Labels and Path Names

For a module label or a path name, you may choose any name so long as it is unique within a job, contains no underscore (`_`) characters and is not one of the names reserved to *art*

(see Section 24.2.

Any name that is a top-level name inside of the `physics` parameter set is either a reserved name or the name of a path.

It is important to recognize which identifiers are module labels and which are path names in a FHiCL file. It is also important to distinguish between a class that is a module and instances of that module class, each uniquely identified by a module label.

art has several rules that were recommended practices in the old framework but which were not strictly enforced by that framework. *art* enforces some of these rules and will, soon, enforce all of them:

- A path may go into either the `trigger_paths` list or into the `end_paths` list, but not both.
- A path that is in the `trigger_paths` list may only contain the module labels of producer modules and filter modules.
- A path that is in the `end_paths` list may only contain the module labels of analyzer modules and output modules.

Analyzer modules and the output modules may be separated into different paths; that might be convenient at some times but it is not necessary. On the other hand, keeping trigger paths separate has real meaning.

24.7 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`; 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.
 - If a module in a trigger path throws, the default behaviour of *art* is to stop all processing and to shut down the job as gracefully as possible. *Art* can be configured, at run time, so that, for selected exceptions, it behaves differently. For example it can be configured to continue with the current trigger path, skip to the next trigger path, skip to the next event, and so on.
- On each event, execute all of the paths listed in the `end_paths`.
 - The module labels listed in `end_paths` are executed exactly once per event, regardless of how many paths there are in the `trigger_paths` and regardless of any filters that failed.

- If a module label appears multiple times among the end paths, it is executed only once. No warning message is given.
 - Even if all trigger_paths have filters that fail, all module labels in the end path will be run.
 - End_path is free to execute the modules in the end_path in any order.
 - If a module in the end_path throws, the default response of *art* is to make a best effort to complete all other modules in the end path and then to shutdown the job in an orderly fashion. This behaviour can be changed at run-time by adding the appropriate parameter set to the top level .fcl file.
- One can ask that an output module be run only for events that pass a given trigger_path; this is done using the SelectEvents parameter set,
 - At present there is no syntax to ask that an analyzer module be run only for events that pass or fail some of the trigger paths. A planned improvement to *art* is to give analyzer modules a SelectEvents parameter that behaves as it does for output modules.
 - If a path appears in neither the trigger_paths nor the end_paths, there is no warning given.
 - If a module label appears in no path, a warning will be given.

In the above there is a lot of focus on which groups of modules are free to be run in an arbitrary order. This is laying the groundwork for module-parallel execution: *art* is capable of identifying which modules may be run in parallel and, on a multi-core machine, *art* could start separate threads for each module. At present both ROOT and G4 are not thread-safe so this is not of immediate interest. But there are efforts underway to make both of these thread-safe and we may one day care about module-parallel execution; our interest in this will depend a great deal on the future evolution of the relative costs of memory and CPU.

For simple cases, in which there is one trigger path with only a few modules in the path, and one end path with only a few modules in the path, the extra level of bookkeeping is just extra typing with no obvious benefit. The benefit comes when many work groups wish to run their modules on the same events during one *art* job; perhaps this is a job skimming off

many different calibration samples or perhaps it is a job selecting many different streams of interesting Monte Carlo events. In such a case, each work group needs only to define their own trigger path and their own end path, without regard for the requirements of other work groups; each work group also needs to ensure that their paths are added to the `end_paths` and `trigger_paths` variables. Art will then automatically, and correctly, schedule the work without redoing any work twice and without skipping work that must be done. This feature came for free with art and, while it imposes a small burden for novice users doing simple jobs, it provides an enormously powerful feature for advanced users. Therefore it was retained in art when some other features were removed.

24.8 Scheduled Reconstruction using Trigger Paths

Consider the following problem. You wish to run a job that has:

- Two producers `MakeA_module.cc` and `MakeB_module.cc`. You want to run both producers on all events.
- One analyzer module that you want to run on all events, `CheckAll_module.cc`.
- You have a filter module, `Filter1_module.cc` that has two modes, 0 and 1; the mode can be selected at run time via the parameter set.
- You wish to write all events that pass mode 0 of the filter to the file `file0.root` and you wish to write all events that pass mode 1 of the filter to `file1.root`

Here is code that would accomplish this:

```
process_name: filter1

source: {
    # Configure some source here.
}

physics: {

    producers : {
        aProducer: { module_type: MakeA }
    }
}
```

```
    bProducer: { module_type: MakeB }
  }

  analyzers : {
    checkAll: { module_type: CheckAll }
  }

  filter : {
    selectMode0: {
      module_type: Filter1
      mode: 0
    }
    selectModel: {
      module_type: Filter1
      mode: 1
    }
  }
}

mode0: [ aProducer, bProducer, selectMode0 ]
model: [ aProducer, bProducer, selectModel ]
analyzermodes: [ checkAll ]
outputFiles: [ out0, out1 ]

trigger_paths : [ mode0, model ]
end_paths : [ analyzermodes, outputFiles ]
}

outputs: {
  out0: {
    module_type: RootOutput
    fileName: "file0.root"
    SelectEvents: { SelectEvents: [ mode0 ] }
  }

  out1: {
```

```

    module_type: RootOutput
    fileName: "file1.root"
    SelectEvents: { SelectEvents: [ model ] }
  }
}

```

Recall that the names `process_name`, `source`, `physics`, `producers`, `analyzers`, `filters`, `trigger_paths`, `end_paths` and `outputs` are reserved to *art*. The names `aProducer`, `bProducer`, `checkAll`, `selectMode0`, `selectMode1`, `out0` and `out1` are module labels, and these are names of paths: `mode0`, `mode1`, `outputFiles`, `analyzermodes`.

24.9 Reconstruction On-Demand

24.10 Bits and Pieces

What variables are known to *art*? `physics` (which has the five reserved identifiers: `filters`, `analyzers`, `producers`, `trigger paths` and `end paths`), what else? input file type `RootInput`

I know that trigger path are // different from end paths, they can contain different types of modules; // event gets frozen after trigger path.

art knows to match the value `defi`

ned by the name `'module_name'` to a C++ object `fi`

le with the name `module_name_module.so` somewhere in the path `defi`

ned by `LD LIBRARY PATH`.

Further information on the FHiCL language and usage can be found at the [mu2e FHiCL page](#).

Part IV

Appendices

A *art* Completion Codes

When *art* completes it prints a message to the log file:

```
Art has completed and will exit with status 0.
```

The number at the end of this line is also the exit status that *art* returns to the shell that issued the *art* command.

Inside the *art* program, the status code is represented by a 4-byte integer. However only the least significant byte is returned to the caller; the other bytes are simply discarded. It is done this way because unix specifies that return codes are only one byte in size. For example, if the last line of the *art* printout is an error code of 8001, the status code seen by the calling shell will be 65.

If you ran *art* interactively you can see the return code by typing the following as the first command after running *art*:

```
echo $?
```

Here the symbol `$?` is a bash variable that holds the status code returned by the previously executed command. If you write a shell script that runs *art*, then you can test *art*'s return code using `$?`.

Listing A.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 A.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 (
8004	68	an <code>std::bad_alloc</code> was caught in main

B Viewing and Printing Figure Files

In many of the exercises in the Workbook you will produce files that contain figures. These files will be in formats such as `.pdf`, `.png`, `.jpg`, and so on.



For those of you who are working on the Fermilab General Purpose Computing Facility (GPCF), this chapter has instructions on how to view these files interactively and how to print them.

If you are logged in to a Fermilab machine from offsite, viewing figures interactively may be too slow to be practical. It depends on the quality of the network connection between Fermilab and your site; and it depends on the details of the file you are viewing. In such cases, the best alternative is to copy the file to a local machine and view it using the tools available there.

B.1 Viewing Figure Files Interactively

On a GPCF machine there are two interactive commands that will allow you to view a figure file:

```
kpdf file.type  
display file.type
```

Despite its name, `kpdf` works on more than just `.pdf` files; it works on most types of graphics formats, as does `display`.

`kpdf` has a more intuitive interface for navigating multipage files: it has a left-hand sidebar with page thumbnails. To navigate a multipage file using `display`:

1. Click anywhere in the image; this will pop up a menu.

2. On this menu click on “File”; this will pop up another menu.
3. On this menu click on “Next” or “Former” to move forward and backward through the pages.

Another interactive command that is present on some Unix machines is,

```
acroread file.pdf
```

Some people prefer `acroread`’s rendering of postscript and PDF files to that of the other tools.

B.2 Printing Figure Files

A future version this document will explain how to print from the GPCF machines.

On most Fermilab machines PDF files can be printed from the command line or from a browser. Other graphics formats need to be viewed in a browser and printed from the browser or saved as PDF and then printed from the command line.

General information about printing at Fermilab is available at fermiprint.fnal.gov.



C CLHEP

C.1 Introduction

The wikipedia entry for CLHEP, <http://en.wikipedia.org/wiki/CLHEP>, describes it as:

CLHEP (short for A Class Library for High Energy Physics) is a C++ library that provides utility classes for general numerical programming, vector arithmetic, geometry, pseudorandom number generation, and linear algebra, specifically targeted for high energy physics simulation and analysis software. The project is hosted by CERN and currently managed by a collaboration of researchers from CERN and other physics research laboratories and academic institutions. According to the project's website, CLHEP is in maintenance mode (accepting bug fixes but no further development is expected).

The *art* Workbook uses CLHEP, as do many of the experiments that use *art*. In both the *art* run-time environment and the *art* development environment CLHEP is made available via UPS and is rooted at `$CLHEP_DIR`. The CLHEP header files can be found at `$CLHEP_INC` and the libraries can be found at `$CLHEP_LIB_DIR`. These environment variables will also be defined in the corresponding environments for your experiment.

This appendix will discuss those parts of CLHEP that are important for the *art* Workbook and will fill in some background information that is assumed by the CLHEP documentation but is not explicitly stated anywhere else.

CLHEP is divided in packages and the *art* Workbook uses classes from four of these packages:

Matrix Support for linear algebra.

Random Support for random engines and random distributions. The distinction between these two ideas will be discussed in *a section to be written in the future*.

Units Support for a standard set of units and for transformations among different units. It also provides the values of many physical constants.

Vector Support for 2-vectors, 3-vectors, 4-vectors.

C.2 Multiple Meanings of *Vector* in CLHEP

CLHEP uses the word *vector* in two different senses, both of which are different from the use of the word in the standard library template `std::vector`. The Matrix package supports linear algebra, by providing classes to represent matrices and vectors of arbitrary dimensions; the package supports operations such as matrix multiplication and the computation of the inverse, transpose, determinant and trace of a matrix. The Vector package, on the other hand, provides classes that represent a point on a plane, a point in 3-space or a point in 4-dimensional space-time; the package supports operations such as dot products, cross products, rotations and Lorentz transformations.



The Vector package does not make a distinction between positions, displacements, velocities and momentum. The same classes are used for all four.



C.3 CLHEP Documentation

The CLHEP home page is <http://proj-clhep.web.cern.ch/proj-clhep>.

The following is a direct link to the CLHEP documentation page:

<http://proj-clhep.web.cern.ch/proj-clhep/index.html#docu>

In many cases the documentation for CLHEP is simply the code or the comments in the code. You can view the header files by looking under `$CLHEP_INC`. You can view the source files by looking under `$CLHEP_DIR/source`. A more convenient format to view the header files is to use the CLHEP Doxygen site:

http://proj-clhep.web.cern.ch/proj-clhep/doc/CLHEP_2_1_3_1/doxygen/html/

Doxygen simply presents the information found in the header file in a format that is easier to view than the header file itself.

To get information about a CLHEP class, go to the Doxygen page, click on the tab named “Classes”, and use your browser’s search function to find the name the class.

C.4 CLHEP Header Files

C.4.1 Naming Conventions and Syntax

The syntax to include a CLHEP header file is:

```
#include "CLHEP/<package-name>/<filename>.h"
```

where `package-name` is the name of the CLHEP package to which the header file belongs and where `filename` is filename of the header file.

Almost all of the class names in CLHEP begin with the prefix `Hep`, for example `HepLorentzVector`, which is the CLHEP representation of a 4-vector. The handful of exceptions to this rule are helper classes used internally by CLHEP.

In most cases, the name of the header file for a class is the name of the class, excluding the leading `Hep`. For example, the header file for `HepLorentzVector` is `CLHEP/Vector/LorentzVector.h`. Some header files also contain the declarations of helper classes that are used by the main class. A few header files contain the declarations of several related classes.

There is one important header file that follows an unusual naming pattern. The header file `CLHEP/Vector/ThreeVector.h` declares the classes `Hep2Vector` and `Hep3Vector` that describe a point in a plane and a point in 3-space, respectively.

C.4.2 `.icc` Files

Following a convention in use during the mid 1990’s when the CLHEP package was developed, CLHEP header files contain only declarations. When inline implementations are required, CLHEP puts them in a file with the same name as the header file but with `.h` replaced with `.icc`, which stands for “inline cc”. The `.icc` file is included near the end of the `.h` file. For example the inline implementation for the class `Hep3Vector` is included in the header file `ThreeVector.h` as:

```
#include "CLHEP/Vector/ThreeVector.icc"
```

The convention of using `.i.cc` files to segregate inline implementations from declarations is no longer recommended, but CLHEP retains it for backwards compatibility. The recommended convention is to simply put inline implementation in the header file.



C.5 The CLHEP Namespace

All identifiers defined by CLHEP are in the `CLHEP` namespace.

C.5.1 `using` Declarations and Directives

An example of a *using* declaration is:

```
1 using CLHEP::Hep3Vector;
```

This tells the compiler that, when ever it sees `Hep3Vector`, it should automatically recognize that it means `CLHEP::Hep3Vector`.

An example of a *using* directive is:

```
1 using namespace CLHEP;
```

This tells the compiler that it should recognize all identifiers from the `CLHEP` namespace without the user having to type the prefix `CLHEP::` in front of every name.

Using declarations and directives are a part of C++ that you can read about in any standard text. Using `using` can reduce the amount of typing you have to do but over use of `using` defeats the entire purpose of namespaces.

One very important rule is that you must never code using directives or using declarations in header files. You should only use them in source files. *A future version of this document will refer to a complete explanation.*



.

For the particular case of CLHEP you should not use `using` declarations or directives even in source files. The remainder of this section explains why.

Consider what happens if you code:

```
1 using namespace CLHEP;
```

The CLHEP Units package defines many identifiers with commonly used short names, `m`, `g` and `s`; in addition there are 17 two character identifiers and 24 three character identifiers, such as `mm`, `m2`, `deg`, `cm3`, `amu`, `rad` and so on. Many of these short identifiers are commonly used in code, `m` for mass, `s` for an arc length, `rad` for a radius and so on. If you give a `using` directive for the namespace CLHEP then all of these short names will be defined with the scope of your code.

A common programming error is to forget to declare a variable before using it. Normally the compiler will recognize this error and issue a diagnostic message. If, on the other hand, one of your undeclared variables matches one of the CLHEP variable names, and if you have used `using namespace CLHEP`, then the compiler will not recognize the error and will not issue a diagnostic. You will need to find the error by tedious debugging.

A similar problem occurs if you code:

```
1  using CLHEP::Hep3Vector;
```

This introduces three one letter identifiers into the scope of your code: `X`, `Y` and `Z`, which are defined as the integers 0, 1, 2.

The bottom line is that CLHEP does not mix with `using` directives and `using` declarations.

C.6 The Vector Package

The *art* Workbook uses the following classes from the Vector package:

`Hep3Vector` A vector in 3-space.

`HepLorentzVector` A 4-vector in 4-dimensional space time.

`HepBoost` Performs Lorentz boosts from one inertial frame to another; it operates on objects of type `HepLorentzVector`.

C.6.1 `CLHEP::Hep3Vector`

The class `Hep3Vector` is used to represent any 3-vector in which each element is represented as `double`. There is no similar class in which each element is represented as a `float`. The header file is found at

Listing C.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 &);

```

```

1  $CLHEP_INC/CLHEP/Vector/ThreeVector.h

```

You can also browse the class header using Doxygen:

http://proj-clhep.web.cern.ch/proj-clhep/doc/CLHEP_2_1_3_1/doxygen/html/classCLHEP_1_1Hep3Vector.html

Because this is a particularly simple class, the header serves as complete documentation.

If you are a beginner to C++ this might not be enough. So this section will show a few of the commonly used features. The goal of this section is to guide you through the header file so that you will learn how to understand a typical CLHEP header file.

`Hep3Vector` has five constructors that are shown in Listing C.1. In each of the first four constructors the missing components are set to 0. If you check the `.icc` file you will see that all five are implemented inline. The `explicit` keyword is beyond the scope of this discussion.

There is a single `Hep3Vector` class that can be used for positions, momenta and velocities. So don't attach special meaning to the names of the arguments; the 4th constructor would have exactly the same behaviour had it been declared as:

```

1  Hep3Vector(double p_x, double p_y, double p_z);

```

or

```

1  Hep3Vector(double v_x, double v_y, double v_z);

```

There are many synonyms for access to the individual elements. If `v` is an object of type `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];

```

similarly for y and z . There is an `enum` near the top of the class declaration that defines the mapping of coordinates to indices

```
1 enum { X=0, Y=1, Z=2, NUM_COORDINATES=3, SIZE=NUM_COORDINATES };
```

The full names `CLHEP::Hep3Vector::X` are so long that it is common to see code in which the integer values are substituted by hand:

```
1 double x = v(0);
```

Why are there so many different ways of doing the same thing? `Hep3Vector` was designed as a common replacement for several other 3-vector packages so its public interface was designed as the union of all of the public interfaces. While you will see all of the above used in code written by others, you should prefer to use the first form. It is preferred over the last two because it executes more quickly. It is preferred over the second form for a stylistic reason; the *art* team recommends that accessors not contain the string “get”.

Why does documentation talk about all of these ways of doing the same thing? Because you are likely to encounter them as you read code written by others.

The header tells you that the above accessors functions exist because it contains:

```
1 double x() const;
2 double getX() const;
3 double & operator () (int);
4 double operator [] (int) const;
```

All of the accessors are implemented inline. The first two accessors are member functions, similar to many you have seen before. You do not need to learn all of the details about the last two accessors — you only need to recognize their use, as illustrated on the previous page.



If you do wish to learn more about the third accessor it will be described in any standard 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[]`, pronounced “operator square brackets”, which may be in a more general section on operator overloading.

If v is an object of type `Hep3Vector`, the following all set the value of the x component to 123.

```
1 v.setX(123.);
2 v(CLHEP::Hep3Vector::X) = 123.;
3 v[CLHEP::Hep3Vector::X] = 123.;
```

Similarly for setting the values of the y and z components. As before, prefer to use the first version because it is faster than the other two.

The header tells you that the above accessors functions exist because it contains:

```
1 void setX(double);
2 double & operator [] (int);
3 double & operator () (int);
```

You can also set the values of all three components at once. For example the following sets the value of v to a unit vector in the z direction:

```
1 v.set(0.,0.,1.);
```

The header tells you that this function exists because it contains:

```
1 void set( double x, double y, double z);
```

Table C.1 lists some of the commonly used member functions of `Hep3Vector`; it shows the declaration of each function and an example of how to use it. Some other member functions are shown in the text below; they do not fit nicely in the table so they are presented separately. In each case you will see a statement in standard mathematical notation, the name of the function, followed by a usage example and the the declaration. In the usage examples, u , v are all objects of type `CLHEP::Hep3Vector` and a is of type `double`.

$\vec{v} = \vec{u}$. The assignment operator:

```
1 v = u;
2 Hep3Vector & operator = (const Hep3Vector &);
```

$\vec{v} = \vec{v} + \vec{u}$. Addition:

```
1 v += u;
2 Hep3Vector& operator += (const Hep3Vector&);
```

Table C.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

$\vec{u} = a\vec{u}$. Multiplication by a scalar:

```
1 u *= a;
2 Hep3Vector& operator *= (double);
```

$\vec{u} = \vec{u}/a$. Division by a scalar:

```
1 u /= a;
2 Hep3Vector& operator /= (double);
```

$\vec{v} = -\vec{u}$. Unary minus:

```
1 v = -u;
2 Hep3Vector operator - () const;
```

$a = \vec{u} \cdot \vec{v}$. Dot product:

```
1 a = u.dot(v);
2 double dot(const Hep3Vector&) const;
```

$\vec{w} = \vec{u} \times \vec{v}$. Cross product:

```
1 w = u.cross(v);
2 Hep3Vector cross(const Hep3Vector&) const;
```

$\vec{v} = \hat{u}$. Computation of a unit vector parallel to \vec{v} .

```
1 v = u.unit();
2 Hep3Vector unit() const;
```

The C++ compiler knows how to use combinations of the above functions to allow you to write:

```
1 w = u + v;  
2 w = u - v;
```

which have the obvious meaning.

Hep3Vector has many other member functions that are not mentioned above. They provide accessors for rapidity and pseudo-rapidity; they provide setters that accept arguments in spherical polar coordinates or in cylindrical coordinates; they provide support for rotations. These functions are not used in the Workbook and will not be discussed here. Consult the header file or the Doxygen site for further information.

C.6.1.1 Some Fragile Member Functions

Hep3Vector has functions to ask if two Hep3Vector's are equal to each other, if they are near to each other, if they are parallel to each other or if they are orthogonal to each other. These functions are fragile and should be used with extreme care. If you think that you have a reason to use one of these functions, speak with an expert to learn if there is a better alternative.

Hep3Vector supports comparison for equality. The function is declared as:

```
1 bool operator == (const Hep3Vector &) const;
```

and is used as follows:

```
1 if ( u == v ){  
2     // do something  
3 }
```

This comparison suffers from the usual dangers about doing comparisons of floating point numbers for equality.

The `isNear` member function is declared as:

```
1 bool isNear (const Hep3Vector &, double epsilon=tolerance) const;
```

and is used as follows:

```

1 double epsilon=1.e-14; // for example
2 if ( u.isNear(v,epsilon ){
3     // do something
4 }

```

The notion of nearness is expressed as,

$$|\vec{v} - \vec{u}|^2 \leq \epsilon^2 |\vec{v} \cdot \vec{u}| \quad (\text{C.1})$$

The fragility comes from the second argument. This argument is optional and has a default value that is set to the value of a static member datum of the class, `CLHEP::Hep3Vector::tolerance`. This member datum can be modified by anyone. Therefore always specify the second argument and never let it take its default value. If you let your code take the default value then, if someone changes the value of tolerance, the behaviour of your code will change.

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;

```

which are used as:

```

1 if ( u.isParallel(v,tolerance );
2 if ( u.isOrthogonal(v,tolerance );

```

C.6.2 CLHEP::HepLorentzVector

The class `HepLorentzVector` is used to represent any 4-vector in 4-dimensional space-time. Each element is represented as `double`; there is no similar class in which each element is represented as a `float`. The header file is found at

```

1 $CLHEP_INC/CLHEP/Vector/LorentzVector.h

```

You can also browse the class header using Doxygen:

http://proj-clhep.web.cern.ch/proj-clhep/doc/CLHEP_2_1_3_1/doxygen/html/classCLHEP_1_1HepLorentzVector.

As for `Hep3Vector`, the only documentation for `HepLorentzVector` is the contents of the header, which can be viewed either as a file or by the Doxygen web page.

If you are not familiar with reading C++ header files you should read Section C.6.2 before reading this section. That section has more examples of how to translate a declaration in a header file into example code; this section presumes that you can do that yourself.

In the `HepLorentzVector` the order of components is given by the enum:

```
1  enum { X=0, Y=1, Z=2, T=3, NUM_COORDINATES=4, SIZE=NUM_COORDINATES };
```

`HepLorentzVector` has eight constructors that may be of interest to users of the Workbook:

```
1  HepLorentzVector();
2  HepLorentzVector(double x, double y, double z, double t);
3  HepLorentzVector(double x, double y, double z);
4  HepLorentzVector(double t);
5  HepLorentzVector(const Hep3Vector & p, double e);
6  HepLorentzVector(double e, const Hep3Vector & p);
7  HepLorentzVector(const HepLorentzVector & v);
8  HepLorentzVector(const Hep3Vector & v);
```

In each constructor, the missing components are set to zero. As for `Hep3Vector`, there is a single `HepLorentzVector` class that can be used for 4-positions, 4-momenta and 4-velocities. So don't attach special meaning to the names of the arguments.

In the following examples the objects `p`, `q`, `s` are of type `HepLorentzVector`, the objects named `u`, `v`, `w` are of `Hep3Vector` and the objects `a`, `b`, `x`, `y`, `z`, `t` are of type `double`.

`HepLorentzVector` has a set of accessor methods that are similar to those of `Hep3Vector`. The following code fragments return the value of the x component:

```
1  double x = p.x();
2  double x = p.px();
3  double x = p.getX();
4  double x = p(CLHEP::LorentzVector::X);
5  double x = p[CLHEP::LorentzVector::X];
```

Similarly for the y , z and t components. Prefer to use one of the first two member functions for same the reasons discussed for `Hep3Vector`.

`HepLorentzVector` has several accessors that return the space part of the 4-vector as a `Hep3Vector`:

```

1 Hep3Vector u = p.vect();
2 Hep3Vector u = p.v();
3 Hep3Vector u = p;
4 Hep3Vector u = p.getV();

```

Prefer to use one of the first two forms. The third form is perfectly correct but it may be difficult for readers of your code to understand its intended function (only real CLHEP experts will understand it). The fourth form executes more slowly than the others.

The following code fragments set the value of the x component to 123.:

```

1 p.setX(123.);
2 p.setPx(123.);
3 p(CLHEP::HepLorentzVector::X) = 123.;
4 p[CLHEP::HepLorentzVector::X] = 123.;

```

Similarly for the y , z and t components. Prefer to use the first two member functions as they execute more quickly than do the last two.

HepLorentzVector has several functions that allow you assign new values to several elements at the same time:

```

1 p = u;
2 p.setVect(u);
3 p.set(x, y, z, t);
4 p.set(u);
5 p.set(u,t);
6 p.setV(u);

```

Two of the above lines in three previous three blocks may be difficult to understand. The two lines:

```

1 Hep3Vector const& u = p; // Return the space part by const reference.
2 p = u;                  // Assignment of a 3-vector to the space part
3                          // of a 4-vector

```

are implemented, respectively, by the following operators of the class HepLorentzVector:

```

1 operator const Hep3Vector & () const;
2 operator Hep3Vector & ();

```

You do not need to understand how these work, just that the above two assignments work. If you do wish to, you can look in any standard C++ reference in a section that might be called type-cast operators or implicit conversions.

`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();
```

A common beginner's mistake involves confusion about the function `mag`. For a `Hep3Vector`, `mag`, returns the magnitude of the 3-vector. For a `HepLorentzVector`, it would make sense that there is a member function to return the magnitude of the space part but there is no such function. Instead you should do the following:

```
1 double momentumMagnitude = p.v().mag();
```

The common mistake is to use `mag` on a `HepLorentzVector` expecting it to return the magnitude of the space part, when it actually returns the invariant mass.



`HepLorentzVector` has member functions that return the invariant mass squared:

```
1 double msq = p.mag2();
2 double msq = p.m2();
3 double msq = p.restMass2();
4 double msq = p.invariantMass2();
```

`HepLorentzVector` has a member function and an operator overload that return the 4-dimensional dot product

$$p \cdot q = p_t q_t - p_x q_x - p_y q_y - p_z q_z \quad (\text{C.2})$$

This function is also provided as an overload of the multiply operator:

```
1 double dot = p.dot(q);
2 double dot = p*q;
```

The `HepLorentzVector` forwards many of the functions of `Hep3Vector`. With one very important exception the member functions described in Table C.1 are also available to a `HepLorentzVector` and perform the same functions. For example the member function `theta` returns the polar angle of the space-part of the 4-vector. The exception is the member function `mag`, which was discussed above.

`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;
```

`HepLorentzVector` has member functions and operators that compare for equality and to ask if two vectors are almost the same, almost parallel or almost orthogonal. These have signatures similar to those for `Hep3Vector` and they have the same fragile behaviour; use them with care.

If you consult the header file or Doxygen documentation you will see that `HepLorentzVector` has many member functions that are not mentioned above. These are not needed for the Workbook and will not be discussed here.

C.6.2.1 HepBoost

This material will be available in a future release.

C.7 The Matrix Package

The *art* Workbook uses the following classes from the Matrix package:

`HepMatrix` A general $n \times m$ matrix class.

`HepSymMatrix` A class that represents symmetric matrices.

`HepVector` A column-vector ($n \times 1$ matrix) class.

This material will be available in a future release.

C.8 The Random Package

The *art* Workbook uses the following classes from the Random package:

`HepRandomEngine` The base class from which all CLHEP random engines must inherit.

`HepJamesRandom` A random engine that implements an algorithm described by F. James of CERN.

`RandFlat` A distribution that returns a random variate that is flat on a specified domain.

`RandGaussQ` A distribution that returns a random variate that is distributed as a Gaussian distribution.

`RandPoissonQ` A distribution that returns a random variate that is distributed as a Poisson distribution.

The two classes with names ending in `Q` have no internal state except for the state of the underlying engine.

This material will be available in a future release.

Part V

Index

Index

- analyzer module, **182**
 - characteristics, 182
 - configure via parameter set, 240
 - events passed by reference, 186
 - flow of execution, 211
 - order in path, 265
 - signature, 186
 - use of override, 183
- API, 12
- art, **7**
 - API, 12
 - applicability, 7
 - artmod, 202
 - as an external product, 23
 - build systems, 199
 - C++, 7
 - command, 14
 - command line options, 134
 - long form, 134
 - short form, 134
 - configuration file, *see* configuration file
 - data persistency, 288
 - data product, *see* data product
 - development environment, 161, 174, 355
 - documentation suite, 9
 - dynamic library usage, 190
 - error conditions, 153
 - error status, 129
 - event, *see* event
 - event ID, *see* event ID
 - event loop, *see* event loop
 - event sharing, 186
 - executable, 123
 - execution syntax, 128
 - flow of execution, 211
 - getting help, 9
 - identifiers, 369
 - input file
 - specify in FHiCL, 136
 - specify on command line, 136
 - job status, 129
 - log file, 129
 - module, *see* module
 - module types, 18
 - operate on multiple modules, 205
 - output file, 129
 - output module, 146
 - paths used in, 40
 - post-initialization steps, 16
 - processing order, 265
 - rerun same module, 140, 262
 - ROOT classes, 288

- ROOT support, 25
- run-time configuration, 123
 - command line options, 130
 - FHiCL file, 130
- run-time environment, 123, 147, 351
- sample output, 128
- services, *see* services
- specify events to process, 135, 138
- specify modules to process, 139
- TFileService, *see* TFileService
- Unix environment, 38
- use as external package, 8
- use of ROOT, 288
- users, 8
- art module, *see* module
- art-users email list, 9
- art::EDAnalyzer, 182
- art::Event, 185, 265
- art::EventID, 269
- art::ServiceHandle, 289
- artdaq, **8**
- artmod, 202
 - options, 205
- auto, 282
- beginJob, 297
- boost, **23**
- browse command, 304
- build system, 22, 161
 - instructions for, 168
- build tools, 355
- building code
 - clean rebuild, 195
 - complete, 195
 - finding dynamic libraries, 198
 - incremental, 195
 - linking, 198
 - saving output files, 195
- buildtool, 172, 174
 - algorithm, 194
 - CMakeLists.txt file, 193
 - error, 196
 - functions, 193
 - verbose mode for experts, 200
- C++, 9
 - Werror, 54
 - .cc files, 47
 - .dylib files, 48
 - .h files, 47
 - .o files, 48
 - .so files, 48
 - automatic type conversion, 66
 - base class, 16
 - build, 47, 64
 - output option, 53
 - build commands, 61
 - c++ command, 61, 63, 65
 - code guards, 59
 - compile, 47, 52
 - declaration, 98
 - definition within declaration, 98
 - exception, 250
 - executable program, 48
 - float, 56
 - free function, 101
 - function
 - argument list, 59
 - declaration, 59
 - definition, 60, 66

- implementation, 60
 - return type, 59
 - function ‘main’, 52, 59, 62
 - header files, 47, 59
 - implementation within declaration, 98
 - include directive, 65
 - include ROOT header syntax, 291
 - inheritance, 16, 162
 - libraries, 48, 58
 - link, 47, 52
 - link list, 49
 - linker, 63
 - linker symbols, 63
 - main program, 52
 - main program, 52, 59
 - module, *see* module
 - object files, 48
 - pointer, 56
 - prerequisites, 52
 - rebuild subset, 58
 - signature, 66
 - Singleton Design Pattern, 21
 - source code files, 47
 - std::vector<T>, 51
 - syntax flexibility, 190
 - temporary object recommendations, 191
 - uninitialized variable, 54
 - unresolved references, 63
 - variable addresses, 54
 - variable type, 56
- C++11, 9
- calibration constants, *see* conditions information
- cetbuildtools, 22, **23**, 161, 355
- CETLIB, **23**
- CINT, 289
- file naming convention, 305
 - relation to C++, 304
 - script, 290
- , 290
- class templates, *see* templates
- CLHEP, **23**
- cmake, 22, 298
- variable, 280, 299
 - definition, 199
- CMakeLists.txt file
- buildtool, 193
- cmsrun, 8
- coding
- best practices, 33
 - conventions, 33
 - rules, 33
 - style, 33
- coding standards, **9**
- C++, 9
 - C++ 11, 9
- collection, 20
- colon initializer syntax, 239, 259
- conditions information, **20**, 123
- configuration file, 14
- constructor
- explicit argument, 183
- data file, 25
- data members
- store parameter values, 252
- data product, **19**, 133, 269
- collection, *see* collection
 - contents, 20
 - distinguish from products, 24

- name, 270, 272
- operations, 25
- persistency, 25
- persistent representation, 25
- specification, 274, 279
- transient representation, 25
- data type
 - friendly name, 273
 - in data product name, 273
- debugging
 - using optional parameters, 255
- defaultExceptions, 284
- development environment, 355
- Doxygen, 12
- dynamic libraries, 21
 - .dylib files, 21
 - .so files, 21
 - build system, 22
 - dependency lists, 198
 - directories, 152
 - file extensions, 123
 - naming rules, 194
 - paths to, 123
 - run-time loading, 183
 - use in instance creation, 190
 - use with buildtool, 194
- dynamic library, 15
 - building, 178
 - matching module name, 178
- dynamic load libraries, *see* dynamic libraries
- EDAnalyzer, 140
- EDM, *see* Event-Data Model
- endJob, 297
- error handling, 250
 - default vs alternate behavior, 251
- event, **13**, 13
 - contents, 185
 - representation, 188
 - representation in art, 185
 - unique identifier, 13
- event ID, **13**, 269
 - event number, 13, 187
 - EventID return type, 187
 - individual parts of, 206
 - run number, 13, 187
 - subRun number, 13, 187
- event loop, 14, **18**, 18
- event-data files, 26
- event-data files for Workbook, 124
- Event-Data Model, **25**
 - ROOT support, 25
- exceptions, 239, 260
- experiment code, *see* user code
- external products, *see* products
- FermiGrid, 24
- Fermilab Hierarchical Configuration Language,
 - see* FHiCL
- FHiCL, 23, 122, 130, **365**
 - fhicl::ParameterSet, 184
 - definition
 - form, 130
 - definitions, 237
 - file extension, 130
 - identifier
 - analyzers, 133
 - physics, 132
 - source, 131
 - numerical forms

- formats, 258
- numerical types, 255
- output module, 146
 - optional parameters, 146
- parameter name, 244
- parameter set, **237**
 - error conditions, 249
 - print, 239, 259
- parameter value
 - store as data member, 252
 - store as local variable, 253
- parameters, 237
 - canonical form, 239, 256, 260
 - default value, 253
 - default values, 239, 260
 - internal representation, 241
 - optional, 253
 - policies, 255
 - properties, 241
 - return type, 242
- paths, 141
- process name, 133
- scope, 131
- source file, 130
- special characters, 132
- specify input files, 136
- syntax, 130
- table, 130, 237
- value, 132
- fhicl::ParameterSet, 237
 - get, 244
- file catalog, 26
- file of Monte Carlo events, *see* event-data files
- file of simulated events, *see* event-data files
- filter module, 18
- framework, **7**
 - boundary with user code, 8
 - infrastructure, 8
- friendly name, 273, 278
 - textbf, 273
- gcc, **23**
- GenParticle, **271**, 276
- GenParticleCollection, **271**
- geometry information, 123
- geometry specification, 20
- getting help, 9
- git, 23, 164
 - about, 165
- handle
 - invalid, 278
- handles, **277**
 - default construct, 278
 - service, 294
 - valid, 278, 281
- header files
 - absence of, 208
 - conflicts, 197
 - finding, 196
 - from UPS product, 197
 - Geant4, 197
 - ROOT, 197
- help with art, 9
- histogram
 - change module label, 310
 - create using TFileService, 295
 - filling, 296
 - formats, 310
 - pointer naming convention, 291

- save, 310
 - structure, 296
 - view interactively, 304
 - view with ROOT, 289
 - view with TBrowse, 301, 310
 - write to PDF, 289
- histogram file, 296, 300
 - change name, 309
 - overwriting, 308
- ifdh_sam, 24
- inheritance, 182
- input tag
 - initialization, 280
- jobsub_tools, 24
- keyword
 - auto, 282
- member function
 - argument names, 185
 - for analyzer module, 182
 - optional, 183
 - override identifier, 183, 185
 - template, 248
- message service, 20
 - purpose of, 134
- MF, **23**
- module, **14**
 - C++ class, 15
 - analyzer, *see* analyzer
 - member function requirement, 245
 - class, 237
 - communication between, 265
 - create with artmod, 202
 - dependencies, 199
 - filter, *see* filter
 - finding, 152
 - header files, 179
 - identify from art output, 178
 - instance, 262, 296
 - in data product name, 273
 - name, 273
 - label, 139, 264, **266**, 296
 - identify parameter set, 266
 - in data product name, 273
 - name instance, 266
 - naming rules, 141
 - parameter set, 237
 - uniqueness, 141
 - naming rules, 194
 - optional member function, 183
 - output, *see* output
 - producer, *see* producer
 - requirements, 15
 - run simultaneously, 205
 - source, *see* source
 - source code, 178
 - three-file style, 209
 - type, **131**, 140, 182
 - types, 18
- module types, **18**
- namespaces
 - fhicl, 184
 - ROOT, 291
 - tex, 181
 - toy experiment, 181
- naming variables, 244
- NTuple, 18

- null pointer, 294
- nullptr, 294
- output directory, 176
- output module, 18
- overload set, 66
- packages, *see* products
- parameter set, 130, 264
 - for module configuration, 140
 - module label, 139
- parameters
 - default value recommendations, 281
- paths
 - art, 141
 - different types of, 144
 - FHiCL, 141
 - module order, 265
- plugins, *see* dynamic libraries
- pointer, 281
 - bare, 282
 - naming convention for histogram, 291
- pointers, 277
 - safe, 278
 - smart, 278
- process name, 133
 - in data product name, 273
- processing loop, *see* event loop
- producer module, 18
 - processing order, 265
- products, **23**, 107
 - access to, 105
 - distinguish from data product, 24
 - distribution via UPS/UPD, 24, 105
 - external, 105
 - product directories, 105
 - PRODUCTS, 105
- reconstruction on demand, 15
- replica manager, 26
- ROOT, 18, 23, 288
 - booking a histogram, 297
 - create histogram, 295
 - cycle numbers, 302
 - deleting a histogram, 297
 - documentation links, 308
 - exit from root program, 302
 - file naming conventions, 297
 - filling a histogram, 296
 - global namespace, 291
 - histogram file, *see* histogram file
 - histogram vs event-data filenames, 297
 - include header syntax, 291
 - input vs output filenames, 297
 - output file structure, 296
 - TBrowserWindow, 301
 - TFileService, *see* TFileService
 - TNtuples, 297
 - treatment of bin edges, 295
 - TTrees, 297
- ROOT files
 - event-data, 289
 - histograms, 289
- RootInput module, 131
- run, 13
- run-time configuration
 - value types, 366
- run-time configuration file, *see* configuration file
- run-time environment, 351
- SAM, 24, 26

- service handle, 294
- services, **20**, 294
 - access via handle, 294
 - message service, *see* message service
 - requesting information from, 21
 - TFileService, *see* TFileService
- set up to run Workbook, 125
 - build window, 168
 - source window, 164
- setup to run Workbook
 - login back in, 217
- shareable libraries, *see* dynamic libraries 21
- site-specific setup, 45
 - procedures, 45
 - Unix environment, 45
- site-specific setup procedure, 148
 - setup git, 165
- smart pointer, 21
- source code
 - compile and link, 179
- source directory, 164
- source module, 18
- std::vector, 254, 269
 - dynamic sizing, 270
- subRun, 13
- TBrowser, 301
 - print histogram, 310
- templates, 239, 242, 260
 - argument, 244, 270, 274
 - dummy, 248
 - type, 244
 - class, 269, 272, 281
 - member function, 248, 269, 279
- templates:collections of objects, use with, 272
- testing
 - using optional parameters, 255
- TFileService, 21, **288**, 289, 294
 - arguments, 295
 - configure, 290, 297
 - create histogram, 295
- toy experiment, 11, **26**
 - namespace, 181
 - setup, 149
- Tree, 18
- typedef, 271
 - recommendation for use of, 188
- Unix
 - art* Workbook environment, 38
 - bash alias, 42
 - bash function, 42
 - bash script, 41
 - bash shell, 36
 - commands, 34
 - computing environment, 37
 - environment, 37
 - environment variables, 37, 38, 45
 - examine environment, 38
 - execute vs source, 41
 - help for commands, 34
 - important concepts, 35
 - login scripts, 43
 - login shell, 36
 - non-standard commands, 34
 - path vs PATH, 39
 - scripts, 36
 - shell variables, 39
 - shells, 36
 - suggested references, 43

- working environment, 37, 45
- UPS
 - product conflicts, 197
 - product header files, 197
 - products area, 123
- UPS/UPD, **23**, 23
 - databases, 105
 - features, 105
- UPS/UPD:initialization, 148
- UPS/UPD:qualifiers, 149
- UPS:product dependency lists, 198
- user code, **7**
- Workbook, **11**
 - build window
 - contents, 170
 - setup, 168
 - disk space, 149
 - event-data files, 124
 - FHiCL files
 - machine-independent, 173
 - multi-site usage, 197
 - obtain code, 164
 - setup to run exercises, 125
 - initial, 125, 127
 - self-managed, 127
 - standard, 125
 - subsequent logins, 128
 - source directory contents, 167
 - toy experiment, *see* toy experiment
 - Unix environment, 38