



Gaudi

LHCb Data Processing Applications Framework

Users Guide

Corresponding to Gaudi release v9

Version: 9
Issue: 0
Edition: 0
Status:
ID:
Date: 19 December 2001



European Laboratory for Particle Physics
Laboratoire Européen pour la Physique des Particules
CH-1211 Genève 23 - Suisse

Document Control Sheet

Document	Title: Gaudi Users Guide
	Version: 9
	Issue: 0
	Edition: 0
	ID:
	Date: 19 December 2001
	Keywords:
Tools	DTP System: Adobe FrameMaker Version: 6.0 Layout Template: Software Documentation V1 - 15 January 1999 Layout Templates
Authorship	Coordinator: M.Cattaneo (from October 1999) P.Maley (until September 1999)

Document Status Sheet

Title: Gaudi Users Guide			
ID:			
Version	Issue	Date	Reason for change
6	0	15/Nov/00	Release of version 6 of the GAUDI framework
6	1	27/Nov/00	Update chapter 8 and section 3.3.1
6	2	5/Dec/00	Corrections to section 6.10, Appendix B
6	3	8/Jan/01	Add first two lines of code listing 76
7	0	20/Apr/01	Release of version 7 of GAUDI framework
8	0	28/Aug/01	Release of version 8 of GAUDI framework
9	0	19/Dec/01	Release of version 9 of GAUDI framework



Table of Contents

Document Control Sheetii
Document Status Sheetii
Chapter 1	
Introduction1
1.1 Purpose of the document1
1.2 Conventions2
1.3 Reporting problems5
1.4 Editor's note5
Chapter 2	
The framework architecture.7
2.1 Overview7
2.2 Why architecture?7
2.3 Data versus code8
2.4 Main components9
2.5 Controlling and Scheduling11
Chapter 3	
Release notes and software installation13
3.1 Release History13
3.2 Current Functionality14
3.3 Changes between releases16
3.4 Availability19
3.5 Using the framework19
3.6 Working with development releases21
3.7 Installation of the framework outside CERN22
Chapter 4	
Getting started25
4.1 Overview25
4.2 Creating a job25
4.3 The main program26
4.4 Configuring the job27
4.5 Algorithms30
4.6 Job execution32
4.7 Examples distributed with Gaudi34
4.8 A additional LHCb specific examples35
Chapter 5	
Writing algorithms37



5.1 Overview 37

5.2 Algorithm base class 37

5.3 Derived algorithm classes 40

5.4 Nesting algorithms 43

5.5 Algorithm sequences, branches and filters 44

Chapter 6

Accessing data 47

6.1 Overview 47

6.2 Using the data stores 47

6.3 Using data objects 49

6.4 Object containers 50

6.5 Using object containers 51

6.6 Data access checklist 53

6.7 Defining Data Objects 53

6.8 The SmartDataPtr/SmartDataLocator utilities 55

6.9 Smart References and Smart Reference Vectors 56

6.10 Persistent storage of data 57

Chapter 7

Modelling Event Data 59

Chapter 8

Detector Description 61

8.1 Overview 61

8.2 Detector Description Database 61

8.3 Detector Data Transient Store 63

8.4 General features of the detector description 69

8.5 Persistent representation based on XML files 73

8.6 Persistent storage in a Conditions Database 87

Chapter 9

Histogram facilities 93

9.1 Overview 93

9.2 The Histogram service 93

9.3 Using histograms and the histogram service 94

9.4 Persistent storage of histograms 95

Chapter 10

N-tuple and Event Collection facilities 97

10.1 Overview 97

10.2 N-tuples and the N-tuple Service 97

10.3 Event Collections 102

10.4 Known Problems 109

Chapter 11

Framework services 111

11.1 Overview 111

11.2 Requesting and accessing services 111

11.3 The Job Options Service 113

11.4 The Standard Message Service 120

11.5 The Particle Properties Service 123

11.6 The Chrono & Stat service 126

11.7 The Auditor Service 129

11.8 The Random Numbers Service 131

11.9 The Incident Service 134

11.10 The GiGa Service 135

11.11 The Gaudi Introspection Service 135

11.12 Developing new services 136

Chapter 12

Tools and ToolSvc 139

12.1 Overview 139

12.2 Tools and Services 139

12.3 The ToolSvc 145

12.4 GaudiTools 147

Chapter 13

Converters 153

13.1 Overview 153

13.2 Persistence converters 153

13.3 Collaborators in the conversion process 154

13.4 The conversion process 156

13.5 Converter implementation - general considerations 158

13.6 Storing Data using the ROOT I/O Engine 158

13.7 The Conversion from Transient Objects to ROOT Objects 159

13.8 Storing Data using other I/O Engines 160

Chapter 14

Scripting and Interactivity 161

14.1 Overview 161

14.2 How to enable Python scripting 161

14.3 Current functionality 163

14.4 Physics Analysis Environment 167

Chapter 15

Visualization Facilities 169

15.1 Overview 169

15.2 The data visualization model 169

15.3 VisSys - the Gaudi visualisation services 170

15.4 Panoramix - the LHCb event display 172

Chapter 16	
Framework packages, interfaces and libraries	173
16.1 Overview	173
16.2 Gaudi Package Structure	173
16.3 Interfaces in Gaudi	176
16.4 Libraries in Gaudi	179
Chapter 17	
Analysis utilities	185
17.1 Overview	185
17.2 CLHEP	185
17.3 HTL	185
17.4 NAG C	186
17.5 ROOT	186
Chapter 18	
Accessing SICB facilities	187
18.1 Overview	187
18.2 Reading tapes	187
18.3 Populating the GAUDI transient data store: SICB Converters	188
18.4 Access to the Magnetic Field	190
18.5 Accessing the SICB detector Geometry from Gaudi	191
18.6 Using FORTRAN code in Gaudi	192
18.7 Handling pile up in Gaudi	193
18.8 Handling Spillover in Gaudi	195
18.9 Increasing the size of the ZEBRA COMMON block	196
Appendix A	
References	197
Appendix B	
Options for standard components	199
Appendix C	
Job Options Grammar and Error Codes	207
Appendix D	
Design considerations	213



Chapter 1

Introduction

1.1 Purpose of the document

This document is intended as a combination of user guide and tutorial for the use of the Gaudi application framework software. It is complemented principally by two other "documents": the Architecture Design Document (ADD) [1], and the online auto-generated reference documentation [2]. A third document [3] lists the User Requirements and Scenarios that were used as input and validation of the architecture design. All these documents and other information about Gaudi, including source code documentation, are available via the Gaudi home page: <http://cern.ch/proj-gaudi>. Documentation of the LHCb extensions to Gaudi, and a copy of this user guide, are available on the LHCb web at: <http://cern.ch/lhcb-comp/Frameworks/Gaudi/>. It is recommended that you also study the material of the LHCb Gaudi tutorials given at CERN, which is available at <http://cern.ch/lhcb-comp/Frameworks/Gaudi/GaudiTutorial.htm>. Should you wish to enrol on a future tutorial, you should send mail to Markus.Frank@cern.ch.

The ADD contains a discussion of the architecture of the framework, the major design choices taken in arriving at this architecture and some of the reasons for these choices. It should be of interest to anyone who wishes to write anything other than private analysis code.

As discussed in the ADD the application framework should be usable for implementing the full range of offline computing tasks: the generation of events, simulation of the detector, event reconstruction, testbeam data analysis, detector alignment, visualisation, etc. . .

In this document we present the main components of the framework which are available in the current release of the software. It is intended to increment the functionality of the software at each release, so this document will also develop as the framework increases in functionality. Having said that, physicist users and developers actually see only a small fraction of the framework code in the form of a number of key interfaces. These interfaces should change very little if at all and the user of the framework cares very little about what goes on in the background.

The document is arranged as follows: Chapter 2 is a short resume of the framework architecture, presented from an "Algorithm-centric" point of view, and re-iterating only a part of what is presented in the ADD.

Chapter 3 contains a summary of the functionality which is present in the current release, and details of how to obtain and install the software.

Chapter 4 discusses in some detail an example which comes with the framework library. It covers the main program, some of the basic aspects of implementing algorithms, the specification of job options and takes a look at how the code is actually executed. The subject of coding algorithms is treated in more depth in Chapter 5.

Chapter 6 discusses the use of the framework data stores and event data. Chapter 7 gives pointers to guidelines for defining the LHCb event data model. Chapters 8, 9, 10 discuss the other types of data accessible via these stores: detector description data (material and geometry), histogram data and n-tuples.

Chapter 11 deals with services available in the framework: job options, messages, particle properties, auditors, chrono & stat, random numbers, incidents, introspection. It also has a section on developing new services. Framework tools are discussed in Chapter 12, the use and implementation of converter classes in Chapter 13.

Chapter 14 discusses scripting as a means of configuring and controlling the application interactively. This is followed by a description in Chapter 15 of how visualisation facilities might be implemented inside Gaudi.

Chapter 16 describes the package structure of Gaudi and discusses the different types of libraries in the distribution.

Chapter 17 gives pointers to the documentation for class libraries which we are recommending to be used within Gaudi.

The use of certain SICB facilities within Gaudi and the wrapping of FORTRAN code are discussed in Chapter 18.

Appendix A contains a list of references. Appendix B lists the options which may be specified for the standard components available in the current release. Appendix C gives the details of the syntax and possible error messages of the job options compiler. Finally, Appendix D is a small guide to designing classes that are to be used in conjunction with the application framework.

1.2 Conventions

1.2.1 Units

We have decided to adopt the same system of units as CLHEP, as used also by GEANT4. This system is fully documented in the CLHEP web pages, at the URL: http://wwwinfo.cern.ch/asd/lhc++/clhep/manual/U_serGuide/Units/units.html.



The list of basic units is reproduced in Table 1.1. Note that this differs from the convention used in SICB, where the basic units of length, time and energy are, respectively, centimetre, GeV, second..

Table 1.1 CLHEP system of units

Quantity	Unit
Length	millimetre
Time	nanosecond
Energy	MeV
Electric charge	positron charge
Temperature	Kelvin
Amount of substance	mole
Plane angle	radian

Users should not actually need to know what units are used in the internal representation of the data, as long as they are consistent throughout the Gaudi data stores. What they care about is that they can define and plot quantities with the correct units. In some specialised algorithms they may also wish to renormalise the data to a different set of units, if the default set would lead to numerical precision problems.

We therefore propose the following rules, which are discussed more fully in reference [5].

1. All dimensioned quantities in the Gaudi data stores shall conform to the CLHEP system of units.
2. All definitions of dimensioned quantities shall be dimensioned by multiplying by the units defined in the `CLHEP/Units/SystemOfUnits.h` header file. For example:

```
const double my_height = 170*cm;
const double my_weight = 75*kg;
```

Note that the user should not care about the numerical value of the numbers `my_height` and `my_weight`. Internally these numbers are represented as 1700. and 4.68e+26, respectively, but the user does not need to know.

3. All output of dimensioned quantities should be converted to the required units by dividing by the units defined in the `CLHEP/Units/SystemOfUnits.h` header file. For example:

```
my_hist = histoSvc()->book( " /stat/diet", "corpulence (kg/m**2)", 30, 10., 40. );
double my_corpulence = my_weight / (my_height*my_height);
my_hist->fill( my_corpulence/(kg/m2), 1. );
```

which, for a healthy person, should plot a number between 19. and 25....

4. Physical constants should not be defined in user code. They should be taken directly from the `CLHEP/Units/PhysicalConstants.h` header file. For example:

```
float my_rest_energy = my_weight * c_squared;
```



- Users may wish to use a different set of units for specific purposes (e.g. when the default units may lead to precision problems). In this case algorithms can renormalise their private copy of the data (as shown in the last line of the rule 3 example) for internal use, but making sure that any data subsequently published to the public data stores is converted back to the CLHEP units.

1.2.2 Coding Conventions

The Gaudi software follows (or should follow!) the LHCb C++ coding conventions described in reference [6]. A tool to check compliance with these coding conventions is described at <http://cern.ch/lhcb-comp/Support/Conventions/RuleChecker.htm>

1.2.2.1 File extensions

One consequence of following the LHCb coding conventions is that the specification of the C++ classes is done in two parts: the header or ".h" file and the implementation or ".cpp" file. We also define file extensions for Gaudi specific files. The recommended file extension for Job Options files is ".opts" (see Section 11.3.3 on page 116). For Python scripts, the extension ".py" is mandatory (see Chapter 14).

1.2.3 Naming Conventions

Histograms In order to avoid clashes in histogram identifiers, we suggest that histograms are placed in named subdirectories of the transient histogram store. The top level subdirectory should be the name of a sub-detector group (e.g. VELO). Below this, users are free to define their own structure. One possibility is to group all histograms produced by a given algorithm into a directory named after the algorithm.

Event data Naming conventions for LHCb data are discussed in reference [8].

1.2.4 Conventions of this document

Angle brackets are used in two contexts. To avoid confusion we outline the difference with an example.

The definition of a templated class uses angle brackets. These are required by C++ syntax, so in the instantiation of a templated class the angle brackets are retained:

```
AlgFactory<UserDefinedAlgorithm> s_factory;
```



This is to be contrasted with the use of angle brackets to denote "replacement" such as in the specification of the string:

```
"<concreteAlgorithmType>/<algorithmName>"
```

which implies that the string should look like:

```
"EmptyAlgorithm/Empty"
```

Hopefully what is intended will be clear from the context.

1.3 Reporting problems

Users of the Gaudi software are encouraged to report problems and requests for new features via the LHCb problem reporting system. This system is integrated in the CERN Problem Report Management System (CPRMS) provided by IT division, based on the Action Request System software from Remedy Corporation.

To report a new problem, go to the LHCb CPRMS home page <http://cern.ch/hep-service-prms/lhcb.html>, click on the **Submit** button, and fill in the form. This will add the report to the system and notify the developers by E-mail. You will receive E-mail notification of any changes in the status of your report.

To view the list of current problems, and their status, click the **Query** button on the LHCb CPRMS home page.

Active developers of the Gaudi software are encouraged to use the gaudi-developers mailing list for discussion of Gaudi features and future developments. This list is not, primarily, intended for bug reports. In order to send mail to gaudi-developers@lhcb.cern.ch, you must first subscribe to the list, using the form at <https://www.lhcb.cern.ch/admin-cg/lhcb-admin?operation=view-list&mail=gaudi-developers@cern.ch>. You need a CERN mailserver account to be able to use this form...

The archive of the mailing list is publicly accessible on the Web, at <http://cern.ch/~major-dom/news/gaudi-developers/index.html>.

1.4 Editor's note

This document is a snapshot of the Gaudi software at the time of the release of version v9. We have made every effort to ensure that the information it contains is correct, but in the event of any discrepancies between this document and information published on the Web, the latter should be regarded as correct, since it is maintained between releases and, in the case of code documentation, it is automatically generated from the code.



We encourage our readers to provide feedback about the structure, contents and correctness of this document and of other Gaudi documentation. Please send your comments to the editor: Marco.Cattaneo@cern.ch



Chapter 2

The framework architecture

2.1 Overview

In this chapter we outline some of the main features of the Gaudi architecture. A (more) complete view of the architecture, along with a discussion of the main design choices and the reasons for these choices may be found in references [1] and [4].

2.2 Why architecture?

The basic "requirement" of the physicists is a set of programs for doing event simulation, reconstruction, visualisation, etc. and a set of tools which facilitate the writing of analysis programs. Additionally a physicist wants something that is easy to use and (though he or she may claim otherwise) is extremely flexible. The purpose of the Gaudi application framework is to provide software which fulfils these requirements, but which additionally addresses a larger set of requirements, including the use of some of the software online.

If the software is to be easy to use it must require a limited amount of learning on the part of the user. In particular, once learned there should be no need to re-learn just because technology has moved on (you do not need to re-take your licence every time you buy a new car). Thus one of the principal design goals was to insulate users (physicist developers and physicist analysts) from irrelevant details such as what software libraries we use for data I/O, or for graphics. We have done this by developing an architecture. An architecture consists of the specification of a number of components and their interactions with each other. A component is a "block" of software which has a well specified interface and functionality. An interface is a collection of methods along with a statement of what each method actually does, i.e. its functionality.

The main components of the Gaudi software architecture can be seen in the object diagram shown in Figure 2.1. Object diagrams are very illustrative for explaining how a system is decomposed. They represent a hypothetical snapshot of the state of the system, showing the



objects (in our case component instances) and their relationships in terms of ownership and usage. They do not illustrate the structure, i.e. class hierarchy, of the software.

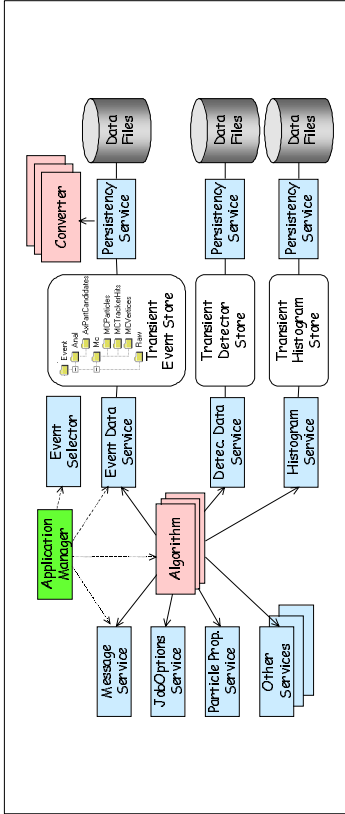


Figure 2.1 Gaudi Architecture Object Diagram

It is intended that almost all software written by physicists, whether for event generation, reconstruction or analysis, will be in the form of specialisations of a few specific components. Here, specialisation means taking a standard component and adding to its functionality while keeping the interface the same. Within the application framework this is done by deriving new classes from one of the base classes:

- DataObject
- Algorithm
- Converter

In this chapter we will briefly consider the first two of these components and in particular the subject of the “separation” of data and algorithms. They will be covered in more depth in chapters 5 and 6. The third base class, Converter, exists more for technical necessity than anything else and will be discussed in Chapter 13. Following this we give a brief outline of the main components that a physicist developer will come into contact with.

2.3 Data versus code

Broadly speaking, tasks such as physics analysis and event reconstruction consist of the manipulation of mathematical or physical quantities; points, vectors, matrices, hits, momenta, etc., by algorithms which are generally specified in terms of equations and natural language. The mapping of this type of task into a programming language such as FORTRAN is very natural, since there is a very clear distinction between “data” and “code”. Data consists of variables such as:

```
integer n
real_p(3)
```



and code which may consist of a simple statement or a set of statements collected together into a function or procedure:

```
real function innerProduct(p1, p2)
real p1(3), p2(3)
innerProduct = p1(1)*p2(1) + p1(2)*p2(2) + p1(3)*p2(3)
end
```

Thus the physical and mathematical quantities map to data and the algorithms map to a collection of functions.

A priori, we see no reason why moving to a language which supports the idea of objects, such as C++, should change the way we think of doing physics analysis. Thus the idea of having essentially mathematical objects such as vectors, points etc. and these being distinct from the more complex beasts which manipulate them, e.g. fitting algorithms etc. is still valid. This is the reason why the Gaudi application framework makes a clear distinction between “data” objects and “algorithm” objects.

Anything which has as its origin a concept such as hit, point, vector, trajectory, i.e. a clear “quantity-like” entity should be implemented by deriving a class from the `DataObject` base class.

On the other hand anything which is essentially a “procedure”, i.e. a set of rules for performing transformations on more data-like objects, or for creating new data-like objects should be designed as a class derived from the `Algorithm` base class.

Further more you should not have objects derived from `DataObject` performing long complex algorithmic procedures. The intention is that these objects are “small”.

Tracks which fit themselves are of course possible: you could have a constructor which took a list of hits as a parameter; but they are silly. Every track object would now have to contain all of the parameters used to perform the track fit, making it far from a simple object.

Track-fitting is an algorithmic procedure; a track is probably best represented by a point and a vector, or perhaps a set of points and vectors. They are different.

2.4 Main components

The principle functionality of an algorithm is to take input data, manipulate it and produce new output data. Figure 2.2 shows how a concrete algorithm object interacts with the rest of the application framework to achieve this.

The figure shows the four main services that algorithm objects use:

- The event data store
- The detector data store
- The histogram service
- The message service






Figure 2.2 The main components of the framework as seen by an algorithm object.

The particle property service is an example of additional services that are available to an algorithm. The job options service (see Chapter 11) is used by the `Algorithm` base class, but is not usually explicitly seen by a concrete algorithm.

Each of these services is provided by a component and the use of these components is via an interface. The interface used by algorithm objects is shown in the figure, e.g. for both the event data and detector data stores it is the `IDataProviderSvc` interface. In general a component implements more than one interface. For example the event data store implements another interface: `IDataManagerSvc` which is used by the application manager to clear the store before a new event is read in.

An algorithm's access to data, whether the data is coming from or going to a persistent store or whether it is coming from or going to another algorithm is always via one of the data store components. The `IDataProviderSvc` interface allows algorithms to access data in the store and to add new data to the store. It is discussed further in Chapter 6 where we consider the data store components in more detail.

The histogram service is another type of data store intended for the storage of histograms and other "statistical" objects, i.e. data objects with a lifetime of longer than a single event. Access is via the `IHistogramSvc` which is an extension to the `IDataProviderSvc` interface, and is discussed in Chapter 9. The n-tuple service is similar, with access via the `INTupleSvc` extension to the `IDataProviderSvc` interface, as discussed in Chapter 10.

In general, an algorithm will be configurable: It will require certain parameters, such as cut-offs, upper limits on the number of iterations, convergence criteria, etc., to be initialised before the algorithm may be executed. These parameters may be specified at run time via the job options mechanism. This is done by the job options service. Though it is not explicitly shown in the figure this component makes use of the `IProperty` interface which is implemented by the `Algorithm` base class.

During its execution an algorithm may wish to make reports on its progress or on errors that occur. All communication with the outside world should go through the message service component via the `IMessageSvc` interface. Use of this interface is discussed in Chapter 11.



As mentioned above, by virtue of its derivation from the `Algorithm` base class, any concrete algorithm class implements the `Algorithm` and `IProperty` interfaces, except for the three methods `initialize()`, `execute()`, and `finalize()` which must be explicitly implemented by the concrete algorithm. `Algorithm` is used by the application manager to control top-level algorithms. `IProperty` is usually used only by the job options service.

The figure also shows that a concrete algorithm may make use of additional objects internally to aid it in its function. These private objects do not need to inherit from any particular base class so long as they are only used internally. These objects are under the complete control of the algorithm object itself and so care is required to avoid memory leaks etc.

We have used the terms "interface" and "implements" quite freely above. Let us be more explicit about what we mean. We use the term interface to describe a pure virtual C++ class, i.e. a class with no data members, and no implementation of the methods that it declares. For example:

```
class PureAbstractClass {
    virtual method1() = 0;
    virtual method2() = 0;
}
```

is a pure abstract class or interface. We say that a class implements such an interface if it is derived from it, for example:

```
class ConcreteComponent: public PureAbstractClass {
    method1() { }
    method2() { }
```

A component which implements more than one interface does so via multiple inheritance, however, since the interfaces are pure abstract classes the usual problems associated with multiple inheritance do not occur. These interfaces are identified by a unique number which is available via a global constant of the form: `IID_InterfaceType`, such as for example `IID_DataProviderSvc`. Interface identifiers are discussed in detail in Chapter 16.

Within the framework every component, e.g. services and algorithms, has two qualities:

- A concrete component class, e.g. `TrackerFinderAlgorithm` or `MessageSvc`.
- Its name, e.g. "KalmanFitAlgorithm" or "MessageService".

2.5 Controlling and Scheduling

2.5.1 Application Bootstrapping

The application is bootstrapped by creating an instance of the `ApplicationMgr` component. The `ApplicationMgr` is in charge of creating an initializing a minimal set of basic and essential services before control is given to specialized scheduling services. These services are shown in



Figure 2.3. The EventLoopMgr is in charge controlling the main physics event¹ loop and scheduling the top algorithms. There will be a number of more or less specialized implementations of EventLoopMgr which will perform the different actions depending on the running environment, and experiment specific policies (clearing stores, saving histograms, etc.). There exists the possibility to give the full control of the application to a component implementing the IRunnable interface. This is needed for interactive applications such as event display, interactive analysis, etc. The Runnable component can interact directly with the EventLoopMgr for triggering the processing of the next physics event.

The essential services that the ApplicationMgr need to instantiate and initialize are the MessageSvc and JobOptionsSvc.




Figure 2.3 Control and Scheduling collaboration

2.5.2 Algorithm Scheduling

The Gaudi architecture foresees explicit invocation of algorithms by the framework or by other algorithms. This latter possibility allows for a hierarchical organization of algorithms, for example, a high level algorithm invoking a number of sub-algorithms.

The EventLoopMgr component is in charge of initializing, finalizing and executing the set of Algorithms that have been declared with the TopAlg property. These Algorithms are executed unconditionally in the order they have been declared. This very basic scheduling is insufficient for many use cases (event filtering, conditional execution, etc.). Therefore, a number of Algorithms have been introduced that perform more sophisticated scheduling and can be configured by some properties. Examples are: Sequencers, Prescatters, etc. and the list can be easily extended. See Section 5.5 for more details on these generic high level Algorithms.

1. We state physics event to differentiate from what is called generally an event in computing.



Chapter 3

Release notes and software installation

3.1 Release History

The Gaudi architecture and framework, which was initially developed by the LHCb collaboration, became a joint development project between several experiments, starting from release v6. At this time the package structure was modified, to split the experiment specific packages from the common packages. The following table reflects the version history since the re-packaging. For the history of earlier releases, please refer to previous versions of the (LHCb) Gaudi Users Guide.

Version	Date	Package List
v9	Dec 2001	GaudiPolicy[v5], GaudiExamples[v9], GaudiKernel[v11], GaudiSvc[v7], GaudiAudi[v5], GaudiAlg[v5], GaudiTools[v5], GaudiNagCl[v6], GaudiDb[v5], GaudiRootDb[v5], GaudiODBCDb[v5], HbookCnv[v11], RootHistCnv[v5], GaudiPython[v2], GaudiObjDesc[v2], GaudiInrospection[v2], GaudiConf[v6], Event/LHCbEvent [v12], SICB/SiebCnv [v14], Det/XmlDDDB [v7], Det/DetDesc [v8], Det/DetCond [v1], Sim/GEANT4 [v3]Op1], Sim/GiGa [v6], Ex/Ex# [v3], Vis/OnXSvc [v3], Vis/SolLHCb [v3], Vis/Panoramix [v3]
v8r1	July 2001	GaudiPolicy[v5], GaudiExamples[v8], GaudiKernel[v10], GaudiSvc[v6], GaudiAudi[v4], GaudiAlg[v4], GaudiTools[v4], GaudiNagCl[v5]p2], GaudiDb[v4], GaudiRootDb[v4], GaudiODBCDb[v3], HbookCnv[v10r1], RootHistCnv[v4], SIPython[v2], GaudiConf[v5], Event/LHCbEvent [v11], Event/DbCnv [v5], SICB/SiebCnv [v13], Det/XmlDDDB [v6], Det/DetDesc [v7], Sim/GEANT4 [v3]Op1], Sim/GiGa [v5]Ex/Ex# [v2]
v7	23/03/2001	GaudiPolicy[v4], GaudiExamples[v7], GaudiKernel[v9], GaudiSvc[v5], GaudiAudi[v3], GaudiAlg[v3], GaudiTools[v3], GaudiNagCl[v5]p1], GaudiDb[v3], HbookCnv[v9], RootHistCnv[v3], SIPython[v1r1], DetDesc[v6], XmlDDDB[v5], XmlEditor[v3], LHCbEvent[v9], DbCnv[v4], GiGa[v4], GaudiConf[v4], SiebCnv[v11], Ex/Ex# [v1]



3.2 Current Functionality

We use an incremental and iterative approach for producing the Gaudi software. We plan to expand its capabilities release by release. The functionality list that follows is organized by categories.

Interfaces A set of interfaces that facilitates the interaction between the different components of the framework. Mainly these are interfaces to services.

Basic framework services This set of services offer the minimal functionality needed for constructing applications. They are described in detail in Chapter 11.

The *message service* is used to send and format messages generated in the code, with an associated severity that is used for filtering and dispatching them.

The *job options service* allows the configuration of the application by end users assigning values to properties defined within the code; properties can be basic types (`Float`, `bool`, `int`, `string`), or extended with bounds checking, hierarchical lists, and immediate callback from string "commands".

The *Random Numbers service* makes available several random number distributions via a standard interface, and ensures that applications use a unique random number engine in a reproducible fashion.

The *Chrono service* offers the functionality for measuring elapsed time and job execution statistics. *Auditors* and *AuditorSvc* provide monitoring of various characteristics of the execution of Algorithms.

Auditors are called before and after invocation of any Algorithm method.

The *Incident service* provides a synchronization between objects within the Application by using named incidents that are communicated to *listener* clients.

The *Tools service*, which provides management of Tools, is discussed in Chapter 12. Tools are lightweight objects which can be requested and used many times by other components to perform well defined tasks. A base class for *associator* tools has been added in this release.

Data services provide the access to the transient data objects (event, detector and statistical data). The *data services* are described in chapters 6 to 10. The basic building blocks for the implementation of the experiment specific data models are also described in Chapter 6.

Event data model The event model has been extended to Velo, L0, L1 and updated for Calo. The current status is presented in Chapter 7.

Event data persistent storage The current version provides a set of generic classes for implementing event data persistence (*GaudiDb* package) and a set of classes supporting persistent I/O to ROOT files (*GaudiRootDb* package). Details can be found in Chapter 13.

Input from event data files (or tapes) produced by SICB in ZEBRA format (RZ) is also supported. The ZEBRA banks are *converted* to C++ objects by specialized fragments of code (SICB converters). This conversion is available in the framework for a number of SICB banks. Adding more converters is possible and will be done on request. The event data produced by the application can be stored in ROOT files and retrieved later. This also requires to write specialized converters.

Histograms & N-tuples The framework provides facilities for creating histograms (1 and 2 dimensional) and n-tuples (row and column wise) from user algorithms. The histogram interface is the AIDA[13] common interface. Saving histograms and n-tuples is currently implemented using the HBOOK and ROOT format. The interface to histograms and n-tuples from the user code should not be affected if the persistency representation is changed later. Details of the histogram and n-tuple facilities can be found in Chapter 9 and 10 respectively.



Event tag collections The framework provides facilities for creating and using collections of event tags (based on an n-tuples implementation) for fast direct access to physics events. The user can specify an event tag collection as input data to an application and perform sophisticated selections using the facilities existing in the data storage technology. This is explained in Chapter 10.

Detector description and geometry The framework provides facilities for accessing detector description and geometry data. The logical structure of the detector is described in terms of a hierarchy of detector elements, the basic geometry in terms of volumes, solids and materials. Facilities for customizing the generic description to many specific detector needs are also provided. This allows the development of detector specific code which can provide geometry answers to questions from the physics algorithms (simulation, reconstruction and analysis). The persistent representation of the detector description is based on text files in XML format. An XML editor that understands the detector description semantics has been developed. A transport service is provided to estimate the amount of material between two arbitrary points in the detector setup. This is described in detail in Chapter 8.

Analysis services A number of facilities and services are included in the current release to facilitate writing physics analysis code. The GaudiAlg package is a collection of general purpose algorithms, including a sequencer which uses the filtering capability of algorithms to manage the execution of algorithm sequences in a filtering application (see Section 5.5). The Particle Properties service (Section 11.5) provides the properties of all the elementary particles. Numerical utilities are available via the CLHEP and NAG C libraries (Chapter 17).

Visualization services The framework provides facilities for the visualisation of event and detector data. These services are currently based on the packages that constitute the Open Scientist suite (OpenGL, OpenInventor(soFree), Lab,...). An event and geometry display application has been built using these facilities. This is briefly described in Chapter 15.

SICB services A number of services are included in the current release to facilitate the access to SICB data (detector description, magnetic field, etc.) and to facilitate re-using existing Fortran code with the framework. These services can be extended to accommodate new concrete requests when integrating big parts of the current legacy code (SICB). In the latest release there full support for event pile-up. It is possible to read events from two input data streams and merge them before any processing is done. Refer to Chapter 18 for more details.

Object Description and Object Introspection The framework provides object modelling and description using XML files. Two code generation back-ends are currently available: to generate the data object header files and to generate the object dictionaries for the object introspection. Refer to Section 6.7 and Section 11.11 for more details.

Scripting services The framework provides a service for interfacing Python with a Gaudi application. The user can interact with a Gaudi application from the Python prompt. The current functionality allows the user to set and get properties from Algorithms and Services, interact with the data stores (event, detector and histogram) using the object introspection capability, and to schedule the execution of the application's algorithms. Refer to Chapter 14 for more details.

Dynamic loading of libraries The framework can be used to implement different data processing applications for different environments. It is important that the services and their concrete implementations are extendable and configurable dynamically at run time and not statically. The latter would imply linking with all the available libraries producing huge executables. And in certain applications and environments some of the libraries will surely never be used. The framework provides support for dynamic libraries for the Windows and Linux platforms.



3.3 Changes between releases

3.3.1 Changes between current release (v9) and previous release (v8)

- **Data-store classes.** Release v9 includes a major re-design of the Data Store classes and interfaces triggered by ATLAS data store implementation.
- **Object description and object introspection.** Two new packages has been added that provide the object description based on XML files and run-time object introspection capability. Refer to Section 6.7 and Section 11.11 for more details.
- **Python service.** The scripting service based on Python has been re-implemented using the Boost library¹. Its functionality has been extended. Refer to Chapter 14 for more details.
- **Algorithms.** Added toolSvc() accessor to Algorithm base class.
- **Algorithm Tools.** Added `initialize()` and `finalize()` methods in `AlgTool` interface. The base class `AlgTool` implements them as dummy but allows an implementation of them on specific Algorithm Tools. Removed the need to implement a `queryInterface()` in specific Algorithm tools. Instead use the expression `declareInterface<I>(this)` in the constructor.
- A number of small internal framework improvements:
 - `ApplicationMgr`. Re-organization to relocate the management of services to `ServiceManager` class. The interfaces `ISvcMgr` and `ISvcLocator` have changed.
 - Introduced a new constructor for `InterfaceID` that uses a name (class name) instead of an interface number.
 - `JobOptions`. Introduced new options `#pragma print on`, `#pragma print off` to switch the printing of job options on and off.
 - Histograms. New job option `HistogramPersistenceSvc::PrintHistos` to steer printing to standard output. Allow RZ directory names up to 16 characters rather than 8.

3.3.1.1 Incompatible changes

In this section we will list changes that users need to make to their code in order to upgrade to the current version of Gaudi from the previous version.

1. In the area of Data Stores many low level base classes (`DataObject`, `DataSvc`, `Converters`, `Registry`, `GenericAddress`, etc.) have changed together with some basic interfaces (`IConverter`, `IDataManagerSvc`, `IDataProviderSvc`, etc.). This implies that some packages, typically converters packages, will need deep changes in the code. Instructions on how to upgrade them can be found in http://cern.ch/lhcb-comp/Frameworks/Gaudi/Gaudi_v9/Changes_cookbook.pdf.
1. <http://www.boost.org/libs/libraries.htm>



End user algorithm packages should not be too affected by these changes.

2. Removed the list of default interfaces in `ApplicationMgr`. Services are late created if needed. This may cause problems if the order of creation played a role. The `Algorithms` and `Tools` that were accessing services using the call `serviceLocator()->service("name", interface)` may require to force the creation of a previously default service by adding a third argument with `true` to force such creation if not existing.
3. The constant `CLID_Event` has been removed from `ClassID.h`. It needs to be defined now in the `Event.h` header file.
4. Algorithm tools are required to implement an interface (pure abstract base class) using the facility provided for declaring it as mentioned above in the list of changes.

3.3.2 Changes between release v8 and release v7

3.3.2.1 Incompatible changes

In this section we will list changes that users need to make to their code in order to upgrade to version v8 of Gaudi from version v7.

1. Location of Histogram Interfaces. Gaudi version v8 uses the standard AIDA interfaces for histograms. These interfaces are located in the AIDA project area. The changes to the end-user code is that the include file should be prefixed with `AIDA/` instead of the current `GaudiKernel/`.
2. Persistent representation of N-tuples. N-tuples saved in HBOOK format no longer have type information in the first row. See the discussion in Section 10.2.3.2 for more details.
3. The output of N-tuples to ODBC (Open DataBase Connectivity) is no longer supported. N-Tuple preselections based on SQL or interpreted C++ are no longer available. If you rely on these features, please contact the Gaudi development team.
4. When saving data objects in a data store, all the sub-directory nodes in the path must already exist or should be explicitly created. In fact this is not a new feature, but a bug fix! Implicit creation of sub-directory nodes will be implemented in a future version.

3.3.3 Changes between release v7 and release v6

- Release v7 includes some re-packaging and changes in location of some include files, see section 3.3.3.1.
- The control of the "physics event" loop has been separated from the `ApplicationMgr` and has become a new component, the event loop manager. A number of subsequent specializations have been provided: `MinimalEventManager`, `EventManager`, and `GaudiEventManager`. These changes have been made to allow the possibility to have other types of event loop processing. These changes are backward compatible.



- The first version of a scripting service based on Python has been released.
- A number of small internal framework improvements:
 - Elimination of the up to now required static libraries.
 - Added version number (major and minor) to the Interface ID to check for interface compatibility at run-time.
 - Re-shuffling of the System class and conversion to a namespace.
 - Handling empty vectors in JobOptions.

3.3.3.1 Incompatible changes

In this section we will list changes that users need to make to their code in order to upgrade to version v7 of Gaudi from version v6.

1. Histogram persistency. In previous versions, the HBOOK histogram persistency service was created by default. From this version there is no default histogram persistency: the ROOT or HBOOK persistency services have to be explicitly declared. See Section 9.4 for details
2. EvtMax. In previous versions it was possible to declare the number of events to be processed through either of the properties ApplicationMgr.EvtMax or EventSelector.EvtMax. In this release, only ApplicationMgr.EvtMax is supported, the default being all the events in the input file
3. The property EventSelector.JobInput has been removed. Use EventSelector.Input instead (note the change in format of the value string).
4. Detector Description. The detector description interfaces (ISolid.h, IDetectorElement.h, IGeometryInfo.h etc) have been moved out of the GaudiKernel package and into the DetDesc package. The detector description related options of the Application Manager (ApplicationMgr.DetDbLocation etc) are now properties of the Detector Data Service (e.g. DetectorDataSvc.DetDbLocation etc).
5. Gaudi has been restructured into a set of experiment independent packages, and a set of LHCB packages. Since many of the packages have changed name, the paths on which to find include files has also changed.

3.3.4 Deprecated features

We list here features of the framework which have become obsolete, because they have been superseded by more recent features. Users are discouraged from using the deprecated features, which may be removed in a future release.

Adding indexed items to N-tuples Use the function `addIndexedItem` instead of `addItem`.

Accessors names in Algorithm Use the service accessors with short names (e.g. `msgSvc()`) instead of the long equivalent ones (e.g. `messageService()`)



Access to extra services from Algorithms. Use the templated method `service()` instead of using the `serviceLocator()` together with the `queryInterface()` to get a reference to a service interface.

User Parameters in detector elements. The XML tag for user parameters in the detector description (detector elements, etc.) is now `<param/>` instead of `<userparameter/>`. The old name will be maintained for a while. The methods in `DetectorElement` and `ConditionClasses` will accordingly be changed to use the word `param` instead of `userParameter`.

3.4 Availability

The application framework is supported on the following platforms:

- Windows NT4 and Windows 2000, using the Developer Studio 6.0 SP2 Visual C++ environment
- RedHat Linux 6.1 (certified CERN Linux distribution with SUE and AFS) with eggs-2.91.66 and gcc-2.95.2.

The code, documentation and installation instructions are available from the LHCB web site at: <http://cern.ch/lhcb-comp/Frameworks/Gaudi/>.

Framework sources and binaries are also available in the CERN AFS cell, at `/afs/cern.ch/sw/Gaudi` for the common Gaudi packages, and at `/afs/cern.ch/lhcb/software/NEW` for the LHCB specific packages.

3.5 Using the framework

3.5.1 CVS repository

The framework sources are stored in CVS and can be accessed using the `getpack` command. The first time you access a Gaudi package, the CVS server will ask for a password, reply `CERNuser`. Note that you only have write access to the LHCB specific part of the repository. For the common Gaudi packages you can send all the CVS commands that don't require write access. If you use a command like `commit`, you will get an error message.

3.5.2 CMT

The framework libraries have been built using the Configuration Management Tool (CMT) [7]. Therefore, using the CMT tool is the recommended way to modify existing packages or re-build the examples included in the release. If CMT is not available in your system, please follow the installation instructions in [7]. The following simple examples are for Unix, but



similar commands exist also for Windows. They assume that the `CMTPEATH` environment variable is set to `$HOME/mycmt:$GAUDIHOME`.

Getting a copy of a package: Suppose you want to build the latest released version of the `GaudiExamples` package:

```
1: cd mycmt
2: cmt checkout GaudiExamples -r v7
```

Building and running an example: Now that you have the code, suppose you want to modify the `AlgSequencer` example, then build it and run it:

```
3: cd GaudiExamples/v7/src/AlgSequencer
4: emacs HelloWorld.cpp
---- Make your modification, then
5: cd ../../cmt
6: source setup.csh
7: emacs requirements
---- Uncomment AlgSequencer and comment all the others
8: gnake
9: cd ../home
10: emacs AlgSequencer.txt
---- Make any modification if needed
11: ../$CMTCONFIG/AlgSequencer.exe
```

Modifying a library and rerunning the example: Suppose now you want to modify one of the `Gaudi` libraries, build it, then rerun the `AlgSequencer` example with it:

```
12: cd $HOME/mycmt
13: cmt checkout GaudiAlg v3
14: cd GaudiAlg/v3/src
15: emacs ...
---- Make your modification...
16: cd ../cmt
17: source setup.csh
18: gnake
19: cd $HOME/mycmt/GaudiExamples/v7/cmt
20: cmt show uses
21: ---- Verify the you are now using the GaudiAlg version from $HOME/mycmt
22: ---- There is no need to relink, since GaudiAlg is a component library
23: cd ../home
24: ../$CMTCONFIG/AlgSequencer.exe
```

3.5.3 Using the framework on Windows with Developer Studio or Nmake

The libraries for Windows are available for download from the web, and in AFS in the `Win32Debug` subdirectory of each package.

Instructions for installing the Gaudi environment on Windows and for customising MS Visual Studio are available at: <http://cern.ch/lhcb-comp/Support/html/DevStudio/default.htm>

The requirements files and CMT commands expect the following environment variables:



- **HOME** Needs to be set to the user's home directory. Typically this is in a network server and will have the form "`\\server\username`" or can also be a local directory like "`C:\home`". This environment variable is used to locate the `.cmtc` file that contains the default `CMTPEATH`.
- **PATH** Should be set up correctly to locate the Developer Studio executables (this is typically the case after installation).
- **TEMP** Location for temporary files. This is set correctly after the operating system installation.
- **STEREOT** This is the root where software is installed. Typically it will point to some share in some server (`\\server\stero`) or to the locally mounted AFS drive (`F:\cern.ch`).
- **CMTPEATH** The first location where CMT is looking for packages. This is typically the local directory `C:/mycmt`
- **CMTSITE** This is your site name. At CERN site and for the Windows platform we use `CERN_WIN32`.

3.5.4 Using the framework in Unix

The libraries for Linux are available for download from the web, and in AFS in the `t386_linux22` and `Linuxdtx` subdirectories of each package (for the optimised and debug versions respectively).

Instructions for installing the Gaudi environment on Linux are available at: http://cern.ch/lhcb-comp/Support/html/start_gaudi.htm

3.6 Working with development releases

This User Guide corresponds to release v9 of the Gaudi software.

For Gaudi as well as for LHCh packages, before they are publicly released and frozen, the development versions are periodically rebuilt from the head revision of the CVS repository in the development release area. These versions are not guaranteed to work and may change without notice; they are intended for integration tests. They should be used with care, mainly if you wish to use new features of the software which have not yet been incorporated in a public release.

3.6.1 The Gaudi common development area

Development releases of Gaudi packages are made periodically into the AFS directory tree below `/afs/cern.ch/sw/Gaudi/dev` (pointed to by the `GAUDIDEV` environment variable on Unix), by checking out and building the CVS head revision of modified packages



3.6.2 The LHCb development release area

Development releases of packages are made periodically into the AFS directory tree below `/afs/cern.ch/lhcb/software/DEV` (pointed to by the `LHCBDEV` environment variable on Unix), by checking out and building the CVS head revision of modified packages. The directory structure of `LHCBDEV` is analogous to `LHCBSOFT` – there is a subdirectory per package, below which there are subdirectories for each development release version of that package. These subdirectories have a name which is constructed from the version number of the package and the date of the build, e.g. `v4r1d010316` corresponds to version `v4r1` and was built on 16th March 2001. The version number is greater than or equal to that of the most recent public release, depending on what has been changed in the package since the release.

3.6.3 Using the development version of packages

The purpose of the development area is to make it easy to work with the development version of packages, without needing to make a private copy of the head revision of a package.

To use packages from the development area, it is sufficient to add the `LHCBDEV` (and/or `$GAUDIDEV`) area to your `CMTPATH`, ahead of the release area. This is done either by editing the file `.cmtc` in your home directory (both Unix and Windows) by typing (Unix only):

```
setenv CMTPATH $(HOME)/mycmt:$(LHCBDEV):$(GAUDIDEV)
```

Note that this latter method is valid only for the current login session, whereas the former method is valid until the next time you modify the file.

3.7 Installation of the framework outside CERN

3.7.1 Package installation

To use the Gaudi framework you also need to have access to installations of some external packages, listed below:

CMT, CLHEP, NAG C, HTL, Python, Xerces, qglhcb, ROOT, BOOST and CERNLIB.

Up to date instructions for installation of these packages and setting of the environment (variables, path...) needed to use the framework can be found on the Web at <http://cern.ch/lhcb-comp/Support/html/Install.htm>.

3.7.2 Event Data access

At CERN, the information about datasets stored in the CASTOR system is obtained from the LHCb bookkeeping database. Other sites may have their own way to access bookkeeping



information and the associated data files. Gaudi reads the bookkeeping database and stages data files via scripts that should be customised for the installation (see <http://cern.ch/lhcb-comp/Support/html/Install.htm>), the only restriction being that they must have the same input and output signatures as the CERN version, which can be found in the package `Tools/Bookkeeping`.



Chapter 4

Getting started

4.1 Overview

In this chapter we walk through one of the example applications (`RandomNumber`) which are distributed with the framework. We look briefly at the different files and go over the steps needed to compile and execute the code. We also outline where various subjects are covered in more detail in the remainder of the document. Finally we cover briefly the other example applications which are distributed and say a few words on what each one is intended to demonstrate.

4.2 Creating a job

Traditionally, a "job" is the running of a program on a specified set of input data to produce a set of output data, usually in batch.

For the example applications supplied this is essentially a two step process. First the executable must be produced, and secondly the necessary environment variables must be set and the required job options specified, as illustrated in Figure 4.1.

The example applications consist of a number of "source code" files which together allow you to generate an executable program. These are:

- The main program.
- Header and implementation files for each of the concrete algorithm classes.
- A CMT requirements file.
- The set of Gaudi libraries.

In order for the job to run as desired you must provide the correct configuration information for the executable. This is done via entries in the job options file.



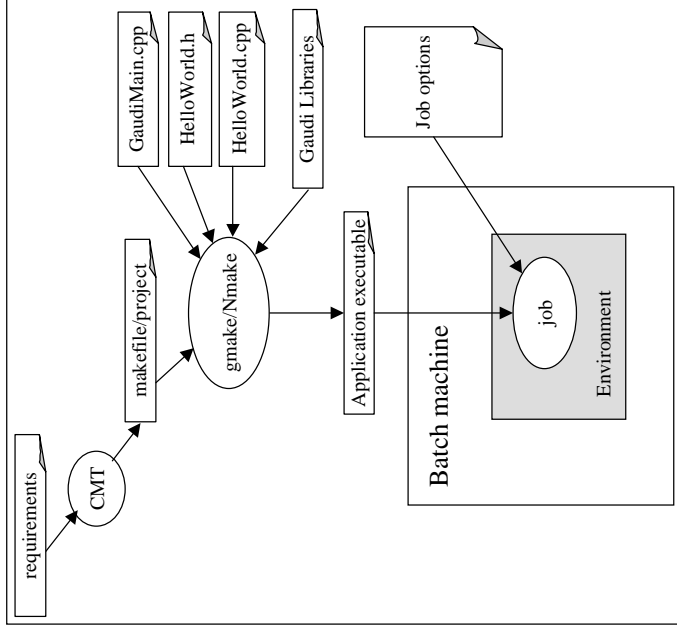


Figure 4.1 Creating a job from the ALGSequencer example application

4.3 The main program

The main program is needed to bootstrap the job. It can be completely general, and can be reused by all Gaudi applications. An example main program, from the package `Gaudi.Examples`, is shown in Listing 4.1. It is constructed as follows:

Include files These are needed for the creation of the application manager and Smart interface pointers.

Application Manager instantiation Line 12 instantiates an `ApplicationMgr` object. The application manager is essentially the job controller. It is responsible for creating and correctly initialising all of the services and algorithms required, for looping over the input data events and executing the algorithms specified in the job options file, and for terminating the job cleanly.

Listing 4.1 The example main program.

```

1: // Include files
2: #include "GaudiKernel/SmartIF.h"
3: #include "GaudiKernel/Bootstrap.h"
4: #include "GaudiKernel/IAppMgrUI.h"
5: #include "GaudiKernel/IProperty.h"
6: #include <iostream>
7:
8: ///-- Example main program
9: int main(int argc, char** argv) {
10:
11: // Create an instance of an application manager
12: IInterface* iface = Gaudi::createApplicationMgr();
13:
14: SmartIF<IProperty> propMgr ( IID_IProperty, iface );
15: SmartIF<IAppMgrUI> appMgr ( IID_IAppMgrUI, iface );
16:
17: if ( !appMgr.isValid() || !propMgr.isValid() ) {
18:   std::cout << "Fatal error creating ApplicationMgr " << std::endl;
19:   return 1;
20: }
21:
22: // Get the input configuration file from arguments
23: std::string opts = (argc>1) ? argv[1] : "../options/job.opts";
24:
25: propMgr->setProperty( "JobOptionsPath", opts );
26:
27: // Run the application manager and process events
28: appMgr->run();
29:
30: // All done - exit
31: return 0;
32: }
  
```

Retrieval of interface pointers The code on lines 14 and 15 retrieves the pointers to the `IProperty` and `IAppMgrUI` interfaces of the application manager.

Setting the application manager's properties The only property which needs to be set explicitly in the main program is the name of the job options file which contains all of the other configuration information needed to run the job. In this example, the name is the first argument of the program and defaults to `"../options/job.opts"` (line 23); it is set on line 25.

Program execution All of the code before line 28 is essentially for setting up the job. Once this is done, a call to `appMgr::run()` is all that is needed to start the job properly! The steps that occur within this method are discussed briefly in section 4.6.

4.4 Configuring the job

The application framework makes use of a job options file for job configuration. Part of the job options file of an example application is shown in Listing 4.2.



Listing 4.2 Part of the job options file for the ROOTRead example application.

```

1: // Include standard option files
2: #include "$SDOPTS/Common.opts"
3:
4: // Private Application Configuration options
5: ApplicationMgr.DLLs += { "GaudiDb", "GaudiRootDb" };
6: ApplicationMgr.ExtSvc += { "DbEventCnvSvc/RootEvtCnvSvc" };
7: ApplicationMgr.TopAlg = { "ReadAlg" };
8:
9: // Set output level threshold (2=DEBUG, 3=INFO, 4=WARNING, 5=ERROR, 6=FATAL)
10: MessageSvc.OutputLevel = 4;
11: EventSelector.OutputLevel = 2;
12:
13: // Input File
14: EventSelector.Input = {"DATAFILE='RootDat.root' TYP='ROOT' OPT='READ'"};
15: EventSelector.FirstEvent = 1;
16: ApplicationMgr.EvtMax = 5;
17:
18: // Persistence service setup:
19: EventPersistenceSvc.CnvServices += { "RootEvtCnvSvc" };
20:
21: // Setup for ROOT I/O system
22: RootEvtCnvSvc.DbType = "ROOT";

```

The format of an options file is discussed fully in Chapter 11. Options may be set both for algorithms and services and the list of available options for standard components is given in Appendix B. Here we look briefly at a few of the more commonly used options.

4.4.1 Defining the algorithms to be executed

The option `ApplicationMgr.TopAlg` (line 7) is a list of algorithms that will be created and controlled directly by the application manager, the so-called top-level algorithms. The syntax is a list of the form:

```
ApplicationMgr.TopAlg = { "Type1/Name1", "Type2/Name2" };
```

The line above instructs the application manager to create two top level algorithms. One of type `Type1` called "Name1" and one of type `Type2` called "Name2".

In the case where the name of the algorithm is the same as the algorithm's type (i.e. class), only the class name is necessary. In the example, an instance of the class "ReadAlg" will be created with name "ReadAlg".

4.4.2 Defining the job input

Event data input is controlled by an `EventSelector`. The `EventSelector` uses a storage technology dependent data persistence service to load the data into the transient event data store, with the help of converters which are able to convert the data from the technology

dependent persistent representation, to the technology independent representation in the transient data store.

In order to set up this mechanism, one needs a number of job options:

- Line 14 defines the input data file, and the persistence technology (ROOT I/O in this example).
- Line 6 tells the application manager to create a new event conversion service, to be called `RootEvtCnvSvc`. Note that this is just a name for our convenience, the service is of type `DbEventCnvSvc` and does not (yet) know that it will deal with ROOT technology. The configuration of `RootEvtCnvSvc` to use the ROOT I/O technology is done in line 22.
- Line 19 tells the event persistence service (`EventPersistenceSvc` created by the application manager by default) to use the `RootEvtCnvSvc` to do the conversion between persistent and transient data representations.
- Line 5 tells the application manager which additional libraries to load in order to find the required conversion service. In this example, the `GaudiDb` library contains the `DbEventCnvSvc` class, the `GaudiRootDb` library contains the ROOT specific database drivers.
- Finally, the options on lines 15 and 16 tell the `EventSelector` to start reading sequentially from the first event in the file, for five events.

In the special case where no event input is required (e.g. for event generation), one can replace the above options by the two options:

```
ApplicationMgr.EvtMax = 20; // events to be processed (default is 10)
ApplicationMgr.EvtSel = "NONE"; // do not use any event input
```

A discussion of event I/O can be found in Chapter 10. Converters and the conversion process are described in Chapter 13.

4.4.3 Defining job output

One can consider three types of job output: event data (including event collections and n-tuples), statistical data (histograms) and `printout`. Here we discuss only the simplest (`printout`); histograms are discussed in Chapter 9, event data in Section 6.10.1, event collections in Section 10.3.1.

`Printout` in Gaudi is handled by the message service (described in Chapter 11), which allows to control the amount of `printout` according to severity level. The global threshold for `printout` is set by the option on line 10 - in this example only messages of severity level `WARNING` or above will be printed. This can be over-ridden for individual algorithms or services, as in line 11, where the threshold for `EventSelector` is set to `DEBUG`.



4.5 Algorithms

The subject of specialising the Algorithm base class to do something useful will be covered in detail in Chapter 5. Here we will limit ourselves to looking at an example `HelloWorld` class.

4.5.1 The `HelloWorld.h` header file

The `HelloWorld` class definition is shown in Listing 4.3.

Listing 4.3 The header file of the class `HelloWorld`.

```

1: // Include files
2: #include "GaudiKernel/Algorithm.h" // Required for inheritance
3: #include "GaudiKernel/Property.h"
4: #include "GaudiKernel/MsgStream.h"
5:
6: class HelloWorld : public Algorithm {
7: public:
8:     // Constructor of this form must be provided
9:     HelloWorld(const std::string& name, ISvcLocator* pSvcLocator);
10:
11:     // Three mandatory member functions of any algorithm
12:     StatusCode initialize();
13:     StatusCode execute();
14:     StatusCode finalize();
15: private:
16:     // These data members are used in the execution of this algorithm
17:     // and are set in the initialisation phase by the job options service
18:     int m_int;
19:     double m_double;
20:     std::string m_string;
21: };

```

Note the following:

- The class is derived from the Algorithm base class as must be all specialised algorithm classes. This implies that the `Algorithm.h` file must be included (line 6).
- All derived algorithm classes must provide a constructor with the parameters shown in line 9. The first parameter is the name of the algorithm and is used amongst other things to locate any options that may have been specified in the job options file.
- The `HiStoAlgorithm` class has three (private) data members, defined in lines 18 to 20. These are properties that can be set via the job options file.
- The three methods on lines 12 to 14 must be implemented, since they are pure virtual in the base class.



4.5.2 The `HelloWorld` implementation file

The implementation file contains the actual code for the constructor and for the methods: `initialize()`, `execute()` and `finalize()`. It also contains two lines of code for the `HelloWorld` factory, which we will discuss in section 5.3.1

The **constructor** must call the base class constructor, passing on its two arguments. As usual, member variables should be initialised. Here we declare and initialise the member variables that we wish to be set by the job options service. This is done by calling the `declareProperty()` method.

```

1: HelloWorld::HelloWorld(const std::string& name, ISvcLocator* pLoc)
2:     : Algorithm(name, pLoc) {
3:     //-----
4:     // Declare the algorithm's properties
5:     declareProperty( "Int", m_int = 100 );
6:     declareProperty( "Double", m_double = 100.);
7:     declareProperty( "String", m_string = std::string("one hundred"));
8: }

```

Initialisation The application manager invokes the `sysInitialize()` method of the algorithm base class which, in turn, invokes the `initialize()` method of the base class, the `setProperty()` method, and finally the `initialize()` method of the concrete algorithm class. As a consequence all of an algorithm's properties will have been set before its `initialize()` method is invoked, and all of the standard services such as the message service are available. This is discussed in more detail in Chapter 5.

Looking at the code in the example (Listing 4.4) we see that we are now able to print out the values of the algorithm's properties, using the message service and the `MsgStream` utility class. A local `MsgStream` object is created (line 3), which uses the Algorithm's standard message service via the `msgSvc()` accessor, and the algorithm's name via the `name()` accessor. The use of these is discussed in more detail in Chapter 11.

Note that the job will stop if the `initialize()` method of any algorithm does not return `StatusCode::SUCCESS`. This is to avoid processing with a badly configured application.:

Listing 4.4 Example of `initialize()` method

```

1: StatusCode HelloWorld::initialize() {
2:     //-----
3:     MsgStream log(msgSvc(), name());
4:     log << MSG::INFO << "initializing....." << endl;
5:     log << MSG::INFO << "Property Int = " << m_int << endl;
6:     log << MSG::INFO << "Property Double = " << m_double << endl;
7:     log << MSG::INFO << "Property String = " << m_string << endl;
8:
9:     m_initialized = true;
10:    return StatusCode::SUCCESS;
11: }

```



execution The `execute()` method is called by the application manager once for every event. This is where most of the real action should take place. The trivial `HelloWorld` class just prints out a message... Note that the method must return `StatusCode::SUCCESS` on successful completion. If a particular algorithm returns a `FAILURE` status code more than a (configurable) maximum number of times, the application manager will decide that this algorithm is badly configured and jump to the finalisation stage before all events have been processed...

```

1: StatusCode HelloWorld::execute() {
2: //-----
3: MsgStream log( msgSvc(), name() );
4: log << MSG << "executing..." << endl;
5:
6: return StatusCode::SUCCESS;
7: }

```

Finalisation The `finalize()` method is called at the end of the job. In this trivial example a message is printed.

```

1: StatusCode HelloWorld::finalize() {
2: //-----
3: MsgStream log(msgSvc(), name());
4: log << MSG << "finalizing..." << endl;
5:
6: return StatusCode::SUCCESS;
7: }

```

4.6 Job execution

From the main program and the CMT requirements file we can make an executable, as explained in section 3.5. This executable together with the file of job options form a job which may be submitted for batch or run interactively. Figure 4.2 shows a trace of an example program execution. The diagram is not intended to be complete, merely to illustrate a few of the points mentioned earlier in the chapter.

1. The application manager instantiates the required services and initialises them. The message service is done first to allow the other services to use it, and the job options service is second so that the other services may be configured at run time.
2. The algorithms which have been declared to the application manager within the job options (via the `TOPALG` option) are created. We denote these algorithms "top-level" as they are the only ones controlled directly by the application manager. For illustration purposes we instantiate an `EmptyAlgorithm` and a `HistoAlgorithm`.
3. The top-level algorithms are initialised. Their properties (if they have any) are set and they may make use of the message service. If any algorithm fails to initialise, the job is stopped.

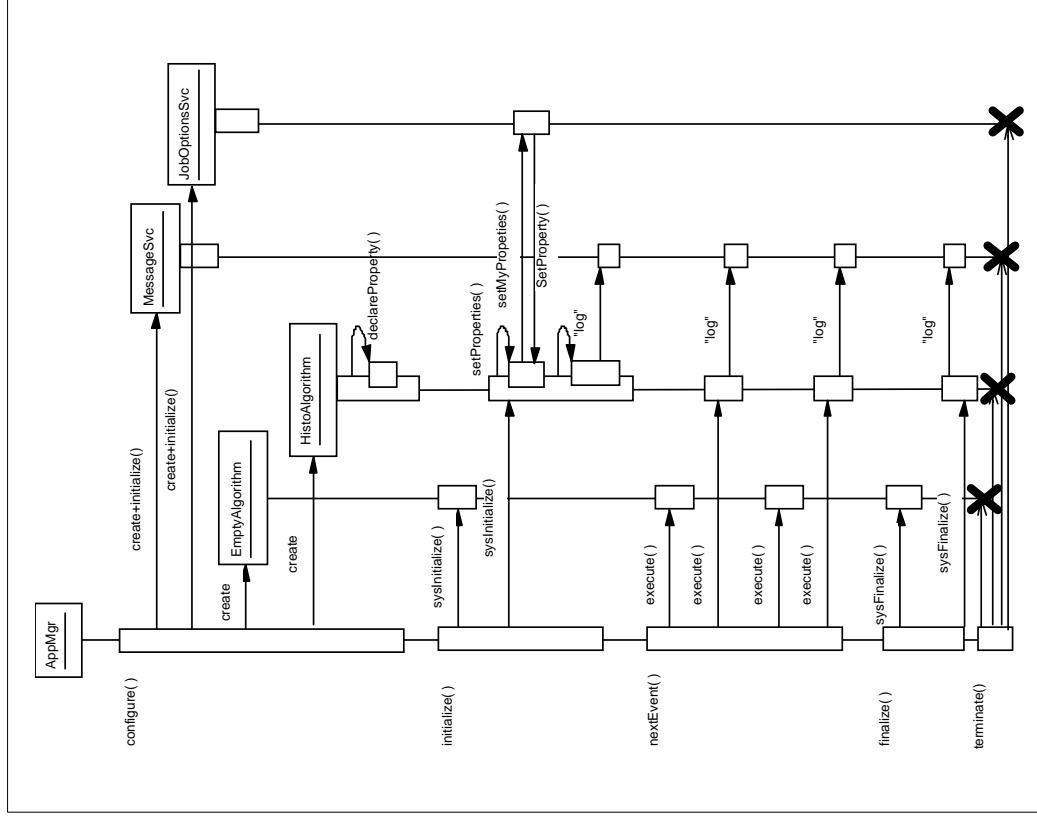


Figure 4.2 A sequence diagram showing a part of the execution of an example program.



4. The application manager now starts to loop over events. After each event is read, it executes each of the top level algorithms in order. The order of execution of the algorithms is the order in which they appear in the `TOPALG` option. This will continue until the required number of events has been processed, unless one or more of the algorithms return a `FAILURE` status code more than the maximum number of times, in which case the application manager will jump to the finalisation stage before all events have been processed.
5. After the required data sample has been read the application manager finalises each top level algorithm.
6. Services are finalised.
7. All objects are deleted and resources freed. The program terminates.

4.7 Examples distributed with Gaudi

A number of examples is included in the current release of the framework, in the `Gaudi.Examples` package. The package has some sub-directories in addition to the standard ones shown in Figure 16.3. The `options` sub-directory contains files of standard job options common to many examples. These files are included in the job options of the specific examples when necessary. The specific job options files can be found in the `home` sub-directory.

The code of the examples is in sub-directories of the `src` directory, one sub-directory per example. The intention is that each example demonstrates how to make use of some part of the functionality of the framework. The list of available examples is shown in Table 4.1.

Table 4.1 List of examples available in Gaudi release v9

Example Name	Target Functionality
AlgSequencer	Illustrating the use of the sequencer algorithm provided in the <code>GaudiAlg</code> package
AlgTool	Example implementation and use of a Gaudi Tool
Common	Actually not a complete example: contains the main program used in the examples
GPYthon	Exercise the Python scripting packages
Properties	Trivial algorithm showing how to set and retrieve Properties
RandomNumber	Example of use of the Random Number service
RootIO	Two examples, reading and writing persistent data with <code>ROOT I/O</code>



4.8 Additional LHCb specific examples

The examples described so far are rather simple and do not contain any specific knowledge about the LHCb event and detector data. A set of LHCb specific examples is provided in the `Ex` group of packages, as listed in Table 4.2.

Table 4.2 List of LHCb specific examples available in the current release

Example Package Name	Target Functionality
Ex/AssociatorExample	Illustrating the use of the <code>AxPart2MCParticleAsct</code> associator tool to navigate between reconstructed tracks and Monte Carlo particles
Ex/DetCondExample	Illustrating the use of the detector conditions database
Ex/DetDescExample	Illustrating the use of the detector description
ExDumpEventExample	Navigation of the LHCb transient event data model
Ex/EvtCollectionExample	Two examples, writing and reading event collections, starting from the vent data file produced by the <code>RootIO</code> example
Ex/FieldGeomExample	Making available existing Svb magnetic field and geometry data to Gaudi algorithms. Example of nested algorithms
ExFortranAlgorithmExample	Wrapping Fortran code in Gaudi
Ex/GiGaExample	Uses <code>GEANT4</code> to visualise the LHCb detector geometry and event data, as an example of using the <code>GiGa</code> service to interface Gaudi to <code>GEANT4</code>
Ex/HistogramsExample	Basic functionality of the framework to execute a simple algorithm, access event data and fill histograms.
Ex/MCPrimaryVertexExample	Retrieve data using <code>SmartDataPtr</code>
Ex/NtuplesExample	Two examples, reading and writing <code>Ntuples</code>
Ex/ParticlePropertiesExample	Access the <code>Particle Properties</code> service to retrieve <code>Particle Properties</code>
Ex/RootIOExample	Two examples, reading and writing persistent data with <code>ROOT I/O</code>
Ex/SimpleAnalysisExample	A realistic example of using the framework for a physics analysis, including access to Monte Carlo data, creation of reconstructed data and filling an n-tuple
Ex/ToolsAnalysisExample	Example of use of framework tools in an analysis

All examples share a single main program and some default job options, which can be found in the `Gaudi1.Conf` package.



4.8.1 Simple Physics Analysis Example

The algorithms in the examples of Table 4.2 use many of the Gaudi Services that someone would want to be able to utilize while doing physics analysis: histograms, ntuples, creating and retrieving private transient data, retrieving particle properties (like mass values), etc. Detailed examples on how to use the specific services are provided in the topical examples but in the `SimpleAnalysisExample` they are combined together. Tools to make physics analysis in a more elegant and complex way are under development and their concrete implementation will be part of DaVinci, the OO Physics Analysis Program. A trivial implementation of an algorithm similar to that of the `SimpleAnalysisAlgorithm` implemented using tools is provided in the `ToolAnalysisExample`.

The `SimpleAnalysisAlgorithm` is an example in which $\pi^+\pi^-\pi^0$ invariant masses are made while requiring the component particles to satisfy some simple kinematic and quality cuts. Private containers of the particles satisfying successive cuts are created and filled (charged particles, detection in the silicon, best particle ID). Invariant masses are made and corresponding histograms are filled for all combinations of the final private containers, for combinations with P_t of both pions greater than a cut value and for combinations with impact parameter of both pions greater than a cut value. The P_t and impact parameter cut values are properties of the algorithm and as such can be specified in the `jobOptions`, where the number is taken in Gaudi Units. CLHEP vectors' classes are used to evaluate transverse momentum and invariant masses as well as to calculate the impact parameter. When nominal mass values are required they are retrieved via the `ParticlePropertySvc`. Since a primary vertex is required a "dummy" algorithm `RecPrimaryVertex` retrieves the Monte Carlo primary vertex and uses the quantities to fill a `MyVertex` object (`/Event/MyVertexes`), which is then retrieved by the `SimpleAnalysisAlgorithm`. Since the `MyVertex` object is created and registered in the Transient Event store by the `RecPrimaryVertex` algorithm, the sequencing of `RecPrimaryVertex` and `SimpleAnalysisAlgorithm` in the `jobOptions` file is very important. A protection is put in place so that the `SimpleAnalysisAlgorithm` will return a failure code if not all of the necessary input data exist in the store.

When doing physics analysis on Monte Carlo data, it is necessary to compare the reconstructed decay with the Monte Carlo truth in order to calculate efficiencies. The `MCDecayFinder` algorithm is an example of how to find any one step decay. The parent of the decay and the list of its direct descendants are properties of the algorithm and can be specified in the `jobOptions` file. If no decay is specified in the `jobOptions` this example will look for a $B^0 \rightarrow \pi^+\pi^0$ decay. The Algorithm will retrieve the particle Geant3 ID from the `ParticlePropertySvc` (the identifying particleID in `MCParticles`) and search the `MCParticles` to find the requested parent and that is has the correct type and number of decay products. If a decay is found kinematic variables are stored in an ntuple that can be accessed by PAW. In addition the Algorithm uses the Message service with `DEBUG` or `INFO` levels to print a summary of its behaviour for each event as well as for the job.



Chapter 5

Writing algorithms

5.1 Overview

As mentioned previously the framework makes use of the inheritance mechanism for specialising the `Algorithm` component. In other words, a concrete algorithm class must inherit from ("be derived from" in C++ parlance, "extend" in Java) the `Algorithm` base class.

In this chapter we first look at the base class itself. We then discuss what is involved in creating concrete algorithms: specifically how to declare properties, what to put into the methods of the `Algorithm` interface, the use of private objects and how to nest algorithms. Finally we look at how to set up sequences of algorithms and how to control processing through the use of branches and filters.

5.2 Algorithm base class

Since a concrete algorithm object is an `Algorithm` object it may use all of the public and protected methods of the `Algorithm` base class. The base class has no protected or public data members, so in fact, these are the only methods that are available. Most of these methods are provided solely to make the implementation of derived algorithms easier. The base class has two main responsibilities: the initialization of certain internal pointers and the management of the properties of derived algorithm classes.

A part of the `Algorithm` base class definition is shown in Listing 5.1. Include directives, forward declarations and private member variables have all been suppressed. It declares a constructor and destructor; some methods of the `Algorithm` interface; several accessors to services that a concrete algorithm will almost certainly require; a method to create a sub algorithm, the two methods of the `IPROPERTY` interface; and a whole series of methods for declaring properties.



Listing 5.1 The definition of the Algorithm base class.

```

1: class Algorithm : virtual public IAlgorithm,
2: public:
3: // Constructor and destructor
4: Algorithm( const std::string& name, ISvcLocator *svcLoc ) ;
5: virtual ~Algorithm();
6:
7: // IAlgorithm interface only partially implemented
8: StatusCode sysInitialize();
9: StatusCode sysExecute();
10: StatusCode sysFinalize();
11: StatusCode beginRun();
12: StatusCode endRun();
13: const std::string& name() const;
14:
15: virtual bool isExecuted() const;
16: virtual StatusCode setExecuted( bool state );
17: virtual StatusCode resetExecuted();
18: virtual bool isEnabled() const;
19: virtual bool filterPassed() const;
20: virtual StatusCode setFilterPassed( bool state );
21:
22: // Service accessors
23: template<class T> StatusCode service( const std::string& name, T*& svc,
24: bool createIf = false );
25: IMessageSvc* msgSvc() const;
26: IAuditorSvc* auditorSvc() const;
27: IDataProviderSvc* eventSvc() const;
28: IConversionSvc* eventChnSvc() const;
29: IDataProviderSvc* getSvc() const;
30: IConversionSvc* detChnSvc() const;
31: IHistogramSvc* histoSvc() const;
32: INTupleSvc* ntupleSvc() const;
33: IChronoStatSvc* chronoSvc() const;
34: IRndmGenSvc* randSvc() const;
35: IToolSvc* toolSvc() const;
36: ISvcLocator* serviceLocator() const;
37:
38: StatusCode createSubAlgorithm( const std::string& type,
39: const std::string& name, Algorithm*& pSubAlg );
40: std::vector<Algorithm*>* subAlgorithms() const;
41:
42: // IProperty interface
43: virtual StatusCode setProperty( const Property& p );
44: virtual StatusCode setProperty( std::istream& s );
45: virtual StatusCode setProperty( const std::string& n,
46: const std::string& v );
47: virtual StatusCode getProperty( Property* p ) const;
48: const Property& getProperty( const std::string& name ) const;
49: virtual StatusCode getProperty( const std::string& n,
50: std::string& v ) const;
51: const std::vector<Property*>* getProperties() const;

```



Listing 5.1 The definition of the Algorithm base class.

```

49: StatusCode setProperties();
50: template <class T>
51: StatusCode declareProperty(const std::string& name, T& property);
52: StatusCode declareRemoteProperty(const std::string& name,
53: IProperty* rsvc, const std::string& name = ""); const;
54: // Methods for IInterface
55: unsigned long addrRef();
56: unsigned long release();
57: StatusCode queryInterface(const IID& riid, void**);
58: protected:
59: bool isInitialized() const;
60: void setInitialized();
61: bool isFinalized() const;
62: void setFinalized();
63: private:
64: // Data members not shown
65: Algorithm(const Algorithm& a); // NO COPY ALLOWED
66: Algorithm& operator=(const Algorithm& rhs); // NO ASSIGNMENT ALLOWED;

```

Constructor and Destructor The base class has a single constructor which takes two

arguments: The first is the name that will identify the algorithm object being instantiated and the second is a pointer to one of the interfaces implemented by the application manager: ISvcLocator. This interface may be used to request special services that an algorithm may wish to use, but which are not available via the standard accessor methods (below).

The IAlgorithm interface The base class only partially implements this interface: the three pure virtual methods initialize(), execute() and finalize() must be implemented by a derived algorithm: these are where the algorithm does its useful work and are discussed in more detail in section 5.3. The base class provides default implementations of the methods beginRun() and endRun(), and the accessor name() which returns the algorithm's identifying name. The methods sysInitialize(), sysFinalize(), sysExecute() are used internally by the framework; they are not virtual and may not be overridden.

Service accessor methods Lines 25 to 35 declare accessor methods which return pointers to key service interfaces. These methods are available for use only after the Algorithm base class has been initialized, i.e. they may not be used from within a concrete algorithm constructor, but may be used from within the initialize() method (see Section 5.3.3). The services and interface types to which they point are self explanatory. Services may be located by name using the templated service() function in line 23 or by using the serviceLocator() accessor method on line 36, as described in Section 11.2. Line 24 declares a facility to modify the message output level from within the code (the message service is described in Section 11.4).

Creation of sub algorithms The methods on lines 38 to 39 are intended to be used by a derived class to manage sub-algorithms, as discussed in section 5.4.

Declaration and setting of properties A concrete algorithm must declare its properties to the framework using the templated declareProperty method (line 50), as discussed in Section 5.3.2 and Section 11.3.1. The Algorithm base class then uses the setProperties() method (line 49) to tell the framework to set these properties to the values defined in the job options file. The methods in lines 42 to 48 can later be used to access and modify the values of specific properties, as explained in Section 11.3.2.



Filtering The methods in lines 14 to 19 are used by sequencers and filters to access the state of the algorithm, as discussed in Section 5.5.

5.3 Derived algorithm classes

In order for an algorithm object to do anything useful it must be specialised, i.e. it must extend (inherit from, be derived from) the `Algorithm` base class. In general it will be necessary to implement the methods of the `Algorithm` interface, and declare the algorithm's properties to the property management machinery of the `Algorithm` base class. Additionally there is one non-obvious technical matter to cover, namely algorithm factories.

5.3.1 Creation (and algorithm factories)

A concrete algorithm class must specify a single constructor with the same parameter signature as the constructor of the base class.

In addition to this, a concrete algorithm factory must be provided. This is a technical matter which permits the application manager to create new algorithm objects without having to include all of the concrete algorithm header files. From the point of view of an algorithm developer it implies adding three lines into the implementation file, of the form:

```
#include "GaudiKernel/AlgFactory.h"
...
static const AlgFactory<ConcreteAlgorithm> s_factory;
const AlgFactory& ConcreteAlgFactory = s_factory;
```

where "`ConcreteAlgorithm`" should be replaced by the name of the derived algorithm class (see for example lines 10 and 11 in Listing 5.2 below).

5.3.2 Declaring properties

In general, a concrete algorithm class will have several data members which are used in the execution of the algorithm proper. These data members should of course be initialized in the constructor, but if this was the only mechanism available to set their value it would be necessary to recompile the code every time you wanted to run with different settings. In order to avoid this, the framework provides a mechanism for setting the values of member variables at run time.

The mechanism comes in two parts: the declaration of properties and the setting of their values. As an example consider the class `TriggerDecision` in Listing 5.2 which has a number of variables whose value we would like to set at run time.

The default values for the variables are set within the constructor (within an initialiser list). To declare them as properties it suffices to call the `declareProperty()` method. This method



Listing 5.2 Declaring member variables as properties.

```
1: //----- In the header file -----//
2: class TriggerDecision : public Algorithm {
3:
4: private:
5:   bool m_passAllMode;
6:   int m_muonCandidateCut;
7:   std::vector m_ECALEnergyCuts;
8: }
9: //----- In the implementation file -----//
10: static const AlgFactory<TriggerDecision> s_factory;
11: const AlgFactory& TriggerDecisionFactory = s_factory;
12:
13: TriggerDecision::TriggerDecision(std::string name, ISvcLocator *pSL) {
14:   Algorithm(name, pSL), m_passAllMode(false), m_muonCandidateCut(0) {
15:     m_ECALEnergyCuts.push_back(0.0);
16:     m_ECALEnergyCuts.push_back(0.6);
17:
18:   declareProperty("PassAllMode", m_passAllMode);
19:   declareProperty("MuonCandidateCut", m_muonCandidateCut);
20:   declareProperty("ECALEnergyCuts", m_ECALEnergyCuts);
21: }
22:
23: StatusCode TriggerDecision::initialize() {
24: }
```

is templated to take an `std::string` as the first parameter and a variety of different types for the second parameter. The first parameter is the name by which this member variable shall be referred to, and the second parameter is a reference to the member variable itself.

In the example we associate the name "PassAllMode" to the member variable `m_passAllMode`, and the name "MuonCandidateCut" to `m_muonCandidateCut`. The first is of type boolean and the second an integer. If the job options service (described in Section 11.3 on page 113) finds an option in the job options file belonging to this algorithm and whose name matches one of the names associated with a member variable, then that member variable will be set to the value specified in the job options file.

5.3.3 Implementing IAlgorithm

Any concrete algorithm must implement the three pure virtual methods `initialize()`, `execute()` and `finalize()` of the `IAlgorithm` interface. For a top level algorithm, i.e. one controlled directly by the application manager, the methods are invoked as is described in section 4.6. This dictates what it is useful to put into each of the methods.

Initialization Figure 5.1 shows an example trace of the initialization phase. In a standard job the application manager will initialize all top level algorithms exactly once before reading any event data. It does this by invoking the `sysInitialize()` method of each top-level algorithm in turn, in which the framework takes care of setting up internal references to standard services and to set the algorithm properties (using the mechanism described in Section 11.3.1 on page 113). At the end, `sysInitialize()` calls the `initialize()` method, which can be used to do such things as creating histograms, or creating sub-algorithms if



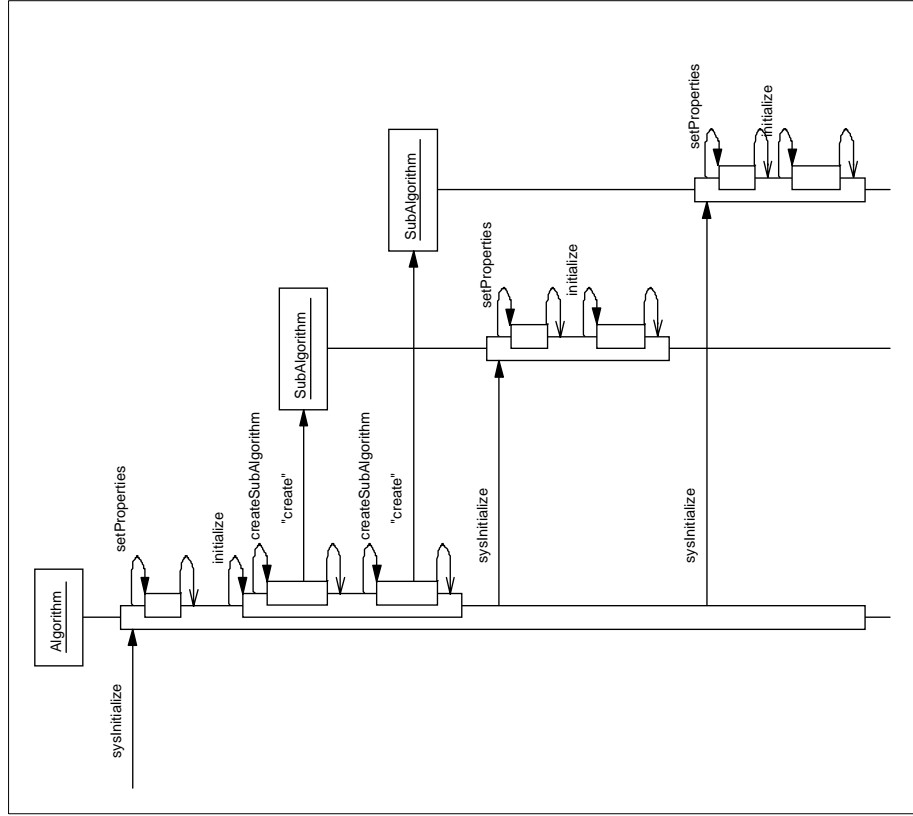


Figure 5.1 Algorithm initialization.

required (sub-algorithms are discussed in Section 5.4). If an algorithm fails to initialize it should return `StatusCode::FAILURE`. This will cause the job to terminate.

Execution The guts of the algorithm class is in the `execute()` method. For top level algorithms this will be called once per event for each algorithm object in the order in which they were declared to the application manager. For sub-algorithms (Section 5.4) the control flow may be as you like; you may call the `execute()` method once, many times or not at all.



Just because an algorithm derives from the `Algorithm` base class does not mean that it is limited to using or overriding only the methods defined by the base class. In general, your code will be much better structured (i.e. understandable, maintainable, etc.) if you do not, for example, implement the `execute()` method as a single block of 100 lines, but instead define your own utility methods and classes to better structure the code.

If an algorithm fails in some manner, e.g. a fit fails to converge, or its data is nonsense it should return from the `execute()` method with `StatusCode::FAILURE`. This will cause the application manager to stop processing events and end the job. This default behaviour can be modified by setting the `<myAlgorithm>.ErrorMaxJob` option to something greater than 1. In this case a message will be printed, but the job will continue as if there had been no error, and just increment an error count. The job will only stop if the error count reaches the `ErrorMaxLimit` set in the job option.

The framework (the `Algorithm` base class) calls the `execute()` method within a try/catch clause. This means that any exception not handled in the execution of an `Algorithm` will be caught at the level of `sysExecute()` implemented in the base class. The behaviour on these exceptions is identical to that described above for errors.

Finalization The `finalize()` method is called at the end of the job. It can be used to analyse statistics, fit histograms, or whatever you like. Similarly to initialization, the framework invokes a `sysFinalize()` method which in turn invokes the `finalize()` method of the algorithm and of any sub-algorithms.

Optionally, the methods `beginRun()` and `endRun()` can also be implemented. These are called at the beginning and the end of the event loop respectively.

Monitoring of the execution (e.g. cpu usage) of each `Algorithm` instance is performed by auditors under control of the Auditor service (described in Section 11.7 on page 129). This monitoring can be turned on or off with the boolean properties `AuditInitialize`, `AuditExecute`, `AuditFinalize`.

The following is a list of things to do when implementing an algorithm.

- Derive your algorithm from the `Algorithm` base class.
- Provide the appropriate constructor and the three methods `initialize()`, `execute()` and `finalize()`.
- Make sure you have implemented a factory by adding the magic two lines of code (see Section 5.3.1).

5.4 Nesting algorithms

The application manager is responsible for initializing, executing once per event, and finalizing the set of top level algorithms, i.e. the set of algorithms specified in the job options file. However such a simple linear structure is very limiting. You may wish to execute some algorithms only for specific types of event, or you may wish to "loop" over an algorithm's `execute` method. Within the Gaudi application framework the way to have such control is via the nesting of algorithms or through algorithm sequences (described in section 5.5). A nested



(or sub-) algorithm is one which is created by, and thus belongs to and is controlled by, another algorithm (its parent) as opposed to the application manager. In this section we discuss a number of points which are specific to sub-algorithms.

In the first place, the parent algorithm will need a member variable of type `Algorithm*` (see the code fragment below) in which to store a pointer to the sub-algorithm.

```
Algorithm* m_pSubAlgorithm; // Pointer to the sub algorithm
// Must be a member variable of the parent class
std::string type; // Type of sub algorithm
std::string name; // Name to be given to subAlgorithm
StatusCode sc; // Status code returned by the call
sc = createSubAlgorithm(type, name, Algorithm* & m_pSubAlgorithm);
```

The sub-algorithm itself is created by invoking the `createSubAlgorithm()` method of the `Algorithm` base class. The parameters passed are the type of the algorithm, its name and a reference to the pointer which will be set to point to the newly created sub-algorithm. Note that the name passed into the `createSubAlgorithm()` method is the same name that should be used within the job options file for specifying algorithm properties.

The algorithm type (i.e. class name) string is used by the application manager to decide which factory should create the algorithm object.

The execution of the sub-algorithm is entirely the responsibility of the parent algorithm whereas the `initialize()` and `finalize()` methods are invoked automatically by the framework as shown in Figure 5.1. Similarly the properties of a sub-algorithm are also automatically set by the framework.

Note that the `createSubAlgorithm()` method returns a pointer to an `Algorithm` object, not an `IAAlgorithm` interface. This means that you have access to the methods of both the `Algorithm` and `IProperty` interfaces, and consequently as well as being able to call `execute()` etc. you may also change the properties of a sub-algorithm during the main event loop as explained in Section 11.3.2. Note also that the vector of pointers to the sub-algorithms is available via the `subAlgorithms()` method.

5.5 Algorithm sequences, branches and filters

A physics application may wish to execute different algorithms depending on the physics signature of each event, which might be determined at run-time as a result of some reconstruction. This capability is supported in Gaudi through sequences, branches and filters. A sequence is a list of Algorithms. Each Algorithm may make a filter decision, based on some characteristics of the event, which can either allow or bypass processing of the downstream algorithms in the sequence. The filter decision may also cause a branch whereby a different downstream sequence of Algorithms will be executed for events that pass the filter decision relative to those that fail it. Eventually the particular set of sequences, filters and branches might be used to determine which of multiple output destinations each event is written to (if at all). This capability is not yet implemented but is planned for a future release of Gaudi.



A `Sequencer` class is available in the `GaudiAlg` package which manages algorithm sequences using filtering and branching protocols which are implemented in the `Algorithm` class itself. The list of Algorithms in a `Sequencer` is specified through the `Members` property. Algorithms can call `setFilterPassed(true/false)` during their `execute()` function. Algorithms in the membership list downstream of one that sets this flag to `false` will not be executed, unless the `StopOverride` property of the `Sequencer` has been set, or the filtering algorithm itself is of type `Sequencer` and its `BranchMembers` property specifies a branch with downstream members. Please note that, if a sub-algorithm is of type `Sequencer`, the parent algorithm must call the `resetExecuted()` method of the sub-algorithm before calling the `execute()` method, otherwise the sequence will only be executed once in the lifetime of the job!

An algorithm instance is executed only once per event, even if it appears in multiple sequences. It may also be enabled or disabled, being enabled by default. This is controlled by the `Enable` property. Enabling and disabling of algorithm instances is a capability that is designed for a future release of Gaudi that will include an interactive scripting language.

The filter passed or failed logic for a particular Algorithm instance in a sequence may be inverted by specifying the `invert` optional flag in the `Members` list for the `Sequencer` in the job options file.

A `Sequencer` will report filter success if either of its main and branch member lists succeed. The two cases may be differentiated using the `Sequencer.branchFilterPassed()` boolean function. If this is set true, then the branch filter was passed, otherwise both it and the main sequence indicated failure.

The following examples illustrate the use of sequences with filtering and branching.

5.5.1 Filtering example

Listing 5.3 is an extract of the job options file of the `AlgSequencer` example; a `Sequencer` instance is created (line 2) with two members (line 5); each member is itself a `Sequencer`, implementing the sequences set up in lines 7 and 8, which consist of `Prescaler`, `EventCounter` and `HelloWorld` algorithms. The `StopOverride` property of the `TopSequence` is set to true, which causes both sequences to be executed, even if the first one indicates a filter failure.

The `Prescaler` and `EventCounter` classes are example algorithms distributed with the `GaudiAlg` package. The `Prescaler` class acts as a filter, passing the fraction of events specified by the `PercentPass` property (as a percentage). The `EventCounter` class just prints each event as it is encountered, and summarizes at the end of job how many events were seen. Thus at the end of job, the `Counter1` instance will report seeing 50% of the events, while the `Counter2` instance will report seeing 10%.

Note the same instance of the `HelloWorld` class appears in both sequences. It will be executed in `Sequence1` if `Prescaler1` passes the event. It will be executed in `Sequence2` if `Prescaler2` passes the event only if `Prescaler1` failed it.



Listing 5.3 Example job options using Sequencers demonstrating filtering

```
1: ApplicationMgr.DLLs += { "GaudiAlg" };
2: ApplicationMgr.TopAlg = { "Sequencer/TopSequence" };
3:
4: // Setup the next level sequencers and their members
5: TopSequence.Members = {"Sequencer/Sequence1", "Sequencer/Sequence2"};
6: TopSequence.StopOverride = true;
7: Sequence1.Members = {"Prescaler/Prescaler1", "HelloWorld",
  "EventCounter/Counter1"};
8: Sequence2.Members = {"Prescaler/Prescaler2", "HelloWorld",
  "EventCounter/Counter2"};
9:
10: Prescaler1.PercentPass = 50.;
11: Prescaler2.PercentPass = 10.;
```

5.5.2 Sequence branching

Listing 5.4 illustrates the use of explicit branching. The BranchMembers property of the Sequencer specifies some algorithms to be executed if the algorithm that is the first member of the branch (which is common to both the main and branch membership lists) indicates a filter failure. In this example the EventCounter instance Counter1 will report seeing 80% of the events, whereas Counter2 will report seeing 20%.

Listing 5.4 Example job options using Sequencers demonstrating branching

```
1: ApplicationMgr.DLLs += { "GaudiAlg" };
2: ApplicationMgr.TopAlg = { "Sequencer" };
3:
4: // Setup the next level sequencers and their members
5: Sequencer.Members = {"HelloWorld", "Prescaler",
  "EventCounter/Counter1"};
6: Sequencer.BranchMembers = {"Prescaler", "EventCounter/Counter2"};
7:
8: Prescaler.PercentPass = 80.;
```

Listing 5.5 illustrates the use of inverted logic. It achieves the same goal as the example in Listing 5.4 through use of two sequences with the same instance of a Prescaler filter, but where the second sequence contains inverted logic for the single instance.

Listing 5.5 Example job options using Sequencers demonstrating inverted logic

```
1: ApplicationMgr.DLLs += { "GaudiAlg" };
2: ApplicationMgr.TopAlg = { "Sequencer/Seq1", "Sequencer/Seq2" };
3:
4: // Setup the next level sequencers and their members
5: Seq1.Members = {"HelloWorld", "Prescaler", "EventCounter/Counter1"};
6: Seq2.Members = {"HelloWorld", "Prescaler:invert",
  "EventCounter/Counter2"};
7:
8: Prescaler.PercentPass = 80.;
```



Chapter 6

Accessing data

6.1 Overview

The data stores are a key component in the application framework. All data which comes from persistent storage, or which is transferred between algorithms, or which is to be made persistent must reside within a data store. In this chapter we use a trivial event data model to look at how to access data within the stores, and also at the `DataObject` base class and some container classes related to it.

We also cover how to define your own data types and the steps necessary to save newly created objects to disk files. The writing of the converters necessary for the latter is covered in Chapter 13.

6.2 Using the data stores

There are four data stores currently implemented within the Gaudi framework: the event data store, the detector data store, the histogram store and the n-tuple store. Event data is the subject of this chapter. The other data stores are described in chapters 8, 9 and 10 respectively. The stores themselves are no more than logical constructs with the actual access to the data being via the corresponding services. Both the event data service and the detector data service implement the same `IDataProviderSvc` interface, which can be used by algorithms to retrieve and store data. The histogram and n-tuple services implement extended versions of this interface (`IHistogramSvc`, `INTupleSvc`) which offer methods for creating and manipulating histograms and n-tuples, in addition to the data access methods provided by the other two stores.

Only objects of a type derived from the `DataObject` base class may be placed directly within a data store. Within the store the objects are arranged in a tree structure, just like a Unix file system. As an example consider Figure 6.1 which shows the trivial transient event data model of the `ROOTIO` example. An object is identified by its position in the tree expressed as a string



such as `"/Event"`, or `"/Event/MyTracks"`. In principle the structure of the tree, i.e. the set of all valid paths, may be deduced at run time by making repeated queries to the event data service, but this is unlikely to be useful in general since the structure will be largely fixed.




Figure 6.4 The structure the event data model of the RootIO example.

Interactions with the data stores are usually via the `IDataProviderSvc` interface, whose key methods are shown in Listing 6.1

Listing 6.1 Some of the key methods of the `IDataProviderSvc` interface.

```

StatusCode findObject(const std::string& path, DataObject& pObject);
StatusCode findObject(DataObject* node, const std::string& path,
                    DataObject* pObject);
StatusCode retrieveObject(const std::string& path, DataObject* pObject);
StatusCode retrieveObject(DataObject* node, const std::string& path,
                    DataObject* pObject);
StatusCode registerObject(const std::string path, DataObject* pObject);
StatusCode registerObject(DataObject* node, DataObject* pObject);
  
```

The first four methods are for retrieving a pointer to an object that is already in the store. How the object got into the store, whether it has been read in from a persistent store or added to the store by an algorithm, is irrelevant.

The `find` and `retrieve` methods come in two versions: one version uses a full path name as an object identifier, the other takes a pointer to a previously retrieved object and the name of the object to look for below that node in the tree.

Additionally the `find` and `retrieve` methods differ in one important respect: the `find` method will look in the store to see if the object is present (i.e. in memory) and if it is not will return a null pointer. The `retrieve` method, however, will attempt to load the object from a persistent store (database or file) if it is not found in memory. Only if it is not found in the persistent data store will the method return a null pointer (and a bad status code of course).

Navigation through the tree structure of the data store is possible via the `IDataManagerSvc` interface of the data service, as described for example in

http://cern.ch/lhcb-comp/Frameworks/Gaudi/Gaudi_v9/Changes_cookbook.pdf.



6.3 Using data objects

Whatever the concrete type of the object you have retrieved from the store the pointer which you have is a pointer to a `DataObject`, so before you can do anything useful with that object you must cast it to the correct type, for example:

```

1: typedef ObjectVector<MyTrack> MyTrackVector;
2: DataObject *pObject;
3:
4: StatusCode sc = eventSvc()->retrieveObject("/Event/MyTracks", pObject);
5: if( sc.IsFailure() )
6:   return sc;
7:
8: MyTrackVector *tv = 0;
9: try {
10:   tv = dynamic_cast<MyTrackVector*>(pObject);
11: } catch(...) {
12:   // Print out an error message and return
13: }
14: // tv may now be manipulated.
  
```

The typedef on line 1 is just to save typing; in what follows we will use the two syntaxes interchangeably. After the `dynamic_cast` on line 10 all of the methods of the `MyTrackVector` class become available. If the object which is returned from the store does not match the type to which you try to cast it, an exception will be thrown. If you do not catch this exception it will be caught by the algorithm base class, and the program will stop, probably with an obscure message. A more elegant way to retrieve the data involves the use of Smart Pointers - this is discussed in section 6.8

The last two methods shown in Listing 6.1 are for registering objects into the store. Suppose that an algorithm creates objects of type `UDO` from, say, objects of type `MyTrack` and wishes to place these into the store for use by other algorithms. Code to do this might look something like:

Listing 6.2 Registering of objects into the event data store

```

1: UDO* pO; // Pointer to an object of type UDO (derived from DataObject)
2: StatusCode sc;
3:
4: pO = new UDO;
5: sc = eventSvc()->registerObject("/Event/tmp", "OK", pO);
6:
7: // THE NEXT LINE IS AN ERROR, THE OBJECT NOW BELONGS TO THE STORE
8: delete pO;
9:
10: UDO autoPO;
11: // ERROR: AUTOMATIC OBJECTS MAY NOT BE REGISTERED
12: sc = eventSvc()->registerObject("/Event/tmp", "notOK", autoPO);
  
```



Once an object is registered into the store, the algorithm which created it relinquishes ownership. In other words the object should not be `deleted`. This is also true for objects which are contained within other objects, such as those derived from or instantiated from the `ObjectVector` class (see the following section). Furthermore objects which are to be registered into the store must be created on the heap, i.e. they must be created with the `new` operator.

6.4 Object containers

As mentioned before, all objects which can be placed directly within one of the stores must be derived from the `DataObject` class. There is, however, another (indirect) way to store objects within a store. This is by putting a set of objects (themselves not derived from `DataObject` and thus not directly storable) into an object which is derived from `DataObject` and which may thus be registered into a store.

An object container base class is implemented within the framework and a number of templated object container classes may be implemented in the future. For the moment, two "concrete" container classes are implemented: `ObjectVector<T>` and `ObjectList<T>`. These classes are based upon the STL classes and provide mostly the same interface. Unlike the STL containers which are essentially designed to hold objects, the container classes within the framework contain only pointers to objects, thus avoiding a lot of memory to memory copying.

A further difference with the STL containers is that the type `T` cannot be anything you like. It must be a type derived from the `ContainedObject` base class, see Figure 6.2. In this way all "contained" objects have a pointer back to their containing object. This is required, in particular, by the converters for dealing with links between objects. A ramification of this is that container objects may not contain other container objects (without the use of multiple inheritance).

As mentioned above, objects which are contained within one of these container objects may not be located, or registered, individually within the store. Only the container object may be located via a call to `FindObject()` or `RetrieveObject()`. Thus with regard to interaction with the data stores a container object and the objects that it contains behave as a single object.

The intention is that "small" objects such as clusters, hits, tracks, etc. are derived from the `ContainedObject` base class and that in general algorithms will take object containers as their input data and produce new object containers of a different type as their output.

The reason behind this is essentially one of optimization. If all objects were treated on an equal footing, then there would be many more accesses to the persistent store to retrieve very small objects. By grouping objects together like this we are able to have fewer accesses, with each access retrieving bigger objects.

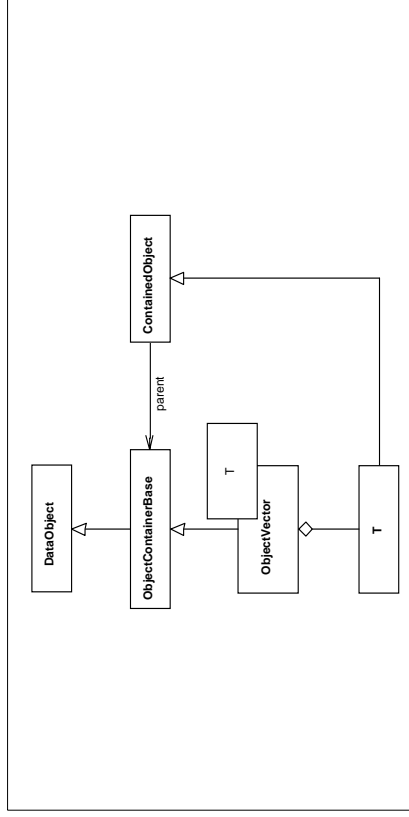


Figure 6.2 The relationship between the `DataObject`, `ObjectVector` and `ContainedObject` classes.

6.5 Using object containers

The code fragment below shows the creation of an object container. This container can contain pointers to objects of type `MyTrack` and only to objects of this type (including derived types). An object of the required type is created on the heap (i.e. via a call to `new`) and is added to the container with the standard STL call.

```

ObjectVector<MyTrack> trackContainer;
MyTrack* h1 = new MyTrack;
trackContainer.push_back(h1);

```

After the call to `push_back()` the `MyTrack` object "belongs" to the container. If the container is registered into the store, the hits that it contains will go with it. Note in particular that if you delete the container you will also delete its contents, i.e. all of the objects pointed to by the pointers in the container.

Removing an object from a container may be done in two semantically different ways. The difference being whether on removal from a container the object is also deleted or not.

Removal with deletion may be achieved in several ways (following previous code fragment):

```

trackContainer.pop_back();
trackContainer.erase( end() );
delete h1;

```



The method `pop_back()` removes the last element in the container, whereas `erase()` may be used to remove any other element via an iterator. In the code fragment above it is used to remove the last element also.

Deleting a contained object, the third option above, will automatically trigger its removal from the container. This is done by the destructor of the `ContainedObject` base class.

If you wish to remove an object from the container without destroying it (the second possible semantic) use the `release()` method:

```
trackContainer.release(h1);
```

Since the fate of a contained object is so closely tied to that of its container, life would become more complex if objects could belong to more than one container. Suppose that an object belonged to two containers, one of which was deleted. Should the object be deleted and removed from the second container, or not deleted? To avoid such issues an object is allowed to belong to a single container only.

If you wish to move an object from one container to another, you must first remove it from one and then add to the other. However, the first operation is done implicitly for you when you try to add an object to a second container:

```
container1.push_back(h1); // Add to first container  
container2.push_back(h1); // Move to second container  
// Internally invokes release().
```

Since the object `h1` has a link back to its container, the `push_back()` method is able to first follow this link and invoke the `release()` method to remove the object from the first container, before adding it into the second.

In general your first exposure to object containers is likely to be when retrieving data from the event data store. The sample code in Listing 6.3 shows how, once you have retrieved an object container from the store you may iterate over its contents, just as with an STL vector.

Listing 6.3 Use of the `ObjectVector` templated class.

```
1: typedef ObjectVector<MyTrack> MyTrackVector;  
2: MyTrackVector *tracks;  
3: MyTrackVector::iterator it;  
4:  
5: for( it = tracks->begin(); it != tracks->end(); it++ ) {  
6:     // Get the energy of the track and histogram it  
7:     double energy = (*it)->fourMomentum().e();  
8:     m_hEnergyDist->fill( energy, 1. );  
9: }
```



The variable `tracks` is set to point to an object in the event data store of type: `ObjectVector<MyTrack>` with a dynamic cast (not shown above). An iterator (i.e. a pointer-like object for looping over the contents of the container) is defined on line 3 and this is used within the loop to point consecutively to each of the contained objects. In this case the objects contained within the `ObjectVector` are of type "pointer to `MyTrack`". The iterator returns each object in turn and in the example, the energy of the object is used to fill a histogram.

6.6 Data access checklist

A little reminder:

- Do not delete objects that you have registered.
- Do not delete objects that are contained within an object that you have registered.
- Do not register local objects, i.e. objects NOT created with the `new` operator.
- Do not delete objects which you got from the store via `findObject()` or `retrieveObject()`.
- Do delete objects which you create on the heap, i.e. by a call to `new`, and which you do not register into a store.

6.7 Defining Data Objects

If you want to create a new data object in the transient or persistent stores of Gaudi, you will have to define the structure of this object. This structure will be defined by C++-classes. These classes in general look very similar to each other; mainly they define the members of the class, which are either data values or which point to another class (eg. via a Smart Reference - see Section 6.9). For each of these members there is usually a `set-` and a `get-`method and some more stuff for the Smart Reference Vectors.

The writing of these classes is a tedious task and having to write this redundant information many times, of course, also bears the risk of many unnecessary typos. To overcome this problem one may use XML in conjunction with the `Gaudi::ObjDesc` package to describe the data-objects. There were two key issues which led to the development of this description language:

- The core information of a data-object lies in the members of the class, most of the rest is redundant information which can be produced automatically around the members.
- There is a lot of information which also must be provided, but which has a default-value in most of the cases.

The information provided in the XML files can be used to produce not only the object information in the classes but also reflection information about the objects (see Section II.11, "The Gaudi Introspection Service"). Future releases may also produce, e.g., converters, or a description of the object in other languages.



As an example, the following XML code describes an `MCParticle` class:

Listing 6.4 Part of the XML description of the `MCParticle` class

```

1: <class name='MCParticle' author='pavel Binko' id='210' desc='The Monte
   Carlo particle kinematics information'>
2:   <base name='ContainedObject' />
3:   <attribute name='subEvtID' type='short' desc='Sub-event ID' />
4:   <relation name='originMCVertex' multiplicity='1' type='MCVertex'
   desc='Pointer to origin vertex' />
5:   <relation name='decayMCVertices' multiplicity='M' type='MCVertex'
   desc='Vector of pointers to decay vertices' />
6: </class>

```

All of the elements in this listing (eg. `<class>`, `<attributes>`, `<relation>`) have several attributes with default values (eg. for relations and attributes " `setMethod='TRUE'` "), which don't have to be mentioned explicitly. If one doesn't want to use the default, the only thing that has to be done is to set the corresponding attribute to another value. There are also several hooks which can be applied eg. to define your own methods if they were not created automatically. The complete syntax of this description language can be found on the web at (<http://cern.ch/lhcb-comp/Frameworks/DataDictionary/>).

Once a set of data-objects is defined, the XML file has to be saved to the `xml` directory of the package. The production of the C++ header files containing the object description can be automated by adding a line to the CMT requirements file of the package, as shown for example below:

```
document obj2doth LHCBEventObj2Doth ../xml/LHCBEvent.xml
```

Another possibility is to produce the information by hand with the tools of the `GaudiObjDesc` package (eg. `GODCpHeaderWriter.exe`) and then compile it.

6.7.1 The class ID

The class definition on line 1 of Listing 6.4 contains an 'id' attribute. This class identifier is required if the objects of this class are to be made persistent. It is used by the data persistency services to make the translation between the transient and persistent representations of the object, using the conversoin mechanism described in Chapter 13. For this mechanism to work, these identifiers must uniquely identify the class and no two classes may have the same identifier. The list of currently allocated class IDs is available from the web at <http://cern.ch/lhcb-comp/Frameworks/EventModel/CLID.htm>. This web page also describes the procedure for allocating new class IDs.

Types which are derived from `ContainedObject` must have a class ID in the range of an unsigned short. Contained objects may only reside in the store when they belong to a container, eg. an `ObjectVector<T>` which is registered into the store. The class identifier of a concrete object container class is calculated (at run time) from the type of the objects which it contains, by setting bit 16.



6.8 The SmartDataPtr/SmartDataLocator utilities

The usage of the data services is simple, but extensive status checking and other things tend to make the code difficult to read. It would be more convenient to access data items in the store in a similar way to accessing objects with a C++ pointer. This is achieved with smart pointers, which hide the internals of the data services.

6.8.1 Using SmartDataPtr/SmartDataLocator objects

The `SmartDataPtr` and a `SmartDataLocator` are smart pointers that differ by the access to the data store. `SmartDataPtr` first checks whether the requested object is present in the transient store and loads it if necessary (similar to the `retrieveObject` method of `IDataProviderSvc`). `SmartDataLocator` only checks for the presence of the object but does not attempt to load it (similar to `findObject`).

Both `SmartDataPtr` and `SmartDataLocator` objects use the data service to get hold of the requested object and deliver it to the user. Since both objects have similar behaviour and the same user interface, in the following only the `SmartDataPtr` is discussed.

An example use of the `SmartDataPtr` class is shown in Listing 6.5.

Listing 6.5 Use of a `SmartDataPtr` object.

```

1: StatusCode myAlgo::execute() {
2:   MsgStream log(msgSvc(), name());
3:   SmartDataPtr<Event> evt(eventSvc(), "/Event");
4:   if ( evt ) {
5:     // Print the event number
6:     log << MSG::INFO << " Run:" << evt->run()
7:     << " Event:" << evt->event() << endl;
8:   }
9:   else {
10:    log << MSG::ERROR << "Error accessing event" << endl;
11:    return StatusCode::FAILURE;
12:   }
13: }

```

The `SmartDataPtr` class can be thought of as a normal C++ pointer having a constructor. It is used in the same way as a normal C++ pointer.

The `SmartDataPtr` and `SmartDataLocator` offer a number of possible constructors and operators to cover a wide range of needs when accessing data stores. Check the online reference documentation [2] for up-to-date information concerning the interface of these utilities.



6.9 Smart References and Smart Reference Vectors

It is foreseen that data objects in the transient data stores can reference other objects in the same data store. This relationship can be described in the XML data description using the 'relation' attribute of the class definition, as shown on line 4 of Listing 6.4

The current implementation of these relationships use 'Smart References' and 'Smart Reference Vectors'. These are similar to smart pointers, they provide safe data access and automate the loading on demand of referenced data, and are used instead of C++ pointers. For example, suppose that `MCParticles` are already loaded but `MCVertices` are not, and that an algorithm dereferences a variable pointing to the origin vertex: if a smart reference is used, the `MCVertices` would be loaded automatically and only after that would the variable be dereferenced. If a C++ plain pointer were used instead, the program would crash. Smart references provide an automatic conversion to a pointer to the object and load the object from the persistent medium during the conversion process.

The XML code in Listing 6.4 will generate Smart Reference and Smart Reference Vector declarations as shown below:

```
#include "/GaudiKernel/SmartRef.h"
#include "/GaudiKernel/SmartRefVector.h"
class MParticle {
private:
    /// Smart reference to origin vertex
    SmartRef<MCVertex> m_originMCVertex;
    /// Vector of smart references to decay vertices
    SmartRefVector<MCVertex> m_decayMCVertices;
public:
    /// Access the origin Vertex
    /// Note: When the smart reference is converted to MCVertex* the object
    /// will be loaded from the persistent medium.
    MCVertex* originMCVertex() { return m_originMCVertex; }
}
```

The syntax of usage of smart references is identical to plain C++ pointers. The Algorithm only sees a pointer to the `MCVertex` object:

```
#include "GaudiKernel/SmartDataPtr.h"
// Use a SmartDataPtr to get the MC particles from the event store
SmartDataPtr<MCParticleVector> particles(eventSvc(), "/Event/MC/MCParticles");
MCParticleVector::const_iterator iter;

// Loop over the particles to access the MCVertex via the SmartRef
for (iter = particles->begin(); iter != particles->end(); iter++) {
    MCVertex* originVtx = (*iter)->originMCVertex();
    if ( 0 != originVtx ) {
        std::cout << "Origin vertex = " << *(*iter) << std::endl; }
}
```

SmartRef offers a number of possible constructors and operators, see the online reference documentation [2].



6.10 Persistent storage of data

6.10.1 Saving event data to a persistent store

Suppose that you have defined your own data type as discussed in section 6.7. Suppose furthermore that you have an algorithm which creates instances of your object type which you then register into the transient event store. How can you save these objects for use at a later date?

You must do the following:

- Write the appropriate converter (see Chapter 13)
- Put some instructions (i.e. options) into the job option file (see Listing 6.6)
- Register your object in the store us usual, typically in the `execute()` method of your algorithm.

```
// myAlg implementation file
StatusCode myAlg::execute() {
    // Create a UDO object and register it into the event data store
    UDO* p = new UDO();
    eventSvc->registerObject( "/Event/myStuff/myUDO", p );
}
```

In order to actually trigger the conversion and saving of the objects at the end of the current event processing it is necessary to inform the application manager. This requires some options to be specified in the job options file:

Listing 6.6 Job options for output to persistent storage

```
ApplicationMgr.OutputStream = { "DstWriter" };

DstWriter.ItemList
= { "/Event#1", "/Event/MyTracks#1" };
DstWriter.EventDataSvc
= "EventDataSvc";
DstWriter.Output
= "DATAFILE= 'RootDst.root' TYP= 'ROOT'";

ApplicationMgr.DLLs
+= { "GaudiDb", "GaudiRootDb" };
ApplicationMgr.ExtSvc
+= { "DbEventChvSvc/RootEvtChvSvc" };
EventPersistenceSvc.ChvServices += { "RootEvtChvSvc" };
RootEvtChvSvc.DbType
= "ROOT";
```

The first option tells the application manager that you wish to create an output stream called "DstWriter". You may create as many output streams as you like and give them whatever name you prefer.

For each output stream object which you create you must set several properties. The `ItemList` option specifies the list of paths to the objects which you wish to write to this output stream. The number after the '#' symbol denotes the number of directory levels below the specified path which should be traversed. The (optional) `EventDataSvc` option specifies in



which transient data service the output stream should search for the objects in the `ItemList`, the default is the standard transient event data service `EventDataSvc`. The `Output` option specifies the name of the output data file and the type of persistence technology, `ROOT` in this example. The last three options are needed to tell the Application manager to instantiate the `ROOTEventSvc` and to associate the `ROOT` persistence type to this service.

An example of saving data to a `ROOT` persistent data store is available in the `ROOTIO` example distributed with the framework.

6.10.2 Reading event data from a persistent store

Suppose you want to read back the file written out in the previous section. To do this, your job options would look something like those described in Section 4.4.2 on page 28.

To read SICB files, please refer to Chapter 18.



Chapter 7

Modelling Event Data

The LHCb event data model is currently under review and discussion. Up to date information on the status of the discussions is available at <http://cern.ch/lhcb-comp/Frameworks/EventModel/>

The current recommendations and conventions for LHCb event data classes are described in reference [8]. Classes containing event data should be defined according to the procedure described in Section 6.7 on page 53.

The table at <http://cern.ch/lhcb-comp/Frameworks/EventModel/EventModelClasses.html> shows the current implementation status of the event data classes.



Chapter 8

Detector Description

8.1 Overview

In this chapter we describe how we make available to the physics application developed using the framework the information related to the detector that resides in the detector description database (DDDB). The DDDB is the persistent storage for all the versions of the detector data needed to describe and qualify the detecting apparatus in order to interpret the event data.

The final clients of the detector description are the algorithms that need this information in order to perform their job (reconstruction, simulation, etc.). To provide this information, there needs to be a sub-detector specific part that understands the sub-detector in question and uses a set of common services. The detector description we are providing in Gaudi is nothing else than a framework for developers to provide the specific detector information to algorithms by using as much as possible common or generic services. A brief introduction to this framework can be found in reference [9].

In the following sections we begin with an overview of the DDDB. We then discuss how to access the detector description data in the Gaudi transient detector data store. This is followed by a discussion of the features of the Gaudi detector description and of its structure. We then describe in detail how the detector description can be built and made persistent using the XML markup language and continue with an introduction to an XML editing tool which allows browsing and modification of the DDDB. Finally we describe some features of the so called Conditions Database, where the time-dependent elements of the DDDB can be stored.

8.2 Detector Description Database

The detector description database (DDDB), see Figure 8.1, includes a physical and a logical description of the detector. The physical description covers dimensions, shape and material of the different types of elements from which the detector is constructed. There is also



information which corresponds to each element which is actually manufactured and assembled into a detector, for example the positioning of each element. Both active and passive elements should be included. The description of active elements should allow for the specification of deficiencies (dead channels), alignment corrections, etc. and also detector response characteristics, e.g. energy normalization in calorimeters, drift velocity in gas chambers.

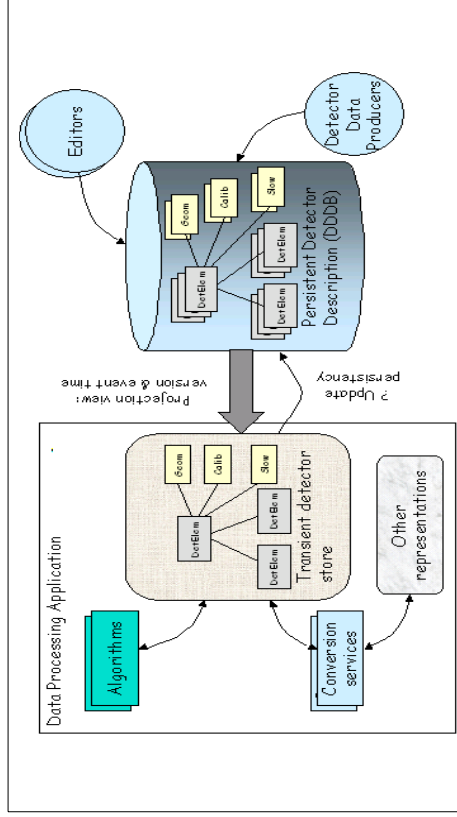


Figure 8.1 Overview of the Detector Description model.

The logical description provides two main functions. The first is a simplified access to particular parts of a physical detector description. This could be a hierarchical description where a given detector setup is composed of various sub-detectors, each of which is made up of a number of stations, modules or layers, etc. and there would be a simple way for a client to use this description to navigate to the information of interest. The second function of the logical description is to provide a means of detector element identification. This allows for different sets of information which are correlated to specific detector elements to be correctly associated with each other.

In a detector description, the definition of the detector elements and of the data associated to their physical description may vary over time, for instance due to real or hypothetical changes to the detector. Each such change should be recorded as a different version of the detector element. Additionally, it should be possible to capture, for an entire description, a version of each of the elements and to associate a name to that set. This is similar to the way CVS allows one to tag a set of files so that one does not need to know the independent version numbers for each file in the set.

The present version of the DDDB includes only the logical description of the detector and the physical description of its geometry and all required materials. This information is described using the XML language and stored in XML text files. The rest of the time varying "detector condition" data (alignment, calibration, readout, slow control and fast control) will eventually



be stored in the so called Conditions Database. This database is meant to record the validity range of each data element, allowing algorithms to process events using the detector description data valid at the given event time. Some features of the Conditions Database are described in Section 8.6.

8.3 Detector Data Transient Store

8.3.1 Structure of the transient store

The transient representation of the detector description is implemented as a standard Gaudi data store. Refer to Chapter 6 for general information on how to access data in stores. The structure of the detector transient store is represented in Figure 8.2.

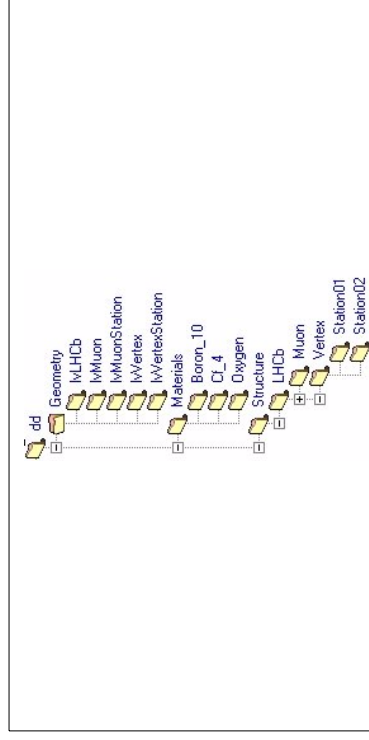


Figure 8.2 The structure of part of the LHCb detector data transient store.

In the present production version of the DDDB (package `Det/Xml/DDDB`) there are three top level catalogs in the transient store. The main catalog, called "Structure" contains the logical structure of the detector identified by the "setup name" i.e. "LHCb" containing the description of the detector and this catalog is used for identification and navigation purposes. Other catalogs are the palette of logical volumes and solids, called "Geometry", used for the geometry description and the palette of materials, called "Materials", used to describe the material properties of the solids needed for the detector simulation, etc.

Five more top level catalogs are foreseen in the transient store as palettes for the additional time varying information about the detector. The "Alignment" catalog will contain the geometrical data coming from regular surveys and necessary to determine the precise position of each subdetector. The "Calibration" catalog should be used to store the list of dead channels and the response parameters of each subdetector, such as energy normalization in calorimeters and drift velocity in gas chambers, while the "Readout" catalog is meant to contain the mapping of active detector elements to readout channels. Finally, the



"SlowControl" and "FastControl" catalogs will be used to store, respectively, online measurements of quantities such as pressures and temperatures, and information about the fast control system. A prototype implementation of this extended transient detector store can be found in package Ex/DetCondExample.

8.3.2 Accessing detector data

An algorithm that needs to access a given detector part uses the detector data service to locate the relevant `DetectorElement`. This operation can be generally done during the initialization phase of the algorithm. Contrary to the Event Data, the Detector Data store is not cleared for each event and the references to detector elements remain valid and are updated automatically during the execution of the program. Locating the relevant detector element is done using the standard `IDataProviderSvc` interface via the `getSvc()` accessor as shown in Listing 8.1

Listing 8.1 Retrieving a detector element by using smart pointers

```
SmartDataPtr<DetectorElement> vertex( detSvc(), "Structure/Velo/VStation01" );
if ( !vertex ) {
    // Error, detector element has not been retrieved
}
```

Remember that data are only loaded on demand, when they are used for the first time. The test on `vertex` not only tests if the data are present but also ensures that the data are loaded.¹

The user can retrieve an array of references to detector elements in a similar way. The code in Listing 8.2 can be used to prepare an array with pointers to all of the Muon stations. Here we use an STL vector of pointers to `DeMuonStation` objects to store the retrieved Muon stations.

Listing 8.2 Retrieving a vector of detector elements using smart references

```
std::vector<DeMuonStation*> d_stations;
SmartDataPtr<DetectorElement> stations( detSvc(),
    "Structure/LHCB/Muon/Stations" );
if ( !stations ) return StatusCode::FAILURE;
// Loop over all the muon stations found in the detector model
for ( DataObject::DirIterator d = stations->dirBegin();
    d != stations->dirEnd(); d++ )
{
    SmartDataPtr<DeMuonStation> s( detSvc(), (*d)->fullpath() );
    if ( !s ) return StatusCode::FAILURE;
    d_stations.push_back( s );
}
```

1. Note that, by comparison with Figure 8.2, the leading `"/dd/"` has been omitted from the full path name `"/dd/Structure/Velo/VStation01"` of the detector element. Either form is accepted; the short form must never begin with a `"/"` character.



Once more, don't forget to test each `DataObject`. Also note that you are not obliged to retrieve `DetectorElements`, you can actually retrieve whatever object inherits from it, provided that the corresponding converter exists (see Section 8.5 for a discussion on the persistent representation and converters). Here some `DeMuonStations` are converted

8.3.3 Using the DetectorElement class

The `DetectorElement` class implements the `IDetectorElement` interface, as shown in Figure 8.3.

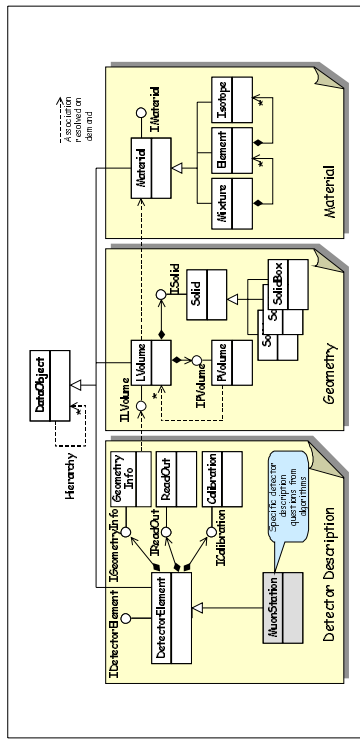


Figure 8.3 Simplified class diagram for the transient detector description

This interface offers six accessor methods to the different types of condition data for each detector element: `geometry()`, `alignment()`, `calibration()`, `readOut()`, `slowControl()` and `fastControl()`. In addition, `DetectorElement` implements the `IValidity` interface. This interface is used to check if the detector element is synchronized with the current event. If the detector element is no longer valid at the time the current event was generated, the corresponding `DataObject` in the Structure tree of the transient store must be updated from the persistent storage. In the current implementation it is not foreseen for end users to use this interface directly.

The accessor method `geometry()` gives access to geometry information offered by the interface of type `IGeometryInfo`. This interface allows the retrieval of a reference to a logical volume associated with the given detector element, its material property, the position in the geometrical hierarchy. In addition to that you can ask it questions like:

1. Transformation matrix from the Global to the Local Reference system
2. Transformation matrix from the Local to the Global Reference system
3. Perform transformation of point from the Global to Local Reference system
4. Perform transformation of point from the Local to Global Reference system
5. Name of daughter volume (of current volume) which contains given (global) point



6. Get a pointer to daughter volume which contains given (global) point
7. Name of daughter volume (at deeper hierarchical level) which contains given point
8. Get a pointer to daughter volume (on deeper level) to which contains given point
9. Get the exact full geometry location¹
10. Whether the given point is inside the associated logical volume or not
11. A pointer to the associated logical volume

As an example, the code fragment in Listing 8.3 shows how to obtain a number of geometrical objects from a given detector element:

Listing 8.3 Getting pointer to a logical volume and retrieving its various properties

```
SmartDataPtr<DetectorElement> vs (detsvc(), "Structure/LHCB/Velo/VStation01");
if ( !vs ) { return StatusCode::FAILURE; }

/// Report the material and its radiation length
IIVolume* stvol = vs->geometry()->slvolume();
log << MSG::INFO << vs->fullpath() << " is made of " << stvol->materialName()
<< " with radiation length " << stvol->material()->radiationlength()
<< endlreg;

/// Retrieve the shape information
const ISolid* stsolid = stvol->solid();

/// Get the rotation and translation
HepTransform3D strans = vs->geometry()->matrixInv();
HepRotation strot = strans.getRotation();
HepVector stvec = strans.getTranslation();
```

The remaining five accessor methods to the condition information of a detector element return pointers to the five corresponding interfaces of type IAlignment, ICalibration, IReadOut, ISlowControl and IFastControl, respectively. In the present implementation of the detector description package Det/DetDesc, these interfaces are much simpler than IGeometryInfo and do not yet offer specialized methods to handle the data they refer to. In particular, they are presently all equivalent to the interface from which they are derived, IConditionInfo, a simple interface which allows to retrieve a pointer to one "Condition" data object, containing all relevant condition data of the given type. Each Condition data object is meant to be loaded in the corresponding tree of the detector transient store: for instance, the slow control condition for Ecal may be stored in

```
/dd/SlowControl/Ecal/scEcal, just like the main logical volume for Ecal is stored in
/dd/Geometry/Ecal/IVeEcal.
```

A Condition is simply a DataObject implementing the IValidity and IUserParameter interfaces. It contains an arbitrary number of user-defined parameters and parameter vectors, taking numerical or string values and identified by their names, which can be retrieved through the IUserParameter interface. The IValidity interface deals with the validity range of a Condition object, allowing it to be updated when it becomes invalid.

1. This operation can be time consuming!



As an example, the code fragment in Listing 8.4 shows how to retrieve and print the parameters of the alignment Condition for a given detector element. Remember that also Conditions, like all other DataObjects, are only loaded on demand, when they are used for the first time: the test on aIEcal not only tests if the data are present but also ensures that the data are loaded. Eventually, the framework will also take care to automatically keep up-to-date all Condition objects used by the algorithms; this is discussed more in detail in Section 8.6. The test on ecalAlignmentInfo, instead, is only necessary because a valid pointer to an IAlignment interface is currently retrieved only for those DetectorElements for which a test version of the alignment information has been defined:

Listing 8.4 Getting a pointer to a Condition for alignment data and retrieving its various properties

```
// Retrieve the Hcal detector element
SmartDataPtr<DetectorElement> hcal ( detsvc(), "Structure/LHCB/Hcal" );
if ( !hcal ) return StatusCode::FAILURE;

// Retrieve the slow control information for the Hcal detector
IAlignment* hcalSCInfo = hcal->slowControl();
// Presently, a test version exists only in package Det/DetCondExample
if ( !hcalSCInfo ) {

// Retrieve slow control Condition for the Hcal detector
Condition* scHcal = hcalSCInfo->condition();
if ( !scHcal ) {

// Retrieve the names of its parameters and parameter vectors
std::vector<std::string> paramNames;
std::vector<std::string> paramVectorNames;
paramNames = scHcal->params();
paramVectorNames = scHcal->paramVectors();

// Print all user parameters
std::vector<std::string>::iterator it;
for ( it = paramNames.begin(); paramNames.end() != it; ++it ) {
log << MSG::INFO << "\n" << *it << "\n" << *it << "\n" << endlreg;
}
}
}
```

8.3.4 Extending the DetectorElement class

As you can see in Listing 8.2, the transient store allows you to retrieve either DetectorElement objects or your own object, extending DetectorElement. The definition of your objects is fully free as you can see in Listing 8.5.

There are a few important things to know here :

- you have to define a new classID for the new object. This one must be unique and will be used to distinguish this object from the default DetectorElement. ClassIDs can be requested using the same procedure as for event data, as described at <http://cern.ch/lhcb-comp/Frameworks/EventModel/CLID.htm>.



Listing 8.5 Extending DetectorElement class

```

class DeMuonStation: public DetectorElement {
public:
    // constructors and destructors
    DeMuonStation();
    DeMuonStation( double t, double x, double y );
    ~DeMuonStation();

    // new classID
    const CLID& clID();

    // Get/Set thickness
    double thickness();
    void setThickness( double t );

    // Get/Set pad dimensions
    double padx();
    double pady();
    void setPadx( double x );
    void setPady( double y );

private:
    double m_thickness; // Aluminium plate thickness
    double m_padx; // Pad X dimension
    double m_pady; // Pad Y dimension
};

```

- the usage of a new object instead of the default `DetectorElement` supposes that the user writes a new converter for this new object. This is described in a dedicated document: "Extending detector elements and implications", by Sebastien Ponce, available at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/detElemExtension.pdf>.
- the new object, as extension of `DetectorElement`, has an `initialize` method. This method is called by the conversion service at the end of the object creation, when every contained object has been computed. This is the place where some computation could be done. Once more, see the dedicated document for further details and examples.

8.3.5 Configuring the Gaudi framework to access the detector description

Applications accessing the detector description have to provide a number of job options to the framework. For convenience, these have been collected together in the file `DetDesc.opts` of the `Gaudi:Conf` package. The user simply needs to add the line below to his job options file:

```
#include "$STDOPIS/DetDesc.opts"
```



In addition, the application must use a number of CMT packages, as shown below.

```
use DetDesc v7 Det
use XmlDDDB v6x1 Det
```

Package `Det/XmlDDDB` contains the XML files for the production version of the persistent DDDB, including all detector elements with their geometry and materials, while package `Det/DetDesc` contains the `DetectorElement` class and all generic XML services and converters. Of course, any sub-detector specific packages containing the sub-detector specific converters should also be used, with the corresponding job options for loading the sub-detector specific DLLs.

8.4 General features of the detector description

As explained in Section 8.3.1, the detector description is divided into three main parts, Structure, Geometry and Materials. We describe here the features and meaning of the elements present in each of these parts. Five other catalogs are also foreseen for Alignment, Calibration, ReadOut, SlowControl and FastControl, although these are not present in the current production version of the DDDB. A brief description of these catalogs is also given.

8.4.1 Structure

The structure of the detector is described as a tree of detector elements. These could also be user defined types that inherit from the class `DetectorElement`.

Each node of the tree, being a detector element, has the following characteristics:

- a class identifier: this tells the framework which converter should be used for this element.
- a geometry info object where all its geometry is described
- five additional condition info objects allowing to access and manipulate its alignment, calibration, readout, slow control and fast control data
- a set of user parameters, each one defined by a name, a comment, a type and a value

8.4.2 Geometry

The geometry in the Gaudi detector description is a bit more than just solids and surfaces. It has to include a way of structuring the geometry in order to avoid repetitions and to facilitate the placement of subparts into their parents.

After giving some general considerations concerning the description of the geometry, we describe here the two classes `LVolume` and `EVVolume` that build the structure of the geometry. Finally, solids and surfaces are quickly described.



You can find more information in the Gaudi web pages, especially at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/default.htm>, in the "Detector geometry model" section.

8.4.2.1 General considerations

The Gaudi geometry description was based on some postulates:

- The geometry is a tree of volumes, where each child describes a subpart of its parent volume.
- There are no "up-links" in the geometry tree. This means that each volume has no information about its parent's volume. Of course, each volume has information about its children.
- No volume has any information about its absolute position. The only spatial information in the whole tree is local: each volume is placed into its parent.
- The geometry (description of shapes and surfaces) and material information is kept by each volume via a `Solid`, a `Surfaces` and a `Material` object. Thus the information can be retrieved at several levels: the top-most volume may contain a global shape and an average material and the leaves of the tree contain the most precise material and geometry.
- All boolean operations on volumes are strictly forbidden¹. Boolean operations should be performed at the level of Solids. This is one of the most essential postulates of the Gaudi geometry structure.

The geometry tree which fulfils all these postulates represents a very effective, simple and convenient tool for description of the geometry. Such a tree is easily formalized. It has many features which are similar to the features of the geometry tree used within the Geant4 toolkit and could easily be transformed to the Geant4 geometry description.

Some consequences of these postulates are:

- The top-level volume (presumably the experimental hall, or cave, or the whole LHCB detector) defines the absolute coordinate reference system. In other words, the null-point (0,0,0) in the so called Global Reference System is just the center of the top volume.
- All geometry calculations, computations, inputs and outputs, performed with the usage of a volume are in the local reference system of this volume.

Sometimes one needs a more efficient way of extracting information from the geometry tree or to compute the unique location of a point in the geometry tree. For these purposes, a simplified detector description tree is introduced into the Gaudi framework².

1. This is equivalent to the absence of the "MANY" flag in the GEANT3 toolkit.
2. Within the Geant4 toolkit there exist two approaches for solving the same problem: Read-Out-Geometry Tree and Navigator. Our approach is quite close to the combined usage of both.



8.4.2.2 LVolume and PVolume

Up to now, we only talked about volumes. This term is mapped to a class called `LVolume` which stands for Logical Volume. In this class, you can find a `Solid` describing the shape of the volume, a list of `Surfaces` associated to it and the `Material` of the volume.

On top of that, `LVolume` is an identifiable object and therefore inherits from `ClassDataObject` and can be identified in the transient data store by a unique name (its "path" in the geometry tree). It implements the `ILVolume` and `IValidity` interfaces.

As we said before, this `LVolume` has a set of children. These are actually `PVolumes`, which stands for Physical Volume but should stand for Placed Volume. The notion of Physical Volume in the Gaudi geometry is extremely primitive, it is just a Logical volume which is positioned inside its mother Logical Volume. It consists of the name of the Logical Volume to be positioned, together with the transformation matrix from the local reference system of the mother Logical Volume to the local reference system of the daughter Logical Volume. `PVolume` is not identifiable and implements the `IPVolume` interface.

The important point here is that a given logical volume can have several times the same logical volume as a child, using several different physical volumes. This is the way volumes can be replicated without replicating the objects in memory.

More detail, especially concerning the implementation can be found in "Volumes", by I. Belyaev at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/volumes.pdf>

8.4.2.3 Solids and surfaces

All solids implement the `ISolid` interface. Currently, seven types of "primitive" solid are implemented: Boxes, Simple Trapezoids, Tube segments, Conical tube segments, Polyconical tube segments, Sphere segments and General trapezoids. Here is a very quick overview of each of them:

- box is obvious.
- `trd` is a simple trapezoid defined by two rectangles, orthogonal to the Z axis and centered on it, and a Z size.
- `trap` is a more generic trapezoid defined by a direction (`theta`, `phi`), that replaces the previous Z axis, a length along this axis and two faces orthogonal to the axis, which are kinds of trapezes. Actually, the faces are defined by two segments parallel to what replaces the X axis, the distance between them and the angle between the orthogonal line to these segments and the line joining the middle of both segments.
- `tub` is a part of a section of tube. The tube is centered on the Z axis and defined by its inner and outer radius. It is then cut to keep a section of a given length centered at the origin. Last, only the part between `phi` and `phi + deltaPhi` is kept.
- `cons` is a part of a section of cone. The cone is defined along the Z axis and its section is defined by the inner and outer radius at each extremity plus the height of the section. Last, only the part between `phi` and `phi + deltaPhi` is kept.
- `polycone` is a part of several cone sections. The cone is defined along Z and the sections are defined by a number of triplets giving for several location in z the inner and outer radiuses of the polycone in these places. At last, only the part between `phi` and `phi + deltaPhi` is kept.



- sphere is a kind of sphere part but with a thickness. The thick sphere is defined by an inner and an outer radius. Then only a portion of the whole sphere is kept, defined by its theta and phi angles.

These solids were chosen from the most frequently used shapes in the GEANT3 and Geant4 toolkits - more shapes can be implemented if necessary. In addition, Boolean Solids have been defined, which allow Subtraction, Union and Intersection operations on solids, to build complex shapes. Details concerning the implementation are available in "Solids" by I. Belayaev at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/Solids.pdf>

The class `Surface` handles surfaces. Details are available in "*Optical properties & Surfaces*" available at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/Optical.pdf>

8.4.3 Materials

The description of materials is actually the description of their chemical composition plus some parameters.

The parameters are the same for all kinds of materials, they are radiation length, density, lambda, temperature, pressure and state. On top of that, tabulated properties can be added.

Concerning the chemical composition, three kinds of materials are distinguished, each of them corresponding to a given class. Here is a short description of each :

- Isotope : this is the definition of a given isotope of a given atom. It has an atomic mass, an atomic number and a number of nucleons.
- Element : this is a real life element, that is in general a mixture of several isotopes with given proportions. It has an effective atomic mass, an effective atomic number and an effective number of nucleons.
- Mixtures : this defines a mixture of several elements or even several other mixtures with given proportions.

8.4.4 Alignment, Calibration, Readout, SlowControl, FastControl

The description of these time varying detector "conditions", as previously stated, has so far only been implemented in a very simple way. For each of these five categories, the corresponding data for a detector element is foreseen to be encapsulated in one Condition object of the given type, containing only the following:

- a set of user defined parameters and parameter vectors, each one defined by a name, a comment, a type and a value (just like those for a detector element).

For all five categories of conditions, it is important to distinguish between the actual measured data and all other quantities which can be derived as a function of it: while the first kind of data must be stored in the persistent DDDb and retrieved by the framework from its persistent storage every time it becomes invalid, the second type of data should not, in general, be stored in the persistent DDDb. For instance, if the gain of a detector was a function of its temperature, only the coefficients of that functional dependence and the temperature of



the detector should be read from the persistent storage, while its actual gain should be computed within the framework at processing time.

Similarly to class `DetectorElement`, it is foreseen that users may define their own types of Condition objects. This is however not recommended unless strictly necessary as, in general, an arbitrary list of parameters and parameter vectors should be sufficient for most uses. Specific questions concerning condition data should be answered by the interfaces like `IAIAlignment` and not by the Conditions themselves; it is then better to identify specific methods to extend these interfaces, rather than extend the Condition class.

8.5 Persistent representation based on XML files

The Gaudi detector description is based on text files whose structure is described in XML (eXtensible Markup Language). XML files are understandable by humans as well as computers. Data in XML are self-descriptive so that by looking at the XML data one can easily guess what the data mean. An advantage of XML is that there exists plenty of software which can be used for processing and manipulating data, as it is an industry standard.

In the future we expect to replace the XML files by an object persistency service, based on a database. However, the data may continue to be described in XML from the user point of view. These issues are described in more detail in Section 8.6.

This part first gives a quick introduction to the XML language and to the document type definitions (DTDs). The DTD used in the particular case of the LHCB detector description is described. Then, the conversion from the persistent to the transient representation is presented as well as its specialization to address particular user cases. Finally, we give some ideas and references concerning the numerical expression parser used during the conversion process and the XMLEditor tool that could be used to edit the persistent representation.

8.5.1 Brief introduction to XML

8.5.1.1 XML Basics

XML is one example of what is called **markup** languages. Another well-know example of it is HTML. The specificity of this class of languages is the usage of "**tags**" to structure the code. A tag is actually a marker (or a pair of them, one for opening and one for closing) that delimits a given part of the code. In the case of both XML and HTML, they look like `<...>` for opening tags or `</...>` for closing tags.

The very basic and mandatory rule of markup languages is that all **tags must nest properly**. Proper nesting means that each opening tag has its corresponding closing tag and that this one must appear before the parent's tag closing tag, as shown in Listing 8.6.

Listing 8.6 Properly nesting in markup languages

```
<ProperNesting> <Something>In</Something> </ProperNesting>  
<WRONGNESTING> <BADTHING>huhu </WRONGNESTING> </BADTHING>
```



The specificity of XML among other markup languages is that it is extensible. This means that a user can define his own tags. For example, Listing 8.7 shows the XML file which describes an e-mail (forget the first two lines and the timeStamp item, we will describe them later).

Listing 8.7 Simple XML file describing an e-mail

```
<?xml version='1.0' encoding='UTF-8'?>
<!-- This is an example of XML -->
<Email>
  <TimeStamp time='11:38:43' date='22/11/1999' />
  <Sender>sender@cern.ch</Sender>
  <Recipient>recipient@cern.ch</Recipient>
  <Subject>Lunch...</Subject>
  <Body> Could we meet at 14:00?
  <Signature>Sender's signature</Signature>
</Body>
</Email>
```

At first sight this markup language looks like screwed-up HTML. This is because HTML is only a subset of XML (with a fixed set of tags). Unlike HTML, you cannot guess from the example above how to present the data described there nor how to visualize them. What is clear however is the meaning of the data items encoded in XML. Thus one can easily recognize the data items and guess what they mean. On the other hand it is relatively easy to instruct a computer program what to do with the given data item according to the XML markup elements it is encapsulated in. Let us analyse the example shown in Listing 8.7.

8.5.1.2 XML components

XML declaration must be at the beginning of each XML document. It is the first line in the example. It says that this file is an XML file conforming to the XML standard version 1.0 and is encoded in UTF-8 encoding. The encoding is very important because XML has been designed to describe data using the Unicode standard for text encoding. This means that all XML documents are treated as 16-bit Unicode characters instead of usual ASCII. So, even if you write your XML files using 7- or 8-bit ASCII, all the XML applications will work with it as with 16-bit Unicode XML data. The encoding information is important, for example when an XML document is transferred over the Internet to some other country where a different encoding is used. If the receiving application can guess the XML encoding from the received file, it can apply transcoding methods to convert the file into proper local encoding, thus preserving readability of the data.

XML comments look like comments in SGML or HTML. They start with `<!--` and end with `-->`. Comments in XML cannot be nested. This means you cannot open comments inside comments, as you can open a tag inside another tag.

XML elements are the tags. In the example we had the following XML elements: Email, TimeStamp, Sender, Recipient, Subject, Body, Signature. The very basic and mandatory rule of XML is that all XML element tags must nest properly (this is a markup language) and that there must be only one root XML element at the top level of each XML document, which contains all the others. An XML document that follows these rules is called **well-formed**. An example is the XML document in Listing 8.7.

XML elements are actually a bit more complex than regular tags. They can contain what is called **attributes**. These are pairs of the form `name=value` that are embedded inside the opening part of the tag. This is the case in the `TimeStamp` tag in Listing 8.7. Two attributes are defined there: `date` and `time`. Attributes usually describe the properties of the given element. The value of an attribute is a string and it has to be enclosed inside quotes (single or double, it doesn't matter). In the content of a given tag, raw text can appear among nested XML elements. This is called **text data**.

8.5.1.3 Document type definitions (DTDs)

The XML documents described up to now are well formed since they follow the syntax rules of XML. This is a first step but one may want to add some grammar inside an XML document so that you are not allowed to use any XML element you could imagine in any place.

This is possible by writing a **document type definition (DTD)**. It can be integrated inside the XML file itself (internal DTD) or can be given in a separated file (external DTD). The external case is the most used since the external file can be reused by many XML files and thus ensure consistency between them.

The DTDs are actually written using a markup language too. The main features of a DTD are:

- the definition of an XML element (ELEMENT tag of the DTD): this specifies the name of the XML element and its possible children
- the definition of the attributes of an XML element (ATTLIST tag of the DTD): this specifies the attributes of a given XML element, as well as their usage (required, default value, ...)

We will not describe here further the syntax of the DTD file since it would be too long but it is pretty obvious to understand.

Listing 8.8 gives an example of the definition and usage of an external DTD for an XML file. Another example is the HTML language. It can be seen (with small exceptions) as the definition of a given DTD, that is used to describe a hypertext document and the way it should be displayed.

Listing 8.8 Document type Definition

```
<?xml version='1.0' encoding='UTF-8'?>
<ELEMENT Email (TimeStamp, Sender, Recipient, Subject, Body)>
<ELEMENT TimeStamp EMPTY>
<!ATTLIST TimeStamp time CDATA #IMPLIED date CDATA #REQUIRED*>
<ELEMENT SENDER #PCDATA>
<ELEMENT RECIPIENT #PCDATA>
<ELEMENT SUBJECT #PCDATA>
<ELEMENT BODY (#PCDATA|Signature)*>
<ELEMENT Signature #PCDATA>
```

Now that the XML file has a grammar, it is called **valid** if it fulfils the grammar requirements. Otherwise, it is not valid, even if well-formed. The process by which the grammar of a given XML file is checked is called "validation process".



8.5.2 DTD of the LHCb detector description

In the case of the LHCb detector description, several DTDs have been defined in order to describe the structure, geometry and materials of the detector. These are detailed in a dedicated document: "The LHCb Detector Description DTD", by Sebastien Ponce, available at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/lhcbtdtd.pdf>. Thus, we will only give here a very short presentation of it.

You will find in the next few sections UML schemas of the different DTDs. For those who don't know this notation, here are the basics:

- every entity here is an XML element, with all the possible attributes enumerated
- following the arrows getting out of an element, you can know which elements could be children of it.

8.5.2.1 Some prerequisites

Some of the elements of the LHCb DTDs are defined in each of the three DTDs. Thus, we describe them here.

Expressions Every numerical value required by an attribute or an element in any of the DTDs of LHCb is an expression. This means that the value will be evaluated by the numerical expressions parser (see Section 8.5.5)

Thus, most of the current units and constants are already known. You can safely use degree, rad or pi for angles for instance. On top of that, many mathematical functions are also known, such as sin, or exp but also arctan and many others.

Parameters A special tag called **parameter** is defined in all DTDs of LHCb. This element allows the user to define his own parameters that can be then reused in any expression or value in the rest of the XML code. It has a name, a type, a value and a comment.

The parameter element is actually a kind of macro since at parsing time it will be replaced by its value wherever it appears.

References In the description of the LHCb DTD, you will see a lot of nodes with names finishing by "ref". Each time there is a corresponding node without the "ref". The "ref" nodes are actually references on the "without ref" ones. All of them have almost the same usage and syntax. The common things they have are attributes for class ID and hyperlink reference. The hyperlink is in general specified using the format

```
protocol://hostname/path/to/the/file.xml#objectID or #objectID
```

The `protocol` and `hostname` parts can be omitted if the file resides on the local host. It is possible to write a hyperlink without the full path name in case one needs to refer to an XML object residing inside another file. In this case the relative path will be appended to the location of the currently parsed XML file.

For example having the current file location `/full/path/to/current.xml` and inside this file a hyperlink as `href="next/file.xml#nextOID"` the hyperlink will be resolved as `/full/path/to/next/file.xml#nextOID`.



If the hyperlink has the form `#ObjectID` this means that the referred object is located in the current file.

Note that relative paths are strongly encouraged for every file except the top-most one, since the whole set of files may be copied to several different locations one day.

8.5.2.2 The structure DTD

The schema of the DTD that should be used for every file describing the structure of the LHCb detector is shown in Figure 8.4. Presently, this DTD is also used for the files containing the description of all kinds of generic condition data such as alignment, although a simpler DTD would be enough.

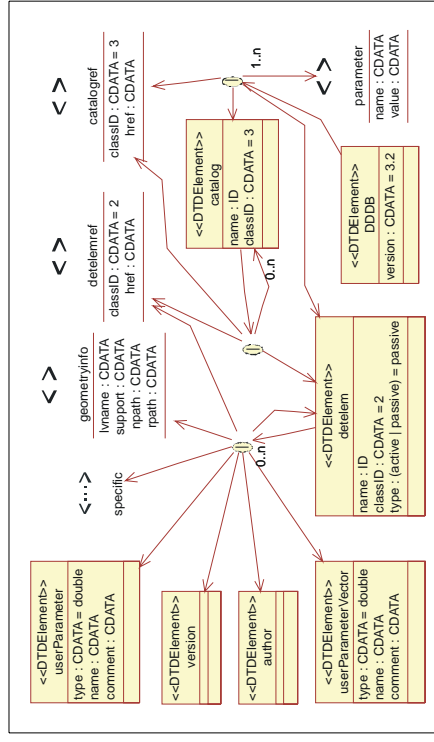


Figure 8.4 The Structure DTD

Let's describe quickly some of the elements (the others are trivial) :

- **DDDBB** : this is to fulfill the XML basic rule that each XML document must have only one root XML element This is the root element.
- **catalog** : this is simply a list of elements, with a given name. This is a way to classify detector elements.
- **detelem** : detector elements are the essential part of the structure of the detector description. They fully describe a given part of the detector by holding data on the geometry of this part as well as on the subparts constituting it.
- **userParameter** and **userParameterVector**: this allows the user to add a parameter or a vector of parameters for a given detector element. This is intended to be used for specific parameters appearing in the subdetector descriptions. Their usage is described deeper in another document: "Extending detector Elements and implications", by Sebastien Ponce, available at



<http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/detElemExtension.pdf>

- **condition**: a condition is simply a container of parameters and parameter vectors.
- **geometry info**: this tag describes the geometry of a given detector element.
- **alignment info, calibration info, readout info, slow control info, fast control info**: these tags, when present, contain the name of the corresponding condition for a given detector element; if absent, no condition info of that type will be available.
- **specific**: this is the place where a user can extend the default detector description language and introduce new tags for his own needs. It is foreseen that the new XML elements be defined in a local DTD section of the XML data file or in a specific DTD file. Its usage is described deeper in "Extending detector Elements and implications", by Sebastien Ponce (See location above).

8.5.2.3 The geometry DTD

We schema of the DTD that should be used for every file describing the geometry of the LHCb detector is shown in Figure 8.5 and Figure 8.6.

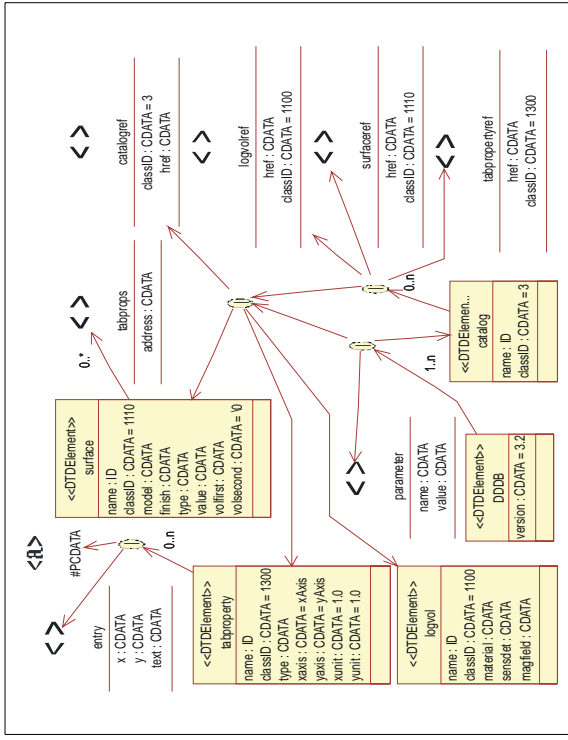


Figure 8.5 Geometry DTD (part 1)

Let's describe quickly some of the elements (the others are trivial or explained in the specialized documentation, see Section 8.5.2) :

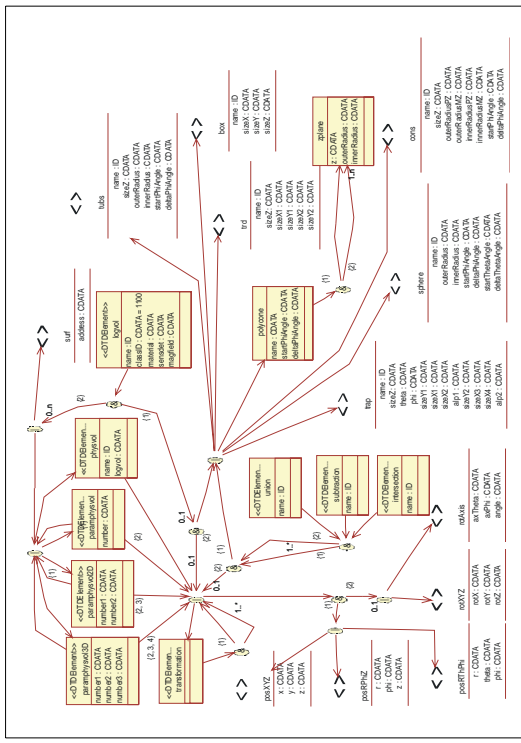


Figure 8.6 Geometry DTD (part 2)

- **DDDB**: this is to fulfill the XML basic rule that each XML document must have only one root XML element This is the root element.
- **catalog**: this is simply a list of elements, with a given name. This is a way to classify the logical volumes.
- **tabproperty**: this defines a tabulated property. This is used to describe optical properties of materials and surfaces. See "Optical properties & Surfaces" available at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/Optical.pdf>.
- **surface**: this defines a surface, see Section 8.4.2.3.
- **transformation**: this defines a new transformation, by composition of several others
- **logvol, physvol**: these are the mapping of the C++ classes described in Section 8.4.2.2
- **paramphysvol2D, paramphysvol3D, paramphysvolBD**: these are ways to define many physical volumes in one shot, by replicating a given volume and applying a given transformation between each replica.
- **posXYZ, posRPhiZ, posRThPhi**: these are 3 ways of defining a translation: cartesian, cylindrical and spherical coordinate systems.
- **rotXYZ, rotAXis**: these are two ways of defining a rotation. Either along X, Y or Z axis, or along a user-defined axis.



- **box**, **trd**, **trap**, **cons**, **polycone**, **tub**, **sphere** : these are all kinds of solids. They are described in Section 8.4.2.3 on page 71. Note that polycone uses a subelement called `zplane` to define the different sections of the polycone.
- **union**, **subtraction**, **intersection** : these are boolean operations on solids.

8.5.2.4 The material DTD

We the schema of the DTD that should be used for every file describing materials of the LHCb detector is shown in Figure 8.7.

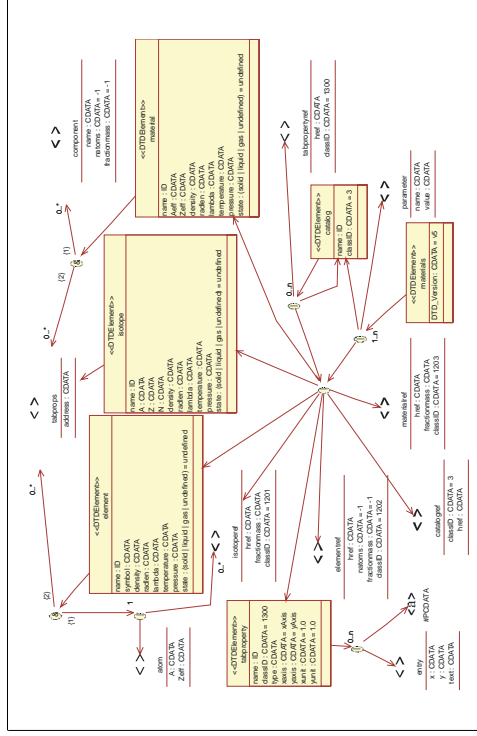


Figure 8.7 Material DTD

Let's describe quickly some of the elements (the others are trivial) :

- **materials** : this is to fulfill the XML basic rule that each XML document must have only the one root XML element This is the root element.
- **catalog** : this is simply a list of elements, with a given name. This is a way to classify materials.
- **tabproperty** : this defines a tabulated property. This is used to describe optical properties of materials and surfaces. See "Optical properties & Surfaces" available at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/Optical.pdf>.
- **isotope**, **element**, **material** : this maps to the three classes `Isotope`, `Element` and `Mixture` described in Section 8.4.3.
- **atom** : this is used to define an element when it is not a mixture of isotopes.



- **component** : this is used to define mixtures. It associates a material and a proportion.

8.5.3 Conversion to transient representation

We describe here the conversion of data from the persistent to the transient world. This matter is actually the main topic of Chapter 13. Please refer to it for an explanation of how the conversion process works and is implemented into Gaudi. Here we give additional information on the special case of detector description data.

8.5.3.1 Overview

The main aim of the conversion process is to populate the transient store with C++ objects, using XML files. If you look at the description of the C++ objects in Section 8.3 and at the description of the XML files in Section 8.5.2, you will notice that there is almost a one to one mapping between the two, where XML elements become objects and XML attributes become members of these objects.

Thus, the conversion service (called `XmlCnvSvc` in our case) uses more or less one converter per XML element, or per kind of C++ object stored in the transient store. Here is the list of existing converters and the list of objects they are creating :

- `XmlTabulatedPropertyCnv` : deals with tabulated properties and creates `TabulatedProperty` objects.
- `XmlSurfaceCnv` : deals with surfaces and creates `Surface` objects.
- `XmlMixtureCnv` : deals with mixtures and creates `Mixture` objects.
- `XmlVolumeCnv` : deals with logical volumes and creates `Volume` objects. This is also the place where all the geometry is parsed and stored, including physical volumes, solids (simple and booleans), transformations and parametrized physical volumes. These last ones do not exist in the C++ world since they are expanded in a set of regular physical volumes.
- `XmlIsotopeCnv` : deals with isotopes and creates `Isotope` objects.
- `XmlElementCnv` : deals with elements and creates `Element` objects.
- `XmlBaseElementCnv` and `XmlUserDetElementCnv` : deal with detector elements and create `DetectorElement` objects.
- `XmlBaseConditionCnv` and `XmlUserConditionCnv` : deal with condition data and create `Condition` objects.
- `XmlCatalogCnv` : deals with catalogs and creates `DataObject` objects containing a list of references on other objects in the store.

8.5.3.2 XmlCnvSvc

The `XmlCnvSvc` is the service responsible for converting XML code into C++ objects. It implements the `IXmlSvc` interface and essentially provides two sets of methods:



- parser related methods : `parse`, `clearCache` and `allowGenericCnv` which respectively parses an XML file, clears the cache of parsed files and says whether user defined detector elements are converted using dedicated converters or the generic converter. This last method is related to the job option property `AllowGenericConversion`.
- expression evaluator related methods : `eval`, `addParameter`, `removeParameter` which respectively computes the value of a given expression, adds a new parameter and remove a given parameter. Parameters can be used in any expression, as soon as they are defined. For more details, the expression evaluator is presented in Section 8.5.5.

8.5.3.3 XmlParserSvc

The `XmlParserSvc` is actually used by the `XmlCnvSvc` for parsing XML files. We will not describe it in detail here but just give some usefull hints.

This service is not only responsible for parsing XML files but also for implementing a cache on parsed files. This cache is very useful due to the behaviour of the transient store. As a matter of fact, the transient store calls the `XmlCnvSvc` and thus, indirectly, the `XmlParserSvc`, every time an object is missing. When this occurs, an XML file is parsed and the object is created and put into the store. But all the objects contained in the XML file are not created and put into the store, only the ones needed. Thus, if the next missing object is defined in the same XML file, this one will be read again and parsed again.

That's where the cache is very useful: it avoids opening and parsing the same file again and again. A recent example showed that the cache could speed up the parsing with a factor 180 under certain conditions (1300 objects in the same file, read one by one).

The important point is that this cache can be configured in order to be efficient for every special case. The two configuration parameters are :

- the maximum number of cached files. This is given by the `MaxDocNbInCache` property in the job options file. The default value is 10.
- the behavior of the cache. This is given by the `CacheBehavior` property in the job options file. The default is 2. The meaning of this parameter is a bit complicated. It is explained below.

8.5.3.3.1 CacheBehavior usage

The quick explanation is :

- set `CacheBehavior` to 0 to have a FIFO (First in, first out) behavior
- put it to more to keep files that are reused. In this second case, the value of the parameter gives you the frequency of usage a given file should have to stay in the cache. Eg. for the default value, a file reused once every two file usages will stay forever in the cache.

The complete explanation is the following :



- for each item in the cache, its `birthdate` and its `utility` are stored. The `birthdate` is the age of the cache at the time the item was added to it and the age grows by 1 each time a file is requested, be it in the cache or not. The `utility` is the number of times the cached item is accessed, once it is in the cache.
- when the cache is full, the following value is computed for every item: `birthdate + CacheBehavior * utility`. The item with the smallest score is removed from the cache.

So, if `CacheBehavior = 0`, the item with the smallest `birthdate` is removed, which means the oldest. So it is a FIFO. If not, the thing is a bit more tricky but imagine that you request the same file every `CacheBehavior` requested files. The `utility` of this file will be `(cacheAge - birthdate) / CacheBehavior` and thus, its score will be `cacheAge`, which is the best possible score for a new arriving file. So this file will be kept. Files that are less often used will have a score further and further from the one of the last arrived file and will eventually be deleted. Files more often accessed have better scores and will be kept.

8.5.3.4 Converters implementation overview

We give here a quick overview of the converters implementation to give an idea of what to do if you want to extend them. This process of extension will be anyway detailed for the case of detector elements in Section 8.5.4.

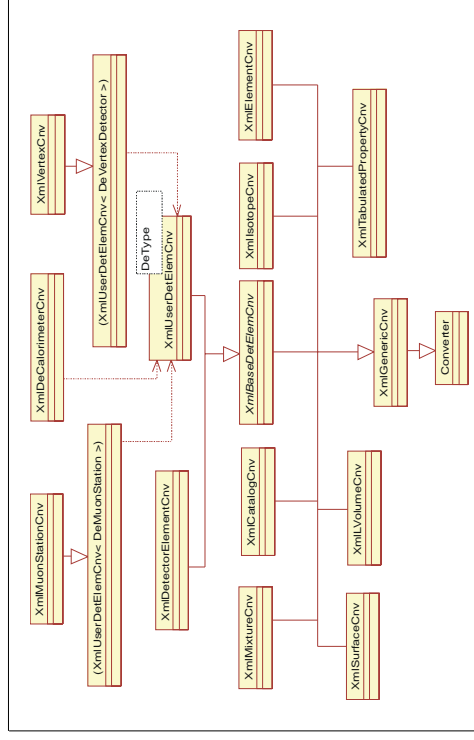


Figure 8.8 C++ class hierarchy for converters

Figure 8.8 gives an overview of the tree of classes defining the whole set of converters. As explained in Chapter 13, every converter inherits from the class `Converter` and implements the methods `initialize`, `finalize`, `createObj`, `updateObj`, `createRep` and



`updateRep`. In the case of the XML converters, the method `createObj` (the only currently implemented one among the `create*` and `update*` methods) is pretty much always the same: retrieve the right XML file, parse it, check the version, find the element to retrieve, create the corresponding C++ object, fill it with its children if needed, make some computation on it and store it.

The only three points where things are converter specific are thus the creation of the object, its filling and the computation made on it. The rest of the code is common and it was put into `XmlGenericCnv`. Only three virtual methods have to be overridden to implement a new converter from the generic one: `i_createObj`, `i_fillObj` and `i_processObj`.

In the case of detector elements, it is a bit more tricky since we want users to be able to extend easily the default converter to add their own stuff. This is done by putting all the default code inside `XmlBasedElementCnv` and defining a new method, called `i_fillSpecifObj`, that has to be implemented by the user in case he adds some stuff. The default converter `XmlDetectorElementCnv` only gives an empty implementation for this method.

See 8.5.4 for more details.

8.5.4 Customizing a detector element

This topic is developed in a dedicated documentation: "Extending detector elements and implications", by Sebastien Ponce, available at <http://cern.ch/lhcb-comp/Frameworks/DetDesc/Documents/defElemExtension.pdf>. Thus, we will only give here a quick overview of the possibilities without any explanation on how to practically do the job.

There are three ways of customizing a detector element, from the simplest but less flexible one to the most complicated, but most flexible one :

- the first one only allows the user to add parameters (or vectors of parameters) to the default `DetectorElement` object in C++.
- the second one allows the user to define a new C++ object, inheriting from the default `DetectorElement`. He can thus do some computation inside the object, using the parameters defined in the first case.
- the last method allows the user to redefine the DTD used in the XML files and to manage structured data instead of just parameters.

8.5.4.1 UserParameters

This is the simplest way to add specific information to a detector element without having to write too much code. It is based on the `UserParameter` and `UserParameterVector` tags defined in the structure DTD of the LHCB detector description files.

The `UserParameter` and `UserParameterVector` tags are children of the `deletelem` tag. They can appear anywhere in the definition of a detector element. Listing 8.9 shows some examples of the usage of user parameters:



Listing 8.9 Example of usage of user parameters

```
<deletelem name="Ecal">
  . . . . .
  <userParameter name="CodingBit" type="int" > 6 </userParameter>
  <userParameterVector name="aVector" type="int">
    10 20 30
    40 50 60
  </userParameterVector>
</deletelem>
```

Once user parameters are defined in XML, they are converted by the regular converter for detector elements and are then reachable in the C++ code using the following methods of the class `DetectorElement`: `UserParameters`, `UserParameterVectors`, `UserParameter`, `UserParameterAsString`, `UserParameterAsInt`, `UserParameterAsDouble`, `UserParameterType`, `UserParameterComment`, `UserParameterVector`, `UserParameterAsString`, `UserParameterAsDouble`, `UserParameterVectorType` and `UserParameterVectorComment`.

The method name should be self-descriptive. For more details, go to the dedicated documentation.

8.5.4.2 Customizing the detector element

The preceding section showed us how to define user parameters inside a regular detector element. We describe here how it is possible to define new C++ classes, inheriting from the original `DetectorElement` class. This has for consequence that the generic converter provided with Gaudi can no longer do the job, since it is not aware of the existence of this new class. Thus, we have to create a new converter, by customizing a bit the default one.

This is done by using the templated class `XmlUserDetectorElementCnv<DeType>` which allows the users to define his own converter for his own type of detector element. It actually inherits the behavior of the default converter and avoids the user to rewrite existing code. The only difference here is that it creates an object of type `DeType` instead of a regular `DetectorElement`.

The parameter class `DeType` has only few constraints: it must inherit from `DetectorElement`, it must have a default constructor with no parameters and it must have a new classID, which has to be unique in whole LHCB software. Apart from that, you have full freedom to use you own class and to add whatever methods or members you may want to the default `DetectorElement`.

In addition, one could override the methods of `DetectorElement`. This is not recommended at all except for the method `initialize()`. This one is a kind of hook for the user. It is called by the converter just after the creation of the object and before its first use. It allows the user to initialize some members using the values parsed from the XML, for example user parameters.



8.5.4.3 Extending the DTD

This last way of extending a detector element is the most flexible one. It allows to define new XML children tags for the `detelem` tag. This is actually the goal of the `specific` tag. It allows the user to extend the default DTD of the XML file where his detector element resides and to use the new tags to add structured data. These data will then be used by a user defined converter to build a detector element.

The extension of the DTD and the use of the `specific` tag are not explained here since they are straightforward for someone who knows XML. More details are available in the dedicated documentation.

Concerning the C++ part, the situation is pretty much the same as before: a new converter must be defined. This looks like the previous one except that it has one more method implemented (actually overridden from `XmlUserDetElemChv`):

```
StatusCode i_fillSpecificObj( DOM_Element childElement,
                             DeType* dataObj );
```

This method uses the DOM interface, which is an interface to XML parsers. If you don't know it and thus don't know what a `DOM_Element` is, please refer to the dedicated documentation.

The principle of the `i_fillSpecificObj` method is pretty simple: it is called by the XML parser service everytime a tag is found directly inside a `specific` tag. The parameters are then the `DOM_Element` corresponding to the tag found and the C++ detector element object being built when this occurred.

Thus, for the XML code in Listing 8.10 the method will be called twice, once for `Al_plate_thickness` and once for `pad_dimensions`. It will not be called for `padX` or `padY`.

Listing 8.10 Extending the DTD

```
<detelem classID="9990" name="MStation01">
<author> Sebastien Ponce </author>
<geometryInfo ... />
<specific>
<Al_plate_thickness value="1.1111*mm"/>
<pad_dimensions>
<padX length="2*mm"/>
<padY length="4*mm"/>
</pad_dimensions>
</specific>
</detelem>
```

This method is thus the place where the XML code can be browsed and data extracted to be put inside the detector element object.



8.5.5 Numerical expressions parser

The framework provides a simple evaluator of numerical expressions based on the CLHEP expression evaluator. It is available for Gaudi framework converters as well as for the user defined converters. The only difference is that in user defined converters the parser is not instantiated explicitly by the user but is accessed through the `IXmlSvc` interface instead.

The numerical expressions recognized by the parser can be composed of integers and floating point numbers assuming one of the formats:

```
100 100. .05 0.1 1.34-e12 -23
```

Supported operations are: `+` `-` `*` `/` `unary +|-` `exponent ^`

Parenthesized expressions: `1.4 * (23.4-e12 / 1.8)`

Operator precedence is: `() unary +|- ^ * / +|-`

In addition, the parser understands CLHEP units. The result is always evaluated to double value. The check for the presence of CLHEP units inside expressions is enabled by default. To suppress this behaviour the call of the `eval()` method must look like:

```
xmlSvc->eval( "2*(34.5 + 1.23-e4) ", false );
```

For further information about this expression parser, please go to the CLHEP documentation at <http://wwwinfo.cern.ch/asd/lhc++/clhep/manual/UserGuide/index.html>

8.5.6 XML Editor

The XML editor is a tool provided to edit the XML files of the detector description database without having to learn the XML syntax. It is provided as a separate Gaudi package and was written in Java. Documentation is available at <http://cern.ch/lhc-comp/Frameworks/DetDesc>.

8.6 Persistent storage in a Conditions Database

While the current implementation of the Gaudi detector description, based on text files in the XML language, is perfectly adequate to describe detector data for which only one or at most a few versions exist (as in the case of the detector logical structure, geometry or materials), a file-based persistent storage is not adequate to describe detector data which is rapidly changing in time, such as slow control measurements of temperature and pressure. It is therefore foreseen that this data will be stored in a database, the so called Conditions Database (or Conddb). Once this storage system is fully functional, it could also be used to store the description of the structure and geometry of the detector, keeping track of its possible time variation due to the installation or removal of subdetector elements between different years.



8.6.1 Generalities on the Conditions Database

The main user requirements and abstract interfaces for the CondDB have been described in the document "Conditions Database for LHCB: Interface Specification Proposal", by Pere Mato, available at <http://cern.ch/lhcb-comp/Frameworks/DetCond/condataspec.pdf>; only a quick overview of the system will be given here. In summary, each block of condition data stored in the CondDB, independently from its data content, has the following three characteristics, as shown in Figure 8.9:

- the name of the data type it refers to (such as "/SlowControl/Hcal");
- the time range during which it is valid, i.e. an interval of the form [since, till];
- a version number (such as the version of the alignment algorithm used).




Figure 8.9 The three axes for identifying uniquely each block of data stored in the conditions database.

The collection of all data blocks referring to the same data type, for any given validity time and version number, will be called a "condition folder" in the following, while a collection of condition folders will be referred to as a "condition folder set"; in the example above, "/SlowControl" is the name of the folderset containing folder "SlowControl/Hcal". Instead of a version number, a "tag" name will most often be used to indicate a consistent set of condition data blocks of a given folder, such that at every given time only one such data block within the folder has the given tag. In addition to "local" tag names defined for a given folder, "global" tag names defined over a collection of many folders of the CondDB can also be used. As a consequence, a block of condition data can be retrieved from the CondDB by specifying three variables: a folder name, a validity time and a (local or global) tag name.

The actual data content of each condition block is required to consist of a byte stream, which may be taken to represent a string. This may be used, for instance, to store XML documents as strings in the CondDB, rather than as text files like in the present implementation of the LHCB DDDDB. Eventually, other possibilities may be investigated, such as storing C++ objects



serialized as byte streams, or storing a string containing an object identifier in an Objectivity database or a table identifier in an Oracle database.

8.6.2 Conditions Database prototype implementation and examples

In collaboration with LHCB and other experiments, a conditions database system with these characteristics has recently been developed by the CERN IT/DB group (see online documentation at <http://wwwinfo.cern.ch/objectivity/docs/conditionsdb/>). The current implementation of this product is based on the Objectivity object database system, but another implementation based on the Oracle relational database is being developed by the same group. It has been decided that the production version of the LHCB conditions database will be based on the Oracle implementation of this package, a first prototype of which might be ready to be used by the next Gaudi software release. In the meantime, the package based on Objectivity has been used to prototype the relevant services of the Gaudi framework.

8.6.2.1 Relevant CMT packages

An example of an algorithm retrieving time-varying condition data from the CondDB exists in package DetCondExample v2r0. This uses a number of CMT packages, as shown below,

```
use DetCond V2r0 Det
use DetDesc V9* Det
```

In addition to the generic detector description code from DetDesc, the example makes use of package DetCond, containing the relevant conversion service for DataObjects stored in the CondDB, class ConditionsDBCnvSvc, described later on. Package DetCond, in turn, depends on package CONDDDB which allows to use Objectivity and the corresponding implementation of the CondDB package from the IT/DB group.

Package DetCondExample does not use the production version of the DDDDB from XmIDDDB, as it contains its own simplified private version of the DDDDB. Two examples are actually contained in this package: one tests the retrieval of Condition objects, stored in XML files, through the specialised interfaces of a detector element such as IAlignment; another tests the retrieval of Condition objects, stored in XML strings in the Objectivity based CondDB, using the ConditionsDBCnvSvc. In this second example, a third job options file allows to populate an Objectivity conditions database using sample data.

While the first example works under both Linux and Windows, the second example only works under Linux, as Objectivity is currently not supported under Windows at CERN. The two examples test different branches of the same DDDDB tree, showing that file-based and CondDB-based persistence systems can be used at the same time. The following will concentrate on the description of the CondDB example.

8.6.2.2 Condition objects in the Gaudi store and data blocks in the CondDB

While user algorithms are only meant to retrieve Condition objects by a request to the corresponding interface of a DetectorElement (for instance, the Hcal slow control Condition is



accessible as `hcal->slowControl()->condition()`, as described in Listing 8.4), it is important to understand the underlying mechanism to retrieve them into the transient store from their persistent representation in the CondDB.

To start with, it is assumed in the following that all parameters and parameter vectors which refer to the same detector element and share the same time validity range should be grouped together into the same Condition object, and that one and only one such Condition should be stored in each data block in the CondDB. While more than one Condition could be stored in the same data block (just like two different elements can be stored in the same XML file), two such Conditions would be forced to share the same validity range, folder name and tag, and are best thought of as two different subsets of parameters of the same Condition. This one-to-one correspondence between the granularity of transient Condition objects and persistent data blocks in the CondDB can be relevant for time performance reasons.

In the design of the framework services for Condition data in the CondDB, it is further assumed that, for every condition folder in the CondDB, only one Condition object from that folder can be loaded in the Gaudi transient data store, i.e. that which is valid at the time of the current event, and for a given global tag specified by the user in the job options. This allows to set a one-to-one correspondence between a folder name, and a path name in the Gaudi store.

In summary, in response to a statement like `hcal->slowControl()->condition()`, the three variables necessary to retrieve a block of condition data from the CondDB, the folder name, the tag name and the validity time, are determined as follows.

- The validity time is a property of the DetectorDataSvc. While in the example it is set by an external "clock" using the IDetDataSvc special interface, eventually the DetectorDataSvc should read it from the current event data whenever it is interrupted by an Incident signalling the start of a new event.
- The default global tag is a property of the ConditionsDBCnvSvc, set in the job options

```
// Production version of detector conditions
ConditionsDBCnvSvc.condDbGlobalTag = "PRODUCTION";
```

- Last, the folder name of the requested data block in the CondDB is dictated by two distinct associations: that of the Hcal interface for slow control information to the path name in the transient store where its Condition must be loaded, and that of this path name to a folder name in the CondDB (which requires the use of a specially defined protocol "condDb:/folder#/entryName"). In the example, both these associations are defined in file-based XML, part of the DDDb, as follows.

```
From file LHCB/structure.xml:
<detElem name="LHCB" type="passive">
  <geometryInfo lvnname="/dd/Geometry/LHCB/lvLHCB"/>
  <slowControlInfo condition="/dd/SlowControl/LHCB/scLHCB"/>
  <detElemref href=" ../Hcal/structure.xml#hcal"/>
</detElem>

From file Hcal/slowcontrol.xml:
<catalog name="Hcal">
  <conditionref href="condDb:/SlowControl/Hcal/scHcal#scHcal"/>
</catalog>
```



It should also be mentioned that special users of the CondDB may need to use condition data in a different way: for instance, one may want to compare two versions of the alignment, accessing at the same time data blocks from the same folder and valid at the same time, but for different global tags. Whenever one needs to access two Condition objects corresponding to the same condition folder, but valid at different times, or differently tagged, this is still possible by direct access to the ConditionsDBCnvSvc through the methods of its special ICConditionsDBCnvSvc interface. In other words, Condition objects may be created for arbitrary folder names, tag names and validity times, but they only exist in C++ memory while they are not loaded into the Gaudi data store.

8.6.2.4 Creation of Condition objects in the ConditionsDBCnvSvc

The conversion of condition data from their persistent storage as XML strings in the CondDB into transient C++ objects of class Condition in the Gaudi data store is more complex than a typical conversion process, such as that of detector elements stored in XML files. The difference is that two persistency mechanisms are used here, the storage of objects as XML strings and the storage of these strings in the CondDB. As a consequence, the conversion service responsible to retrieve Condition objects from the CondDB, the ConditionsDBCnvSvc, is a special type of conversion service. While it has no converters of its own, it is responsible to retrieve XML strings from the CondDB, then send them to another conversion service, the XmlCnvSvc described in Section 8.5.3.2, which in turn sends them for conversion to the appropriate converter (an XmlUserConditionCnv, in this case). After the creation of a Condition object, the XmlCnvSvc returns it to the ConditionsDBCnvSvc, which before returning control sets the appropriate validity range of the Condition, using the information retrieved from the CondDB.

While the example in the DetCondExample package stores and retrieves XML strings from the CondDB, the behaviour of the ConditionsDBCnvSvc does not assume that the data strings retrieved are written in XML. The "secondary" storage type for data blocks in the CondDB, which in this case is XML, is also discovered at runtime in the CondDB, by reading it off from the "description" of the corresponding CondDB folder; a single string which is meant to refer to all data blocks in the folder, independently of their validity ranges and version numbers.

8.6.2.4 Update of Condition objects

Eventually, it is foreseen that the DetectorDataSvc would automatically update all Condition objects loaded in the transient store it manages, whenever they become invalid. In the present implementation, this functionality is not yet provided, and Condition objects must be explicitly updated inside the algorithms before they are used, as in the following example. This will be modified by the time of the next release..

```
Condition* scHcal = hcal->slowControl()->condition();
StatusCode sc = detSvc()->updateObject( scHcal );
if( !sc.isSuccess() ) {
  Log << MSG::ERROR << "Can't update Condition" << endl;
  return StatusCode::FAILURE;
}
```



When requested to update a DataObject, the DetectorDataSvc does nothing if the requested object does not implement IValidity or if it is still valid at the time of the current event. Otherwise, the request is propagated through the DetPersistenceSvc to the appropriate conversion service, i.e. the ConditionsDBCnvSvc for data stored in the CondDB.

The ConditionsDBCnvSvc handles update requests for data objects in the following way. In principle, it should propagate an update request to the secondary conversion service, from where in turn it would be dispatched to the appropriate converter. However, this requires that the `updateObject()` method of the appropriate converter (XmiUserConditionCnv for a Condition object) is properly implemented, while currently this is not the case. Eventually, all converters of data that implements the IValidity interface (including DetectorElements, if these are to be stored in the CondDB) should implement this method: for Xml converters, this could be done by moving the parsing of Xml strings from the `createObject()` to the `updateObject()` method, calling the latter from inside the former. In the present implementation, update requests are handled by creating a new object of the given type, then copying its contents into the location referred to by the pointer to the old object. This "deep" copy of a Condition is actually not a complete deep copy, as the data members of the old DataObject which concern its location in the Gaudi store should not be changed. Rather than by an overloaded `operator=`, the update of the contents of the old Condition using selected contents of the new Condition is performed using the virtual `void` method `Condition::update(Condition& obj)`. This method needs to be properly overloaded by all users desiring to write derived Conditions, to perform a deep copy of any additional data members.



Chapter 9

Histogram facilities

9.1 Overview

The histogram data store is one of the data stores discussed in Chapter 2. Its purpose is to store statistics based data and user created objects that have a lifetime of more than a single event (e.g. histograms).

As with the other data stores, all access to data is via a service interface. In this case it is via the `IHistogramSvc` interface, which is derived from the `IDataProviderSvc` interface discussed in Chapter 6. The user asks the Histogram Service to book a histogram and register it in the histogram data store. The service returns a pointer to the histogram, which can then be used to fill and manipulate the histogram, using the methods defined in the `IHistogram1D` and `IHistogram2D` interfaces and documented on the AIDA (Abstract Interfaces for Data Analysis) project web pages: <http://www.inf.cern.ch/asd/lhc++/AIDA/>.

Internally, Gaudi uses the transient part of HTL (Histogram Template Library, <http://www.inf.cern.ch/asd/lhc++/HTL/>) to implement histograms.

Examples of histogram usage are available in `$LHCBSOFT/Ex/HistogramsExample`, and in section 5 of the Basic Gaudi Tutorial (http://cern.ch/lhc-comp/Frameworks/Caudi/Tutorial/5_Histogramming_and_N_tuples.pdf)

9.2 The Histogram service.

An instance of the histogram data service is created by the application manager. After the service has been initialised, the histogram data store will contain a root directory, always called `"/stat"`, in which users may book histograms and/or create sub-directories (for example, in the code fragment below, the histogram is stored in the subdirectory `"/stat/simple"`). A suggested naming convention for the sub-directories is given in

Section 1.2.3. Note that the string `"/stat/"` can be omitted when referring to a histogram in the data store: `"/stat/simple"` is equivalent to `"simple"`, **without** a leading `"/"`.

As discussed in Section 5.2, the `Algorithm` base class defines a member function which returns a pointer to the `IHistogramSvc` interface of the standard histogram data service.

```
IHistogramSvc* histoSvc()
```

Access to any other non-standard histogram data service (if one exists) must be sought via the `ISvcLocator` interface of the application manager as discussed in section 11.2.

9.3 Using histograms and the histogram service

The code fragment below shows how to book a 1D histogram and place it in a directory within the histogram data store, followed by a simple statement which fills the histogram.

```
#include "AIDA/IHistogramId.h"
...
// Book 1D histogram in the histogram data store
IHistogramId* m_hTrackCount= histoSvc()->
book( "simple", 1, "TrackCount", 100, 0., 3000. );
SmartDataPtr<MyTrackVector> particles( eventSvc()."/Event/MyTracks" )
if ( 0 != particles ) {
// Filling the track count histogram
m_hTrackCount->fill(particles->size(), 1.);
}
```

The parameters of the book function are the directory in which to store the histogram in the data store, the histogram identifier, the histogram title, the number of bins and the lower and upper limits of the X axis. 1D histograms with fixed and variable binning are available. In the case of 2D histograms, the book method requires in addition the number of bins and lower and upper limits of the Y axis.

If using HBOOK for persistency, the histogram identifier should be a valid HBOOK histogram identifier (number) and must be unique within the RZ directory the histogram is assigned to. The name of the RZ directory is given by the directory and parent directories in the transient histogram store. Please note that HBOOK accepts only directory names, which are shorter than 16 characters and that HBOOK internally converts any directory name into upper case. Even if using another persistency solution (e.g. ROOT) it is recommended to comply with the HBOOK constraints in order to make the code independent of the persistency choice.

The call to `histoSvc()->book(...)` returns a pointer to an object of type `IHistogramId` (or `IHistogram2D` in the case of a 2D histogram). All the methods of this interface can be used to further manipulate the histogram, and in particular to fill it, as shown in the example. Note that this pointer is guaranteed to be non-null, the algorithm would have failed the initialisation step if the histogram data service could not be found. On the contrary the user variable `particles` may be null (in case of absence of tracks in the transient data store and in

the persistent storage), and the fill statement would fail - so the value of `particles` must be checked before using it.

Algorithms that create histograms will in general keep pointers to those histograms, which they may use for filling operations. However it may be that you wish to share histograms between different algorithms. Maybe one algorithm is responsible for filling the histogram and another algorithm is responsible for fitting it at the end of the job. In this case it may be necessary to look for histograms within the store. The mechanism for doing this is identical to the method for locating event data objects within the event data store, namely via the use of smart pointers, as discussed in section 6.8.

```
SmartDataPtr<IHistogramId> histID( histoSvc(), "simple/1" );
if ( 0 != histID ) {
// Print the found histogram
histoSvc()->print( histID );
}
```

9.4 Persistent storage of histograms

By default, Gaudi does not produce a persistent histogram output. The options exist to write out histograms either in HBOOK or in ROOT format. The choice is made by giving the job option `ApplicationMgr.HistogramPersistency`, which can take the values "NONE" (no histograms saved, default), "HBOOK" or "ROOT". Depending on the choice, additional job options are needed, as described below.

9.4.1 HBOOK persistency

The HBOOK conversion service converts objects of types `IHistogramId` and `IHistogram2D` into a form suitable for storage in a standard HBOOK file. In order to use it you first need to tell Gaudi where to find the `HbookCnv` shared library. This is done by adding the following line to the `CMT Requirements` file:

```
use HbookCnv v*
```

You then have to tell the application manager to load this shared library and to create the HBOOK conversion service, by adding the following line to your job options file:

```
#include "$TDOPTS/Hbook.opts"
```

Finally, you have to tell the histogram persistency service the name of the output file:

```
HistogramPersistencySvc.OutputFile = "histo.hbook";
```



Note that it is also possible to print the histograms to the standard output destination (HISTDO) by setting the following job option (default is false).

```
HistogramPersistenceSvc.PrintHistos = true;
```

9.4.2 ROOT persistency

The ROOT conversion service converts objects of types THistogramLD and THistogram2D into a form suitable for storage in a standard ROOT file. In order to use it you first need to tell Gaudi where to find the RootHistCnv shared library. This is done by adding the following line to the CMT requirements file:

```
use RootHistCnv v*
```

You then have to tell the application manager to load this shared library and to create the ROOT histograms conversion service, by adding the following line to your job options file:

```
#include "$STDOPPTS/RootHist.opts"
```

Finally, you have to tell the histogram persistency service the name of the output file:

```
HistogramPersistenceSvc.OutputFile = "histo.root";
```



Chapter 10

N-tuple and Event Collection facilities

10.1 Overview

In this chapter we describe facilities available in Gaudi to create and retrieve N-tuples. We discuss how Event Collections, which can be considered an extension of N-tuples, can be used to make preselections of event data. Finally, we explore some possible tools for the interactive analysis of N-tuples.

10.2 N-tuples and the N-tuple Service

User data - so called N-tuples - are very similar to event data. Of course, the scope may be different: a row of an N-tuple may correspond to a track, an event or complete runs. Nevertheless, user data must be accessible by interactive tools such as PAW or ROOT.

Gaudi N-tuples allow to freely format structures. Later, during the running phase of the program, data are accumulated and written to disk.

The transient image of an N-tuple is stored in a Gaudi data store which is connected to the N-tuple service. Its purpose is to store user created objects that have a lifetime of more than a single event.

As with the other data stores, all access to data is via a service interface. In this case it is via the `INTupleSvc` interface which extends the `IDataProvider` interface. In addition the interface to the N-tuple service provides methods for creating N-tuples, saving the current row of an N-tuple or retrieving N-tuples from a file. The N-tuples are derived from `DataObject` in order to be storable, and are stored in the same type of tree structure as the event data. This inheritance allows to load and locate N-tuples on the store with the same smart pointer mechanism as is available for event data items (c.f. Chapter 6).

10.2.1 Access to the N-tuple Service from an Algorithm.

The `Algorithm` base class defines a member function which returns a pointer to the `INTupleSvc` interface .

```
INTupleSvc* ntupleSvc()
```

The N-tuple service provides methods for the creation and manipulation of N-tuples and the location of N-tuples within the persistent store.

The top level directory of the N-tuple transient data store is always called `"/NTUPLES"`. The next directory layer is connected to the different output streams: e.g. `"/NTUPLES/FILE1"`, where `FILE1` is the logical name of the requested output file for a given stream. There can be several output streams connected to the service. In case of persistency using `HBOOK`, `"FILE1"` corresponds to the top level `RZ` directory of the file (...the name given to `HROOPEN`). From then on the tree structure is reflected with normal `RZ` directories (caveat: `HBOOK` only accepts directory names with less than 8 characters! It is recommended to keep directory names to less than 8 characters even when using another technology (e.g. `ROOT`) for persistency, to make the code independent of the persistency choice.). Note that the top level directory name `"/NTUPLES/"` can be omitted when referring to an N-tuple in the transient data store - in the example above the name could start with `"FILE1"` (without a leading `"/"`).

10.2.2 Using the N-tuple Service.

This section explains the steps to be performed when defining an N-tuple:

- The N-tuple tags must be defined.
- The N-tuple must be booked and the tags must be declared to the N-tuple.
- The N-tuple entries have to be filled.
- The filled row of the N-tuple must be committed.
- Persistent aspects are steered by the job options.

10.2.2.1 Defining N-tuple tags

When creating an N-tuple it is necessary to first define the tags to be filled in the N-tuple, as shown for example in Listing 10.1:

Listing 10.1 Definition of N-tuple tags from the `NTuples.WriteAlg.h` example header file.

```
1: NTuple::Item<long>          m_ntrk; // A scalar item (number)
2: NTuple::Array<bool>       m_flag; // Vector items
3: NTuple::Array<long>       m_index;
4: NTuple::Array<float>      m_px, m_py, m_pz;
5: NTuple::Matrix<long>      m_hits; // Two dimensional tag
```



Typically the tags belong to the filling algorithm and hence should be provided in the Algorithm's header file. Currently the supported data types are: `bool`, `long`, `float` and `double`. `double` types (fortran `REAL*8`) are not recommended if using `HBOOK` for persistency. `HBOOK` will complain if the N-tuple structure is not defined in a way that aligns `double` types to 8 byte boundaries. In addition `PAW` cannot understand `double` types.

10.2.2.2 Booking and Declaring Tags to the N-tuple

Listing 10.2 shows how to book a column-wise N-tuple. The first directory specifier (`FILE1` in the example) must correspond to an open output stream (see Section 10.2.3.2); lower directory levels are created automatically. After booking, the previously defined tags must be declared to the N-tuple; if not, they are invalid and will cause an access violation at run-time.

Listing 10.2 Creation of a column-wise N-tuple in a specified directory and file.

```
1: #include "GaudiKernel/NTuple.h"
2: ..
3: NTuplePtr ntl(intupleSvc(), "FILE1/MC/1");
4: if ( !ntl ) { // Check if already booked
5:   ntl=ntupleSvc()->book("FILE1/MC/1", CLID_ColumnWiseTuple, "Hello World");
6:   if ( 0 != ntl ) {
7:     // Add an index column
8:     status = ntl->addItem("Ntrck", m_ntrk, 0, 5000 );
9:     // Add a variable size column, type float (length=length of index col)
10:    status = ntl->addIndexedItem("px", m_ntrk, m_px);
11:    status = ntl->addIndexedItem("py", m_ntrk, m_py);
12:    status = ntl->addIndexedItem("pz", m_ntrk, m_pz);
13:    // Another one, but this time of type bool
14:    status = ntl->addIndexedItem("flg", m_ntrk, m_flg);
15:    // Another one, type integer, numbers must be within [0, 5000]
16:    status = ntl->addIndexedItem("idx", m_ntrk, m_index, 0, 5000 );
17:    // Add 2-dim column: [0:m_ntrk][0:2]; numerical numbers within [0, 8]
18:    status = ntl->addIndexedItem("hit", m_ntrk, 2, m_hits, 0, 8 );
19:  }
20: } else { // did not manage to book the N tuple.....
21:   return StatusCode::FAILURE;
22: }
23: }
```

In previous versions of Gaudi (up to v8), indexed items were added with the `addItem` function, causing confusion for users. For this reason the calls to add indexed arrays and matrices were changed, these should now be added using the member function `addIndexedItem`. Please consult the doxygen code documentation for further details. The old calls still exist, however they are deprecated.

Row wise N-tuples are booked in the same way, but giving the type `CLID_RowWiseTuple`. However, only individual items (class `NTuple::Item`) can be filled, no arrays and no matrices. Clearly this excludes the usage of indexed items. For row-wise N-tuples to be saved in `HBOOK` format, it is recommended to use only `float` type, for the reasons explained in Section 10.2.3.3.

When using `HBOOK` for persistency, the N-tuple identifier ("1" in this example) must be a number and must be unique in a given directory. This is a limitation imposed by `HBOOK RZ`.



directories. It is recommended to keep this number unique even when using another technology (e.g. ROOT) for persistency, to make the code independent of the persistency choice.

10.2.2.3 Filling the N-tuple

Tags are usable just like normal data items, where

- Items<TYPE> are the equivalent of numbers: bool, long, float.
- Array<TYPE> are equivalent to 1 dimensional arrays: bool [size], long [size], float [size]
- Matrix<TYPE> are equivalent to an array of arrays or matrix: bool [dim1] [dim2].

Implicit bounds checking is not possible without a rather big overhead at run-time. Hence it is up to the user to ensure the arrays do not overflow.

When all entries are filled, the row must be committed, i.e. the record of the 7N-tuple must be written.

Listing 10.3 Filling an N-tuple.

```

1: m_ntrk = 0;
2: for( MyTrackVector::iterator i = mytracks->begin(); i !=
   mytracks->end(); i++) {
3:   const HepLorentzVector& mom4 = (*i)->fourMomentum();
4:   m_px[m_ntrk] = mom4.px();
5:   m_py[m_ntrk] = mom4.py();
6:   m_pz[m_ntrk] = mom4.pz();
7:   m_index[m_ntrk] = cnt;
8:   m_flag[m_ntrk] = (m_ntrk%2 == 0) ? true : false;
9:   m_hits[m_ntrk][0] = 0;
10:  m_hits[m_ntrk][1] = 1;
11:  m_ntrk++;
12:  // Make sure the array(s) do not overflow.
13:  if ( m_ntrk > m_ntrk->range().distance() ) break;
14: }
15: // Commit N tuple row. See Listing 10.2 for initialisation of m_ntuple
16: status = m_ntuple->write();
17: if ( !status.isSuccess() ) {
18:   log << MSG::ERROR << "Cannot fill id 1" << endl;
19: }

```

10.2.2.4 Reading N-tuples

Although N-tuples intended for interactive analysis, they can also be read by a regular program. An example of reading back such an N-tuple is given in Listing 10.4.

Listing 10.4 Reading an N-tuple.

```

1: NTuplePtr nt(ntuplesvc(), "FILE1/ROW_WISE/2");
2: if ( nt ) {
3:   long count = 0;
4:   ntuple::Item<float> px, py, pz;
5:   status = nt->item("px", px);
6:   status = nt->item("py", py);
7:   status = nt->item("pz", pz);
8:   // Access the N tuple row by row and print the first 10 tracks
9:   while ( nt->read().isSuccess() ) {
10:    log << MSG::INFO << " Entry [ " << count++ << " ]:";
11:    log << " Px=" << px << " Py=" << py << " Pz=" << pz << endl;
12:  }
13: }

```

10.2.3 N-tuple Persistency

10.2.3.1 Choice of persistency technology

N-tuples are of special interest to the end-user, because they can be accessed using commonly known tools such as PAW, ROOT or Java Analysis Studio (JAS). In the past it was not a particular strength of the software used in HEP to plug into many possible persistent data representations. Except for JAS, only proprietary data formats are understood. For this reason the choice of the output format of the data depends on the preferred analysis tool/viewer.

HBOOK This data format is used by PAW. PAW can understand this and only this data format. Files of this type can be converted to the ROOT format using the h2root data conversion program. The use of PAW in the long term is deprecated.

ROOT This data format is used by the interactive ROOT program.

In the current implementation, N-tuples must use the same persistency technology as histograms. The choice of technology is therefore made in the same way as for histograms, as described in Section 9.4. Obviously the options have to be given only once and are valid for both histograms and N-tuples. The only difference is that histograms are saved to a different output file (defined by the job option HistogramPersistencySvc.OutputFile), a different output file (or set of output files) must be defined for the N-tuples.

10.2.3.2 Input and Output File Specification

Conversion services exist to convert N-tuple objects into a form suitable for persistent storage in a number of storage technologies. In order to use this facility it is necessary to add the following line in the job options file:

```

NTupleSvc.Output = { "FILE1 DATAFILE='tuples.hbook' OPT='NEW'",
  "FILE2 ...",
  ...
  "FILEN ..." };

```



where `<tuplees.hbook>` should be replaced by the name of the file to which you wish to write the N-tuple. FILE1 is the logical name of the output file, which must be the same as the data store directory name described in Section 10.2.1. Several files are possible, corresponding to different data store directories whose name can be chosen at will.

The detailed syntax of the options is as follows. In each case only the three leading characters are significant: `DATAFILE=<...>`, `DATABASE=<...>` or simply `DATA=<...>` would lead to the same result.

- `DATAFILE='<file-spec>'`
Specifies the datafile (file name) of the output stream.
- `OPT='<opt-spec>'`
 - `NEW, CREATE, WRITE`: Create a new data file. Not all implementations allow to over-write existing files.
 - `OLD, READ`: Access an existing file for read purposes
 - `UPDATE`: Open an existing file and add records. It is not possible to update already existing records.

A similar option `NTupleSvc . Input` exists for N-tuple input.

10.2.3.3 Saving row wise N-tuples in HBOOK

Since the persistent representation of row wise N-tuples in HBOOK is done by floats only, a convention is needed to access the proper data type. By default the `Float` type is assumed, i.e. all data members are of `Float` type. This is the recommended format.

It is possible to define row wise N-tuples in Gaudi containing data types other than `Float`. This was the default in Gaudi versions previous to v8, where the first row of the N-tuple contained the type information. This possibility can be switched on by using the option

```
HistogramPersistencySvc . RowWiseNTuplePolicy = "USE_DATA_TYPES";
```

which also provides backwards compatibility for reading back old N-tuples produced with old Gaudi versions. Remember however that when using PAW to read N-tuples produced using this option, you must skip the first row and start with the second event.

10.3 Event Collections

Event collections or, to be more precise, event tag collections, are used to minimize data access by performing preselections based on small amounts of data. Event tag data contain flexible event classification information according to the physics needs. This information could either be stored as flags indicating that the particular event has passed some preselection criteria, or as a small set of parameters which describe basic attributes of the event. Fast access is required for this type of event data.

Event tag collections can exist in several versions:

- Collections recorded during event processing stages from the online, reconstruction, reprocessing etc.
- Event collections defined by analysis groups with pre-computed items of special interest to a given group.
- Private user defined event collections.

Starting from this definition an event tag collection can be interpreted as an N-tuple which allows to access the data used to create the N-tuple. Using this approach any N-tuple which allows access to the data is an event collection.

Event collections allow pre-selections of event data. These pre-selections depend on the underlying storage technology.

First stage pre-selections based on scalar components of the event collection. First stage preselection is not necessarily executed on your computer but on a database server e.g. when using ORACLE. Only the accessed columns are read from the event collection. If the criteria are fulfilled, the N-tuple data are returned to the user process. Preselection criteria are set through a job options, as described in section 10.3.2.3.

The **second stage pre-selection** is triggered for all items which passed the first stage pre-selection criteria. For this pre-selection, which is performed on the client computer, all data in the N-tuple can be used. The further preselection is implemented in a user defined function object (functor) as described in section 10.3.2.3. Gaudi algorithms are called only when this pre-selector also accepts the event, and normal event processing can start.

10.3.1 Writing Event Collections

Event collections are written to the data file using a Gaudi sequencer. A sequencer calls a series of algorithms, as discussed in section 5.2. The execution of these algorithms may terminate at any point of the series (and the event not selected for the collection) if one of the algorithms in the sequence fails to pass a filter.

10.3.1.1 Defining the Address Tag

The event data is accessed using a special N-tuple tag of the type

```
NTuple::Item<IOpaqueAddress*> m_evtAddress
```

It is defined in the algorithm's header file in addition to any other ordinary N-tuple tags, as described in section 10.2.2.1. When booking the N-tuple, the address tag must be declared like any other tag, as shown in Listing 10.5. It is recommended to use the name "Address" for this tag.

The usage of this tag is identical to any other tag except that it only accepts variables of type `IOpaqueAddress` - the information necessary to retrieve the event data.



Listing 10.5 Connecting an address tag to an event collection N-tuple.

```

1: StatusCode status = service("EvtTupleSvc", m_evtTupleSvc);
2: if ( status.isSuccess() ) {
3:   NtuplePtr nt(m_evtTupleSvc, "/NTUPLES/EvtColl/Collection");
4:   ... Book N-tuple
5:   // Add an event address column
6:   status = nt->addItem("Address", m_evtAddress);

```

Please note that the event tuple service necessary for writing event collections is not instantiated by default and hence must be specified in the job options file:

Listing 10.6 Adding the event tag collection service to the job options.

```

1: ApplicationMgr.ExtSvc += { "TagCollectionsSvc/EvtTupleSvc" };

```

It is up to the user to locally remember within the algorithm writing the event collection tuple the reference to the corresponding service. Although the TagCollectionsSvc looks like an N-tuple service, the implementation is different.

10.3.1.2 Filling the Event Collection

At fill time the address of the event must be supplied to the Address item. Otherwise the N-tuple may be written, but the information to retrieve the corresponding event data later will be lost. Listing 10.7 also demonstrates the setting of a filter to steer whether the event is written out to the event collection.

Listing 10.7 Fill the address tag of an N-tuple at execution time:

```

1: SmartDataPtr<Event> evt(eventSvc(), "/Event");
2: if ( evt ) {
3:   ... Some data analysis deciding whether to keep the event or not
4:   // keep_event=true if event should be written to event collection
5:   setFilterPassed( keep_event );
6:   m_evtAddrColl = evt->address();
7: }

```

10.3.2 Event Collection Persistency

10.3.2.1 Output File Specification

Conversion services exist to convert event collection objects into a form suitable for persistent storage in a number of storage technologies. In order to use this facility it is necessary to add the following line in the job options file:

```

EvtTupleSvc.Output = { "FILE1 DATAFILE='coll.root' TYP='ROOT' OPT='NEW'",
                      "FILE2 ...",
                      ...
                      "FILEN ..." };

```



where `<coll.root>` should be replaced by the name of the file to which you wish to write the event collection. FILE1 is the logical name of the output file - it could be any other string.

The options are the same as for N-tuples (see Section 10.2.3.2) with the following additions:

- **TYP=<typ-spec>**
Specifies the type of the output stream. Currently supported types are:
 - ROOT: Write as a ROOT tree.
 - MS Access: Write as a Microsoft Access database.
 - There is also weak support for the following database types:
 - SQL Server
 - MySQL
 - Oracle ODBC

These database technologies are supported through their ODBC interface. They were tested privately on local installations. However all these types need special setup to grant access to the database.

You need to specify the use of the technology specific persistency package (i.e. GaudiRootDb) in your CMT requirements file and to load explicitly in the job options the DLLs containing the generic (GaudiDb) and technology specific (e.g. GaudiRootDb) implementations of the database access drivers:

ApplicationMgr.DLLs += { "GaudiDb", "GaudiRootDb" };

- **SVC=<service-spec> (optional)**

Connect this file directly to an existing conversion service. This option however needs special care. It should only be used to replace default services.

- **AUTHENTICATION=<authentication-specs> (optional)**

For protected datafiles (e.g. Microsoft Access) it can happen that the file is password protected. In this case the authentication string allows to connect to these databases. The connection string in this case is the string that must be passed to ODBC, for example: AUTH=SERVER=server_host;UID=user_name;PWD=my_password;

- All other options are passed without any interpretation directly to the conversion service responsible to handle the specified output file.

For all options at most three leading characters are significant: DATAFILE=<...>, DATABASE=<...> or simply DATA=<...> would lead to the same result. Additional options are available when accessing events using an event tag collection.

10.3.2.2 Writing out the Event Collection

The event collection is written out by an EvtCollectionStream, which is the last member of the event collection Sequencer. Listing 10.8 (which is taken from the job options of EvtCollection example), shows how to set up such a sequence consisting of a user written Selector algorithm (which could for example contain the code in Listing 10.7), and of the EvtCollectionStream.



Listing 10.8 Job options for writing out an event collection

```

1: ApplicationMgr.OutStream = { "Sequencer/EvtCollection" };
2: ApplicationMgr.ExtSvc += { "TagCollectionSvc/EvtTupleSvc" };
3: EvtCollection.Members = { "EvtCollectionWrite/Selector",
    "EvtCollectionStream/Writer" };
4: Writer.ItemList = { "/NTUPLES/EvtColl/Collection" };
5: Writer.EvtDataSvc = "EvtTupleSvc";
6: EvtTupleSvc.Output = { "EvtColl DATAFILE="MyEvtCollection.root",
    OPT="NEW" TYP="ROOT" " " };

```

10.3.2.3 Reading Events using Event Collections

Reading event collections as the input for further event processing in Gaudi is transparent. The main change is the specification of the input data to the event selector:

Listing 10.9 Connecting an address tag to an N-tuple.

```

1: EventSelector.Input = {
2:   "COLLECTION='Collection' ADDRESS='Address'
   DATAFILE='MyEvtCollection.root' TYP='ROOT' SEL='(Ntrack>80)'
   FUN='EvtCollectionsSelector' "
3: };

```

The tags that were not already introduced earlier are:

- **COLLECTION**
Specifies the sub-path of the N-tuple used to write the collection. If the N-tuple which was written was called e.g. "/NTUPLES/FILE1/Collection", the value of this tag must be "Collection".
- **ADDRESS (optional)**
Specifies the name of the N-tuple tag which was used to store the opaque address to be used to retrieve the event data later. This is an optional tag, the default value is "address". Please use this default value when writing, conventions are useful!
- **SELECTION (optional):**
Specifies the selection string used for the first stage pre-selection. The syntax depends on the database implementation; it can be:
 - SQL like, if the event collection was written using ODBC.
Example: (NTrack>200 AND Energy>200)
 - C++ like, if the event collection was written using ROOT.
Example: (NTrack>200 && Energy>200).
Note that event collections written with ROOT also accept the SQL operators 'AND' instead of '&&' as well as 'OR' instead of '||'. Other SQL operators are not supported.
- **FUNCTION (optional)**
Specifies the name of a function object used for the second-stage preselection. An example of a such a function object is shown in Listing 10.10. Note that the factory declaration on line 16 is mandatory in order to allow Gaudi to instantiate the function object.



- The **DATAFILE** and **TYP** tags, as well as additional optional tags, have the same meaning and syntax as for N-tuples, as described in section 10.2.3.2.

Listing 10.10 Example of a function object for second stage pre-selections.

```

1: class EvtCollectionsSelector : public NTuple::Selector {
2:   NTuple::Item<long> m_ntrack;
3: public:
4:   EvtCollectionsSelector(Interface* svc) : NTuple::Selector(svc) { }
5:   virtual ~EvtCollectionsSelector() { }
6:   // Initialization
7:   virtual StatusCode initialize(NTuple::Tuple* nt) {
8:     return nt->Item("Ntrack", m_ntrack);
9:   }
10:  // Specialized callback for NTuples
11:  virtual bool operator()(NTuple::Tuple* nt) {
12:    return m_ntrack>cut;
13:  }
14: };
15:
16: ObjectFactory<EvtCollectionsSelector> EvtCollectionsSelectorFactory

```

10.3.3 Interactive Analysis using Event Tag Collections

Event tag collections are very similar to N-tuples and hence allow within limits some interactive analysis.

10.3.3.1 Interactive Access to Event Tag Collections written with ROOT

This data format is used by the interactive ROOT program. Using the helper library TB1ob located in the package GaudiRootDb it is possible to interactively analyse the N-tuples written in ROOT format. However, access is only possible to scalar items (int, float, ...) not to arrays.

Analysis is possible through directly plotting variables:

```

root [1] gSystem->Load("D:/mycmt/GaudiRootDb/v3/Win32Debug/TB1ob");
root [2] TFile* f = new TFile("tuple.root");
root [3] TTree* t = (TTree*) f->Get("<local>_MC_ROW_WISE_2");
root [4] t->Draw("pz");

```

or using a ROOT macro interpreted by ROOT's C/C++ interpreter (see for example the code fragment interactive.C shown in Listing 10.11):

```

root [0] gSystem->Load("D:/mycmt/GaudiRootDb/v3/Win32Debug/TB1ob");
root [1] L ./v8/NTuples/interactive.C
root [2] interactive("./v8/NTuples/tuple.root");

```

More detailed explanations can be found in the ROOT tutorials (<http://root.cern.ch>).



Listing 10.11 Interactive analysis of ROOT N-tuples: interactive.C

```
1: void interactive(const char* fname) {
2:   TFile *input = new TFile(fname);
3:   TTree *tree = (TTree*)input->Get("<local>_MC_ROW_WISE_2");
4:   if ( 0 == tree ) {
5:     printf("Cannot find the requested tree in the root file!\n");
6:     return;
7:   }
8:   Int_t ID, OBJSIZE, NUMLINK, NUMSYMB;
9:   TBlob *BUFFER = 0;
10:  Float_t px, py, pz;
11:  tree->SetBranchAddresses("ID",&ID);
12:  tree->SetBranchAddresses("OBJSIZE",&OBJSIZE);
13:  tree->SetBranchAddresses("NUMLINK",&NUMLINK);
14:  tree->SetBranchAddresses("NUMSYMB",&NUMSYMB);
15:  tree->SetBranchAddresses("BUFFER",&BUFFER);
16:  tree->SetBranchAddresses("px",&px);
17:  tree->SetBranchAddresses("py",&py);
18:  tree->SetBranchAddresses("pz",&pz);
19:  Int_t nbytes = 0;
20:  for (Int_t i = 0, nentries = tree->GetEntries(); i<nentries;i++) {
21:    nbytes += tree->GetEntry(i);
22:    printf("Trk#=%d PX=%f PY=%f PZ=%f\n",i,px,py,pz);
23:  }
24:  printf("I have read a total of %d Bytes.\n", nbytes);
25:  delete input;
26: }
```



Listing 10.12 Feed event tag data from MS Access into an Excel spreadsheet using Visual Basic:

```
27: Sub FillSpreadSheet ()
28: Dim dbs As Database, rst As Recordset
29: Dim strSQL as String
30: Const conPath = "D:\mycmt\GaudiExamples\vb8\Visual\tuple.mdb"
31: the_sheet_name = ActiveSheet.Name
32: Sheets(the_sheet_name).Select
33: If (IsEmpty(Selection)) Then
34:   GoTo Done
35: End If
36: ' Open database and return reference to Database object.
37: Set dbs = DBEngine.Workspaces(0).OpenDatabase(conPath)
38: strSQL = "SELECT px, py, pz FROM tblMCROWS2 ORDER BY ID;"
39: ' Open dynaset-type recordset.
40: Set rst = dbs.OpenRecordset(strSQL, dbOpenDynaset)
41: bin = 1
42: Do Until rst.EOF
43:   Cells(bin, "A:A") = rst!px
44:   Cells(bin, "B:B") = rst!py
45:   Cells(bin, "C:C") = rst!pz
46:   bin = bin + 1
47:   rst.MoveNext
48: Loop
49: rst.Close ' Close recordset and database.
50: dbs.Close
51: Done:
52: End Sub
```

10.3.3.2 Interactive Access to Event Tag Collections written with ODBC

Open DataBase Connectivity (ODBC) developed by Microsoft allows to access a very wide range of relational databases using the same callable interface. A Gaudi interface to store and retrieve data from ODBC tables was developed and offers the entire range of MS Office applications to access these data. The small Visual Basic program in Listing 10.12 shows how to fill an Excel spreadsheet using n-tuple data from an Access database. Apparently access to ODBC compliant databases using ROOT is also possible, but this was not tested.

10.4 Known Problems

Nothing is perfect and there are always things to be sorted out....

- When building the Gaudi.RootDb package on Linux using CMT you must first set up the ROOT environment, by sourcing the `setup.csh` file





Chapter 11

Framework services

11.1 Overview

Services are generally sizeable components that are setup and initialized once at the beginning of the job by the framework and used by many algorithms as often as they are needed. It is not desirable in general to require more than one instance of each service. Services cannot have a "state" because there are many potential users of them so it would not be possible to guarantee that the state is preserved in between calls.

In this chapter we describe how services are created and accessed, and then give an overview of the various services, other than the data access services, which are available for use within the Gaudi framework. The `JobOptions` service, the `Message` service, the `ParticleProperties` service, the `Chrono & Stat` service, the `Auditor` service, the `RandomNumbers` service, the `Incident` service and the `Introspection` service are available in this release. The `Tools` service is described in Chapter 12.

We also describe how to implement new services for use within the Gaudi environment. We look at how to code a service, what facilities the `Service` base class provides and how a service is managed by the application manager.

11.2 Requesting and accessing services

The Application manager only creates by default the `JobOptionsSvc` and `MessageSvc`. Other services are created on demand the first time they are accessed, provided the corresponding DLL has been loaded. The services in the GaudiSvc package are accessible in this way by default - these are the default data store services (`EventDataSvc`, `DetectorDataSvc`, `HistogramDataSvc`, `NTupleSvc`) and many of the framework services described in this chapter and in Chapter 12 (`ToolSvc`, `ParticlePropertySvc`, `ChronoStatSvc`, `AuditorSvc`, `RndmGenSvc`, `IncidentSvc`).



Additional services can be made accessible by loading the appropriate DLL, using the property `ApplicationMgr.DLLs` in the job options file, as shown for example in Listing 6.6 on page 57.

Sometimes it may be necessary to force the Application Manager to create a service at initialisation (for example if the order of creation is important). This can be done using the property `ApplicationMgr.ExtSvc`. In the example below this option is used to create a specific type of persistence service.:

Listing 11.1 Job Option to create additional services

```
ApplicationMgr.ExtSvc += { "DbEventChvSvc/RootEvtChvSvc" };
```

Once created, services must be accessed via their interface. The Algorithm base class provides a number of accessor methods for the standard framework services, listed on lines 25 to 36 of Listing 5.1 on page 38. Other services can be located using the templated `service` function. In the example below we use this function to return the `IParticlePropertySvc` interface of the Particle Properties Service. The third argument is optional: when set to true,

Listing 11.2 Code to access the IParticlePropertySvc interface from an Algorithm

```
#include "GaudiKernel/IParticlePropertySvc.h"
...
IParticlePropertySvc* m_ppSvc;
StatusCode sc = service( "ParticlePropertySvc", m_ppSvc, true );
if ( sc.isFailure() )
...

```

the service will be created if it does not already exist; if it is missing, or set to `false`, the service will not be created if it is not found, and an error is returned.

In components other than Algorithms and Services (e.g. Tools, Converters), which do not provide the `service` function, you can locate a service using the `serviceLocator` function:

```
#include "GaudiKernel/IParticlePropertySvc.h"
...
IParticlePropertySvc* m_ppSvc;
IService* theSvc;
StatusCode sc=serviceLocator()->getService("ParticlePropertySvc",theSvc,true);
if ( sc.isSuccess() ) {
  sc = theSvc->queryInterface(IID_IParticlePropertySvc, (void**)&m_ppSvc);
}
if ( sc.isFailure() ) {
...

```



11.3 The Job Options Service

The Job Options Service is a mechanism which allows to configure an application at run time, without the need to recompile or relink. The options, or properties, are set via a job options file, which is read in when the Job Options Service is initialised by the Application Manager. In what follows we describe the format of the job options file, including some examples.

11.3.1 Algorithm, Tool and Service Properties

In general a concrete Algorithm, Service or Tool will have several data members which are used to control execution. These data members (properties) can be of a basic data type (`int`, `float`, etc.) or class (`Property`) encapsulating some common behaviour and higher level of functionality. Each concrete Algorithm, Service, Tool declares its properties to the framework using the `declareProperty` templated method as shown for example on line 12 of Listing 11.4 (see also Section 5.3.2 on page 40). The method `setProperty()` is called by the framework in the initialization phase; this causes the job options service to make repeated calls to the `setProperty()` method of the Algorithm, Service or Tool (once for each property in the job options file), which actually assigns values to the data members.

11.3.1.1 SimpleProperties

Simple properties are a set of classes that act as properties directly in their associated Algorithm, Tool or Service, replacing the corresponding basic data type instance. The primary motivation for this is to allow optional bounds checking to be applied, and to ensure that the Algorithm, Tool or Service itself doesn't violate those bounds. Available SimpleProperties are:

- `int` `==> IntegerProperty or SimpleProperty<int>`
 - `double` `==> DoubleProperty or SimpleProperty<double>`
 - `bool` `==> BooleanProperty or SimpleProperty<bool>`
 - `std::string` `==> StringProperty or SimpleProperty<std::string>`
- and the equivalent vector classes
- `std::vector<int>` `==> IntegerArrayProperty or SimpleProperty<std::vector<int>>`
 - etc.

The use of these classes is illustrated by the `EventCounter` class (Listings 11.3 and 11.4).

In the Algorithm constructor, when calling `declareProperty`, you can optionally set the bounds using any of:

```
setBounds( const T& lower, const T& upper );
setLower ( const T& lower );
setUpper ( const T& upper );
```



Listing 11.3 EventCounter.h

```

1: #include "GaudiKernel/Algorithm.h"
2: #include "GaudiKernel/Property.h"
3: class EventCounter : public Algorithm {
4: public:
5:   EventCounter( const std::string& name, ISvcLocator* pSvcLocator );
6:   ~EventCounter( );
7:   StatusCode initialize();
8:   StatusCode execute();
9:   StatusCode finalize();
10: private:
11:   IntegerProperty m_frequency;
12:   int m_skip;
13:   int m_total;
14: };

```

Listing 11.4 EventCounter.cpp

```

1: #include "GaudiAlg/EventCounter.h"
2: #include "GaudiKernel/MsgStream.h"
3: #include "GaudiKernel/AlgFactory.h"
4:
5: static const AlgFactory<EventCounter> Factory;
6: const IAlgFactory& EventCounterFactory = Factory;
7:
8: EventCounter::EventCounter( const std::string& name, ISvcLocator*
9:   pSvcLocator ) :
10:   Algorithm( name, pSvcLocator ),
11:   m_skip ( 0 ), m_total ( 0 ) {
12:   declareProperty( "Frequency", m_frequency=1 ); // [1]
13:   m_frequency.setBounds( 0, 1000 ); // [2]
14: }
15:
16: StatusCode EventCounter::initialize() {
17:   MsgStream log( msgSvc(), name() );
18:   log << MSG::INFO << "Frequency: " << m_frequency << endl; // [3]
19:   return StatusCode::SUCCESS;
20: }

```

Notes:

1. A default value may be specified when the property is declared.
2. Optional upper and lower bounds may be set (see later).
3. The value of the property is accessible directly using the property itself.

There are similar selectors and modifiers to determine whether a bound has been set etc., or to clear a bound.

```

bool hasLower ( )
bool hasUpper ( )
T lower ( )
T upper ( )
void clearBounds ( )
void clearLower ( )
void clearUpper ( )

```



You can set the value using the "=" operator or the set functions

```

bool set( const T& value )
bool setValue( const T& value )

```

The function value indicates whether the new value was within any bounds and was therefore successfully updated. In order to access the value of the property, use:

```
m_property.value( );
```

In addition there's a cast operator, so you can also use m_property directly instead of m_property.value().

11.3.1.2 CommandProperty

CommandProperty is a subclass of StringProperty that has a handler that is called whenever the value of the property is changed. Currently that can happen only during the job initialization so it is not terribly useful. Alternatively, an Algorithm could set the property of one of its sub-algorithms. However, it is envisaged that Gaudi will be extended with a scripting language such that properties can be modified during the course of execution.

The relevant portion of the interface to CommandProperty is:

```

class CommandProperty : public StringProperty {
public:
  [...]
  virtual void handler( const std::string& value ) = 0;
  [...]
};

```

Thus subclasses should override the handler() member function, which will be called whenever the property value changes. A future development is expected to be a ParsableProperty (or something similar) that would offer support for parsing the string.

11.3.2 Accessing and modifying properties

Properties are private data which are initialised by the framework using the default values given when they are declared in constructors, or the values read from the job options file. On occasions it may be necessary for components to access (or even modify) the values of properties of other components. This can be done by using the getProperty() and setProperty() methods of the IProperty interface. In the example below, an algorithm stores the default value of a cut of its sub-algorithm, then executes the sub-algorithm with a different cut, before resetting the cut back to its default value. Note that in the example we begin with a pointer to an Algorithm object, not an IAlgorithm interface. This means that we have access to the methods of both the IAlgorithm and IProperty interfaces and can therefore call the methods of the IProperty interface. In the general one may need to navigate to the IProperty interface first, as explained in Section 16.3.2.



```

Algorithm* myAlg;
...
std::string dfltCut;
StatusCode sc = myAlg->getProperty( "TheCut", dfltCut );
if ( sc.isSuccess () ) {
    myAlg->setProperty( "TheCut", "0.8" );
    StatusCode sc1 = myAlg->execute();
    ...
}
if ( sc1.isSuccess () ) msgProp->setProperty( "The Cut", dfltCut );

```

11.3.3 Job options file format

An example of a job options file was shown in Listing 4.2 on page 28. The job options file has a well-defined syntax (similar to a simplified C++-Syntax) without data types. The data types are recognised by the "Job Options Compiler", which interprets the job options file according to the syntax (described in Appendix C together with possible compiler error codes).

The job options file is an ASCII-File, composed logically of a series of statements. The end of a statement is signalled by a semicolon ";" - as in C++.

Comments are the same as in C++, with "//" until the end of the line, or between "*" and "*//".

There are four constructs which can be used in a job options file:

- Assignment statement
- Append statement
- Include directive
- Platform dependent execution directive

11.3.3.1 Assignment statement

An assignment statement assigns a certain value (or a vector of values) to a property of an object or identifier. An assignment statement has the following structure:

```
<Object / Identifier> . < Propertyname > = < value >;
```

The first token (Object / Identifier) specifies the name of the object whose property is to be set. This must be followed by a dot (.)

The next token (Propertyname) is the name of the option to be set, as declared in the declareProperty() method of the IProperty interface. This must be followed by an assign symbol (=).

The final token (value) is the value to be assigned to the property. It can be a vector of values, in which case the values are enclosed in array brackets ({}), and separated by commas (,). The token must be terminated by a semicolon (;).



The type of the value(s) must match that of the variable whose value is to be set, as declared in declareProperty(). The following types are recognised:

Boolean-type, written as true or false.

```
e.g. true; false;
```

Integer-type, written as an integer value (containing one or more of the digits '0', '1', '2', '3', '4', '5', '6', '7', '8', '9')

```
e.g.: 123; -923; or in scientific notation, e.g.: 1.2e2;
```

Real-type (similar to double in C++), written as a real value (containing one or more of the digits '0', '1', '2', '3', '4', '5', '6', '7', '8', '9' followed by a dot '.' and optionally one or more of digits again)

```
e.g.: 123.; -123.45; or in scientific notation, e.g. 12.5e7;
```

String type, written within a pair of double quotes (" ")

```
e.g.: "I am a string"; (Note: strings without double quotes are not allowed!)
```

Vector of the types above, within array-brackets ({}), separated by a comma (,)

```
e.g.: {true, false, true};
```

```
e.g.: {124, -124, 135e2};
```

```
e.g.: {123.53, -23.53, 123., 12.5e2};
```

```
e.g.: {"String 1", "String 2", "String 3"};
```

A single element which should be stored in a vector must be within array-brackets without a comma

```
e.g. {true};
```

```
e.g. {"String"};
```

A vector which has already been defined earlier in the file (or in included files) can be reset to an empty vector

```
e.g. {};
```

11.3.3.2 Append Statement

Because of the possibility of including other job option files (see below), it is sometimes necessary to extend a vector of values already defined in the other job option file. This functionality is provided by the append statement.

An append statement has the following syntax:

```
<Object / Identifier> . < Propertyname > += < value >;
```

The only difference from the assignment statement is that the append statement requires the '+' symbol instead of the '=' symbol to separate the Propertyname and value tokens.



The value must be an array of one or more values

```
eg. {true};  
eg. {"String"};  
eg. {true, false, true};  
eg. {124, -124, 135e2};  
eg. {123.53, -23.53, 123., 12.5e2};  
eg. {"String 1", "String 2", "String 3"};
```

The job options compiler itself tests if the object or identifier already exists (i.e. has already been defined in an included file) and the type of the existing property. If the type is compatible and the object exists the compiler appends the value to the existing property. If the property does not exist then the append operation "+=" behaves as assignment operation "=".

11.3.3.3 Including other Job Option Files

It is possible to include other job option files in order to use pre-defined options for certain objects. This is done using the `#include` directive:

```
#include "filename.opts"
```

The "`filename`" can also contain the path where this file is located. By convention we use ".opts" as the file extension for job options. The include directive can be placed anywhere in the job option file, usually at the top (as in C++). Note that the value of a property defined earlier in the file may be over-ridden by assigning a new value to the same property; the last value assigned is the valid value! This makes it possible to over-ride the value of a property defined in a previously included file without changing the include file.

It is possible to use environment variables in the `#include` statement, either standalone or as part of a string. Both Unix style ("`$environmentvariable`") and Windows style ("`%environmentvariable%`") are understood (on both platforms). For example, in line 2: of Listing 4.2 the logical name `$STDOPTS`, which is defined in the `Gaudi.Examples` package, points to a directory containing a number of standard job options include files that can be used by applications.

As mentioned above, you can append values to vectors defined in an included job option file. The interpreter creates these vectors at the moment he interprets the included file, so you can only append elements defined in a file included before the append-statement!

As in C/C++, an included job option file can include other job option files. The compiler checks itself whether the include file has already been included, so there is no need for `#ifndef` statements as in C or C++ to check for multiple inclusion.



11.3.3.4 Platform dependent execution

The possibility exists to execute statements only according to the used platform. Statements within platform dependent clauses are only executed if they are asserted to the current used platform.:

```
#ifdef WIN32  
(Platform-Dependent Statement)  
#else (optional)  
(Platform-Dependent Statement)  
#endif
```

Only the variable `WIN32` is defined! An `#ifdef WIN32` will check if the used platform is a Windows platform. If so, it will execute the statements until an `#endif` or an optional `#else`. On non-Windows platforms it will execute the code within `#else` and `#endif`. Alternatively one directly can check for a non-Windows platform by using the `#ifndef WIN32` clause.

11.3.3.5 Switching on/off printing

By default, the Job Options Service prints out the contents of the Job Options files to the standard output destination. The possibility exists to switch off this printing, and to toggle between the two states, as shown below:

```
1: // Switch off printing  
2: #pragma print off  
3: ..(some job options)  
4: //Switch printing back on  
5: #pragma print on
```

In the example above, all lines between line 2 and line 5 will not be printed.



11.4 The Standard Message Service

One of the components directly visible to an algorithm object is the message service. The purpose of this service is to provide facilities for the logging of information, warnings, errors etc. The advantage of introducing such a component, as opposed to using the standard `std::cout` and `std::cerr` streams available in C++ is that we have more control over what is printed and where it is printed. These considerations are particularly important in an online environment.

The Message Service is configurable via the job options file to only output messages if their "activation level" is equal to or above a given "output level". The output level can be configured with a global default for the whole application:

```
// Set output level threshold
// (1=VERBOSE, 2=DEBUG, 3=INFO, 4=WARNING, 5=ERROR, 6=FATAL, 7=ALWAYS)
MessageSvc.OutputLevel = 4;
```

and/or locally for a given client object (e.g. `myAlgorithm`):

```
myAlgorithm.OutputLevel = 2;
```

Any object wishing to print some output should (must) use the message service. A pointer to the `IMessageSvc` interface of the message service is available to an algorithm via the accessor method `msgSvc()`, see section 5.2. It is of course possible to use this interface directly, but a utility class called `MsgStream` is provided which should be used instead.

11.4.1 The MsgStream utility

The `MsgStream` class is responsible for constructing a `Message` object which it then passes onto the message service. Where the message is ultimately sent to is decided by the message service.

In order to avoid formatting messages which will not be sent because the verbosity level is too high, a `MsgStream` object first checks to see that a message will be printed before actually constructing it. However the threshold for a `MsgStream` object is not dynamic, i.e. it is set at creation time and remains the same. Thus in order to keep synchronized with the message service, which in principle could change its printout level at any time, `MsgStream` objects should be made locally on the stack when needed. For example, if you look at the listing of the `HelloWorld` class (see also Listing 11.5 below) you will note that `MsgStream` objects are instantiated locally (i.e. not using `new`) in all three of the `IAAlgorithm` methods and thus are destructed when the methods return. If this is not done messages may be lost, or too many messages may be printed.

The `MsgStream` class has been designed to resemble closely a normal stream class such as `std::cout`, and in fact internally uses an `ostream` object. All of the `MsgStream` member functions write unformatted data; formatted output is handled by the insertion operators.



An example use of the `MsgStream` class is shown below.

Listing 11.5 Use of a `MsgStream` object.

```
1: #include "GaudiKernel/MsgStream.h"
2:
3: StatusCode myAlgo::finalize() {
4:   StatusCode status = Algorithm::finalize();
5:   MsgStream log(msgSvc(), name());
6:   if ( status.isFailure() ) {
7:     // Print a two line message in case of failure.
8:     log << MSG::ERROR << " Finalize failed" << endl
9:         << "Error initializing Base class." << endl;
10:  }
11:  else {
12:    log << MSG::DEBUG << "Finalize completed successfully" << endl;
13:  }
14:  return status;
15: }
```

When using the `MsgStream` class just think of it as a configurable output stream whose activation is actually controlled by the first word (message level) and which actually prints only when "endreq" is supplied. For all other functionality simply refer to the C++ `ostream` class.

The "activation level" of the `MsgStream` object is controlled by the first expression, e.g. `MSG::ERROR` or `MSG::DEBUG` in the example above. Possible values are given by the enumeration below:

```
enum MSG::Level { VERBOSE, DEBUG, INFO, WARNING, ERROR, FATAL, ALWAYS };
```

Thus the code in Listing 11.5 will produce NO output if the print level of the message service is set higher than `MSG::ERROR`. In addition if the service's print level is lower than or equal to `MSG::DEBUG` the "Finalize completed successfully" message will be printed (assuming of course it was successful).

User interface

What follows is a technical description of the part of the `MsgStream` user interface most often seen by application developers. Please refer to the header file for the complete interface.

Insertion Operator

The `MsgStream` class overloads the '<<' operator as described below.

```
MsgStream& operator <<(TYPE arg);
```

Insertion operator for various types. The argument is only formatted by the stream object if the print level is sufficiently high and the stream is active. Otherwise the insertion operators simply return. Through this mechanism extensive debug printout does not cause large run-time overheads. All common base types such as `char`, `unsigned char`, `int`, `float`, etc. are supported



```
MsgStream& operator << (MSG::Level level);
```

This insertion operator does not format any output, but rather (de)activates the stream's formatting and forwarding engine depending on the value of level.

Accepted Stream Manipulators

The `MsgStream`-specific manipulators are presented below, e.g. `endl`: `MsgStream& endl`, `MsgStream& stream`. Besides these, the common `ostream` and `ios` manipulators such as `std::ends`, `std::endl`,... are also accepted.

endl Inserts a newline sequence. Opposite to the `ostream` behaviour this manipulator does not flush the buffer. Full name: `MsgStream& endl` (`MsgStream& s`)

ends Inserts a null character to terminate a string. Full name: `MsgStream& ends` (`MsgStream& s`)

flush Flushes the stream's buffer but does not produce any output! Full name: `MsgStream& flush` (`MsgStream& s`)

endreq Terminates the current message formatting and forwards the message to the message service. If no message service is assigned the output is sent to `std::cout`. Full name: `MsgStream& endreq` (`MsgStream& s`)

endmsg Same as `endreq`

11.5 The Particle Properties Service

The Particle Property service is a utility to find information about a named particle's Geant3 ID, Jetset/Pythia ID, Geant3 tracking type, charge, mass or lifetime. The database used by the service can be changed, but by default is the same as that used by the LHCb S/CB program. Note that the units conform to the CLHEP convention, in particular MeV for masses and ns for lifetimes. Any comment to the contrary in the code is just a leftover which has been overlooked!

11.5.1 Initialising and Accessing the Service

This service is created by adding the following line in the Job Options file:

```
// Create the particle properties service
ApplicationMgr::ExtSvc += { "ParticlePropertySvc" };
```

Listing 11.2 on page 112 shows how to access this service from within an algorithm.

11.5.2 Service Properties

The Particle Property Service currently only has one property: `ParticlePropertiesFile`. This string property is the name of the database file that should be used by the service to build up its list of particle properties. The default value of this property, on all platforms, is `$LHCDBASE/cdf/particle.cdf`

11.5.3 Service Interface

The service implements the `IParticlePropertySvc` interface. In order to use it, clients must include the file `GaudiKernel/IParticlePropertySvc.h`.

The service itself consists of one STL vector to access all of the existing particle properties, and three STL maps, one to map particles by name, one to map particles by Geant3 ID and one to map particles by stdHep ID.

Although there are three maps, there is only one copy of each particle property and thus each property must have a unique particle name and a unique Geant3 ID. Particles that are known to Geant but not to stdHep, such as Deuteron, have an artificial stdHep ID using unreserved (>7) digits. Although retrieving particles by name should be sufficient, the second and third maps are there because most often generated data stores a particle's Geant3 ID or stdHep ID, and not the particle's name. These maps speed up searches using the IDs.

1. This is an LHCb specific file. A generic implementation will be available in a future release of Gaudi



The IParticlePropertySvc interface provides the following functions:

Listing 11.6 The IParticlePropertySvc interface.

```
// IParticlePropertySvc interface:
// Create a new particle property.
// Input: particle, String name of the particle.
// Input: geantId, Geant ID of the particle.
// Input: jetsetId, Jetset ID of the particle.
// Input: type, Particle type.
// Input: charge, Particle charge (/e).
// Input: mass, Particle mass (MeV).
// Input: tlife, Particle lifetime (ns).
// Return: StatusCode - SUCCESS if the particle property was added.
virtual StatusCode push_back( const std::string& particle, int geantId, int
jetsetId, int type, double charge, double mass, double tlife );

// Create a new particle property.
// Input: pp, a Particle property class.
// Return: StatusCode - SUCCESS if the particle property was added.
virtual StatusCode push_back( ParticleProperty* pp );

// Get a const reference to the beginning of the map.
virtual const_iterator begin() const;

// Get a const reference to the end of the map.
virtual const_iterator end() const;

// Get the number of properties in the map.
virtual int size() const;

// Retrieve a property by geant id.
// Pointer is 0 if no property found.
virtual ParticleProperty* find( int geantId );

// Retrieve a property by particle name.
// Pointer is 0 if no property found.
virtual ParticleProperty* find( const std::string& name );

// Retrieve a property by StdHep id
// Pointer is 0 if no property found.
virtual ParticleProperty* findByStdHepID( int stdHepId );

// Erase a property by geant id.
virtual StatusCode erase( int geantId );

// Erase a property by particle name.
virtual StatusCode erase( const std::string& name );

// Erase a property by StdHep id
virtual StatusCode eraseByStdHepID( int stdHepId );
```



The IParticlePropertySvc interface also provides some typedefs for easier coding:

```
typedef ParticleProperty* mapped_type;
typedef std::map< int, mapped_type, std::less<int> > MapID;
typedef std::map< std::string, mapped_type, std::less<std::string> > MapName;
typedef std::map< int, mapped_type, std::less<int> > MapStdHepID;
typedef IParticlePropertySvc::VectPP VectPP;
typedef IParticlePropertySvc::const_iterator const_iterator;
typedef IParticlePropertySvc::iterator iterator;
```

11.5.4 Examples

Below are some extracts of code from the LHCb ParticleProperties example to show how one might use the service:

Listing 11.7 Code fragment to find particle properties by particle name.

```
// Try finding particles by the different methods
log << MSG::INFO << "Trying to find properties by Geant3 ID..." << endl;
ParticleProperty* pp1 = m_ppSvc->find( 1 );
if ( pp1 ) log << MSG::INFO << *pp1 << endl;
log << MSG::INFO << "Trying to find properties by name..." << endl;
ParticleProperty* pp2 = m_ppSvc->find( "e+" );
if ( pp2 ) log << MSG::INFO << *pp2 << endl;
log << MSG::INFO << "Trying to find properties by StdHep ID..." << endl;
ParticleProperty* pp3 = m_ppSvc->findByStdHepID( 521 );
if ( pp3 ) log << MSG::INFO << *pp3 << endl;
```

Listing 11.8 Code fragment showing how to use the map iterators to access particle properties.

```
// List all properties
log << MSG::DEBUG << "Listing all properties..." << endl;
for( IParticlePropertySvc::const_iterator i = m_ppSvc->begin();
    i != m_ppSvc->end(); i++ ) {
    if ( *i ) log << *( *i ) << endl;
}
```



11.6 The Chrono & Stat service

The Chrono & Stat service provides a facility to do time profiling of code (Chrono part) and to do some statistical monitoring of simple quantities (Stat part). The service is created by default by the Application Manager, with the name "ChronoStatSvc" and service ID `extern const CLID& IID_IChronoStatSvc`. To access the service from inside an algorithm, the member function `ChronoSvc()` is provided. The job options to configure this service are described in Appendix B, Table B.28.

11.6.1 Code profiling

Profiling is performed by using the `ChronoStart()` and `ChronoStop()` methods inside the codes to be profiled, e.g:

```

1:  /// ...
2:  IChronoStatSvc* svc = ChronoSvc();
3:  /// start
4:  svc->ChronoStart( "Some Tag" );
5:  /// here some user code are placed:
6:  ...
7:  /// stop
8:  svc->ChronoStop( "SomeTag" );

```

The profiling information accumulates under the tag name given as argument to these methods. The service measures the time elapsed between subsequent calls of `ChronoStart()` and `ChronoStop()` with the same tag. The latter is important, since in the sequence of calls below, only the elapsed time between lines 3 and 5 lines and between lines 7 and 9 lines would be accumulated.:

```

1:  svc->ChronoStop("Tag");
2:  svc->ChronoStop("Tag");
3:  svc->ChronoStart("Tag");
4:  svc->ChronoStart("Tag");
5:  svc->ChronoStop("Tag");
6:  svc->ChronoStop("Tag");
7:  svc->ChronoStart("Tag");
8:  svc->ChronoStart("Tag");
9:  svc->ChronoStop("Tag");

```

The profiling information could be printed either directly using the `ChronoPrint()` method of the service, or in the summary table of profiling information at the end of the job.

Note that this method of code profiling should be used only for fine grained monitoring inside algorithms. To profile a complete algorithm you should use the Auditor service, as described in section 11.7.



11.6.2 Statistical monitoring

Statistical monitoring is performed by using the `stat()` method inside user code:

```

1:  /// ... Flag and Weight to be accumulated:
2:  svc->stat( " Number of Tracks ", Flag , Weight );

```

The statistical information contains the "accumulated" flag, which is the sum of all Flags for the given tag, and the "accumulated" weight, which is the product of all Weights for the given tag. The information is printed in the final table of statistics.

In some sense the profiling could be considered as statistical monitoring, where the variable Flag equals the elapsed time of the process.

11.6.3 Chrono and stat helper classes

To simplify the usage of the Chrono & Stat Service, two helper classes were developed: `Class Chrono` and `class Stat`. Using these utilities, one hides the communications with Chrono & Stat Service and provides a more friendly environment.

11.6.3.1 Chrono

Chrono is a small helper class which invokes the `ChronoStart()` method in the constructor and the `ChronoStop()` method in the destructor. It must be used as an automatic local object.

It performs the profiling of the code between its own creation and the end of the current scope, e.g:

```

1:  #include GaudiKernel/Chrono.h
2:  /// ...
3:  { // begin of the scope
4:  Chrono chrono( ChronoSvc() , "ChronoTag" ) ;
5:  // some codes:
6:  ...
7:  } // end of the scope
8:  } // end of the scope
9:  /// ...

```

If the Chrono & Stat Service is not accessible, the Chrono object does nothing



11.6.3.2 Stat

Stat is a small helper class, which invokes the `stat()` method in the constructor:

```
1: GaudiKernel/Stat.h
2: /// ...
3: Stat stat ( chronoSvc(), "statTag", Flag, weight );
4: /// ...
```

If the Chrono & Stat Service is not accessible, the Stat object does nothing.

11.6.4 Performance considerations

The implementation of the Chrono & Stat Service uses two `std::map` containers and could generate a performance penalty for very frequent calls. Usually the penalty is small relative to the elapsed time of algorithms, but it is worth avoiding both the direct usage of the Chrono & Stat Service as well as the usage of it through the Chrono or Stat utilities inside internal loops:

```
1: /// ...
2: { /// begin of the scope
3: Chrono chrono( chronoSvc(), "Good Chrono"); /// OK
4: long double a = 0 ;
5: for( long i = 0 ; i < 1000000 ; ++i )
6: {
7: Chrono chrono( svc , "Bad Chrono"); /// not OK
8: /// some codes ;
9: a += sin( cos( sin( cos( long double i ) ) ) );
10: /// end of codes
11: Stat stat ( svc , "Bad Stat", a ); /// not OK
12: }
13: Stat stat ( svc , "Good Stat", a); /// OK
14: } /// end of the scope!
15: /// ...
```

11.7 The Auditor Service

The Auditor Service provides a set of auditors that can be used to provide monitoring of various characteristics of the execution of Algorithms. Each auditor is called immediately before and after each call to each Algorithm instance, and can track some resource usage of the Algorithm. Calls that are thus monitored are `initialize()`, `execute()` and `finalize()`, although monitoring can be disabled for any of these for particular Algorithm instances. Only the `execute()` function monitoring is enabled by default.

Several examples of auditors are provided. These are:

- **NameAuditor.** This just emits the name of the Algorithm to the Standard Message Service immediately before and after each call. It therefore acts as a diagnostic tool to trace program execution.
- **ChronoAuditor.** This monitors the cpu usage of each algorithm and reports both the total and per event average at the end of job.
- **MemoryAuditor.** This monitors the state of memory usage during execution of each Algorithm, and will warn when memory is allocated within a call without being released on exit. Unfortunately this will in fact be the general case for Algorithms that are creating new data and registering them with the various transient stores. Such Algorithms will therefore cause warning messages to be emitted. However, for Algorithms that are just reading data from the transient stores, these warnings will provide an indication of a possible memory leak. Note that currently the **MemoryAuditor** is only available for Linux.
- **MemStatAuditor.** The same as **MemoryAuditor**, but prints a table of memory usage statistics at the end.

11.7.1 Enabling the Auditor Service and specifying the enabled Auditors

The Auditor Service is enabled by the following line in the Job Options file:

```
/// Enable the Auditor Service
ApplicationMgr.DLLs += { "GaudiAud" };
```

Specifying which auditors are enabled is illustrated by the following example:

```
/// Enable the NameAuditor and ChronoAuditor
AuditorSvc.Auditors = { "NameAuditor", "ChronoAuditor" };
```



11.7.2 Overriding the default Algorithm monitoring

By default, only monitoring of the Algorithm execute () function is enabled by default. This default can be overridden for individual Algorithms by use of the following Algorithm properties:

```
// Enable initialize and finalize auditing & disable execute auditing
// for the myAlgorithm Algorithm
myAlgorithm.AuditInitialize = true;
myAlgorithm.AuditExecute   = false;
myAlgorithm.AuditFinalize  = true;
```

11.7.3 Implementing new Auditors

The relevant portion of the IAuditor abstract interface is shown below:

```
virtual StatusCode beforeInitialize( IAlgorithm* theAlg ) = 0;
virtual StatusCode afterInitialize ( IAlgorithm* theAlg ) = 0;

virtual StatusCode beforeExecute   ( IAlgorithm* theAlg ) = 0;
virtual StatusCode afterExecute    ( IAlgorithm* theAlg ) = 0;

virtual StatusCode beforeFinalize  ( IAlgorithm* theAlg ) = 0;
virtual StatusCode afterFinalize   ( IAlgorithm* theAlg ) = 0;
```

A new Auditor should inherit from the Auditor base class and override the appropriate functions from the IAuditor abstract interface. The following code fragment is taken from the ChronoAuditor:

```
#include "GaudiKernel/Auditor.h"
class ChronoAuditor : virtual public Auditor {
public:
    ChronoAuditor(const std::string& name, ISvcLocator* pSvcLocator);
    virtual ~ChronoAuditor();
    virtual StatusCode beforeInitialize( IAlgorithm* alg);
    virtual StatusCode beforeExecute( IAlgorithm* alg);
    virtual StatusCode afterExecute( IAlgorithm* alg);
    virtual StatusCode beforeFinalize( IAlgorithm* alg);
    virtual StatusCode afterFinalize( IAlgorithm* alg);
};
```

11.8 The Random Numbers Service

When generating random numbers two issues must be considered:

- reproducibility and
- randomness of the generated numbers.

In order to ensure both, Gaudi implements a single service ensuring that these criteria are met. The encapsulation of the actual random generator into a service has several advantages:

- Random seeds are set by the framework. When debugging the detector simulation, the program could start at any event independent of the events simulated before. Unlike the random number generators that were known from CERNLIB, the state of modern generators is no longer defined by one or two numbers, but rather by a fairly large set of numbers. To ensure reproducibility the random number generator must be initialized for every event.
- The distribution of the random numbers generated is independent of the random number engine behind. Any distribution can be generated starting from a flat distribution.
- The actual number generator can easily be replaced if at some time in the future better generators become available, without affecting any user code.

The implementation of both generators and random number engines are taken from CLHEP. The default random number engine used by Gaudi is the RanLux engine of CLHEP with a luxury level of 3, which is also the default for Geant4, so as to use the same mechanism to generate random numbers as the detector simulation.

Figure 11.1 shows the general architecture of the Gaudi random number service. The client interacts with the service in the following way:

- The client requests a generator from the service, which is able to produce a generator according to a requested distribution. The client then retrieves the requested generator.
- Behind the scenes, the generator service creates the requested generator and initializes the object according to the parameters. The service also supplies the shared random number engine to the generator.
- After the client has finished using the generator, the object must be released in order to inhibit resource leaks




Figure 11.1 The architecture of the random number service. The client requests from the service a random number generator satisfying certain criteria



There are many different distributions available. The shape of the distribution must be supplied as a parameter when the generator is requested by the user.

Currently implemented distributions include the following. See also the header file `GaudiKernel/RndmGenerators.h` for a description of the parameters to be supplied.

- Generate random bit patterns with parameters `Rndm::Bit()`
- Generate a flat distribution with boundaries `[min, max]` with parameters:
`Rndm::Flat(double min, double max)`
- Generate a gaussian distribution with parameters: `Rndm::Gauss(double mean, double sigma)`
- Generate a poissonian distribution with parameters: `Rndm::Poisson(double mean)`
- Generate a binomial distribution according to n tests with a probability p with parameters: `Rndm::Binomial(long n, double p)`
- Generate an exponential distribution with parameters:
`Rndm::Exponential(double mean)`
- Generate a Chi^2 distribution with n_{dof} degrees of freedom with parameters:
`Rndm::Chi2(long n_dof)`
- Generate a Breit-Wigner distribution with parameters:
`Rndm::BreitWigner(double mean, double gamma)`
- Generate a Breit-Wigner distribution with a cut-off with parameters:
`Rndm::BreitWignerCutOff(mean, gamma, cut-off)`
- Generate a Landau distribution with parameters:
`Rndm::Landau(double mean, double sigma)`
- Generate a user defined distribution. The probability density function is given by a set of discrete points passed as a vector of doubles:
`Rndm::DefinedPdf(const std::vector<double>& pdf, long intpol)`

Clearly the supplied list of possible parameters is not exhaustive, but probably represents most needs. The list only represents the present content of generators available in CLHEP and can be updated in case other distributions will be implemented.

Since there is a danger that the interfaces are not released, a wrapper is provided that automatically releases all resources once the object goes out of scope. This wrapper allows the use of the random number service in a simple way. Typically there are two different usages of this wrapper:



- Within the user code a series of numbers is required only once, i.e. not every event. In this case the object is used locally and resources are released immediately after use. This example is shown in Listing 11.9.

Listing 11.9 Example of the use of the random number generator to fill a histogram with a Gaussian distribution within a standard Gaudi algorithm

```
1: Rndm::Numbers gauss(randSvc(), Rndm::Gauss(0.5,0.2));
2: if ( gauss ) {
3:   IHistogramI* his = histoSvc()->book("/stat/2", "Gaussian", 40, 0., 3.);
4:   for ( long i = 0; i < 5000; i++ )
5:     his->fill(gauss(), 1.0);
6: }
```

- One or several random numbers are required for the processing of every event. An example is shown in Listing 11.10.

Listing 11.10 Example of the use of the random number generator within a standard Gaudi algorithm, for use at every event. The wrapper to the generator is part of the Algorithm itself and must be initialized before being used. Afterwards the usage is identical to the example described in Listing 11.9

```
1: #include "GaudiKernel/RndmGenerators.h"
2: // Constructor
3: // class myAlgorithm : public Algorithm {
4:   Rndm::Numbers m_gaussDist;
5:   ...
6: };
7: // Initialization
8: // StatusCode myAlgorithm::initialize() {
9:   ...
10:  StatusCode sc=m_gaussDist.initialize(randSvc(), Rndm::Gauss(0.5,0.2));
11:  if ( !status.isSuccess() ) {
12:    // put error handling code here...
13:  }
14:  ...
15: }
```

There are a few points to be mentioned in order to ensure the reproducibility:

- Do not keep numbers across events. If you need a random number ask for it. Usually caching does more harm than good. If there is a performance penalty, it is better to find a more generic solution.
- Do not access the `RndmEngine` directly.
- Do not manipulate the engine. The random seeds should only be set by the framework on an event by event basis.



11.9 The Incident Service

The Incident service provides synchronization facilities to components in a Gaudi application. Incidents are named software events that are generated by software components and that are delivered to other components that have requested to be informed when that incident happens. The Gaudi components that want to use this service need to implement the `IIncidentListener` interface, which has only one method: `handle(IIncident&)`, and they need to add themselves as Listeners to the `IncidentSvc`. The following code fragment works inside Algorithms.

```
#include "GaudiKernel/IIncidentListener.h"
#include "GaudiKernel/IIncidentSvc.h"
...
};

class MyAlgorithm : public Algorithm, virtual public IIncidentListener {
...
};

MyAlgorithm::Initialize() {
    IncidentSvc* incsvc;
    StatusCode sc = service("IncidentSvc", incsvc);
    int priority = 100;
    if( sc.isSuccess() ) {
        incsvc->addListener( this, "BeginEvent", priority);
        incsvc->addListener( this, "EndEvent");
    }
}

MyAlgorithm::handle(Incident& inc) {
    log << "Got informed of incident: " << inc.type()
    << " generated by: " << inc.source() << endl;
}
```

The third argument in method `addListener()` is for specifying the priority by which the component will be informed of the incident in case several components are listeners of the same named incident. This parameter is used by the `IncidentSvc` to sort the listeners in order of priority.

11.9.1 Known Incidents

Table 11.1 Table of known named incidents

Incident Type	Source	Description
BeginEvent	ApplicationMgr	The ApplicationMgr is starting processing of a new physics event. This incident can be used to clear caches of the previous event in Services and Tools.
EndEvent	ApplicationMgr	The ApplicationMgr has finished processing the physics event. The Event data store is not yet purged at this moment.



11.10 The GiGa Service

The GiGa service allows Gaudi applications to use the GEANT4 toolkit as a black box, without detailed knowledge of the GEANT4 internal features.

An overview and a detailed description of GiGa are available in references [10] and [11] respectively. The `GiGaExample` application demonstrates the usage of GiGa to display LHCb events and detector geometry.

11.11 The Gaudi Introspection Service

Introspection is the ability of a programming language to interact with objects from a meta-level. The Gaudi Introspection package defines a meta-model which gives the layout of this meta-information.

The data to fill this meta-information (i.e. the dictionary) can be generated by the Gaudi Object Description package (described in Section 6.7 on page 53) by adding a few lines to the CMT requirements file, as shown for example in Listing 11.11.

Listing 11.11 CMT requirements for generation of data dictionary of the LHCbEvent package

```
#---- dictionary
document obj2dict LHCbEventObj2Dict -group=dict ../xml/LHCbEvent.xml
library LHCbEventDict -group=dict ../dict/*.cpp
macro LHCbEventDict_shlibflags "$(use_linkopts) $(libraryshr_linkopts)"
```

The C++ code generated in this way is compiled into a dll and loaded into the Gaudi Introspection Model at runtime.

To get a reference to information about a real object, clients have to use the Gaudi Introspection Service (`IntrospectionSvc`). The service can also be used to load the meta-information into the model. The Gaudi Introspection Service is already used in several places in the framework (e.g. Interface to Python, Data Store Browser).

Further information about this service is available at <http://cern.ch/lhcb-comp/Frameworks/DataDictionary/default.htm>.



11.12 Developing new services

11.12.1 The Service base class

Within Gaudi we use the term "Service" to refer to a class whose job is to provide a set of facilities or utilities to be used by other components. In fact we mean more than this because a concrete service must derive from the Service base class and thus has a certain amount of predefined behaviour; for example it has `initialize()` and `finalize()` methods which are invoked by the application manager at well defined times.

Figure 11.2 shows the inheritance structure for an example service called `SpecificService`. The key idea is that a service should derive from the Service base class and additionally implement one or more pure abstract classes (interfaces) such as `IConcreteSvcType1` and `IConcreteSvcType2` in the figure.

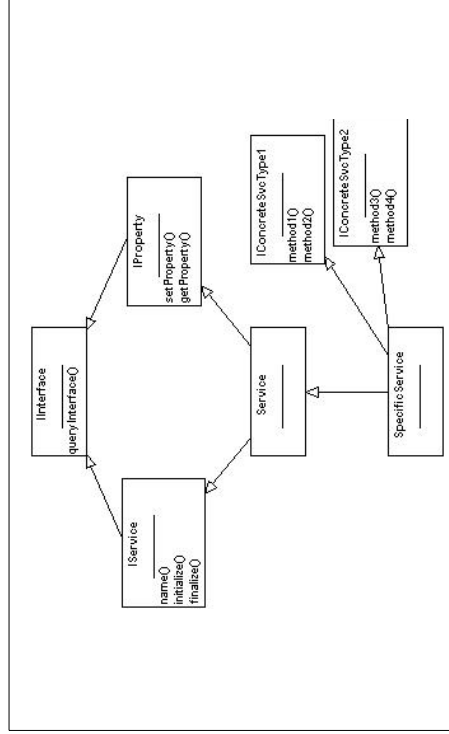


Figure 11.2 Implementation of a concrete service class. Though not shown in the figure, both of the `IConcreteSvcType` interfaces are derived from `IInterface`.

As discussed above, it is necessary to derive from the Service base class so that the concrete service may be made accessible to other Gaudi components. The actual facilities provided by the service are available via the interfaces that it provides. For example the `ParticLeProperties` service implements an interface which provides methods for retrieving, for example, the mass of a given particle. In figure 11.2 the service implements two interfaces each of two methods.

A component which wishes to make use of a service makes a request to the application manager. Services are requested by a combination of name, and interface type, i.e. an algorithm would request specifically either `IConcreteSvcType1` or `IConcreteSvcType2`.



The identification of what interface types are implemented by a particular class is done via the `queryInterface` method of the `IInterface` interface. This method must be implemented in the concrete service class. In addition the `initialize()` and `finalize()` methods should be implemented. After initialization the service should be in a state where it may be used by other components.

The service base class offers a number of facilities itself which may be used by derived concrete service classes:

- Properties are provided for services just as for algorithms. Thus concrete services may be fine tuned by setting options in the job options file.
- A `ServiceLocator` method is provided which allows a component to request the use of other services which it may need.
- A message service.

11.12.2 Implementation details

The following is essentially a checklist of the minimal code required for a service.

1. Define the interfaces
2. Derive the concrete service class from the Service base class.
3. Implement the `queryInterface()` method.
4. Implement the `initialize()` method. Within this method you should make a call to `Service::initialize()` as the first statement in the method and also make an explicit call to `setProperties()` in order to read the service's properties from the job options (note that this is different from Algorithms, where the call to `setProperties()` is done in the base class).

Listing 11.12 An interface class

```
#include "GaudiKernel/IInterface.h"
class IConcreteSvcType1 : virtual public IInterface {
public:
    void method1() = 0;
    int method2() = 0;
};

#include "IConcreteSvcType1.h"
const IID& IID_IConcreteSvcType1 = 143; // UNIQUE within LHCB !!
```



Listing 11.13 A minimal service implementation

```
#include "GaudiKernel/Service.h"
#include "IConcreteSvcType1.h"
#include "IConcreteSvcType2.h"

class SpecificService : public Service,
    virtual public IConcreteSvcType1,
    virtual public IConcreteSvcType2 {
public:
    // Constructor of this form required:
    SpecificService(const std::string& name, ISvcLocator* sl);

    queryInterface(const IID& riid, void** ppvIF);
};

// Factory for instantiation of service objects
static SvcFactory<SpecificService> s_factory;
const ISvcFactory& SpecificServiceFactory = s_factory;

// UNIQUE Interface identifiers defined elsewhere
extern const IID& IID_IConcreteSvcType1;
extern const IID& IID_IConcreteSvcType2;

// queryInterface
StatusCode SpecificService::queryInterface(const IID& riid, void** ppvIF) {
    if (IID_IConcreteSvcType1 == riid) {
        *ppvIF = dynamic_cast<IConcreteSvcType1*> (this);
        return StatusCode::SUCCESS;
    } else if (IID_IConcreteSvcType2 == riid) {
        *ppvIF = dynamic_cast<IConcreteSvcType2*> (this);
        return StatusCode::SUCCESS;
    } else {
        return Service::queryInterface(riid, ppvIF);
    }
}

// Implement the specifics ...
SpecificService::method1() {...}
SpecificService::method2() {...}
SpecificService::method3() {...}
SpecificService::method4() {...}
```



Chapter 12

Tools and ToolSvc

12.1 Overview

Tools are light weight objects whose purpose is to help other components perform their work. A framework service, the `ToolSvc`, is responsible for creating and managing Tools. An `Algorithm` requests the Tools it needs to the `ToolSvc`, specifying if requesting a private instance by declaring itself as the parent. Since Tools are managed by the `ToolSvc`, any component¹ can request a tool. `Algorithms`, `Services` and other Tools can declare themselves as `Tools` parents.

In this chapter we first describe these objects and the difference between “private” and “shared” tools. We then look at the `AlgTool` base class and how to write concrete Tools.

In section 12.3 we describe the `ToolSvc` and show how a component can retrieve `Tools` via the service.

Finally we describe `Associators`, common utility `GaudiTools` for which we provide the interface and base class.

12.2 Tools and Services

As mentioned elsewhere **Algorithms** make use of framework services to perform their work. In general the same instance of a service is used by many algorithms and **Services** are setup and initialized once at the beginning of the job by the framework. `Algorithms` also delegate some of their work to sub-algorithms. Creation and execution of sub-algorithms are the responsibilities of the parent algorithm whereas the `initialize()` and `finalize()` methods are invoked automatically by the framework while initializing the parent algorithm.

1. In this chapter we will use an `Algorithm` as example component requesting tools.

The properties of a sub-algorithm are automatically set by the framework but the parent algorithm can change them during execution. Sharing of data between nested algorithms is done via the Transient Event Store.

Both Services and Algorithms are created during the initialization stage of a job and live until the jobs ends.

Sometimes an encapsulated piece of code needs to be executed only for specific events, in which case it is desirable to create it only when necessary. On other occasions the same piece of code needs to be executed many times per event. Moreover it can be necessary to execute a sub-algorithm on specific contained objects that are selected by the parent algorithm or have the sub-algorithm produce new contained objects that may or may not be put in the Transient Store. Finally different algorithms may wish to configure the same piece of code slightly differently or share it as-is with other algorithms.

To provide this kind of functionality we have introduced a category of processing objects that encapsulate these “light” algorithms. We have called this category **Tools**.

Some examples of possible tools are single track fitters, association to Monte Carlo truth information, vertexing between particles, smearing of Monte Carlo quantities.

12.2.1 “Private” and “Shared” Tools

Algorithms can share instances of Tools with other Algorithms if the configuration of the tool is suitable. In some cases however an Algorithm will need to customize a tool in a specific way in order to use it. This is possible by requesting the ToolSvc to provide a “private” instance of a tool.

If an Algorithm passes a pointer to itself when it asks the ToolSvc to provide it with a tool, it is declaring itself as the parent and a “private” instance is supplied. Private instances can be configured according to the needs of each particular Algorithm.

As mentioned above many Algorithms can use a tool as-is, in which case only one instance of a Tool is created, configured and passed by the ToolSvc to the different algorithms. This is called a “shared” instance. The parent of “shared” tools is the ToolSvc.

12.2.2 The Tool classes

12.2.2.1 The AlgTool base class

The main responsibilities of the AlgTool base class (see Listing 12.1) are the identification of the tools instances, the initialisation of certain internal pointers when the tool is created and the management of the tools properties. The AlgTool base class also offers some facilities to help in the implementation of derived tools and management of the additional tools interfaces.



Listing 12.1 The definition of the AlgTool Base class. Highlighted in bold are methods relevant for the implementation of concrete tools.

```

1: class AlgTool : public virtual IAlgTool,
2:                 public virtual IProperty {
3:
4: public:
5: // Standard Constructor.
6: AlgTool( const std::string& type, const std::string& name,
7:         const IInterface* parent);
8: ISvLocator* serviceLocator() const;
9: IMessageSvc* msgSvc() const;
10:
11: virtual StatusCode setProperty( const Property& p );
12: virtual StatusCode setProperty( std::istream& s );
13: virtual StatusCode setProperty( const std::string& n,
14:                               const std::string& v );
15: virtual StatusCode getProperty(Property* p) const;
16: virtual StatusCode getProperty( const std::string& name ) const;
17: virtual StatusCode getProperty( const std::string& n, std::string& v )
18:                               const;
19: virtual const std::vector<Property*>& getProperties( ) const;
20:
21: StatusCode setProperties( );
22: template <class T>
23: StatusCodes declareProperty(const std::string& name, T& property) const
24: virtual const std::string& name() const;
25: virtual const std::string& type() const;
26: virtual const IInterface* parent() const;
27:
28: virtual StatusCode initialize();
29: virtual StatusCode finalize();
30:
31: virtual StatusCode queryInterface(const IID& riid, void** ppvUnknown);
32: void queryInterface( const IID&, void* );
33: template <class I> class declareInterface {
34:     public:
35:         template <class T> declareInterface(I* tool)
36:     }
37:     protected:
38:         // Standard destructor.
39:     virtual ~AlgTool();

```

Constructor - The base class has a single constructor which takes three arguments. The first is the type (i.e. the class) of the Tool object being instantiated, the second is the full name of the object and the third is a pointer to the IInterface of the parent component. The name is used for the identification of the tool instance as described below. The parent interface is used by the tool to access for example the outputLevel of the parent.



Access to Services - A `ServiceLocator()` method is provided to enable the derived tools to locate the services necessary to perform their jobs. Since concrete `Tools` are instantiated by the `ToolSvc` upon request, all `Services` created by the framework prior to the creation of a tool are available. In addition access to the message service is provided via the `msgSvc()` method. Both pointers are retrieved from the parent of the tool.

Properties - A template method for declaring properties similarly to `Algorithms` is provided. This allows tuning of data members used by the `Tools` via `JobOptions` files. The `ToolSvc` takes care of calling the `setProperties()` method of the `AlgTool` base class after having instantiated a tool. Properties need to be declared in the constructor of a `Tool`. The property `outputLevel` is declared in the base class and is identically set to that of the parent component, unless specified otherwise in the `JobOptions`. For details on Properties see section 11.3.1.

IAlgTool Interface - It consists of three accessor methods for the identification and management of the tools: `type()`, `name()` and `parent()`. These methods are all implemented by the base class and should not be overridden. Two additional methods, `initialize()` and `finalize()`, allow concrete tools to be configured after creation and orderly terminated before deletion. An empty implementation is provided by the `AlgTool` base class and concrete tools need to implement these methods only when relevant for their purpose. The `ToolSvc` is responsible for calling these methods at the appropriate time.

Tools interfaces - Concrete tools must implement additional interfaces that will inherit from `IAlgTool`. When a component implements more than one interface it is necessary to "recognize" the various interfaces. This is taken care of by the `AlgTool` base class once the additional interface is declared by a concrete tool (or 'tools' base class). The declaration of the additional interface must be done in the constructor of a concrete tool and is done via the template method `declareInterface`.

12.2.2.2 Tools identification

A tool instance is identified by its full name. The name consist of the concatenation of the parent name, a dot, and a tool dependent part. The tool dependent part can be specified by the user, when not specified the tool type (i.e. the class) is automatically taken as the tool dependent part of the name. Examples of tool names are `RecPrimaryVertex.VertexSmearer` (a private tool) and `ToolSvc.AddFourMom` (a shared tool). The full name of the tool has to be used in the `JobOptions` file to set its properties.

12.2.2.3 Concrete tools classes

Operational functionalities of tools must be provided in the derived tool classes. A concrete tool class must inherit directly or indirectly from the `AlgTool` base class to ensure that it has the predefined behaviour needed for management by the `ToolSvc`.

Concrete tools must implement additional interfaces, specific to the task a tool is designed to perform. Specialised tools intended to perform similar tasks can be derived from a common base class that will provide the common functionality and implement the common interface. Consider as example the vertexing of particles, where separate tools can implement different algorithms but the arguments passed are the same. The `ToolSvc` interacts with specialized



tools only through the additional tools interface, therefore the interface itself must inherit from the `IAlgTool` interface in order for the tool to be correctly managed by the `ToolSvc`. The inheritance structure of derived tools is shown in Figure 12.1. `ConcreteTool1` implements one additional abstract interface while `ConcreteTool2` and `ConcreteTool3` derive from a base class `SubTool` that provides them with additional common functionality.

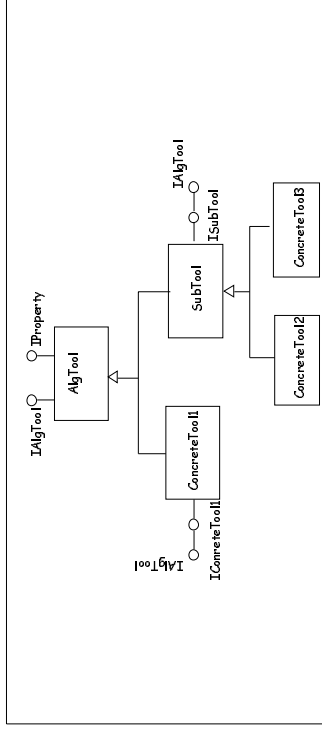


Figure 12.1 Tools class hierarchy

12.2.2.4 Implementation of concrete tools

An example minimal implementation of a concrete tool is shown in Listings 12.2, 12.3 and 12.4, taken from the `LHCbToolsAnalysis` example application .

Listing 12.2 Example of a concrete tool additional interface

```

1: static const InterfaceID IID_IVertexSmearer("IVertexSmearer", 1, 0);
2:
3: class IVertexSmearer : virtual public IAlgTool {
4: public:
5:     // Retrieve interface ID
6:     static const InterfaceID& interfaceID() { return IID_IVertexSmearer; }
7:     // Actual operator function
8:     virtual StatusCode smear( MyAxVertex* ) = 0;
9: };

```

Listing 12.3 Example of a concrete tool minimal implementation header file

```

1: #include "GaudiKernel/AlgTool.h"
2: class VertexSmearer : public AlgTool, virtual public IVertexSmearer {
3: public:
4:     // Constructor
5:     VertexSmearer( const std::string& type, const std::string& name,
6:                   const Interface* parent );
7:     virtual ~VertexSmearer() { }
8:     // specific method of this tool
9:     StatusCode smear( MyAxVertex* pvertex );

```



Listing 12.4 Example of a concrete tool minimal implementation file

```

1: #include "GaudiKernel/ToolFactory.h"
2: // Static factory for instantiation of algoTool objects
3: static ToolFactory<VertexSmearer> s_factory;
4: const IToolFactory& VertexSmearerFactory = s_factory;
5:
6: // Standard constructor
7: VertexSmearer::VertexSmearer(const std::string& type,
    const std::string& name,
    const IInterface* parent) {
8:     : AlgoTool( type, name, parent ) {
9:
10:    // Locate service needed by the specific tool
11:    m_randSvc = 0;
12:    if( serviceLocator() ) {
13:        StatusCode sc=StatusCode::FAILURE;
14:        sc = serviceLocator()->service( "RandGenSvc", m_randSvc, true );
15:    }
16:    // Declare additional interface
17:    declareInterface<IVertexSmearer>(this);
18:
19:    // Declare properties of the specific tool
20:    declareProperty("dxVtx", m_dxVtx = 9 * micrometer);
21:    declareProperty("dyVtx", m_dyVtx = 9 * micrometer);
22:    declareProperty("dzVtx", m_dzVtx = 38 * micrometer);
23: }
24: // Implement the specific method ....
25: StatusCode VertexSmearer::smear( MyAxVertex* pvertex ) { ...}

```

The creation of concrete tools is similar to that of Algorithms, making use of a Factory Method. As for Algorithms, Tool factories enable their creator to instantiate new tools without having to include any of the concrete tools header files. A template factory is provided and a tool developer will only need to add the concrete factory in the implementation file as shown in lines 1 to 4 of Listing 12.4

In addition a concrete tool class must specify a single constructor with the same parameter signatures as the constructor of the `AlgoTool` base class as shown in line 5 of Listing 12.3.

Below is the minimal checklist of the code necessary when developing a Tool:

1. Define the specific interface (inheriting from the `IAlgoTool` interface).
2. Derive the tool class from the `AlgoTool` base class
3. Provide the constructor
4. Declare the additional interface in the constructor.
5. Implement the factory adding the lines of code shown in Listing 12.4
6. Implement the specific interface methods.

In addition if a tool requires special initialization and termination you can implement the initialize and finalize methods.



12.3 The ToolSvc

The `ToolSvc` manages `Tools`. It is its responsibility to create tools, configure them, make them available to `Algorithms` or `Services` and terminate them in an orderly fashion before deleting them.

The `ToolSvc` verifies if a tool type is available and creates the necessary instance after having verified if it doesn't already exist. If a tool instance exists the `ToolSvc` will not create a new identical one but pass to the algorithm the existing instance. Tools are created on a "first request" basis: the first `Algorithm` requesting a tool will prompt its creation. The relationship between an algorithm, the `ToolSvc` and `Tools` is shown in Figure 12.2.

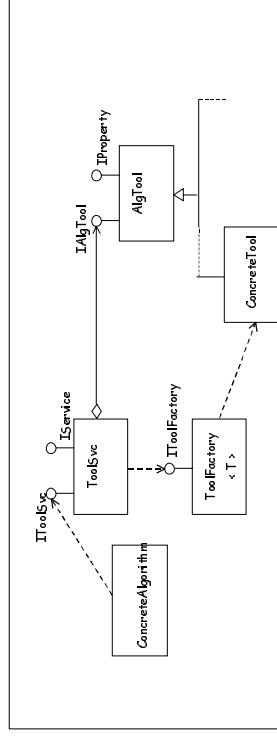


Figure 12.2 ToolSvc design diagram

Immediately after having created a tool, the `ToolSvc` will configure it by setting its properties and calling the tool `initialize()` method.

The `ToolSvc` will "hold" a tool until it is no longer used by any component or until the `finalize()` method of the tool service is called. Algorithms can inform the `ToolSvc` they are not going to use a tool previously requested via the `releaseTool` method of the `IToolSvc` interface. Before deleting the tools the `ToolSvc` will cleanly terminate them by calling their `finalize()` method.

The `ToolSvc` is created by default by the `ApplicationMgr` and algorithms wishing to use the service can do so via the algorithm `toolSvc()` accessor method. `Services` and `AlgoTools` need to retrieve it using the `serviceLocator()` method of their respective base classes.

12.3.1 Retrieval of tools via the IToolSvc interface

The `IToolSvc` interface is the `ToolSvc` specific interface providing methods to retrieve tools. The interface has two retrieve methods that differ in their parameters signature, as shown in Listing 12.5

The arguments of the method shown in Listing 12.5, line 1, are the tool type (i.e. the class), the tool additional interface ID and the `IAlgoTool` interface of the returned tool. In addition there are two arguments with default values: one is the `IInterface` of the component



Listing 12.5 The IToolSvc interface methods

```

1: virtual StatusCode retrieve(const std::string& type,
    const IID&,
    IAlgTool* & tool,
    const Interface* parent=0,
    bool createIf=true ) = 0;
2: virtual StatusCode retrieve(const std::string& type,
    const IID&,
    const std::string& name,
    IAlgTool* & tool,
    const Interface* parent=0,
    bool createIf=true ) = 0;

```

requesting the tool, the other a boolean creation flag. If the component requesting a tool passes a pointer to itself as the third argument, it declares to the ToolSvc that it is asking for a "private" instance of the tool. By default a "shared" instance is provided. In general if the requested instance of a Tool does not exist the ToolSvc will create it. This behaviour can be changed by setting to `false` the last argument of the method.

The method shown in Listing 12.5, line 2 differs from the one shown in line 1 by an extra argument, a string specifying the tool dependent part of the full tool name. This enables a component to request two separately configurable instances of the same tool.

When retrieving concrete tools, it is recommended to use the two templated functions provided in the IToolSvc interface file which are shown in Listing 12.6.

Listing 12.6 The IToolSvc template methods

```

1: template <class T>
2:   StatusCode retrieveTool( const std::string& type,
    T* & tool,
    const Interface* parent=0,
    bool createIf=true ) {...}
3: template <class T>
4:   StatusCode retrieveTool( const std::string& type,
    T* & tool,
    const Interface* parent=0,
    bool createIf=true ) {...}

```

The two template methods correspond to the IToolSvc retrieve methods but have the tool returned as a template parameter. Using these methods the component retrieving a tool avoids explicit dynamic-casting to specific additional interfaces or to derived classes.

Listing 12.7 shows an example of retrieval of a shared and of a common tool.



Listing 12.7 Example of retrieval by an algorithm of a shared tool in line 4: and of a private tool in line 10:

```

1: // Example of tool belonging to the ToolSvc and shared between
2: // algorithms
3: StatusCode sc;
4: sc = toolSvc()->retrieveTool("AddFourMom", m_sum4p );
5: if ( sc.isFailure() ) {
6:   log << MSG::FATAL << " Unable to create AddFourMom tool" << endl;
7:   return sc;
8: }
9: // Example of private tool
10: sc = toolSvc()->retrieveTool("ImpactPar", m_ip, this );
11: if ( sc.isFailure() ) {
12:   log << MSG::FATAL << " Unable to create ImpactPar tool" << endl;
13:   return sc;
14: }

```

12.4 GaudiTools

In general concrete tools are specific to applications or detectors' code but there are some tools of common utility for which interfaces and base classes can be provided. The Associators described below and contained in the Gaudi.Tools package are one of such tools.

12.4.1 Associators

When working with Monte Carlo data it is often necessary to compare the results of reconstruction or physics analysis with the original corresponding Monte Carlo quantities on an event-by-event basis as well as on a statistical level.

Various approaches are possible to implement navigation from reconstructed simulated data back to the Monte Carlo truth information. Each of the approaches has its advantages and could be more suited for a given type of event data or data-sets. In addition the reconstruction and physics analysis code should treat simulated data in an identical way to real data.

In order to shield the code from the details of the navigation procedure, and to provide a uniform interface to the user code, a set of Gaudi Tools, called Associators, has been introduced. The user can navigate between any two arbitrary classes in the Event Model using the same interface as long as a corresponding associator has been implemented. Since an Associator retrieves existing navigational information, its actual implementation depends on the Event Model and how the navigational information is stored. For some specific Associators, in addition, it can depend on some algorithmic choices: consider as an example a physics analysis particle and a possible originating Monte Carlo particle where the associating discriminant could be the fractional number of hits used in the reconstruction of the tracks. An advantage of this approach is that the implementation of the navigation can be modified without affecting the reconstruction and analysis algorithms because it would affect only the associators. In addition short-cuts or complete navigational information can be provided to the user in a transparent way. By limiting the use of such associators to dedicated monitoring algorithms where the comparison between raw/reconstructed data and MC truth is done, one



could ensure that the reconstruction and analysis code treat simulated and real data in an identical way.

Associators must implement a common interface called `IAssociator`. An `Associator` base class providing at the same time common functionality and some facilities to help in the implementation of concrete Associators is provided. A prototype version of these classes is provided in the current release of Gaudi.

12.4.1.1 The `IAssociator` Interface

As already mentioned Associators must implement the `IAssociator` interface.

In order for Associators to be retrieved from the `ToolSvc` only via the `IAssociator` interface, the interface itself inherits from the `IAlgTool` interface. While the implementation of the `IAlgTool` interface is done in the `AlgTool` base class, the implementation of the `IAssociator` interface is the full responsibility of concrete associators.

The four methods of the `IAssociator` interface that a concrete Associator must implement are show in Listing 12.8

Listing 12.8 Methods of the `IAssociator` Interface that must be implemented by concrete associators

```

1: virtual StatusCode i_retrieveDirect( ContainedObject* objFrom,
    ContainedObject*& objTo,
    const CLID idFrom,
    const CLID idTo ) = 0;
2: virtual StatusCode i_retrieveDirect( ContainedObject* objFrom,
    ContainedObject*& objTo,
    std::vector<ContainedObject*>& vObjTo,
    const CLID idFrom,
    const CLID idTo ) = 0;
3: virtual StatusCode i_retrieveInverse( ContainedObject* objFrom,
    ContainedObject*& objTo,
    const CLID idFrom,
    const CLID idTo ) = 0;
4: virtual StatusCode i_retrieveInverse( ContainedObject* objFrom,
    ContainedObject*& objTo,
    std::vector<ContainedObject*>& vObjTo,
    const CLID idFrom,
    const CLID idTo ) = 0;

```

Two `i_retrieveDirect` methods must be implemented for retrieving associated classes following the same direction as the links in the data: for example from reconstructed particles to Monte Carlo particles. The first parameter is a pointer to the object for which the associated Monte Carlo quantity(ies) is requested. The second parameter, the discriminating signature between the two methods, is one or a vector of pointers to the associated Monte Carlo objects of the type requested. Some reconstructed quantities will have only one possible Monte Carlo associated object of a certain type, some will have many, others will have many out of which a "best" associated object can be extracted. If one of the two methods is not valid for a concrete associator, such method must return a failure. The third and fourth parameters are the class IDs of the objects for which the association is requested. This allows to verify at run time if the objects' types are those the concrete associator has been implemented for.



The two `i_retrieveInverse` methods are complementary and are for retrieving the association between the same two classes but in the opposite direction to that of the links in the data: for example from Monte Carlo particles to reconstructed particles. The different name is intended to alert the user that navigation in this direction may be a costly operation

Four corresponding template methods are implemented in `IAssociator` to facilitate the use of Associators by Algorithms (see Listing 12.9). Using these methods the component retrieving a tool avoids some explicit dynamic-casting as well as the setting of class IDs. An example of how to use such methods is described in section 12.4.1.3.

Listing 12.9 Template methods of the `IAssociator` interface

```

1: template <class T1, class T2>
    StatusCode retrieveDirect( T1* from, T2*& to ) {...}
2: template <class T1>
    StatusCode retrieveDirect( T1* from,
        std::vector<ContainedObject*>& objVTo,
        const CLID idTo ) {...}
3: template <class T1, class T2>
    StatusCode retrieveInverse( T1* from, T2*& to ) {...}
4: template <class T1>
    StatusCode retrieveInverse( T1* from,
        std::vector<ContainedObject*>& objVTo,
        const CLID idTo ) {...}

```

12.4.1.2 The `Associator` base class

An `associator` is a type of `AlgTool`, so the `Associator` base class inherits from the `AlgTool` base class. Thus, Associators can be created and managed as `AlgTools` by the `ToolSvc`. Since all the methods of the `AlgTool` base class (as described in section 12.2.2.1) are available in the `Associator` base class, only the additional functionality is described here.

Access to Event Data Service - An `eventSvc()` method is provided to access the Event Data Service since most concrete associators will need to access data, in particular if accessing navigational short-cuts.

Associator Properties - Two properties are declared in the constructor and can be set in the `jobOptions`: "FollowLinks" and "DataLocation". They are respectively a `bool` with initial value `true` and a `std::string` with initial value set to ". ". The first is foreseen to be used by an associator when it is possible to either follow links between classes or retrieve navigational short cuts from the data. A user can choose to set either behaviour at run time. The second property contains the location in the data where the stored navigational information is located. Currently it must be set via the `jobOptions` when necessary, as shown in Listing 12.10 for a particular implementation provided in the `Associator` example. Two corresponding methods are provided for using the information from these properties: `followLinks()` and `whichTable()`.

Inverse Association - Retrieving information in the direction opposite to that of the links in the data is in general a time consuming operation, that implies checking all the direct associations to access the inverse relation for a specified object. For this reason Associators should keep a local copy of the inverse associations after receiving the first request for an event. A few methods are provided to facilitate the work of Associators in this case. The



methods `inverseExist()` and `setInverseFlag(bool)` help in keeping track of the status of the locally kept inverse information. The method `buildInverse()` has to be overridden by concrete associators since they choose in which form to keep the information and should be called by the associator when receiving the first request during the processing of an event.

Locally kept information - When a new event is processed, the associator needs to reset its status to the same conditions as those after having been created. In order to be notified of such an incident happening the `Associator` base class implements the `ILiStener` interface and, in the constructor, registers itself with the `IncidentService` (see section 11.9 for details of the `IncidentService`). The associator's `FlushCache()` method is called in the implementation of the `ILiStener` interface in the `Associator` base class. This method must be overridden by concrete associators wanting to do a meaningful reset of their initial status.

12.4.1.3 A concrete example

In this section we look at an example implementation of a specific associator. The code is taken from the `LHCbAssociator` example, but the points illustrated should be clear even without a knowledge of the `LHCb` data model.

The `AxPart2MCParticleAsct` provides association between physics analysis particles (`AxPartCandidate`) and the corresponding Monte Carlo particles (`MCParticle`). The direct navigational information is stored in the persistent data as short-cuts, and is retrieved in the form of a `SmartRefTable` in the `TransientEventStore`. This choice is specific to `AxPart2MCParticleAsct`, any associator can use internally a different navigational mechanism. The location in the `EventStore` where the navigational information can be found is set in the job options via the "DataLocation" property, as shown in Listing 12.10.

Listing 12.10 Example of setting properties for an associator via `jobOptions`

```
ToolSvc.AxPart2MCParticleAsct.DataLocation = "/Event/Anal/AxPart2MCParticle";
```

In the current `LHCb` data model only a single `MCParticle` can be associated to one `AxPartCandidate` and vice-versa only one or no `AxPartCandidate` can be associated to one `MCParticle`. For this reason only the `i_retrieveDirect` and `i_retrieveInverse` methods providing one-to-one association are meaningful. Both methods verify that the objects passed are of the correct type before attempting to retrieve the information, as shown in Listing 12.11. When no association is found, a `StatusCode::FAILURE` is returned.

Listing 12.11 Checking if objects to be associated are of the correct type

```
1: if ( idFrom != AxPartCandidate::classID() ) {
2:   objTo = 0;
3:   return StatusCode::FAILURE;
4: }
5: if ( idTo != MCParticle::classID() ) {
6:   objTo = 0;
7:   return StatusCode::FAILURE;
8: }
```

The `i_retrieveInverse` method providing the one-to-many association returns a failure, while a fake implementation of the one-to-many `i_retrieveDirect` method is



implemented in the example, to show how an Algorithm can use such a method. In the `AxPart2MCParticleAsct` example the inverse table is kept locally and both the `buildInverse()` and `FlushCache()` methods are overridden. In the example the choice has been made to implement an additional method `buildDirect()` to retrieve the direct navigational information on a first request per event basis.

Listing 12.12 shows how a monitoring Algorithm can get an associator from the `ToolSvc` and use it to retrieve associated objects through the template interfaces.

Listing 12.12 Extracted code from the `AscExampleAlgorithm`

```
1: #include "GaudiTools/IAsociator.h"
2: // Example of retrieving an associator
3: IAsociator
4: StatusCode sc = toolSvc()->retrieveTool("AxPart2MCParticleAsct",
                                     m_pAsct);
5: if ( sc.isFailure() ) {
6:   log << MSG::FATAL << "Unable to create Associator tool" << endl;
7:   return sc;
8: }
9: // Example of retrieving inverse one-to-one information from an
10: // associator
11: SmartDataPtr<MCParticleVector> vmcparts (evt, "/MC/MCParticles");
12: for( MCParticleVector::iterator itm = vmcparts->begin();
      vmcparts->end() != itm; itm++) {
13:   AxPartCandidate* mptr = 0;
14:   StatusCode sc = m_pAsct->retrieveInverse( *itm, mptr );
15:   if ( sc.isSuccess() ) {...}
16:   else {...}
17: }
18: // Example of retrieving direct one-to-many information from an
19: // associator
20: SmartDataPtr<AxPartCandidateVector> candidates(evt,
                                                "/Anal/AxPartCandidates");
21: std::vector<ContainedObject*> pptry;
22: AxPartCandidate* itp = *(candidates->begin());
23: StatusCode sa =
24:   m_pAsct->retrieveDirect(itp, pptry, MCParticle::classID());
25: else {
26:   for (std::vector<ContainedObject*>::iterator it = pptry.begin();
        pptry.end() != it; it++) {
27:     MCParticle* imc = dynamic_cast<MCParticle*>( *it );
28:   }
29: }
```





Chapter 13 Converters

13.1 Overview

Consider a small piece of detector, a silicon wafer for example. This "object" will appear in many contexts: it may be drawn in an event display, it may be traversed by particles in a Geant4 simulation, its position and orientation may be stored in a database, the layout of its strips may be queried in an analysis program, etc. All of these uses or views of the silicon wafer will require code.

One of the key issues in the design of the framework was how to encompass the need for these different views within Gaudi. In this chapter we outline the design adopted for the framework and look at how the conversion process works. This is followed by sections which deal with the technicalities of writing converters for reading SICB data and writing to ROOT files.

13.2 Persistency converters

Gaudi gives the possibility to read event data from either ZEBRA (SICB) or ROOT files, and to write data back to ROOT files. The use of ODBC compliant databases is also possible.

Figure 13.1 is a schematic illustrating how converters fit into the transient-persistent translation of event data. We will not discuss in detail how the transient data store (e.g. the event data service) or the persistency service work, but simply look at the flow of data in order to understand how converters are used. An introduction to the persistency mechanism of Gaudi can be found in reference [12].

One of the issues considered when designing the Gaudi framework was the capability for users to "create their own data types and save objects of those types along with references to already existing objects". A related issue was the possibility of having links between objects






Figure 13.1 Persistence conversion services in Gaudi

which reside in different stores (i.e. files and databases) and even between objects in different types of store.

Figure 13.1 shows that data may be read from an ODBC database and /or ROOT files into the transient event data store and that data may be written out again to the same media. It is the job of the persistency service to orchestrate this transfer of data between memory and disk.

The figure shows two "slave" services: the ODBC conversion service and the ROOT I/O service. These services are responsible for managing the conversion of objects between their transient and persistent representations. Each one has a number of converter objects which are actually responsible for the conversion itself. As illustrated by the figure a particular converter object converts between the transient representation and one other form, here either MS Access or ROOT. The mechanism is identical when reading data from a SICB data file, via the SicbCnv conversion service.

13.3 Collaborators in the conversion process

In general the conversion process occurs between the transient representation of an object and some other representation. In this chapter we will be using persistent forms, but it should be borne in mind that this could be any other "transient" form such as those required for visualisation or those which serve as input into other packages (e.g. Geant4).

Figure 13.2 shows the interfaces (classes whose name begins with "I") which must be implemented in order for the conversion process to function.

The conversion process is essentially a collaboration between the following types:

- IConversionSvc
- IConverter

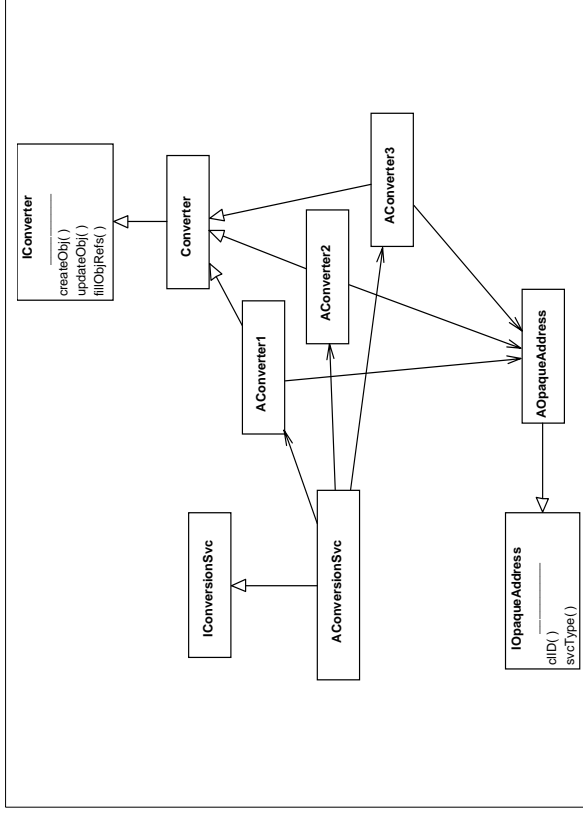


Figure 13.2 The classes (and interfaces) collaborating in the conversion process.

- IOpaqueAddress

For each persistent technology, or "non-transient" representation, a specific conversion service is required. This is illustrated in the figure by the class `AConversionSvc` which implements the `IConversionSvc` interface.

A given conversion service will have at its disposal a set of converters. These converters are both type and technology specific. In other words a converter knows how to convert a single transient type (e.g. `MuonHit`) into a single persistent type (e.g. `RootMuonHit`) and vice versa. Specific converters implement the `IConverter` interface, possibly by extending an existing converter base class.

A third collaborator in this process are the opaque address objects. A concrete opaque address class must implement the `IOpaqueAddress` interface. This interface allows the address to be passed around between the transient data service, the persistency service, and the conversion services without any of them being able to actually decode the address. Opaque address objects are also technology specific. The internals of an `OpaqueAddress` object are different from those of a `RootAddress` object.

Only the converters themselves know how to decode an opaque address. In other words only converters are permitted to invoke those methods of an opaque address object which do not form a part of the `IOpaqueAddress` interface.



Converter objects must be "registered" with the conversion service in order to be usable. For the "standard" converters this will be done automatically. For user defined converters (for user defined types) this registration must be done at initialisation time (see Section 6.10).

13.4 The conversion process

As an example (see Figure 13.3) we consider a request from the event data service to the persistence service for an object to be loaded from a data file.

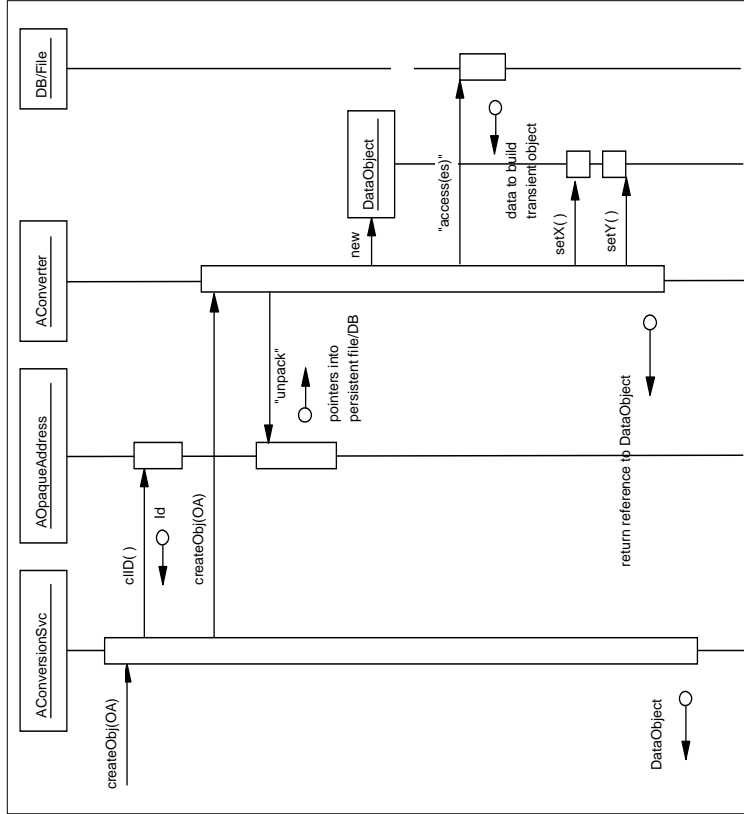


Figure 13.3 A trace of the creation of a new transient object.

As we saw previously, the persistence service has one conversion service slave for each persistent technology in use. The persistence service receives the request in the form of an opaque address object. The svcType () method of the IOpaqueAddress interface is

invoked to decide which conversion service the request should be passed onto. This returns a "technology identifier" which allows the persistence service to choose a conversion service.

The request to load an object (or objects) is then passed onto a specific conversion service. This service then invokes another method of the IOpaqueAddress interface, cid () , in order to decide which converter will actually perform the conversion. The opaque address is then passed onto the concrete converter who knows how to decode it and create the appropriate transient object.

The converter is specific to a specific type, thus it may immediately create an object of that type with the new operator. The converter must now "unpack" the opaque address, i.e. make use of accessor methods specific to the address type in order to get the necessary information from the persistent store.

For example, a SICB converter might get the name of a bank from the address and use that to locate the required information in the ZEBRA common block. On the other hand a ROOT converter may extract a file name, the names of a ROOT Tree and an index from the address and use these to load an object from a ROOT file. The converter would then use the accessor methods of this "persistent" object in order to extract the information necessary to build the transient object.

We can see that the detailed steps performed within a converter depend very much on the nature of the non-transient data and (to a lesser extent) on the type of the object being built.

If all transient objects were independent, i.e. if there were no references between objects then the job would be finished. However in general objects in the transient store do contain references to other objects.

These references can be of two kinds:

- i. "Macroscopic" references appear as separate "leaves" in the data store. They have to be registered with a separate opaque address structure in the data directory of the object being converted. This must be done after the object was registered in the data store in the method fillObjRefs () .
- ii. Internal references must be handled differently. There are two possibilities for resolving internal references:

1. Load on demand. If the object the reference points to should only be loaded when accessed, the pointer must no longer be a raw C++ pointer, but rather a smart pointer object containing itself the information for later resolution of the reference. This is the preferred solution for references to objects within the same data store (e.g. references from Monte-Carlo tracks to Monte-Carlo vertices) and is generated by the Object Description Tools when a relation tag is found in the XML class description (see Section 6.9).
2. Filling of raw C++ pointers. This is only necessary if the object points to an object in another store, e.g. the detector data store, and should be avoided in classes foreseen to be made persistent. To resolve the reference a converter has to retrieve the other object and set the raw pointer. These references should be set in the fillObjRefs () method. This of course is more complicated, because it must be ensured that both objects are present at the time the reference is accessed (i.e. when the pointer is actually used).



13.5 Converter implementation - general considerations

After covering the ground work in the preceding sections, let us look exactly what needs to be implemented in a specific converter class. The starting point is the `Converter` base class from which a user converter should be derived.

Listing 13.1 An example converter class

```
// Converter for class UDO.
extern const CLID& CLID_UDO;
extern unsigned char OBJV_StorageType;

static ChvFactory<UDOChv> s_factory;
const IChvFactory& UDOChvFactory = s_factory;

class UDOChv : public Converter {
public:
    UDOChv(ISvLocator* svcLoc) :
        Converter(Objectivity_StorageType, CLID_UDO, svcLoc) {}

    createRep(DataObject* po, IOpaqueAddress* a); // transient->persistent
    createObj(IOpaqueAddress* pa, DataObject* po); // persistent->transient

    fillObjRefs( ... ); // transient->persistent
    fillRepRefs( ... ); // persistent->transient
}
```

The converter shown in Listing 13.1 is responsible for the conversion of UDO type objects into objects that may be stored into an Objectivity database and vice-versa. The `UDOChv` constructor calls the `Converter` base class constructor with arguments which contain this information. These are the values `CLID_UDO`, defined in the UDO class, and `Objectivity_StorageType` which is also defined elsewhere. The first two external statements simply state that these two identifiers are defined elsewhere.

All of the “book-keeping” can now be done by the `Converter` base class. It only remains to fill in the guts of the converter. If objects of type `UDO` have no links to other objects, then it suffices to implement the methods `createRep()` for conversion from the transient form (to `Objectivity` in this case) and `createObj()` for the conversion to the transient form.

If the object contains links to other objects then it is also necessary to implement the methods `fillRepRefs()` and `fillObjRefs()`.

13.6 Storing Data using the ROOT I/O Engine

One possibility for storing data is to use the `ROOT I/O` engine to write `ROOT` files. Although `ROOT` by itself is not an object oriented database, with modest effort a structure can be built on top to allow the Converters to emulate this behaviour. In particular, the issue of object linking had to be solved in order to resolve pointers in the transient world.



The concept of `ROOT` supporting paged tuples called trees and branches is adequate for storing bulk event data. Trees split into one or several branches containing individual leaves with data.

The data structure within the Gaudi data store is also tree like. In the transient world Gaudi objects are sub-class instances of the “`DataObject`”. The `DataObject` offers some basic functionality like the implicit data directory which allows e.g. to browse a data store. This tree structure will be mapped to a flat structure in the `ROOT` file resulting in a separate tree representing each leaf of the data store. Each data tree contains a single branch containing objects of the same type. The Gaudi tree is split up into individual `ROOT` trees in order to give easy access to individual items represented in the transient model without the need of loading complete events from the root file i.e. to allow for selective data retrieval. The feature of `ROOT` supporting selective data reading using split trees did not seem too attractive since, generally, complete nodes in the transient store should be made available in one go.

However, `ROOT` expects “`ROOT`” objects, they must inherit from `TObject`. Therefore the objects from the transient store have to be converted to objects understandable by `ROOT`.

The following sections are an introduction to the machinery provided by the Gaudi framework to achieve the migration of transient objects to persistent objects. The `ROOT` specific aspects are not discussed here; the `ROOT I/O` engine is documented on the `ROOT` web site <http://root.cern.ch>. Note that Gaudi only uses the `I/O` engine, not all `ROOT` classes are available. Within Gaudi the `ROOT I/O` engine is implemented in the `GaudiROOTDb` package.

13.7 The Conversion from Transient Objects to ROOT Objects

As for any conversion of data from one representation to another within the Gaudi framework, conversion to/from `ROOT` objects is based on Converters. The support of a “generic” Converter accesses pre-defined entry points in each object. The transient object converts itself to an abstract byte stream. However, for specialized objects specific converters must be built.

Whenever objects must change their representation within Gaudi, data converters are involved. For the `ROOT` case, the converters must have some knowledge of `ROOT` internals and of the service finally used to migrate `ROOT` objects (`->TObject`) to a file. They must be able to translate the functionality of the `DataObject` component to/from the `ROOT` storage. Within `ROOT` itself the object is stored as a `Binary Large Object (BLOB)`.

The generic data conversion mechanism relies on two functionalities, which must be present:

- When writing or reading objects, the object’s data must be “serializable”. The corresponding persistent type is of a generic type, the data are stored as a machine independent byte stream. This method is implemented automatically if the class is described using the Gaudi Object Description tools (described in Section 6.7 on page 53). Some examples of classes implemented automatically in this way can be found in the `Event/LHCbEvent` package.



- When reading objects, an empty object must be created before any de-serialization can take place. The constructor must be called. This functionality does not imply any knowledge of the conversion mechanism itself and hence can be encapsulated into an object factory simply returning a DataObject. These data object factories are distinguished within Gaudi through the persistent data type information, the class ID. For this reason the class ID of objects, which are written must only depend on the object type, i.e. every class needs it's own class ID. The instantiation of the appropriate factory is done by a macro. Please see the RootIO example for details how to instantiate the factory.

13.8 Storing Data using other I/O Engines

Once objects are stored as BLOBs, it is possible to adopt any storage technology supporting this datatype. This is the case not only for ROOT, but also for

- Objectivity/DB
- most relational databases, which support an ODBC interface like
 - Microsoft Access,
 - Microsoft SQL Server,
 - MySQL,
 - ORACLE and others.

Note that although storing objects using these technologies is possible, there is currently no experiment wide policy on how to use Objectivity or other client server based technologies. For this reason only the example to store data using Microsoft Access is described in the example `ROOTIOExample1.c`. All other technologies are currently not supported. If you desperately want to use SQL Server, MySQL or Objectivity, please contact Markus Frank (Markus.Frank@cern.ch).



Chapter 14

Scripting and Interactivity

14.1 Overview

A scripting capability has been added to the Gaudi framework. The current functionality is likely to change rapidly, so users should check with the latest release notes for changes or new functionality that might not be documented here.

In keeping with the design philosophy of the Gaudi architecture, scripting is defined by an abstract scripting service interface, with the possibility of there being several different implementations. The first implementation available is based on Python, a public-domain programming language. Python is ideal both as a scripting interface for modern systems, and as a standalone rapid-development language. Its object-oriented nature mixes well with frameworks written in C++.

The Python scripting language will not be described in detail here. There are many Python books available, among them we recommend:

- Learning Python, by M. Lutz & D. Ascher, O'Reilly, 1999
- Programming Python (2nd ed.), by M. Lutz, O'Reilly, 2001

14.2 How to enable Python scripting

Three different mechanisms are available for enabling Python scripting.

1. Replace the job options text file by a Python script that is specified on the command line.
2. Use a job options text file which hands control over to the Python shell once the initial configuration has been established.
3. Load and start a Gaudi application from a Python shell.



14.2.1 Using a Python script for configuration and control

One can avoid using a job options text file for configuration by specifying a Python script as a command line argument, as shown in Listing 14.1.

Listing 14.1 Using a Python script for job configuration

```
myJob MyPythonScript.py [1]
```

Note:

1. The file extension `.py` is used to identify the job options file as a Python script. All other extensions are assumed to be job options text files.

This approach may be used in two modes. The first uses such a script to establish the configuration, but results in the job being left at the Python shell prompt. This supports interactive sessions. The second specifies a complete configuration and control sequence and thus supports a batch style of processing. The particular mode is controlled by the presence or absence of Gaudi-specific Python commands described in Section 14.3.6.

14.2.2 Using a text JobOptions file and giving control to the Python interactive shell

Python scripting is enabled when using a job options text file for job configuration by adding the lines shown in Listing 14.2 to the job options file.

Listing 14.2 Job Options text file entries to enable Python scripting

```
ApplicationMgr.Dlls += { "GaudiPython" }; [1]  
ApplicationMgr.Runnable = "PythonScriptingSvc"; [2]  
PythonScriptingSvc.StartupScript = "../options/AnalysisTest.py"; [3]
```

Notes:

1. This entry specifies the component library that implements Python scripting.
2. This entry specifies that the Python scripting should take the control (runable) of the application.
3. Optional startup python script.

Once the initial configuration has been established by the job options text file, control will be handed over to the Python shell when the startup script, if specified, will be executed. The user can then issue interactive commands.

14.2.3 Starting a Gaudi application from the Python shell

It is also possible to bootstrap a Gaudi application directly from a Python shell. The user needs to import the Python extension module called `gaudimodule`, which allows the interaction with Gaudi from Python. Listing 14.3 shows a small Python program that instantiates a Gaudi



application, configure it and runs for a number of events. This program would work from the Python shell a long as the environment (`LD_LIBRARY_PATH/PATH`) is properly set up.

Listing 14.3 Example of a Python program that executes a Gaudi program

```
from gaudi.module import *  
theApp = AppMgr ()  
theApp.JobOptionsType = 'NONE'  
theApp.EvtSel = 'NONE'  
theApp.config ()  
theApp.Dlls = ['GaudiAlg']  
mySeq = theApp.algorithm ('Sequencer/MySeq')  
mySeq.Members = ['EventCounter/Count1', 'EventCounter/Count2']  
theApp.topAlg = ['MySeq']  
theApp.initialize ()  
theApp.run (10)  
theApp.exit ()
```

14.3 Current functionality

The current functionality is limited to the following capabilities:

1. The ability to set and get basic properties for all framework components (Algorithms, Services, Auditors etc.) and the main `ApplicationMgr` that controls the application. Arrays of simple properties are mapped into Python Lists.
2. The ability to interact with the transient data stores. Browsing store contents, registering, unregistering and retrieving objects, getting and setting object data members (with the help of the `IntrospectionSvc`) and limited method invocation.
3. The ability to interact with the Histograms (1D and 2D) in the transient store. This includes booking, filling, dumping contents, etc.
4. The ability to add new services and component libraries and access their capabilities.
5. The ability to control the execution of the application by adding Algorithms into the list of top level Algorithms, executing single events or a set of events, executing single Algorithms, etc.
6. The ability to define Python Algorithms that will be managed and scheduled as normal Gaudi Algorithms.

14.3.1 Property manipulation

An example of the use of the scripting language to display and set component properties is shown in Listing 14.4:



Listing 14.4 Property manipulation from the Python interactive shell

```

>>> theApp
<AppMgr object at 00AD22B8>
>>> theApp.ExtSvc
[ 'IntrospectionSvc', 'ParticlePropertySvc' ]
>>> theApp.ExtSvc = theApp.ExtSvc + [ 'AnotherSvc' ]
>>> theApp.ExtSvc
[ 'IntrospectionSvc', 'ParticlePropertySvc', 'AnotherSvc' ]
>>> theApp.EventMax = 100
>>> theApp.properties()
{ 'EvtMax': 100, 'JobOptionsType': 'NONE', 'TopAlg': [ 'seql', 'PhysAnaAlg' ],
  'Go': 0, 'Exit': 0, 'Dills': [ 'GaudiAlg', 'GaudiIntrospection' ],
  'JobOptionsPath': '\\home\\test-py', 'OutputStream': [], 'OutputLevel': 3,
  'EventLoop': 'EventLoopMgr', 'HistogramPersistence': 'NONE', 'EvtSel':
  'NONE', 'ExtSvc': [ 'IntrospectionSvc', 'ParticlePropertySvc', 'AnotherSvc' ],
  'Runnable': 'PythonScriptingSvc' }
>>> theApp.algorithms()
[ 'seql', 'WriteAlg', 'PhysAnaAlg' ]
>>> alg = Algorithm('WriteAlg')
>>> alg.properties()
{ 'ErrorCount': 0, 'OutputLevel': 0, 'AuditExecute': 1, 'AuditInitialize': 0,
  'Enable': 1, 'AuditFinalize': 0, 'ErrorMax': 1 }

```

Notes:

1. The ">>>" is the Python shell prompt. Typing the name of a variable, Python prints its value in textual form.
2. The variable theApp is always defined and it represents the ApplicationMgr. An alias g has also been defined.
3. The name of the property is used as a data member in Python. It returns the correct type directly.
4. You can use the properties in normal Python expressions
5. The list of all properties (as a Python Dictionary) of a component can be obtained with the method properties().
6. The list of algorithms can be obtained with the method algorithms().
7. To access an Algorithm by name (creating it if it does not exist) the constructor Algorithm() is used. Similarly for services with Service().

14.3.2 Creating Algorithms and Services

It is possible to create new Algorithms or Services as a result of a scripting command. Examples of this are shown in Listing 14.5.

If the specified Algorithm or Service does not exist, it is created. Its properties can immediately be accessed for read and write. They will be initialized when the application will start processing events.



Listing 14.5 Examples of Python commands that create new Algorithms or Services

```

>>> myseq = theApp.algorithm('Sequencer/MySeq')
>>> myseq.members = ['HelloWorld', 'WriteAlg']
MySeq      INFO HelloWorld doesn't exist - created and
           INFO HelloWorld already exists - appended to member list
>>> theApp.TopAlg = ['MySeq']
>>> g.run(1)
HelloWorld INFO initializing....
HelloWorld INFO INFO executing....
WriteAlg   INFO INFO Generated event 5
StatusCode::SUCCESS

```

14.3.3 Interaction with Transient Data Stores

It is possible to get the list of data stores and to interact with them. The following commands are available:

- ```

theApp.datastores()
Returns the list of all available data stores (all services that implement the
IDataProviderSvc interface).

theApp.datastore(name)
Returns a generic data store.

theApp.histogram()
Returns the standard histogram data service ("HistogramDataSvc")

theApp.eventSvc()
Returns the standard event data service ("EventDataSvc")

theApp.detSvc()
Returns the standard detector data service ("DetectorDataSvc")

theApp.ntuplesSvc()
Returns the standard ntuple data service ("NTupleSvc")

datastore.dump()
Prints the contents of the transient data store (names and types)

datastore.clear()
Clears the contents of the transient data store

datastore.object(name), datastore[name], datastore.retrieve(name)
Retrieves the named object from the transient data store. If the IntrospectionSvc
is loaded and the dictionaries are available for the requested object then it creates
an object that can be introspected.

datastore.register(name, obj), datastore.unregister(name)
Registers and unregisters an object to/from the data store

```

1. See Section 11.11 on page 135



```
datastore.clear(), datastore.clear(name)
Clears the whole store or a sub-tree.
```

A complete example of the capabilities on the interaction with the event transient store is shown in Listing 14.6.

**Listing 14.6** Example of interaction with the event transient store taken from LHCB

```
evt = theApp.evtSvc()
evt.dump()
parts = evt['/Event/MC/MCParticles']
for p in parts:
 print P.particleID.id
```

### 14.3.4 Interaction with Histograms

The data store commands described in the previous section are also available for the histogram data store. The following histogram specific commands are available in addition:

```
histosvc.histo(name), histosvc[name]
Retrieves the histogram from the histogram data store. It returns either an ID or 2D histogram.
histosvc.book(id, title, xbin, xmin, xmax [,ybin, ymin, ymax])
Books ID or 2D histogram and registers it in the histogram store.
histo.title(), dim(), mean(), rms(), maxbin(), minbin()
Returns the title, dimensions, mean, rms, maximum bin contents, minimum bin contents of the histogram
histo.fill(x [,y,w]), histo.reset()
Fills 1D or 2D histogram, resets the contents
histo.heights(), entries(), errors(), edges()
Returns in a Python list the heights, entries, errors and edges of the 1D histogram
histo.projections()
Returns X and Y projections (tuple) of a 2D histogram
```

### 14.3.5 Interaction with Data Objects

If the dictionaries of the classes have been loaded by the object introspection service, then it is possible to browse and interact with the data objects in the transient store. Loading the dictionaries is done by configuring the application as it is shown in Listing 14.7

**Listing 14.7** Configuring the Gaudi Introspection Service using the JobOptions text file

```
ApplicationMgr.Dlls += { "GaudiIntrospection"};
ApplicationMgr.ExtSvc += { "IntrospectionSvc" };
IntrospectionSvc.Dictionaryes = { "LHCBEventDict", "PhysEventDict" };
```



The following commands are currently available:

```
obj.<attribute>
Returns the value of the attribute for the object. If the attribute is of a complex type it returns a reference to it, such that it can be browsed recursively.
obj.<attribute> = value
Sets the value of attribute for the object if this is a simple type.
obj.<method>()
Invokes the class method. This is currently only available for methods without arguments.
obj.values()
Returns a Python dictionary with all the attributes and their values.
obj.names()
Returns a Python list with all the available attribute names.
obj.desc()
Prints the description of the class of the object.
obj.type()
Returns the object type (C++ class)
```

### 14.3.6 Controlling job execution

There exist a few commands to control the job execution interactively:

```
theApp.run(number)
The control is returned from the Python shell to the Gaudi environment with this command. The argument is the number (-1 for infinite) of events to be processed, after which control will be returned to the Python shell.
theApp.exit()
Typing Ctrl-D (or Ctrl-Z in Windows) at the Python shell prompt will cause an orderly termination of the job. Alternatively, this command will also cause an orderly application termination.
```

## 14.4 Physics Analysis Environment

It is possible to declare an Algorithm in Python that can be declared to the list of top level algorithms to be executed for each event by the ApplicationMgr. This can be useful for constructing an interactive physics analysis environment. An example is shown in Listing 14.8.



Listing 14.8 Example Analysis

```
-- User analysis algorithm class
class PhysAnalAlg(PyAlgorithm):
 def initialize(self):
 global hi
 hi = his.book('hi', 'Histogram Test', 10, 0., 10.)
 print '....User Analysis Initialized'
 return 1
 def finalize(self):
 print 'Finalizing User Analysis...'
 return 1
 def execute(self):
 cand = evt['Anal/AxPartCandidates']
 print 'Found '+ `len(cand)` + ' candidates'
 for c in cand :
 hi.fill(c.momentum)
 return 1

-- Initialization and Configuration
his = theApp.histSvc()
evt = theApp.evtSvc()
det = theApp.detSvc()
pdt = PartSvc()
physalg = PhysAnalAlg(theApp, 'PhysAnalAlg')
theApp.topAlg = theApp.topAlg + ['PhysAnalAlg']
```

- Notes:
1. The analysis algorithm must inherit from the class PyAlgorithm
  2. Useful variables to avoid long typing
  3. An instance of the new class needs to be instantiated and declared in the list of top level algorithms.





## Chapter 15

# Visualization Facilities

### 15.1 Overview

In this chapter we describe how visualization facilities are provided to the applications based on the Gaudi framework. We present how we interface the physics event data objects, detector components or statistical objects to their graphical representation. One example of an application that uses the visualization services is an event display program. With this program we display graphically the event data from files or being acquired by the data acquisition. Another example could be an interactive analysis program that combines in the same application histogramming or manipulation of statistical entities, event display, and full interactive control by the end user of the data objects and algorithms of the application.

In the current release we have implemented the mechanism of converting event and detector objects into their graphical representation and built an event display application (Panoramix). This application can be used to help in the development of physics algorithms (e.g. pattern recognition) or in the verification of the detector geometry.

### 15.2 The data visualization model

The Gaudi architecture envisaged implementing data visualization using a similar pattern to data persistency. We do not want to implement visualization methods in each data object. In other words, we do not want to tell an object to "draw" itself. Instead we would implement converters as separate entities that are able to create specific graphical representations for each type of data object and for each graphical package that we would like to use. In that way, as for the persistency case, we decouple the definition and behaviour of the data objects from the various technologies for graphics. We could configure at run time to have 2D or 3D graphics depending on the needs of the end-user at that moment.

Figure 15.1 illustrates the components that need to be included in an application to make it capable of visualizing data objects. The interactive user interface is a Service which allows the



end-user to interact with all the components of the application. The user could select which objects to display, which algorithms to run, what properties of which algorithm to inspect and modify, etc. This interaction can be implemented using a graphical user interface or by using a scripting language.

The User interface service is also in charge of managing one or more GUI windows where views of the graphical representations are going to be displayed.

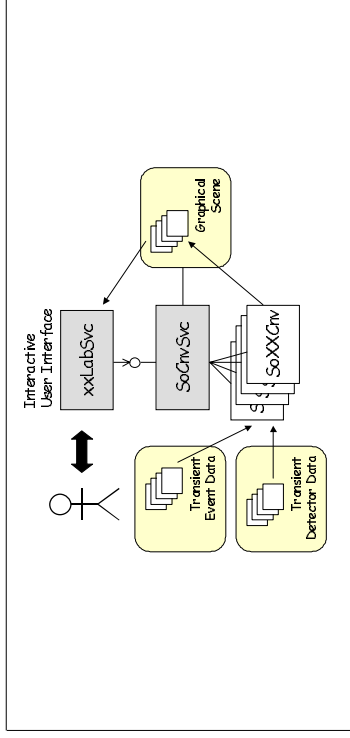


Figure 15.1 Components for visualization

The other main component is a Conversion Service that handles the conversion of objects into their graphical representation. This service requires the help of a number of specialized converters, one for each type of data object that needs to be graphically displayed. The transient store of graphical representations is shared by the conversion service, together with the converters, and the user interface component. The form of this transient store depends on the choice of graphics package. Typically it is the user interface component that would trigger the conversion service to start the conversion of a number of objects (next event), but this service can also be triggered by any algorithm that would like to display some objects.

## 15.3 VisSys - the Gaudi visualisation services

### 15.3.1 Design layout

The Gaudi Visualisation service (OnXSvc) uses the OnX(ML) package to drive interactivity. From an XML description of the graphical user interface (GUI), OnX creates an application GUI by using various "toolkits" like Motif, Win32, GTK+. It permits also to handle various scripting languages (currently tcl, CINT, KUIP) in the GUI XML files, to describe the behaviour of the GUI pieces.

The graphics are handled by Open Inventor. The Inventor "viewers", being part of the GUI, are created by OnX from a placement in the application GUI XML files.

Connection to Gaudi is done through the wrapping of Gaudi C++ code to scripting languages (Panoramix package).

Inventor modeling of the LHCB detector and event data uses Gaudi Inventor converters (the So<xxx>Cnv classes, SolHcb package) to produce Inventor scene graphs.

The SoDetectorCnv and SolVolumeCnv build the Inventor scene graph for the detector. Others (such as SoMCParticleCnv) build Inventor scene graphs for event data.

### 15.3.2 Writing graphic converters

The role of each converter So<xxx>Cnv is to produce an Open Inventor node that represents the object. The following fragment of code shows how this is done for the geometry of a detector element. The code has been simplified to be more illustrative. The 3D graphical objects that are created are standard OpenInventor objects (in bold).

Listing 15.1 Fragment of SoDetectorElementCnv

```

30: StatusCode SoDetElemCnv::createRep(DataObject* aObject, IOpaqueAddress&&)
31: {
32: DetectorElement* de = dynamic_cast<DetectorElement*>(aObject);
33: IIVolume* lv = de->geometry()->lvolume();
34: SolidBox* box = dynamic_cast<SolidBox*>(lv->solid()->coverTop());
35:
36: SoSeparator* separator = new SoSeparator;
37: SoDrawStyle* drawStyle = new SoDrawStyle;
38: SoMaterial* material = new SoMaterial;
39: separator->addChild(drawStyle);
40: separator->addChild(material);
41: // set drawing styles
42: drawStyle->style.setValue(SoDrawStyle::LINES);
43: drawStyle->linePattern.setValue(0xFFFF);
44: material->diffuseColor.setValue(SbColor(0.,1.,0.));
45:
46: // Code related to the transformation
47: SoTransform* trans = new SoTransform;
48: ...
49: separator->addChild(trans);
50:
51: SoCube* cube = new SoCube();
52: cube->width = box->xHalfLength() * 2;
53: cube->height = box->yHalfLength() * 2;
54: cube->depth = box->zHalfLength() * 2;
55:
56: separator->addChild(cube);
57: m_pSo->addChild(separator);
58: return StatusCode::SUCCESS;
59: }

```



## 15.4 Panoramix - the LHCb event display

Panoramix is the association of the three packages `Vis/OmXSvc`, `Vis/SolHcb`, `Vis/Panoramix`, and of sub-detector converter packages (`Vis/So<Det>`) to produce an interactive event and geometry display for LHCb. Detailed instructions for its installation and use can be found in the Visualisation pages of the LHCb web, at <http://cern.ch/lhcb-comp/Frameworks/Visualization/>.



## Chapter 16

# Framework packages, interfaces and libraries

---

### 16.1 Overview

It is clearly important to decompose large software systems into hierarchies of smaller and more manageable entities. This decomposition can have important consequences for implementation related issues, such as compile-time and link dependencies, configuration management, etc. A package is the grouping of related components into a cohesive physical entity. A package is also the minimal unit of software release.

In this chapter we describe the Gaudi package structure, and how these packages are implemented in libraries. We also discuss abstract interfaces, which are one of the main design features of Gaudi

### 16.2 Gaudi Package Structure

The Gaudi software is decomposed into the packages shown in Figure 16.1.

At the lower level we find `GaudiKernel`, which is the framework itself, and whose only dependency is on the `GaudiPolicy` package, which contains the various flags defining the CMT [7] configuration management environment needed to build the Gaudi software. At the next level are the packages containing standard framework components (`GaudiSvc`, `GaudiDb`, `GaudiTools`, `GaudiAlg`, `GaudiAud`, `GaudiIntrospection`), which depend on the framework and on widely available foundation libraries such as `CLHEP` and `HTL`. These external libraries are accessed via CMT interface packages which use environment variables defined in the `ExternalLibs` package, which should be tailored to the software installation at a given site. All the above packages are grouped into the `GaudiSys` set of packages which are the minimal set required for a complete Gaudi installation



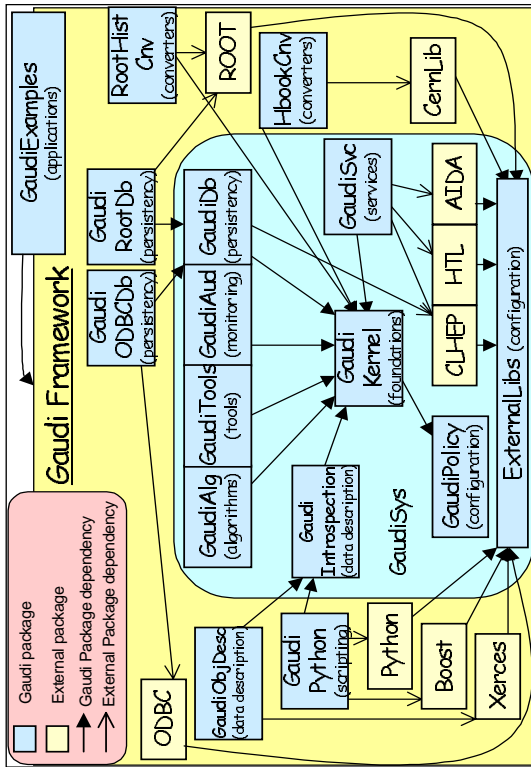


Figure 16.1 Package structure of the Gaudi software

The remaining packages are optional packages which can be used according to the specific technology choices for a given application. In this distribution, there are two specific implementations of the histogram persistence service, based on HBOOK (HbookCnv) and ROOT (RootHistCnv) and two implementations of the event data persistence service (GaudiRootDb and GaudiODBCdb) which understand ROOT and ODBC compliant databases respectively. There is also a scripting service (GaudiPython) depending on the Python scripting language and a data description service (GaudiObjDesc) based on the Xerces XML parser. Finally, at the top level we find the applications (Gaudi.Examp.Les) which depend on Gaudi.Sys and the scripting and persistence services.

In addition to the Gaudi packages, there are a number of LHCb specific packages built on top of Gaudi, as shown in Figure 16.2. DetDesc gives access to the detector description data stored in XML files in the Xml1DDb package; these files can be edited with the XmlEditor. DetEvent is a prototype implementation of a conditions database based on Objectivity. The LHCbEvent package contains the LHCb event data model. The event data is populated from SICB files via converters (SicbCnv package). The SicbCnv converters depend on some of the Fortran packages of SICB (in particular FortIO and FortInclude) and give access to certain SICB facilities, such as the magnetic field. The GiGa package interfaces Gaudi to GEANT4. Finally, there is a large number of examples, each of which has its own package.

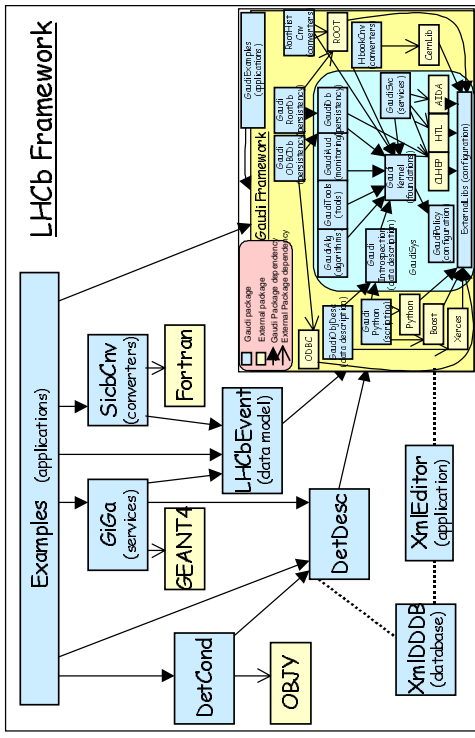


Figure 16.2 Package structure of the LHCb Gaudi software

### 16.2.1 Gaudi Package Layout

Figure 16.3 shows the layout for Gaudi packages. Note that the binaries directories are not in CVS, they are created by CMT when building a package.

### 16.2.2 Packaging Guidelines

Packaging is an important architectural issue for the Gaudi framework, but also for the experiment specific software packages based on Gaudi. Typically, experiment packages consist of:

- Specific event model
- Specific detector description
- Sets of algorithms (digitisation, reconstruction, etc.)

The packaging should be such as to minimise the dependencies between packages, and must absolutely avoid cyclic dependencies. The granularity should not be too small or too big. Care should be taken to identify the external interfaces of packages; if the same interfaces are shared by many packages, they should be promoted to a more basic package that the others would then depend on. It is a good idea to discuss your packaging with the librarian and/or architect.



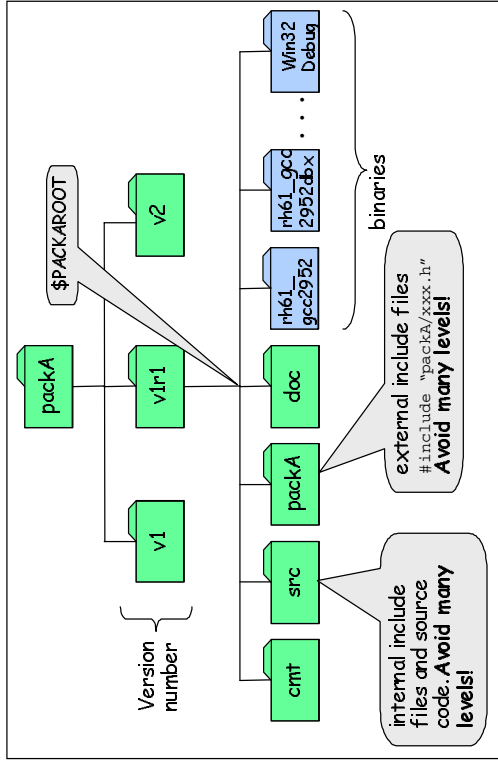


Figure 16.3 Layout of Gaudi software packages

## 16.3 Interfaces in Gaudi

One of the main design choices at the architecture level in Gaudi was to favour abstract interfaces when building collaborations of various classes. This is the way we best decouple the client of a class from its real implementation.

An abstract interface in C++ is a class where all the methods are pure virtual. We have defined some practical guidelines for defining interfaces. An example is shown in Listing 16.1:

From this example we can make the following observations:

- **Interface Naming.** The name of the class has to start with capital "I" to denote that it is an interface.
- **Derived from Interface.** We follow the convention that all interfaces should be derived from a basic interface `IInterface`. This interface defined 3 methods: `address() , release() and queryInterface()`. This methods allow the framework to manage the reference counting of the framework components and the possibility to obtain a different interface of a component using any interface (see Section 16.3.2).



Listing 16.1 Example of an abstract interface (IService)

```

1: // $Header: $
2: #ifndef GAUDIKERNEL_ISERVICE_H
3: #define GAUDIKERNEL_ISERVICE_H
4:
5: // Include files
6: #include "GaudiKernel/IInterface.h"
7: #include <string>
8:
9: // Declaration of the interface ID. (id, major, minor)
10: static const InterfaceID IID_IService(2, 1, 0);
11:
12: /** @class IService IService.h GaudiKernel/IService.h
13:
14: General service interface definition
15:
16: @author Pere Mato
17: */
18: class IService : virtual public IInterface {
19: public:
20: /// Retrieve name of the service
21: virtual const std::string& name() const = 0;
22: /// Retrieve ID of the Service. Not really used.
23: virtual const IID& type() const = 0;
24: /// Initialize Service
25: virtual StatusCode initialize() = 0;
26: /// Finalize Service
27: virtual StatusCode finalize() = 0;
28: /// Retrieve interface ID
29: static const InterfaceID& interfaceID() { return IID_IService; }
30: };
31:
32: #endif // GAUDIKERNEL_ISERVICE_H

```

- **Pure Abstract Methods.** All the methods should be pure abstract (`virtual Returntype method(...) = 0;`) With the exception of the static method `interfaceID()` (see later) and some inline templated methods to facilitate the use of the interface by the end-user.
- **Interface ID.** Each interface should have a unique identification (see Section 16.3.1) used by the query interface mechanism.

### 16.3.1 Interface ID

We needed to introduce an interface ID for identifying interfaces for the queryInterface functionality. The interface ID is made of a numerical identifier (generated from the interface name by a hash function) and major and minor version numbers. The version number is used to decide if the interface the service provider is returning is compatible with the interface the client is expecting. The rules for deciding if the interface request is compatible are:

- The interface identifier is the same
- The major version is the same



- The minor version of the client is less than or equal to the one of the service provider. This allows the service provider to add functionality (incrementing minor version number) keeping old clients still compatible.

The interface ID is defined in the same header file as the rest of the interface. Care should be taken of globally allocating the interface identifier (by giving a unique name to the constructor), and of modifying the version whenever a change of the interface is required, according to the rules. Of course changes to interfaces should be minimized.

```
static const InterfaceID IID_Ixxx("Ixxx" / *id*/, 1 / *major*/, 0 / *minor*/);
class Ixxx : public Interface {
 ...
 static const InterfaceID& interfaceID() { return IID_Ixxx; }
};
```

The static method `Ixxx::interfaceID()` is useful for the implementation of templated methods and classes using an interface as template parameter. The construct `T::interfaceID()` returns the interface ID of interface T.

### 16.3.2 Query Interface

The method `queryInterface()` is used to request a reference to an interface implemented by a component within the Gaudi framework. This method is implemented by each component class of the framework and allows us to navigate from one interface of a component to another, as shown for example in Listing 16.2, where we navigate from the `IMessageSvc` interface of the message service to its `IProperty` interface, in order to discover the value of its "outputLevel" property.

**Listing 16.2** Example usage of `queryInterface` to navigate between interfaces

```
1: IMessageSvc* msgSvc ();
2: ...
3: IProperty* msgProp;
4: msgSvc()->queryInterface(IID_IProperty, (void**) &msgProp);
5: std::string dfItrLevel;
6: StatusCode scl = msgProp->getProperty("outputLevel", dfItrLevel);
```

The implementation of `queryInterface()` is usually not very visible since it is done in the base class from which you inherit. A typical implementation is shown in Listing 16.3:

The implementation returns the corresponding interface pointer if there is a match between the received `InterfaceID` and the implemented one. The method `versionMatch()` takes into account the rules mentioned in Section 16.3.1.

If the requested interface is not recognized at this level (line 9), the call can be forwarded to the inherited base class or possible sub-components of this component.



**Listing 16.3** Example implementation of `queryInterface()`

```
1: StatusCode DataSvc::queryInterface(const InterfaceID& riid,
2: void** ppvInterface) {
3: if (IID_DataProviderSvc.versionMatch(riid)) {
4: *ppvInterface = (IDataProviderSvc*)this;
5: }
6: else if (IID_DataManagerSvc.versionMatch(riid)) {
7: *ppvInterface = (IDataManagerSvc*)this;
8: }
9: else {
10: return Service::queryInterface(riid, ppvInterface);
11: }
12: address();
13: return SUCCESS;
14: }
```

## 16.4 Libraries in Gaudi

Two different sorts of library can be identified that are relevant to the framework. These are *component* libraries, and *linker* libraries. These libraries are used for different purposes and are built in different ways.

### 16.4.1 Component Libraries

Component libraries are shared libraries that contain standard framework components which implement abstract interfaces. Such components are *Algorithms*, *Auditors*, *Services*, *Tools* or *Converters*. These libraries do not export their symbols apart from one which is used by the framework to discover what components are contained by the library. Thus component libraries should not be linked against, they are used purely at run-time, being loaded dynamically upon request, the configuration being specified by the job options file. Changes in the implementation of a component library do not require the application to be relinked.

Component libraries contain factories for their components, and it is important that the factory entries are declared and loaded correctly. The following sections describe how this is done.

When a component library is loaded, the framework attempts to locate a single entrypoint, called `getFactoryEntries()`. This is expected to declare and load the component factories from the library. Several macros are available to simplify the declaration and loading of the components via this function.

Consider a simple package `MyComponent_s`, that declares and defines the `MyAlgorithm` class, being a subclass of `Algorithm`, and the `MyService` class, being a subclass of `Service`. Thus the package will contain the header and implementation files for these classes (`MyAlgorithm.h`, `MyAlgorithm.cpp`, `MyService.h` and `MyService.cpp`) in addition to whatever other files are necessary for the correct functioning of these components.

In order to satisfy the requirements of a component library, two additional files must also be present in the package. One is used to declare the components, the other to load them. Because of the technical



limitations inherent in the use of shared libraries, it is important that these two files remain separate, and that no attempt is made to combine their contents into a single file.

The names of these files and their contents are described in the following sections.

#### 16.4.1.1 Declaring Components

Components within the component library are declared in a file `MyComponents_load.cpp`. By convention, the name of this file is the package name concatenated with `_load`. The contents of this file are shown below:

**Listing 16.4** The `MyComponents_load.cpp` file

```
#include "GaudiKernel/DeclareFactoryEntries.h"
DECLARE_FACTORY_ENTRIES (MyComponents) {
 DECLARE_ALGORITHM (MyAlgorithm); [1]
 DECLARE_SERVICE (MyService); [2]
}
```

Notes:

1. The argument to the `DECLARE_FACTORY_ENTRIES` statement is the name of the component library.
2. Each component within the library should be declared using one of the `DECLARE_XXX` statements discussed in detail in the next Section.

#### 16.4.1.2 Component declaration statements

The complete set of statements that are available for declaring components is given below. They include those that support C++ classes in different namespaces, as well as for `DataObjects` or `ContainedObjects` using the generic converters.

**Listing 16.5** The available component declaration statements

```
DECLARE_ALGORITHM(X)
DECLARE_AUDITOR(X)
DECLARE_CONVERTER(X)
DECLARE_GENERIC_CONVERTER(X) [1]
DECLARE_OBJECT(X)
DECLARE_SERVICE(X)

DECLARE_NAMESPACE_ALGORITHM(N,X)
DECLARE_NAMESPACE_AUDITOR(N,X)
DECLARE_NAMESPACE_CONVERTER(N,X)
DECLARE_NAMESPACE_GENERIC_CONVERTER(N,X)
DECLARE_NAMESPACE_OBJECT(N,X)
DECLARE_NAMESPACE_SERVICE(N,X)
```

**Listing 16.5** The available component declaration statements

Notes:

1. Declarations of the form `DECLARE_GENERIC_CONVERTER(X)` are used to declare the generic converters for `DataObject` and `ContainedObject` classes. For `DataObject` classes, the argument should be the class name itself (e.g. `EventHeader`), whereas for `ContainedObject` classes, the argument should be the class name concatenated with either `List` or `Vector` (e.g. `CellVector`) depending on whether the objects are associated with an `ObjectList` or `ObjectVector`.
2. Declarations of this form are used to declare components from explicit C++ namespaces. The first argument is the namespace (e.g. `At1Fast`), the second is the class name (e.g. `CellMarker`).

#### 16.4.1.3 Loading Components

Components within the component library are loaded in a file `MyComponents_dll.cpp`. By convention, the name of this file is the package name concatenated with `_dll`. The contents of this file are shown below:

**Listing 16.6** The `MyComponents_dll.cpp` file

```
#include "GaudiKernel/LoadFactoryEntries.h"
LOAD_FACTORY_ENTRIES (MyComponents) [1]
```

Notes:

1. The argument of `LOAD_FACTORY_ENTRIES` is the name of the component library.

#### 16.4.1.4 Specifying component libraries at run-time

The fragment of the job options file that specifies the component library at run-time is shown below.

**Listing 16.7** Selecting and running the desired tutorial example

```
ApplicationMgr.DLLs += { "MyComponents" }; [1]
```

Notes:

1. This is a list property, allowing multiple such libraries to be specified in a single line.
2. It is important to use the "+=" syntax to append the new component library or libraries to any that might already have been configured.

The convention in Gaudi is that component libraries have the same name as the package they belong to (prefixed by "lib" on Linux). When trying to load a component library, the framework will look for it in various places following this sequence:





- Look for an environment variable with the name of the package, suffixed by "Shr" (e.g. `{MyComponent}sShr`). If it exists, it should translate to the full name of the library, without the file type suffix (e.g. `{MyComponent}sShr`)  
`=$MYSOFT/MyComponents/v1/i386_linux22/libMyComponents` ).
- Try to locate the file `libMyComponents.so` using the `LD_LIBRARY_PATH` (on Linux), or `MyComponents.dll` using the `PATH` (on Windows).

### 16.4.2 Linker libraries

These are libraries containing implementation classes. For example, libraries containing code of a number of base classes or specific classes without abstract interfaces, etc. These libraries, contrary to the component libraries, export all the symbols and are needed during the linking phase in the application building. These libraries can be linked to the application "statically" or "dynamically", requiring a different file format. In the first case the code is added physically to the executable file. In this case, changes in these libraries require the application to be re-linked, even if these changes do not affect the interfaces. In the second case, the linker only adds into the executable minimal information required for loading the library and resolving the symbols at run time. Locating and loading the proper shareable library at run time is done exclusively using the `LD_LIBRARY_PATH` for Linux and `PATH` for Windows. The convention in Gaudi is that linker libraries have the same name as the package, suffixed by "Lib" (and prefixed by "lib" on Linux, e.g. `libMyComponentsLib.so`).

### 16.4.3 Library strategy and dual purpose libraries

Because component libraries are not designed to be linked against, it is important to separate the functionalities of these libraries from linker libraries. For example, consider the case of a `DataProvider` service that provides `DataObjects` for clients. It is important that the declarations and definitions of the `DataObjects` be handled by a different shared library than that handling the service itself. This implies the presence of two different packages - one for the component library, the other for the `DataObjects`. Clients should only depend on the second of these packages. Obviously the package handling the component library will in general also depend on the second package.

It is possible to have dual purpose libraries - ones which are simultaneously component and linker libraries. In general such libraries will contain `DataObjects` and `ContainedObjects`, together with their converters and associated factories. It is recommended that such dual purpose libraries be separated from single purpose component or linker libraries. Consider the case where several Algorithms share the use of several `DataObjects` (e.g. where one Algorithm creates them and registers them with the transient event store, and another Algorithm locates them), and also share the use of some helper classes in order to decode and manipulate the contents of the `DataObjects`. It is recommended that three different packages be used for this - one pure component package for the Algorithms, one dual-purpose for the `DataObjects`, and one pure linker package for the helper classes.



### 16.4.4 Building and linking with the libraries

Gaudi libraries and applications are built using CMT taking advantage of the CMT macros defined in the `GaudiPolicy` package. As an example, the CMT requirements file of the `GaudiTools` package is shown in Listing 16.8. The linker and component libraries are defined on lines 23 and 26 respectively - the linker library is defined first because it must be built ahead of the component library. Lines 28 and 34 set up the generic linker options and flags for the linker library, which are suffixed by the package specific flags set up by line 35. Line 31 tells CMT to generate the symbols needed for the component library, while line 33 sets up the corresponding linker flags for the component library. Finally, line 30 updates `LD_LIBRARY_PATH` (or `PATH` on Windows) for this package. In packages with only a component library and no linker library, line 30 could be replaced by "apply\_pattern packageShr", which would create the logical name required to access the component library by the first of the two methods described in Section 16.4.1.4.

Listing 16.8 CMT requirements file for the `GaudiTools` package

```
15: package GaudiTools
16: version v1
17:
18: branches GaudiTools cmt doc src
19: use GaudiKernel v8*
20: include_dirs "${GAUDIIToolsROOT}"
21:
22: #linker library
23: library GaudiToolsLib ../src/Associaator.cpp ../src/IInterface.cpp
24:
25: #component library
26: library GaudiTools ../src/GaudiTools_load.cpp ../src/GaudiTools_dll.cpp
27:
28: apply_pattern package_LlLinkopts
29:
30: apply_pattern ld_library_path
31: macro_append GaudiTools_stamps "${GaudiToolsDir}/GaudiToolsLib.stamp"
32:
33: apply_pattern package_CshlibFlags
34: apply_pattern package_LshlibFlags
35: macro_append GaudiToolsLib_shlibFlags ${GaudiKernel_linkopts}
```

### 16.4.5 Linking FORTRAN code

Any library containing FORTRAN code (more specifically, code that references COMMON blocks) must be linked statically. This is because COMMON blocks are, by definition, static entities. When mixing C++ code with FORTRAN, it is recommended to build separate libraries for the C++ and FORTRAN, and to write the code in such a way that communication between the C++ and FORTRAN worlds is done exclusively via wrappers. This makes it possible to build shareable libraries for the C++ code, even if it calls FORTRAN code internally. An example of a wrapper is the class `SiCbFunctions` in the file `src/Static/SiCbFortran.cpp` (`SiCb/SiCbCnv` package).



In cases where the FORTRAN calls are confined to a few well identified C++ algorithms, it may be more convenient to include the algorithms in the static library, removing the need for the wrapper class. A good example of this is the `SICE/CaloSiCbBckCnv` package. The `/src` directory contains two sub directories: `/static` for code going in the static library (including all FORTRAN code), and `/component` for code going in the C++ component library. Note that in the `/cmt/requirements` file, the link options explicitly request static linking of the `CaloSiCbBckCnv_LoadRef` file, in order to statically link the `CaloSiCbBckCnv` algorithm and all the FORTRAN functions which it calls.



## Chapter 17

# Analysis utilities

---

### 17.1 Overview

In this chapter we give pointers to some of the third party software libraries that we use within Gaudi or recommend for use by algorithms implemented in Gaudi.

### 17.2 CLHEP

CLHEP ("Class Library for High Energy Physics") is a set of HEP-specific foundation and utility classes such as random generators, physics vectors, geometry and linear algebra. It is structured in a set of packages independent of any external package. The documentation for CLHEP can be found on WWW at <http://wwwinfo.cern.ch/asd/lhc++/clhep/index.html>

CLHEP is used extensively inside Gaudi, in the `GaudiSvc`, `GaudiDb` and `LHCbEvent` packages.

### 17.3 HTL

HTL ("Histogram Template Library") is used internally in Gaudi (GaudiSvc package) to provide histogramming functionality. It is accessed through its abstract AIDA [13] compliant interfaces. Gaudi uses only the transient part of HTL. Histogram persistency is available with ROOT or HBOOK.

The documentation on HTL is available at <http://cern.ch/anaphe/documentation.html>.



## 17.4 NAG C

The NAG C library is a commercial mathematical library providing a similar functionality to the FORTRAN mathlib (part of CERNLIB). It is organised into chapters, each chapter devoted to a branch of numerical or statistical computation. A full list of the functions is available at [http://cern.ch/anaphe/documentation/Nag\\_C/NAGdoc/c/html/mark6.html](http://cern.ch/anaphe/documentation/Nag_C/NAGdoc/c/html/mark6.html)

NAG C is not explicitly used in the Gaudi framework, but developers are encouraged to use it for mathematical computations. Instructions for linking NAG C with Gaudi can be found at <http://cern.ch/lhcb-comp/Support/NagC/nagC.html>

Some NAG C functions print error messages to `stdout` by default, without any information about the calling algorithm and without filtering on severity level. A facility is provided by Gaudi to redirect these messages to the Gaudi MessageSvc. This is documented at <http://cern.ch/lhcb-comp/Support/NagC/GaudiNagC.html>

## 17.5 ROOT

ROOT is used by Gaudi for I/O and as a persistency solution for event data, histograms and n-tuples. In addition, it can be used for interactive analysis, as discussed in Chapter 10. Information about ROOT can be found at <http://root.cern.ch/>



## Chapter 18

# Accessing SICB facilities

---

### 18.1 Overview

In order to facilitate the transition towards C++ based code we have implemented a number of features into the Gaudi framework whose purpose is to allow access to certain SICB facilities. Gaudi can read data from SICB ZEBRA banks.

In this chapter we cover: staging data from tapes, converting SICB data for use by Gaudi algorithms, accessing the magnetic field map, accessing the SICB geometry description, the use of the SUNFIT, SUANAL and SULAST routines from within the Gaudi framework and handling event pileup and spillover in Gaudi. We also explain how to modify the size of the ZEBRA COMMON block.

When using the geometry and magnetic field descriptions described here, remember that they are a temporary solution for backwards compatibility with SICB. Gaudi includes the machinery to provide any algorithm with the detector data stored in XML format. Investing some time now to describe your detector in XML may be, in many cases, more convenient than using the old SICB routines to access detector description data. If you do use the old tools, use them only to populate some pre-defined class which can then be used in your algorithms. In this way, when you will move to the new detector description tools in Gaudi, the required changes in your code will be confined to the parts which access the geometry.

### 18.2 Reading tapes

There are three ways to specify a SICB input data file in the Gaudi job options:



1. Specify one or more file names. The name can be either the complete path on disk, or a "Castor file name" as returned by the book-keeping database. Using the Castor file name is the recommended method

```
EventSelector.Input = { "FILE = 'C:/data/mydata.dst,
/castor/ceern.ch/lhcb/mc/L251367'" };
```

2. Specify a data set by giving the job identification numbers (JOBIDs), as was done in SICB. The framework will make an http query to the book-keeping database in order to find the corresponding dataset, and then issue a stage request. This database query can be avoided by using Castor names instead of JOBIDs, as recommended above:

```
EventSelector.Input = { "JOBID='16434, 16435'" };
```

3. Specify explicitly a tape and file sequence number. This option is available just for completeness, tapes should never be accessed directly, use Castor instead.

```
EventSelector.Input = { "TAPE='Y21221-7, Y21223-24'" ;
```

The format is <Volume serial number of the tape><File sequence number>

When a Gaudi job requires to stage more than one file, the program waits for the first file to be staged. The rest of the required files are staged while the first file is processed. Once the program ends reading the first file, it will check if the next file is ready before continuing. If the file is not staged yet the program writes a message and waits until the file is in the staging disk pools.

**Skipping events** When reading SICB data, one may wish to skip some events. This can be done by setting the property `EventSelector.FirstEvent`. For example, with the following job option, the algorithms will only start processing at event 109:

```
EventSelector.FirstEvent = 109;
```

Note that Gaudi will take a few seconds to reach the requested event as it has to read all the records in the Zebra file before the one requested.

## 18.3 Populating the GAUDI transient data store: SICB Converters

### 18.3.1 General considerations

Access to the SICB data sets is basically via wrappers to the FORTRAN code. A complete event is read into the Zebra common block, and then the conversion to transient objects is done on request by specialised converters.



As mentioned in Chapter 13, a converter must implement the `IConverter` interface, by deriving from a specific base class. In this way any actions which are in common to all converters of a specific technology may be implemented in a single place.

In the following section we give detailed instructions on how to implement converters within the `SicbCnv` package. These are intended primarily for Gaudi developers themselves.

### 18.3.2 Implementing converters in the `SicbCnv` package

SICB converters are available for reading most SICB banks, the full list of converted banks is available at <http://cern.ch/lhcb-comp/Support/html/ConvertedSICBBanks.htm>. Writing back into persistent storage (ZEBRA files) is possible for a few banks.

Typically, `GAUDI` `DataObjects` can be of two types:

- Simple classes, which contain the data of a single SICB bank. These classes are of type `DataObject`. An example is the `Event` class containing the data of the PASS bank.
- Container classes, which contain data from multiple SICB banks. An example is the `ObjectVector<MCParticle>`, which contains Monte-Carlo particles with data from the ATMC bank.

Template files exist in the directory `$LHCBSOFT/SicbCnv/<version>/doc` for both types of converters, to ease the creation of user converters:

- `SicbCnv.Class.Template.cpp` and `SicbCnv.Class.Template.h` to be used when writing a converter for a single class.
- `SicbCnv.Class.Container.cpp` and `SicbCnv.Container.Template.h` to be used when writing a container of an object container.

If you intend to write your own SICB converter, follow the instructions below:

- Copy `SicbCnv.xxxxx.Template.h` to `Sicb<your-class>Cnv.h`, where `<your-class>` is the name of your persistent class.
- Copy `SicbCnv.xxxxx.Template.cpp` to `Sicb<your-class>Cnv.cpp`
- Now customize the header and the implementation file

- Follow `TODO` instructions in `Sicb<your-class>Cnv.h`
- Follow `TODO` instructions in `Sicb<your-class>Cnv.cpp`

- The converter factory must be made known to the system. This in fact depends on the linking mechanism: If the converter is linked into the executable as an object file, no action is necessary. However, usually the converter code resides in a shared or archive library. In this case the library must have an initialisation routine which creates an artificial reference to the created converter and forces the linker to include the code in the executable. An example of creating such a reference can be found in the file `$LHCBSOFT/SicbCnv/<version>/SicbCnv/SicbCnvDll/SicbCnv_load.cpp`. The convention for these initialization files is the following: for any other package replace the string "`SicbCnv`" with "OtherPackage".



- Compile link, debug
- Once the converter works, remove unnecessary TODO comments.

### 18.3.3 Back Converters

In some cases it may be necessary to convert event data from the Gaudi Transient Event Data store back to SICB banks in the ZEBRA COMMON block, to make this data available to FORTRAN algorithms. This "back" conversion is done by dedicated converter algorithms. An example is the `CaLoDigitToSicbConverter` algorithm in the `SICB/CaLoSicbCnv` package

## 18.4 Access to the Magnetic Field

The magnetic field map will be accessible in the future via the transient detector store. For the time being, as this is not implemented and as access to the magnetic field has been requested, we have provided a magnetic field service. Again this is effectively just a wrapper which uses SICB routines to read the information from a `.cdf` file.

The location of the `field.cdf` file is provided by the `standard.stream` file which is read in by a SICB routine called from Gaudi. This file is in the data base area, `$LHCDBASE/standard.stream` in AFS. For every version the file used in the production is read in. The location of the `standard.stream` file will be taken from an environment variable as per normal SICB operation.

To use the Magnetic field service one should modify the `jobOptions.txt` file to include the following:

```
ApplicationMgr.ExtSvc += { "MagneticFieldSvc" };
```

Any algorithm which requires the use of the service makes a request via the `service()` method of the `Algorithm` base class:

```
IMagneticFieldSvc* pIMF= 0;
StatusCode sc = service("MagneticFieldSvc", pIMF);
```

The service provides a method:

```
StatusCode fieldVector(HepPoint3D& Pos, HepVector3D& field)
```

which gives a magnetic field vector at a given point in space, for example:

The magnetic field service uses a new version of the SICB routine `GUFIELD`. In this new version the best possible description of the magnet geometry and the field are assumed in order to eliminate dependencies with other parts of SICB. Technically in SICB this corresponds to:

```
HepPoint3D P(10.*cm, 10.*cm, 120.*cm);
HepVector3D B;
pIMF->fieldVector(P, B);
```

```
IMAGLEV = IUVERS('GEBM','MAGN') = 4
IFLDLEV = IUVERS('GEBM','MFLD') = 4
```

These two parameters have been fixed to 4 in the production since a few months before the Technical Proposal: version 111 of SICB. Thus this seems to be a reasonable restriction until the field map is provided in the detector description database.

For an example of the use of this service see the sub-algorithm `readMagField` in the example `FieldGeom` distributed with the release.

## 18.5 Accessing the SICB detector Geometry from Gaudi

As discussed previously, the detector geometry will be included along with the field map in the XML detector description database. Currently only a part of the LHCb detector is in the new database. However, the detector geometry used in the production of the data can be accessed by calling a function in the `SicbFortran` name space (this function is just a wrapper for the FORTRAN function of similar name):

```
void SicbFortran::utdget(const std::string& a1, const std::string& a2,
 int& nParam, int* data);
```

`nParam` should be set to the number of data words required and on return from the function will contain the number of data words actually copied into the array `data`. The first string contains the name of the sub detector whose geometry is being requested and the second string is a list of options:

```
V'- version;
G'- geometry description parameters (default);
C'- calculated geometry (not in *.cdf);
H'- hit description parameters;
D'- digitization parameters;
R'- reconstruction parameters;
F'- floating point parameters (default);
I'- Integer parameters;
N'- take parameters from the *.cdf file
L'- ZEBRA pointer to the beginning of the parameters storage is returned in IARRAY(I)
```

An algorithm requiring this access should include the header file:

```
#include "SicbCnv/TopLevel/SicbFortran.h"
```



and call it so:

```
float rpar[300];
SicbFortran::utdget("WDRFN", "D", npar, (int *) rpar);
log << MSG::INFO << " wdrpar(" << j << ") = " << vF[j] << endl;
```

Note that the data returned by the function is written into an integer array. However we can also read floating point numbers as in the code fragment above by casting a float array.

One should notice that the geometry returned by this function is that which was used in the production of the data and not that which is in the current version of the `.cdf` files. Only if the option 'N' is specified are the `.cdf` files read in from the standard location. In order to be able to use the array of parameters returned one has to know in advance the organization of these data in the `cdf` file since the data are stored in an array and not in a common block with named variables.

The sub-algorithm `readTrackerGeom` in the example `FieldGeom` extracts and writes out some digitization and geometry parameters of the outer tracker.

## 18.6 Using FORTRAN code in Gaudi

Existing FORTRAN code can be used within a Gaudi application by using the `FortranAlgorithm` class. This is a standard Gaudi algorithm which calls the FORTRAN routines `SUINIT` in the `initialize()` method, `SUANA1` in the `execute()` method for each event, and `SUFAST` in `finalize()`. Implementing these three routines allows you to write code in FORTRAN and have it called from within Gaudi, in particular to import routines already written in SICB.

Note, however that there are some points that should be kept in mind when importing SICB code into Gaudi. The following list is far from being complete but we hope that it will help anyone using the `FortranAlgorithm`. It may be updated in the future with the user's experiences and our own findings.

- Gaudi is linked with only two sub-packages of SICB: `Finclue` and `Futio`. This means that taking some code from SICB and running it successfully in Gaudi will not always be straight forward. Every case will have to be studied separately. In some cases you may find that the effort required to make the code compatible with Gaudi is comparable to writing it directly in C++. For some pieces of code the integration into Gaudi may be very simple. The difficulties will come mainly from dependencies of the code you want to include on other parts of SICB. For instance there may be common blocks which your code needs but which are never initialized.
- As most of SICB is not executed, you will have either to include and execute the required initialization or to try to eliminate those dependencies.



- Gaudi uses SICB to access the data in the DST's, to extract the information from the banks and to convert the data into the appropriate classes. This needs some initialization but not every SICB initialization step is followed. The standard stream file and setup `cdf` are read in from the standard locations. The program will have access to the data in the `cdf` files which were used to produce the DST.
- `Sicb.dat` is not read in. If you need to change the running conditions of your program do it using the `jobOptions.txt` file or write your own cards file and read it in `SUINIT`.
- In `Finclue` you will find all the `NAME_BANK.INC` and `NAME_FUNC.INC` include files. This means that you have access to the bank data with the usual utilities provided by SICB. For example, to access the momentum of the reconstructed track one can use (as in SICB):

```
P = AXTKSP(NOBI)
```

- `Futio` includes most of the `UT*` and `UB` routines which can therefore be used by FORTRAN code within Gaudi.
- Initialize the histogram files yourself in `SUINIT`. Gaudi initializes the histogram service but this can be accessed only from C++ code.

## 18.7 Handling pile up in Gaudi.

In this section we explain the pile-up structure implemented in the Gaudi framework (see Figure 18.1). Please note that this possibility has been rendered somewhat redundant by the implementation of pile-up at generator level in SICBMC v244.

Pile-up in Gaudi is performed by a pile-up algorithm. An example (`PileUpAlg`) can be found in the package `SicbCrv`. The pile-up algorithm creates a second instance of the event selector which has to be configured in the job options file. The leading event will be read in from the `EventSelector` created by the `ApplicationMgr`. The piled-up events are read in from the second instance created by the pile-up algorithm. There are two different iterators, one for each `EventSelector` instance, which loop over these two data streams.

When reading the ZEBRA files produced by SICB the events are merged at the level of the ZEBRA common blocks by a call to the `Futio` routine `REMERGE`. Every C++ algorithm requesting data from one of the converted banks will get the merged data. Every FORTRAN algorithm using the SICB banks will also read the merged data as SICBDST does.

`PileUpAlg` must be the first algorithm called (unless spillover is also enabled, in which case it should be the second algorithm after the spillover algorithm), other algorithms will access the merged events when they retrieve some data from the store. The pile-up algorithm controls the number of events which have to be piled-up to every leading event. `PileUpAlg` uses a C++ version of the SICB routine `RELUMI` to get the number of pile-up events as a function of the luminosity and some other beam parameters. Those parameters are currently read in from the beam `cdf` file. The C++ version of `RELUMI` uses the random number service in Gaudi. Other implementations of the algorithm, for instance to return a fix number of pile-up events every time, may be implemented if they are needed.







event structure from the Spillover input file. This approach is an approximation: if `num_inter > 0`, the Spillover algorithm always reads one (and only one) event from the input file, regardless of the value of `num_inter`. In theory the event read from the input file should contain `num_inter` piled up events; in practice one just reads the next event. It would be possible to do the correct thing by skipping events until one with the right PileUp multiplicity is found, or to open several input files, each containing events with a fixed PileUp multiplicity. Neither of these possibilities is currently implemented.

- Since Spillover events will be combined by digitisers into a single Raw event, it is only foreseen to provide the /MC part of the event. Furthermore, only those parts of the /MC subevent whose converter foresees Spillover are available. Listing 18.2 for example shows how the SICB `MCParticle`s converter has been extended to allow spillover. Any SICB converter can be extended in a similar way if required

Listing 18.2 Adding provision for spillover to a SICB converter

```
1: // Next line declares path for main event
2: declareObject("/Event/MC/MCParticles", objType(), "ATMC");
3: // Next 4 lines declare path for spillover events
4: declareObject("/Event/Prev/MC/MCParticles", objType(), "ATMC");
5: declareObject("/Event/Next/MC/MCParticles", objType(), "ATMC");
6: declareObject("/Event/PrevPrev/MC/MCParticles", objType(), "ATMC");
7: declareObject("/Event/NextNext/MC/MCParticles", objType(), "ATMC");
```

## 18.9 Increasing the size of the ZEBRA COMMON block

The default Gaudi implementation initializes the ZEBRA COMMON block to a rather small size (3M words) in order to avoid a large memory overhead for applications that need to access SICB facilities but not large events. This limited size is however insufficient to read all but the simplest of events. Data processing applications need to override this size. One way to do this is to provide a private version of the `GETZEBRASIZE` routine, to be linked ahead of the one in the `SICBConv` library. An example is shown in Listing 18.3, where the ZEBRA COMMON size is set to 12M words

Listing 18.3 Code to redefine the size of the ZEBRA COMMON block

```
1: SUBROUTINE GETZEBRASIZE(size)
2:
3: IMPLICIT NONE
4: #define NRPW 1000000
5: #define NWGEAN 12000000
6: REAL GSTORE
7: COMMON / GCBANK / GSTORE(NWGEAN)
8: REAL HPW
9: COMMON / PAWC / HPW(NRPW)
10:
11: INTEGER size
12: size = NWGEAN
13:
14: RETURN
15: END
```



## Appendix A

# References

- 1 GAUDI - Architecture Design Report [LHCb 98-064 COMP]
- 2 GAUDI online code documentation (<http://cern.ch/proj-gaudi/Doxygen/v9/>)
- 3 GAUDI - User Requirements Document [LHCb 98-065 COMP]
- 4 G. Barrand et al., GAUDI - A software architecture and framework for building LHCb data processing applications. [Proc CHEP 2000, Computer Physics Communications 140 (2001) 45-55]  
(<http://cern.ch/lhcb-comp/General/Publications/longpap-a152.pdf>)
- 5 LHCb Physical Units Convention  
(<http://cern.ch/lhcb-comp/Reconstruction/Conventions/units.pdf>)
- 6 Revised LHCb coding conventions [LHCb 2001-054 COMP]
- 7 CMT configuration management environment  
<http://www.lal.in2p3.fr/technique/si/SI/CMT/CMT.htm>
- 8 The new LHCb Event Data Model [LHCb 2001-142 COMP]
- 9 R. Chytráček et al., The LHCb Detector Description Framework. [Proc CHEP 2000]  
| <http://cern.ch/lhcb-comp/General/Publications/pap-a155.pdf>
- 10 I. Belyaev et al., Integration of GEANT4 with GAUDI. [Proc CHEP 2001]  
<http://cern.ch/lhcb-doc/presentations/conferencetalks/BelyaevProceedingsCHEP01.pdf>
- 11 GiGa: Geant4 Interface for Gaudi Applications  
<http://cern.ch/lhcb-comp/Frameworks/Gaudi/Documents/GiGa.pdf>
- 12 M. Frank et al., Data Persistency Solution for LHCb. [Proc CHEP 2000],  
<http://cern.ch/lhcb-comp/General/Publications/pap-c153.pdf>
- 13 AIDA: Abstract Interfaces for Data Analysis <http://aida.freehep.org/>



## Appendix B

# Options for standard components

The following is a list of options that may be set for the standard components: e.g. data files for input, print-out level for the message service, etc. The options are listed in tabular form for each component along with the default value and a short explanation. The component name is given in the table caption thus: [ComponentName].

**Table B.1** Standard Options for the Application manager [ApplicationMgr]

| Option name          | Default value    | Meaning                                                                                                                                  |
|----------------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| EvtSel               | ""               | If "NONE", no event input                                                                                                                |
| EvtMax               | -1               | Maximum number of events to process. The default is -1 (infinite) unless EvtSel = "NONE"; in which case it is 10.                        |
| TopAlg               | {}               | List of top level algorithms. Format: {<Type>/<Name> <Type2>/<Name2>...};                                                                |
| ExtSvc               | {}               | List of external services to be explicitly created by the ApplicationMgr (see section 11.2). Format: {<Type>/<Name> <Type2>/<Name2>...}; |
| OutStream            | {}               | Declares an output stream object for writing data to a persistent store, e.g. {"DstWriter"}; See also Table B.10                         |
| DLLs                 | {}               | Search list of libraries for dynamic loading. Format: {<dll> <dll2>...};                                                                 |
| HistogramPersistence | "NONE"           | Histogram and N-tuple persistence mechanism. Available options are "HBOOK", "ROOT", "NONE".                                              |
| Runnable             | "AppMgrRunnable" | Type of runnable object to be created by Application manager                                                                             |
| EventLoop            | "EventLoopMgr"   | Type of event loop:<br>"EventLoopMgr" is standard event loop<br>"MinimalEventLoop" executes algorithms but does not read events          |
| OutputLevel          | MSG::INFO        | Same as MessageSvc.OutputLevel.<br>See Table B.2 for possible values                                                                     |

**Table B.1** Standard Options for the Application manager [ApplicationMgr]

| Option name                                                                                                                                                                                                                                                                                                                                                                   | Default value    | Meaning                           |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|-----------------------------------|
| The last two options define the source of the job options file and so they cannot be defined in the job options file itself. There are two possibilities to set these options, the first one is using a environment variable called JOBPATH or setting the option to the application manager directly from the main program <sup>a</sup> . The coded option takes precedence. |                  |                                   |
| JobOptionsType                                                                                                                                                                                                                                                                                                                                                                | "FILE"           | Type of file (FILE implies ascii) |
| JobOptionsPath                                                                                                                                                                                                                                                                                                                                                                | "jobOptions.txt" | Path for job options source       |

- a. A basic DataObject object is created as event root (" /Event")
- b. The setting of properties from the main program is discussed in Chapter 4.

**Table B.2** Standard Options for the message service [MessageSvc]

| Option name | Default value            | Meaning                                                                                                     |
|-------------|--------------------------|-------------------------------------------------------------------------------------------------------------|
| OutputLevel | 0                        | Verboseness threshold level:<br>0=NIL, 1=VERBOSE, 2=DEBUG, 3=INFO,<br>4=WARNING, 5=ERROR, 6=FATAL, 7=ALWAYS |
| Format      | "% F%18W%S%7W%R%T %0W%M" | Format string.                                                                                              |

**Table B.3** Standard Options for all algorithms [<myAlgorithm>]

| Option name     | Default value | Meaning                                                                        |
|-----------------|---------------|--------------------------------------------------------------------------------|
| OutputLevel     | 0             | Message Service Verboseness threshold level. See Table B.2 for possible values |
| Enable          | true          | If false, application manager skips execution of this algorithm                |
| ErrorMax        | 1             | Job stops when this number of errors is reached                                |
| ErrorCount      | 0             | Current error count                                                            |
| AuditInitialize | false         | Enable/Disable auditing of Algorithm initialisation                            |
| AuditExecute    | true          | Enable/Disable auditing of Algorithm execution                                 |
| AuditFinalize   | false         | Enable/Disable auditing of Algorithm finalisation                              |

**Table B.4** Standard Options for all services [<myService>]

| Option name | Default value | Meaning                                                                        |
|-------------|---------------|--------------------------------------------------------------------------------|
| OutputLevel | 0             | Message Service Verboseness threshold level. See Table B.2 for possible values |

**Table B.5** Standard Options for all Tools [<myTools>]

| Option name | Default value | Meaning                                                                        |
|-------------|---------------|--------------------------------------------------------------------------------|
| OutputLevel | 0             | Message Service Verboseness threshold level. See Table B.2 for possible values |

**Table B.6** Standard Options for all Associators [<myAssociator>]

| Option name  | Default value | Meaning                                                                         |
|--------------|---------------|---------------------------------------------------------------------------------|
| FollowLinks  | true          | Instruct the associator to follow the links instead of using cached information |
| DataLocation | ""            | Location where to get association information in the data store                 |

**Table B.7** Standard Options for Auditor service [AuditorSvc]

| Option name | Default value | Meaning                                                                                         |
|-------------|---------------|-------------------------------------------------------------------------------------------------|
| Auditors    | {}            | List of Auditors to be loaded and to be used.<br>See section 11.7 for list of possible auditors |

**Table B.8** Standard Options for all Auditors [<myAuditor>]

| Option name | Default value | Meaning                                                                        |
|-------------|---------------|--------------------------------------------------------------------------------|
| OutputLevel | 0             | Message Service Verboseness threshold level. See Table B.2 for possible values |
| Enable      | true          | If false, application manager skips execution of the auditor                   |

**Table B.9** Options of Algorithms in GaudiAlg package (see Section 5.5)

| Algorithm name | Option Name   | Default value | Meaning                                                  |
|----------------|---------------|---------------|----------------------------------------------------------|
| EventCounter   | Frequency     | 1;            | Frequency with which number of events should be reported |
|                | PercentPass   | 100.0;        | Percentage of events that should be passed               |
| Sequencer      | Members       |               | Names of algorithms in the sequence                      |
|                | BranchMembers |               | Names of algorithms on the branch                        |
| Sequencer      | StopOverride  | false;        | If true, do not stop sequence if a filter fails          |



**Table B.10** Options available for output streams (e.g. DstWriter)

Output stream objects are used for writing user created data into data files or databases. They are created and named by setting the option `ApplicationMgr.OutStream`. For each output stream the following options are available

| Option name      | Default value       | Meaning                                                                                                     |
|------------------|---------------------|-------------------------------------------------------------------------------------------------------------|
| ItemList         | {}                  | The list of data objects to be written to this stream, e.g. <code>{"/Event#1","Event/MyTracks/#1"}</code> ; |
| Preload          | true;               | Preload items in ItemList                                                                                   |
| Output           | ""                  | Output data stream specification. Format: <code>["DATAFILE=mydstroot'TYP=ROOT"]</code> ;                    |
| OutputFile       | ""                  | Output file specification - same as DATAFILE in previous option                                             |
| EventDataSvc     | "EventDataSvc"      | The service from which to retrieve objects.                                                                 |
| EvtConversionSvc | "EvtPersistenceSvc" | The persistence service to be used                                                                          |
| AcceptAlgs       | {}                  | If any of these algorithms sets <code>filterflag=true</code> ; the event is accepted                        |
| RequireAlgs      | {}                  | If any of these algorithms is not executed, the event is rejected                                           |
| VetoAlgs         | {}                  | If any of these algorithms does not set <code>filterflag = true</code> ; the event is rejected              |

**Table B.11** Standard Options for persistence services (e.g. EventPersistenceSvc)

| Option name | Default value | Meaning                                                                                                           |
|-------------|---------------|-------------------------------------------------------------------------------------------------------------------|
| CnvServices | {}            | Conversion services to be used by the service to load or store persistent data (e.g. <code>RootEvtCnvSvc</code> ) |

**Table B.12** Standard Options for conversion services (e.g. RootEvtCnvSvc)

| Option name | Default value | Meaning                                          |
|-------------|---------------|--------------------------------------------------|
| DbType      | ""            | Persistence technology (e.g. <code>ROOT</code> ) |

**Table B.13** Standard Options for the histogram service [HistogramPersistenceSvc]

| Option name         | Default value | Meaning                                                                                                                                  |
|---------------------|---------------|------------------------------------------------------------------------------------------------------------------------------------------|
| OutputFile          | ""            | Output file for histograms. Histograms not saved if not given.                                                                           |
| RowWiseNTuplePolicy | "FLOAT_ONLY"  | Persistent representation of NTuple data types. Other possible value is <code>"USE_DATA_TYPES"</code> . See Section 10.2.3.2 for details |
| PrintHistos         | false         | Print the histograms also to standard output (HBOOK only)                                                                                |



**Table B.14** Standard Options for the N-tuple service [NTupleSvc] (see Section 10.2.3.2)

| Option name | Default value | Meaning                                                                                                                                   |
|-------------|---------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Input       | {}            | Input file(s) for n-tuples. Format: <code>["FILE1 DATAFILE=tuple1.typ' OPT=OLD' ", ["FILE2 DATAFILE=tuple2.typ' OPT=OLD' ", ...]]</code>  |
| Output      | {}            | Output file(s) for n-tuples. Format: <code>["FILE1 DATAFILE=tuple1.typ' OPT=NEW' ", ["FILE2 DATAFILE=tuple2.typ' OPT=NEW' ", ...]]</code> |
| StoreName   | " /NTUPLES"   | Name of top level entry                                                                                                                   |

**Table B.15** Standard Options for the Event Collection service [TagCollectionSvc] (see Section 10.3.2)

| Option name | Default value | Meaning                                                   |
|-------------|---------------|-----------------------------------------------------------|
| Output      | {}            | Output file specification. See Section 10.3.2 for details |
| StoreName   | " /NTUPLES"   | Name of top level entry                                   |

**Table B.16** Standard Options for the standard event selector [EventSelector]

| Option name | Default value | Meaning                                                                                                                                                                                                                                                                                                               |
|-------------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input       | {}            | Input data stream specification.<br>Format: <code>" &lt;tagname&gt; = '&lt;tagvalue&gt;' &lt;opt&gt;"</code><br><br>Possible tags are different depending on input data type.<br>For SICB data, see Section 18.2<br>For other (e.g. ROOT) Event data, see Section 6.10.2<br>For Event Collections, see Section 10.3.2 |
| FirstEvent  | 1             | First event to process (allows skipping of preceding events)                                                                                                                                                                                                                                                          |
| PrintFreq   | 10            | Frequency with which event number is reported                                                                                                                                                                                                                                                                         |

**Table B.17** Standard Options for Sicb Event Conversion service [SicbEventCnvSvc]

| Option name     | Default value | Meaning                                            |
|-----------------|---------------|----------------------------------------------------|
| enableSpillover | " false;"     | If true, spillover is enabled. See also Table B.19 |

**Table B.18** Standard Options for Sicb Pileup Algorithm [PileUpAlg]

| Option name | Default value | Meaning                                                                                                                                                              |
|-------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PileUpMode  | ""            | Pileup mode. Possible values:<br>"LUMISIGNAL": Luminosity weighted for signal leading event<br>"LUMINBIAS": LUMI: Luminosity weighted for minimum bias leading event |



**Table B.19** Standard Options for Sicb Spillover Algorithm [SpilloverAlg]

| Option name   | Default value | Meaning                                                                                    |
|---------------|---------------|--------------------------------------------------------------------------------------------|
| SpilloverMode | "LUMI"        | Spillover mode. Possible value: "LUMI": Luminosity weighted                                |
| SpilloverPrev | 1             | Number of preceding bunch crossings to be considered for spillover. Possible values: 0,1,2 |
| SpilloverNext | 0             | Number of following bunch crossings to be considered for spillover. Possible values: 0,1,2 |
| SpilloverData | ""            | List of event data containers to be loaded into transient data store for spillover events. |

**Table B.20** Event Tag Collection Selector [EventCollectionSelector]

The following options are used internally by the EventCollectionsSelector. They should not normally be used directly by users, who should set them via the "tags" of the EventSelector.Input option

| Option name    | Corresponding tag of EventSelector.Input | Default value     | Meaning                       |
|----------------|------------------------------------------|-------------------|-------------------------------|
| ChvService     | SVC                                      | "EvtTupleSvc"     | Conversion service to be used |
| Authentication | AUTH                                     | ""                | Authentication to be used     |
| Container      |                                          | "B2PPI"           | Container name                |
| Item           |                                          | "Address"         | Item name                     |
| Criteria       | SEL                                      | ""                | Selection criteria            |
| DB             | DATAFILE                                 | ""                | Database name                 |
| DbType         | TYP                                      | ""                | Database type                 |
| Function       | FUN                                      | "NTuple:Selector" | Selection function            |

**Table B.21** Standard Options for Python scripting service [PythonScriptingSvc]

| Option name   | Default value | Meaning                                                                       |
|---------------|---------------|-------------------------------------------------------------------------------|
| startupScript | ""            | Script to be executed after reading the job options file (see Section 14.2.2) |

**Table B.22** Standard Options for the Introspection service [IntrospectionSvc]

| Option name  | Default value | Meaning                                   |
|--------------|---------------|-------------------------------------------|
| Dictionaries | ""            | Initial list of dictionaries to be loaded |



**Table B.23** Standard Options for Detector Data service [DetDataSvc]

| Option name    | Default value | Meaning                                                                                               |
|----------------|---------------|-------------------------------------------------------------------------------------------------------|
| DetStorageType | 7             | Detector database storage type: 7=XML                                                                 |
| DetDbLocation  | "empty"       | Detector database location (filename,URL). If empty then the environment variable XMLDDBROOT is used. |
| DetDbRootName  | "id"          | Name of the root node of the detector transient store                                                 |
| UsePersistency | false;        | If true, load from persistent medium.                                                                 |

**Table B.24** Standard Options for XML conversion service [XmlConvSvc]

| Option name            | Default value | Meaning                                                     |
|------------------------|---------------|-------------------------------------------------------------|
| AllowGenericConversion | false;        | Switch for generic detector element conversion <sup>a</sup> |

a. The XML conversion service allows the possibility to convert user defined detector elements to generic detector elements. This means that only the generic part of the detector element, and its associated geometry, will be converted but not the user defined detector element data. This feature can be used when the user defined detector element data are not needed (e.g. visualization) or when the corresponding user defined XML converters are not available (testing). When this feature is switched ON, the Gaudi application will run successfully with an information message saying that this feature is enabled, and will print out the information about all the detector elements to which this generic conversion is applied. The limitation of this feature is that after the generic conversion the returned reference points to a `DetElement` object and not to the user defined class. This means that `SmartDataPtr` class can be parameterized only by `DetectorElement` class and not by the user defined class.

**Table B.25** Standard Options for XML parser service [XmlParserSvc]

| Option name     | Default value | Meaning                         |
|-----------------|---------------|---------------------------------|
| MaxDocNbInCache | 5;            | Maximum number of cached files  |
| CacheBehavior   |               | See Section 8.3.3.3.1 for usage |

**Table B.26** Standard Options for Random Numbers Generator Service [RndmGenSvc]

| Option name | Default value                   | Meaning                            |
|-------------|---------------------------------|------------------------------------|
| Engine      | "HepRndm::Engine<RanluxEngine>" | Random number generator engine     |
| Seeds       |                                 | Table of generator seeds           |
| Column      | 0                               | Number of columns in seed table -1 |
| Row         | 1                               | Number of rows in seed table -1    |
| Luxury      | 3                               | Luxury value for the generator     |
| UseTable    | false                           | Switch to use seeds table          |





**Table B.27** Standard Options for Particle Property Service [ParticlePropertySvc]

| Option name            | Default value                  | Meaning                               |
|------------------------|--------------------------------|---------------------------------------|
| ParticlePropertiesFile | "(SLHCDBASE)/cdf/particle.cdf" | Particle properties database location |

**Table B.28** Standard Options for Chrono and Stat Service [ChronoStatSvc]

| Option name            | Default value | Meaning                                                  |
|------------------------|---------------|----------------------------------------------------------|
| ChronoPrintOutTable    | true          | Global switch for profiling printout                     |
| PrintUserTime          | true          | Switch to print User Time                                |
| PrintSystemTime        | false         | Switch to print System Time                              |
| PrintElapsedTime       | false         | Switch to print Elapsed time (Note typo in option name!) |
| ChronoDestinationCout  | false         | If true, printout goes to cout rather than MessageSvc    |
| ChronoPrintLevel       | 3             | Print level for profiling (values as for MessageSvc)     |
| ChronoTableToBeOrdered | true          | Switch to order printed table                            |
| StatPrintOutTable      | true          | Global switch for statistics printout                    |
| StatDestinationCout    | false         | If true, printout goes to cout rather than MessageSvc    |
| StatPrintLevel         | 3             | Print level for profiling (values as for MessageSvc)     |
| StatTableToBeOrdered   | true          | Switch to order printed table                            |

## B.1 Obsolete options

The following options are obsolete and should not be used. They are documented here for completeness and may be removed in a future release.

**Table B.29** Obsolete Options

| Obsolete Option      | Replacement                       |
|----------------------|-----------------------------------|
| EventSelector:EvtMax | ApplicationMgr:EvtMax (Table B.1) |



## Appendix C

# Job Options Grammar and Error Codes

---

### C.1 The EBNF grammar of the Job Options files

The syntax of the Job-Options-File is defined through the following EBNF-Grammar.

```
Job-Options-File =
 {Statements} .

Statements =
 {Include-Statement} | {Assign-Statement} | {Append-Statement} |
 {Platform-Dependency} .

AssertableStatements =
 {Include-Statement} | {Assign-Statement} | {Append-Statement} .

AssertionStatement =
 '#ifdef' | '#ifndef' .

Platform-Dependency =
 AssertionStatement 'WIN32' <AssertableStatements> [#else <Asserta-
bleStatements>] #endif

Include-Statement =
 '#include' string .

Assign-Statement =
 Identifier '\,' Identifier '=' value ';' .

Append-Statement =
 Identifier '\,' Identifier '+=' value ';' .
```

```

Identifier =
 letter {letter | digit} .

value =
 boolean | integer | double | string | vector .

vector =
 { ' vectorvalue { ',' vectorvalue } ' } .

vectorvalue =
 boolean | integer | double | string .

boolean =
 'true' | 'false' .

integer =
 prefix scientificdigit .

double =
 (prefix <digit> '.' [scientificdigit]) |
 (prefix '.' scientificdigit) .

string =
 ' ' {char} ' ' .

scientificdigit =
 <digit> [('e' | 'E') <digit>] .

digit =
 <figure> .

prefix =
 ['+' | '-'] .

figure =
 '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9' .

char =
 any character from the ASCII-Code

letter =
 set of all capital- and non-capital letter

```



## C.2 Job Options Error Codes and Error Messages

The table below lists the error codes and error messages that the Job Options compiler may generate, their reason and how to avoid them.

Table 18.1 Possible Error-Codes

| Error-Code                                                            | Reason                                                                                                                                                                               | How to avoid it                                                                                                                                             |
|-----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Error #000 Internal compiler error                                    | -                                                                                                                                                                                    | This code normally should never appear. If this code is shown there is maybe a problem with your memory, your disk-space or the property-file is corrupted. |
| Error #001 Included property-file does not exist or can not be opened | * wrong path in #include-directive<br>* wrong file or mistyped filename<br>* file is exclusively locked by another application<br>* no memory available to open this file            | Please check if any of the listed reasons occurred in your case.                                                                                            |
| Warning #001 File already included by another file                    | The file was already included by another file and will not be included a second time.<br>The compiler will ignore this #include-directive and will continue with the next statement. | Remove the #include-directive                                                                                                                               |
| Error #002 syntax error: Object expected                              | The compiler expected an object at the given position.                                                                                                                               | Maybe you mistyped the name of the object or the object contains unknown characters or does not fit the given rules.                                        |
| Error #003 syntax error: Missing dot between Object and Propertyname  | The compiler expect a dot between the Object and the Propertyname.                                                                                                                   | Check if the dot between the Object and the Propertyname is missing.                                                                                        |
| Error #004 syntax error: Identifier expected                          | The compiler expected an identifier at the given position.                                                                                                                           | Maybe you mistyped the name of the identifier or the identifier contains unknown characters or does not fit the given rules.                                |
| Error #005 syntax error: Missing operator '+' or '='                  | The compiler expected an operator between the Propertyname and the value.                                                                                                            | Check if there is a valid operator after the Propertyname. Note that a blank or tab is not allowed between '+'='!                                           |



Table 18.1 Possible Error-Codes

| Error-Code                                                                  | Reason                                                                                                                | How to avoid it                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Error #006 String is not terminated by a "                                  | A string (value) was not terminated by a "                                                                            | Check if all your strings are beginning and ending with ". Note that the position given by the compiler can be wrong because the compiler may thought that following statements are part of the string!                                                                                               |
| Error #007 syntax error: #include-statement is not correct                  | The next token after the #include is not a string.                                                                    | Make sure that after the #include-directive there is specified the file to include. The file must be defined as a string!                                                                                                                                                                             |
| Error #008 syntax error: #include does not end with a ;                     | The include-directive was terminated by a ;                                                                           | Remove the ; after the #include-directive.                                                                                                                                                                                                                                                            |
| Error #009 syntax error: Values must be separated with , ;                  | One or more values within a vector were not separated with a , ; or one ore more values within a vector are mistyped. | Check if every value in the vector is separated by a , ;. If so the reason for this message may result in mistyped values in the vector (maybe there is a blank or tab between numbers).                                                                                                              |
| Error #010 syntax error: Vector must end with };                            | The closing bracket is missing or the vector is not terminated correctly.                                             | Check, if the vector ends with a } and if there is no semicolon before the ending-bracket.                                                                                                                                                                                                            |
| Error #011 syntax error: Statement must end with a ;                        | The statement is not terminated correctly.                                                                            | Check if the statement ends with a semicolon ;.                                                                                                                                                                                                                                                       |
| Runtime-Error #012: Cannot append to object because it does not exists      | The compiler cannot append the values to the object-propertyname because the object does not exist.                   | Check if the referred object is defined in one of the included files, if so check if you writed the object-name exactly like in the include-file.                                                                                                                                                     |
| Runtime-Error #013 Cannot append to object because Property does not exists | The compiler cannot append the values to the object-propertyname because the property does not exist.                 | Check if there was already something assigned to the referred property (in the include-file or in the current file). If not then modify the append-statement into a assign-statement.<br>If there was already something assigned, check if the object-name and the property-name are typed correctly. |



Table 18.1 Possible Error-Codes

| Error-Code                                                               | Reason                                                                                                                                                                                                                 | How to avoid it                                                                                                                                                   |
|--------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Error #014 Elements in the vector are not of the same type               | One or more elements in the vector have a different type than the first element in the vector. All elements must have the same type like the first declared element.                                                   | Check declaration of vector, check the types and check, if maybe a value is mistyped.                                                                             |
| Error #015 Value(s) expected                                             | The compiler didn't find values to append or assign                                                                                                                                                                    | Check the statement if there exists values and if they are written correctly.<br>Maybe this error is a result of a previous error!                                |
| Error #016 Specified property-file does not exist or can not be resolved | The compiler was not able to include a property-file or didn't found the file.<br>A reason can be that the compiler was not able to resolve an environment-variable which points to the location of the property-file. | Check, if you are using environment-variables to resolve the file, if they are mistyped (wether in the system or in the #include-directive) or not set correctly. |
| Error #017 #fdef not followed by an identifier                           | The #fdef-statement is not followed by the assertion-identifier (WIN32).                                                                                                                                               | Add WIN32 after the #fdef-statement.                                                                                                                              |
| Error #018 Identifier in #fdef / #fndef not known                        | The assertion-identifier used in the #fdef. / #fndef-statement is not known. At the moment there can only be used WIN32!                                                                                               | Change identifier to WIN32.                                                                                                                                       |
| Error #019 #fdef / #fndef / #else / #endif doesn't end with a ; ;        | A semicolon was found after the #fdef. / #fndef. / #else. / #endif-statement. These statements don't end with a semicolon.                                                                                             | Remove the semicolon after the #fdef / #fndef / #else / #endif-statement.                                                                                         |





## Appendix D

# Design considerations

---

### D.1 Generalities

In this chapter we look at how you might actually go about designing and implementing a real physics algorithm. It includes points covering various aspects of software development process and in particular:

- The need for more “thinking before coding” when using an OO language like C++.
- Emphasis on the specification and analysis of an algorithm in mathematical and natural language, rather than trying to force it into (unnatural?) object orientated thinking.
- The use of OO in the design phase, i.e. how to map the concepts identified in the analysis phase into data objects and algorithm objects.
- The identification of classes which are of general use. These could be implemented by the computing group, thus saving you work!
- The structuring of your code by defining private utility methods within concrete classes.

When designing and implementing your code we suggest that your priorities should be as follows: (1) Correctness, (2) Clarity, (3) Efficiency and, very low in the scale, OOness

Tips about specific use of the C++ language can be found in the coding rules document [6] or specialized literature.

### D.2 Designing within the Framework

A physicist designing a real physics algorithm does not start with a white sheet of paper. The fact that he or she is using a framework imposes some constraints on the possible or allowed



designs. The framework defines some of the basic components of an application and their interfaces and therefore it also specifies the places where concrete physics algorithms and concrete data types will fit in with the rest of the program. The consequences of this are: on one hand, that the physicists designing the algorithms do not have complete freedom in the way algorithms may be implemented; but on the other hand, neither do they need worry about some of the basic functionalities, such as getting end-user options, reporting messages, accessing event and detector data independently of the underlying storage technology, etc. In other words, the framework imposes some constraints in terms of interfaces to basic services, and the interfaces the algorithm itself is implementing towards the rest of the application. The definition of these interfaces establishes the so called "master walls" of the data processing application in which the concrete physics code will be deployed. Besides some general services provided by the framework, this approach also guarantees that later integration will be possible of many small algorithms into a much larger program, for example a reconstruction program. In any case, there is still a lot of room for design creativity when developing physics code within the framework and this is what we want to illustrate in the next sections.

To design a physics algorithm within the framework you need to know very clearly what it should do (the requirements). In particular you need to know the following:

- What is the input data to the algorithm? What is the relationship of these data to other data (e.g. event or detector data)?
- What new data is going to be produced by the algorithm?
- What's the purpose of the algorithm and how is it going function? Document this in terms of mathematical expressions and plain english.<sup>1</sup>
- What does the algorithm need in terms of configuration parameters?
- How can the algorithm be partitioned (structured) into smaller "algorithm chunks" that make it easier to develop (design, code, test) and maintain?
- What data is passed between the different chunks? How do they communicate?
- How do these chunks collaborate together to produce the desired final behaviour? Is there a controlling object? Are they self-organizing? Are they triggered by the existence of some data?
- How is the execution of the algorithm and its performance monitored (messages, histograms, etc.)?
- Who takes the responsibility of bootstrapping the various algorithm chunks.

For didactic purposes we would like to illustrate some of these design considerations using a hypothetical example. Imagine that we would like to design a tracking algorithm based on a Kalman-filter algorithm.

1. Catalan is also acceptable.



### D.3 Analysis Phase

As mentioned before we need to understand in detail what the algorithm is supposed to do before we start designing it and of course before we start producing lines of C++ code. One old technique for that, is to think in terms of data flow diagrams, as illustrated in Figure A.1, where we have tried to decompose the tracking algorithm into various processes or steps.

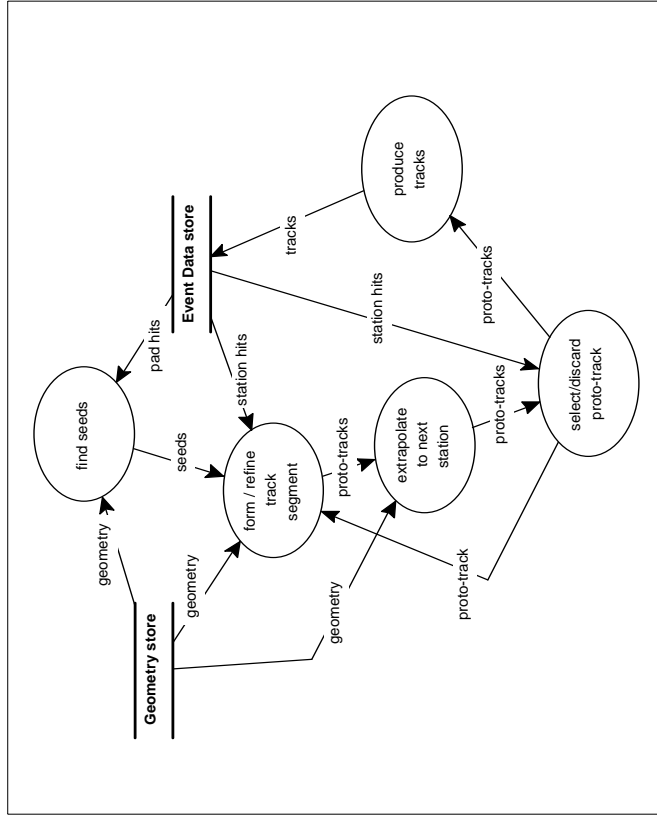


Figure A.1 Hypothetical decomposition of a tracking algorithm based on a Kalman filter using a Data flow Diagram

In the analysis phase we identify the data which is needed as input (event data, geometry data, configuration parameters, etc.) and the data which is produced as output. We also need to think about the intermediate data. Perhaps this data may need to be saved in the persistency store to allow us to run a part of the algorithm without starting always from the beginning.

We need to understand precisely what each of the steps of the algorithm is supposed to do. In case a step becomes too complex we need to sub-divide it into several ones. Writing in plain english and using mathematics whenever possible is extremely useful. The more we understand about what the algorithm has to do the better we are prepared to implement it.



## D.4 Design Phase

We now need to decompose our physics algorithm into one or more Algorithms (as framework components) and define the way in which they will collaborate. After that we need to specify the data types which will be needed by the various Algorithms and their relationships. Then, we need to understand if these new data types will be required to be stored in the persistency store and how they will map to the existing possibilities given by the object persistency technology. This is done by designing the appropriate set of Converters. Finally, we need to identify utility classes which will help to implement the various algorithm chunks.

### D.4.1 Defining Algorithms

Most of the steps of the algorithm have been identified in the analysis phase. We need at this moment to see if those steps can be realized as framework Algorithms. Remember that an Algorithm from the view point of the framework is basically a quite simple interface (initialize, execute, finalize) with a few facilities to access the basic services. In the case of our hypothetical algorithm we could decide to have a "master" Algorithm which will orchestrate the work of a number of sub-Algorithms. This master Algorithm will be also in charge of bootstrapping them. Then, we could have an Algorithm in charge of finding the tracking seeds, plus a set of others, each one associated to a different tracking station in charge of propagating a proto-track to the next station and deciding whether the proto-track needs to be kept or not. Finally, we could introduce another Algorithm in charge of producing the final tracks from the surviving proto-tracks.

It is interesting perhaps in this type of algorithm to distribute parts of the calculations (extrapolations, etc.) to more sophisticated "hits" than just the unintelligent original ones. This could be done by instantiating new data types (clever hits) for each event having references to the original hits. For that, it would be required to have another Algorithm whose role is to prepare these new data objects, see Figure A.2.

The master Algorithm (TrackingAlg) is in charge of setting up the other algorithms and scheduling their execution. It is the only one that has a global view but it does not need to know the details of how the different parts of the algorithm have been implemented. The application manager of the framework only interacts with the master algorithm and does not need to know that in fact the tracking algorithm is implemented by a collaboration of Algorithms.

### D.4.2 Defining Data Objects

The input, output and intermediate data objects need to be specified. Typically, the input and output are specified in a more general way (algorithm independent) and basically are pure data objects. This is because they can be used by a range of different algorithms. We could have various types of tracking algorithm all using the same data as input and producing similar data as output. On the contrary, the intermediate data types can be designed to be very algorithm dependent.

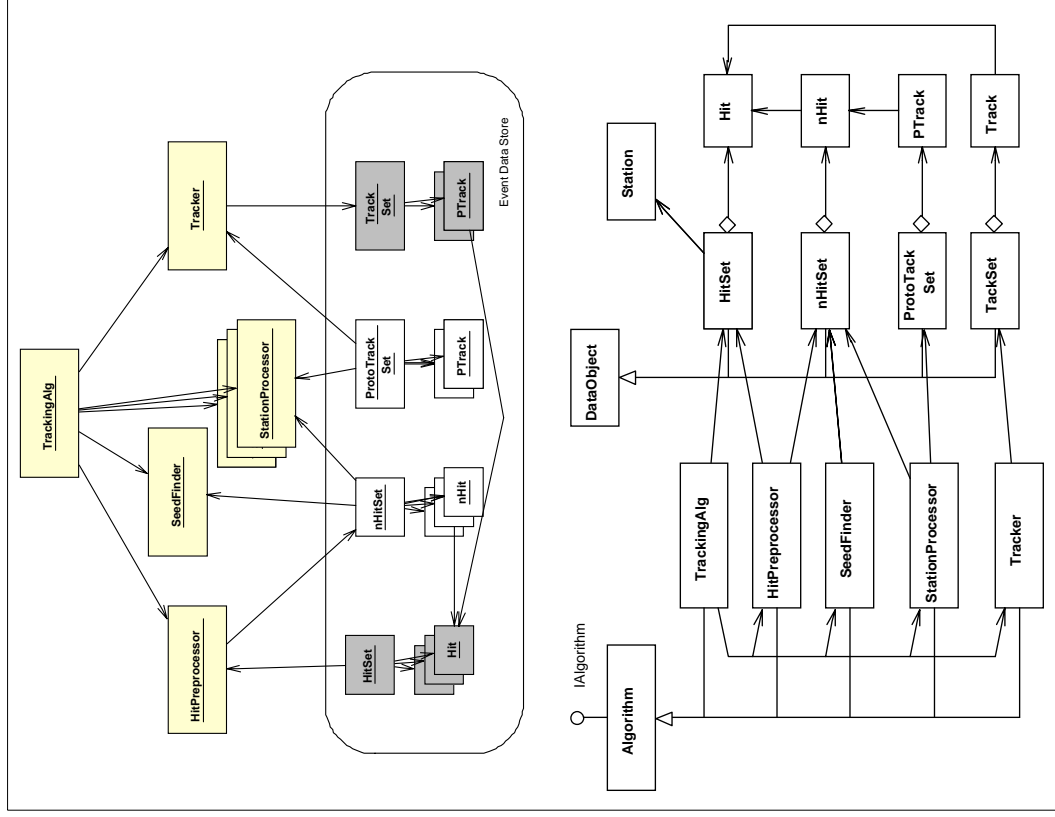


Figure A.2 Object diagram (a) and class diagram (b) showing how the complete tracking algorithm could be decomposed into a set of specific algorithms that collaborate to perform the complete task.





The way we have chosen to communicate between the different Algorithms which constitute our physics algorithm is by using the transient event data store. This allows us to have low coupling between them, but other ways could be envisaged. For instance, we could implement specific methods in the algorithms and allow other "friend" algorithms to use them directly.

Concerning the relationships between data objects, it is strongly discouraged to have links from the input data objects to the newly produced ones (i.e. links from hits to tracks). In the other direction this should not be a problem (i.e. from tracks to constituent hits).

For data types that we would like to save permanently we need to implement a specific Converter. One converter is required for each type of data and each kind of persistence technology that we wish to use. This is not the case for the data types that are used as intermediate data, since these data are completely transient.

### D.4.3 Mathematics and other utilities

It is clear that to implement any algorithm we will need the help of a series of utility classes. Some of these classes are very generic and they can be found in common class libraries. For example the standard template library. Other utilities will be more high energy physics specific, especially in cases like fitting, error treatment, etc. We envisage making as much use of these kinds of utility classes as possible.

Some algorithms or algorithm-parts could be designed in a way that allows them to be reused in other similar physics algorithms. For example, perhaps fitting or clustering algorithms could be designed in a generic way such that they can be used in various concrete algorithms. During design is the moment to identify this kind of re-usable component or to identify existing ones that could be used instead and adapt the design to make possible their usage.

# Index

## Symbols

, 16, 191

## A

AIDA, 185

see Interfaces

Algorithm, 8

Base class, 9, 37

branches, 44

Concrete, 37, 40

Constructor, 39, 40

Declaring properties, 39

Execution, 32, 42

Filters, 44

Finalisation, 32, 43

Initialisation, 31, 39, 41, 43

Nested, 43

sequences, 44

Setting properties, 39

Algorithms

  EventCounter, 45, 113

  Prescaler, 45

  Sequencer, 45

  Application Manager, 10

  instantiation, 26

  ApplicationMgr. See Application Manager

  Architecture, 7

  Associators, 147

  Example, 150

## B

Branches, 44

## C

Casting  
  of DataObjects, 49

Changes  
  in the new release, 16  
  incompatible in release v7, 18  
  incompatible in release v8, 17  
  incompatible in release v9, 16  
  see also Deprecated Features  
  Checklist  
  for implementing algorithms, 43  
Class  
  identifier (CLID), 54  
CLHEP, 185  
Units, 2  
CMT  
  Building libraries with, 183  
CMTPATH  
  to get development version of packages, 22  
Component, 7, 179  
libraries, 179, 181  
Condition objects, 89  
  Creation of, 91  
  update of, 91  
Conditions data  
  Accessing from Detector Element, 66  
Conditions Database, 72  
  Persistence solution, 87  
ContainedObject, 50  
Conventions, 2  
  Coding, 4  
  Naming, 4  
  Units, 2  
  used in this this document, 4  
Converters, 153  
  Graphic (for visualisation), 171  
  SICB Back Converters, 190  
  SICB Converters, 188  
  XML, 83  
CVS  
  password, 19

**D**  
Data  
  Reading, 187  
  Data Store, 47  
  Detector Data, 63  
  Detector data, 63  
  finding objects in, 48, 55  
  Histograms, 93  
  registering objects into, 49  
  DataObject, 9, 47, 49, 50  
  Defining, 53  
  ownership, 50



DECLARE\_ALGORITHM, 180  
DECLARE\_FACTORY\_ENTRIES, 180  
Deprecated Features, 18  
Detector  
  Geometry, 69  
  Materials, 72  
  Structure, 69  
  Detector Data  
    Accessing, 64  
  Detector Description, 61  
    Conditions, 72  
    DTD, 76  
  General features, 69  
  Job options for accessing, 68  
  See also Geometry  
  DetectorElement  
    Customizing the, 84  
    Extending the, 67  
    Using the, 65  
  DTD  
    of LHCb detector description, 76

**E**  
endreq, MsgStream manipulator, 122  
Event Collections, 102  
  Filling, 104  
  Reading Events with, 106  
  Writing, 103  
Event Data Model, 59  
Event Display  
  Panoramix, 172  
EventCounter algorithm. See Algorithms  
Example Application  
  Main program, 26  
  Trace of execution, 27  
Examples  
  Associator, 150  
  distributed with Gaudi, 34  
  HelloWorld, 30  
  Simple Analysis, 36  
  when casting, 49

**F**  
Factory  
  for a concrete algorithm, 40  
Filters, 44  
FORTRAN, 8  
  and shareable libraries, 183  
  How to use in Gaudi, 192



**G**

- GEANT4
  - interfacing to. See GiGa
  - units, 2
- Geometry
  - from SICB, 191
  - Solids, 71
  - Surfaces, 71
- getFactoryEntries, 179
- GiGa, 135
- Guidelines
  - for software packaging, 175

**H**

- HBOOK
  - Constraints on histograms, 94
  - For histogram persistency, 95
  - Limitations on N-tuples, 98, 99, 102
- Histograms
  - data service, 93
- HTL, 185
  - Naming convention for, 4
  - Persistency service, 95
- HTL, 185

**I**

- Inheritance, 37
- Installation
  - of the framework, 13
  - Outside CERN, 22
- Interactive Analysis
  - of N-tuples, 107
- Interface, 7
  - and multiple inheritance, 11
  - Identifier, 11, 177
  - In C++, 11

## Interfaces

- AIDA, 93, 185
- IAlgorithm, 11, 37, 39, 41
- IAlgTool, 142
- IAlignment, 66
- IAppMgrUI, 27
- IAssociator, 148
- IAuditor, 130
- ICalibration, 66
- IConditionInfo, 66
- IConversionSvc, 154
- IConverter, 154, 189
- IDataManagerSvc, 10, 48
- IDataProviderSvc, 10, 47, 48, 64, 97
- IDataProvideSvc, 93
- IDetectorElement, 65
- IFastControl, 66
- IGeometryInfo, 65
- IHistogram1D, 93
- IHistogram2D, 93
- IHistogramSvc, 10, 47, 93
- IncidentListener, 134
- ILVolume, 71
- IMessageSvc, 10
- in Gaudi, 176
- INTupleSvc, 47, 97
- INTupleSvc, 10
- IOpaqueAddress, 155
- IParticlePropertySvc, 123
- IProperty, 10, 27, 37
- IPVolume, 71
- IReadOut, 66
- IRunnable, 12
- ISlowControl, 66
- ISolid, 71
- ISvcLocator, 39
- IToolSvc, 145
- IUserParameter, 66
- IValidity, 65, 66, 71
- IXmlSvc, 87
  - Navigating between, 178
- Introspection, 135

**J**

- Job Options
  - see also Properties
  - Job options, 113



**L**

## Libraries

- Building, 183
  - Building with CMT, 183
  - Component, 179, 181
  - containing FORTRAN code, 183
  - Linker, 182
  - Linux, 19
- LOAD\_FACTORY\_ENTRIES, 181

**M**

## Magnetic Field. See Services

- Message service, 120
- Monitoring
  - of algorithm calls, with the Auditor service, 129
  - statistical, using the Chrono&stat service, 127
- Monte Carlo truth
  - navigation using Associators, 147

**N**

## NAG C, 186

- N-tuples, 97
  - Booking and declaring tags, 99
  - filling, 100
  - Interactive Analysis of, 107
- Limitations imposed by HBOOK, 98, 99, 102
- persistence, 101, 104
- reading, 100
- Service, 97

**O**

- Object Container, 50
  - and STL, 50
- ObjectList, 50
- ObjectVector, 50
- ODBC
  - For N-Tuple analysis, 108

**P**

- Package, 173
  - Internal layout, 175
  - structure of LHCb software, 173
- Packages
  - Dependencies of Gaudi, 173
  - Guidelines, 175
- Panoramix, 172
- Parsers
  - Numerical expressions, 87

**P**

- for N-Tuple analysis, 101
- Persistence
  - of histograms, 95
  - of N-tuples, 101, 104
- Persistent store
  - saving data to, 57
- Pile-up, 193
- Platform
  - Available platforms, 19
- Platforms
  - on which Gaudi is supported, 19
- Prescaler algorithm. See Algorithms
- Problems
- Profiling
  - of execution time, using the Chrono&Stat service, 126
  - of execution time, with the Auditor service, 129
  - of memory usage, with the Auditor service, 129
- Properties
  - Accessing and Modifying, 115
- Python, 161
- Python scripts
  - file extension for, 162

**R**

- Random numbers
  - generating, 131
  - Service, 131
- Reading Data, 187
- Release notes, 13
- Reporting problems, 5
- Retrieval, 145
- ROOT, 186
  - for histogram persistence, 96
  - for N-Tuple analysis, 101, 107

**S**

- Saving data, 57
- Sequencer algorithm. See Algorithms
- Sequences, 44



**S**

Services, 9  
Auditor Service, 129  
Chrono&Stat service, 126  
Histogram data service, 93  
Histogram Persistence Services, 95  
Incident service, 134  
Introspection service, 135  
Job Options service, 113  
Magnetic Field Service, 190  
Message Service, 120  
N-tuples Service, 97  
Particle Properties Service, 123  
Random numbers service, 131  
requesting and accessing, 111  
ToolSvc, 139, 145  
vs. Tools, 139  
XmlCnvSvc, 81  
XmlParserSvc, 82

**SICB**

Converters, 188  
geometry, 191  
SmartDataLocator, 55  
SmartDataPtr, 55  
SmartRef, 56  
SpillOver, 195  
Stripping, 161  
StatusCode, 42

**T**

Tapes  
Reading, 187  
Tools, 139  
Associators, 147  
provided in Gaudi, 147  
vs. Services, 139  
ToolSvc, see Services  
Tutorials, 1

**U**

Units, 2  
Convention, 2

**V**

Visualization, 169  
Volume  
Logical, 71  
Physical, 71

**W**

Windows NT, 19

**X**

XML  
Conversion Service, 81  
Converters, 83  
introduction to, 73  
Parser Service, 82

**Z**

ZEBRA  
COMMON block size, 196

