

AthenaMT: Upgrading the ATLAS Software Framework for the Many-Core World with Multi-Threading

Charles Leggett¹, John Baines², Tomasz Bold³, Paolo Calafiura¹,
Steven Farrell¹, Peter van Gemmeren⁴, David Malon⁴, Elmar Ritsch⁵,
Graeme Stewart⁶, Scott Snyder⁷, Vakhtang Tsulaia¹, Benjamin
Wynne⁸ on behalf of the ATLAS Collaboration

¹Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA

²STFC Rutherford Appleton Laboratory, Harwell Oxford, Didcot OX110QX, United Kingdom

³AGH, University of Science and Technology al. Mickiewicza 30, PL-30-059 Cracow, Poland

⁴Argonne National Laboratory, 9700 S. Cass Ave, Argonne, IL 60439, USA

⁵European Organization for Nuclear Research, CH-1211 Geneva 23, Switzerland

⁶SUPA - School of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom

⁷Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973, USA

⁸SUPA - School of Physics and Astronomy, University of Edinburgh, Edinburgh EH9 3JZ, United Kingdom

E-mail: cgleggett@lbl.gov

Abstract. ATLAS's current software framework, Gaudi/Athena, has been very successful for the experiment in LHC Runs 1 and 2. However, its single threaded design has been recognized for some time to be increasingly problematic as CPUs have increased core counts and decreased available memory per core. Even the multi-process version of Athena, AthenaMP, will not scale to the range of architectures we expect to use beyond Run2.

After concluding a rigorous requirements phase, where many design components were examined in detail, ATLAS has begun the migration to a new data-flow driven, multi-threaded framework, which enables the simultaneous processing of singleton, thread unsafe legacy Algorithms, cloned Algorithms that execute concurrently in their own threads with different Event contexts, and fully re-entrant, thread safe Algorithms.

In this paper we report on the process of modifying the framework to safely process multiple concurrent events in different threads, which entails significant changes in the underlying handling of features such as event and time dependent data, asynchronous callbacks, metadata, integration with the Online High Level Trigger for partial processing in certain regions of interest, concurrent I/O, as well as ensuring thread safety of core services. We also report on upgrading the framework to handle Algorithms that are fully re-entrant.

1. Introduction

ATLAS's[1] framework (Gaudi[2]/Athena[3]) was designed to process serially one event at a time. Limitations of existing and emerging computing technology, as well as the requirements of the ATLAS reconstruction environment, have forced us to examine concurrent, multi-threaded implementations[5].



40 In 2014, the ATLAS *Future Framework Requirements Task Force* created a report that
41 summarized a list of requirements and recommendations for such a framework, and in 2015
42 ATLAS began the process of migrating its software to the new design[6]. For a large experiment
43 like ATLAS, with a massive software code base, a complete rewrite of the software was
44 deemed untenable, and instead, capitalizing on certain fundamental features of the GaudiHive[7]
45 framework, which allow the localization of enforced thread safe code to only a limited set of
46 components, an *evolutionary* migration strategy was developed.

47 The most difficult part of this migration is the requirement that all shared software Services,
48 *i.e.* components that can be accessed concurrently by clients from multiple threads, must not
49 only be thread safe, but must also be able to process requests from different events. These core
50 Services must be fixed before any realistic migration of user analysis algorithms can commence.
51 ATLAS has created an aggressive migration schedule in order to be ready in time for Run 3,
52 and in 2016, we have focused our attention on the migration of core Services. We will detail our
53 efforts in this paper.

54 **2. Enabling Concurrency in Core Services**

55 In order to function properly in AthenaMT, shared Services must be thread safe, and also able
56 to process requests from various clients that may be executing in different concurrent events.
57 These requirements have a broad range of impacts on the migration of the software. Some
58 Services are already event agnostic, and can be made thread safe with simple mutexes or thread
59 safe data structures.

60 Some Services need more intrusive modifications to be able to handle state information that
61 is associated with multiple concurrent events. For example, the MagFieldSvc, which is used to
62 calculate the value of the detector's magnetic field at any given location, was caching localized
63 data internally, in order to speed up sequential requests which are usually for very similar
64 physical locations. Not only was this thread unsafe, but this localization was broken if multiple
65 clients queried the Service simultaneously, requesting information about very different regions.
66 The Service was re-designed to carry a cache object along with each request, localizing the cache
67 to the client, and not the Service, thus maintaining both thread safety and performance benefits.

68 Another example is the THistSvc, which is used to manage histograms and ntuples. Since the
69 objects it manages are identified either by name, or a pointer to the object, multiple concurrent
70 clients that updated these objects could interfere with each other. The Service was upgraded so
71 that clients could either request a shared object, with an enforced locking policy that prevented
72 simultaneous updates to the object, or their own copy of the object that would not be shared.
73 At the end of the job, these copied objects could be automatically merged.

74 And some Services, such as those that manage Asynchronous Data, (*e.g.* the IOVSvc which
75 manages detector Conditions or the GeoModelSvc which deals with detector alignments), need
76 a complete redesign. These will be discussed in further detail below.

77 **3. Concurrent Processing of Asynchronous Data**

78 One of the challenges in the migration process has been the handling of Asynchronous Data, *i.e.*
79 data which can have a lifetime of more than one event. The period of time for which any piece
80 of such data is valid is referred to as an Interval of Validity (IOV). While we do have a solution
81 for managing multiple concurrent Event Stores belonging to different events, Asynchronous data
82 cannot be stored there, as the contents of the Event Store are erased at the end of each event,
83 so a different solution must be found.

84 We can classify Asynchronous Data into two, somewhat interrelated, categories: Conditions,
85 such as high voltages, calibrations, etc., and Detector Geometry and Alignments. Closely related
86 to these are Asynchronous Callbacks (Incidents), which are functions that need to be executed

87 at non-predetermined intervals, such as in response to the opening of a file, or the signaling of
88 the beginning of a new run.

89 *3.1. Conditions*

90 In serial Athena, Conditions were managed by the Interval of Validity Service (IOVSvc). At
91 the start of a job, the IOVSvc is configured to manage a number of objects in an associated
92 Conditions Database, which stores the value of each object for each IOV. At the start of each
93 event, the IOVSvc examines the validity of each registered object. Objects that are no longer
94 valid are re-read from the database, and any required post-processing of the data is performed
95 by an associated callback function. The processed objects are then placed in a conditions store,
96 from whence they can be retrieved by a user Algorithm.

97 This workflow fails when multiple events are processed concurrently. Since only a single
98 instance of the conditions data can be held at any one time in the conditions store, if two
99 events are processed concurrently, with associated conditions data from different IOVs, one
100 will overwrite the other. Furthermore, neither the IOVSvc itself nor any of the conditions
101 callback functions were designed to be thread safe, and since these are shared instances, threading
102 problems are bound to occur. A major rewrite of the entire infrastructure is required.

103 Several different designs for the condition handling were examined, with two key requirements
104 in mind: minimize changes to client code (as there is so much of it), and minimize memory usage
105 (as an overall memory shortage is one of the main reasons we need to use a multi-threaded
106 framework). Designs considered and discarded were the use of a processing barrier, and the use
107 of multiple conditions stores. With a processing barrier the framework would only concurrently
108 process events that had the same set of conditions, draining the event scheduler until all these
109 events had finished processing, then loading a new set of conditions data for the next set of events,
110 before resuming processing events concurrently. The problem with this design is that it assumes
111 that conditions boundaries are infrequent, so that the loss of concurrency when the scheduler is
112 drained and refilled is minimal. On ATLAS, however, Conditions changes can sometimes occur
113 very rapidly, for example as frequently as once per event in the Muon subsystem. This would
114 have the effect of serializing event processing, with complete loss of concurrency. Processing out
115 of order events could also result in a loss of concurrency.

116 The other rejected design used multiple conditions stores, one per concurrent event, in the
117 same manner as the Event Stores are duplicated for each concurrent event. The mechanism
118 by which data is retrieved from the conditions store would be modified, such that clients would
119 associate with the correct Store. Impact on client code would be small - only the conditions data
120 retrieval syntax would need to be updated. However, beyond merely ensuring thread safety of
121 the IOVSvc and the callback functions, there are two significant problems with this design: the
122 memory usage would balloon, as objects would be duplicated between each Store instance; and
123 also the execution of the callback functions that are used to process data would be duplicated,
124 resulting in extra CPU overhead.

125 The chosen solution was to implement a single conditions store in the form of a multi-
126 cache. Instead of holding individual Condition Objects, it holds containers of them, where the
127 elements in each container correspond to individual IOVs. Clients access Condition Objects via
128 smart references, called ConditionHandles, which implement logic to determine which element
129 in any ConditionContainer is appropriate for a given event. The callback functions are migrated
130 to fully- fledged Condition Algorithms, which are managed by the framework like any other
131 Algorithm, but only executed on demand when the Conditions Objects they create need to be
132 updated.

133 One of the fundamental changes in the client code needed for the migration to AthenaMT
134 is that all access to event data must be done via smart references, called DataHandles.
135 DataHandles are declared as member variables of Algorithms, and provide two fundamental

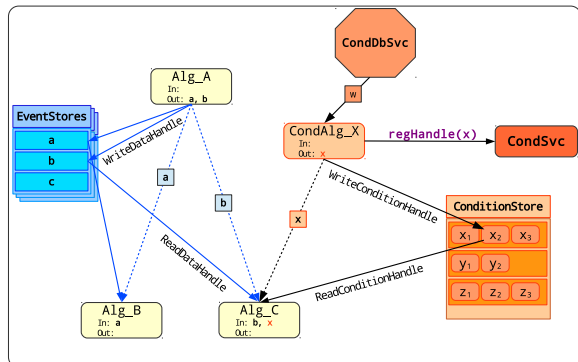


Figure 1. ConditionHandles

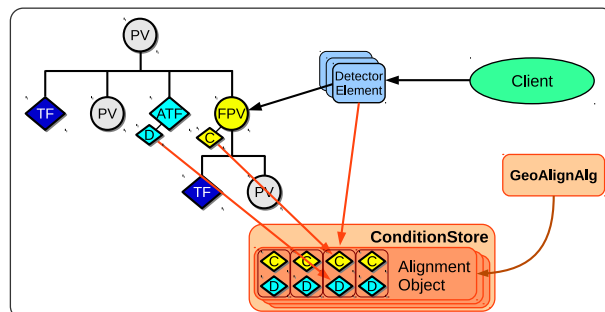


Figure 2. Detector Geometry Alignments

136 functions: to perform the recording and retrieval of event data, and to automatically declare the
 137 data dependencies of the Algorithms to the framework, so that the Algorithms can be executed
 138 by the Scheduler as the data becomes available. We capitalized on the migration to DataHandles
 139 by requiring that all access to Conditions data be done via related ConditionHandles. By
 140 using ConditionHandles in the Condition Algorithms to write data to the conditions store,
 141 the framework solves the problem of Algorithm ordering for us, ensuring that the Condition
 142 Algorithm is executed, and the updated Condition Objects are written to the store before any
 143 downstream Algorithm which needs to use them are executed.

144 When a ConditionHandle is initialized, it will look in the conditions store for its associated
 145 container, identified by a unique key. This container holds a set of objects of the same type and
 146 their associated IOVs. Upon dereferencing, the ConditionHandle will use the current event and
 147 run numbers to look inside this container, and determine what action needs to be taken. At the
 148 start of the event, the Condition Service analyzes the subset of the objects held in the condition
 149 store that have been registered with it at the start of the job by the Condition Algorithms, and
 150 determines which are valid or invalid for the current event. If an object is found to be invalid,
 151 the Condition Algorithm that produces that object will be scheduled. If an object is found to be
 152 valid, then the Scheduler will be informed that this object is present, and placed in the registry
 153 of existing objects. In this case the Condition Algorithm will not execute.

154 When a Condition Algorithm is executed, it queries the Conditions Database for data
 155 corresponding to the current event, as well as its associated IOV, creates the new object for
 156 which it is responsible, and adds a new entry in the ConditionContainer that is associated
 157 with a ConditionHandle (see Fig. 1). When a downstream Algorithm that needs to read
 158 a ConditionHandle from the store is executed, the data is guaranteed to be present. The
 159 ConditionHandle uses the current event number to identify which element in the container is
 160 the appropriate one, and returns its value.

161 By using basic features of the new framework, namely the use of ConditionHandles and data
 162 flow dependencies to automatically schedule Algorithms as needed, we are able to minimize
 163 changes to client code, and delegate the majority of the work to the framework. The use of
 164 collections of Condition Objects inside of a single conditions store allows us to minimize the
 165 total memory footprint.

166 3.2. Detector Geometry and Alignments

167 The detector geometry model used in ATLAS (GeoModel), is a hierarchical tree that is built
 168 from several components (see Fig. 2): a Physical Volume (PV) which are the basic building
 169 blocks; a Transform (TF) that is fixed at construction; and an Alignable Transform (ATF),
 170 which accounts for the movement of the detector component as a function of time, reading

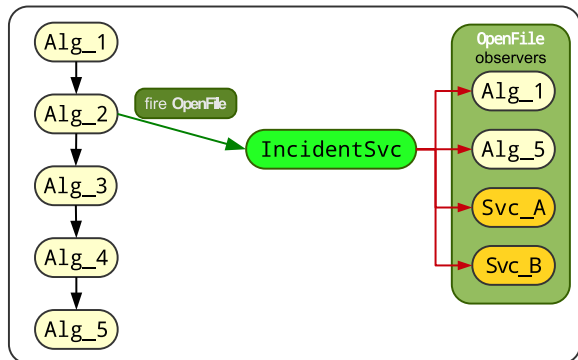


Figure 3. The IncidentSvc in Serial Athena

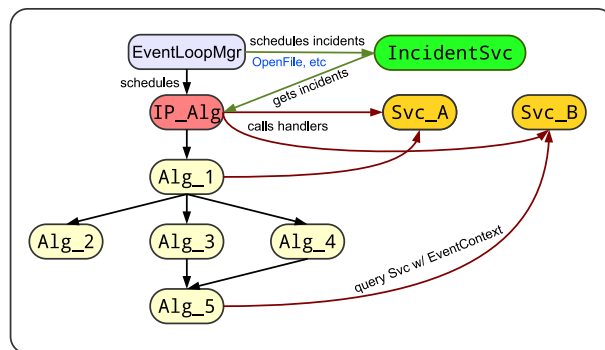


Figure 4. Handling Incidents in AthenaMT

171 Deltas (D) from a database. When a client requests the position of a Detector Element, the Full
 172 Physical Volume (FPV) is assembled, and the position is cached (C). As the detector alignment
 173 changes, new Deltas are read in by the ATF, and the cache held by the FPV is invalidated, until
 174 the position of the element is again requested, recomputed, and cached.

175 When multiple concurrent events are processed, this design will fail, as there is only a single
 176 shared instance of the GeoModel tree, and the ATF and FPV can only keep track of single Delta
 177 or cache at any one time. We can solve this problem in the same way as for the conditions.
 178 The time dependent information (*i.e.* the Deltas and cache) held by the GeoModel is decoupled
 179 from the static entries, and held in a new AlignmentObject located inside the ConditionStore.
 180 The ATF and FPV use ConditionHandles to access this data, and they are updated by a new
 181 GeoAlignAlg which is scheduled on demand by the framework. Clients of the DetectorElements
 182 are entirely blind to this change, and the only code that needs to be modified are base classes
 183 inside the GeoModel structure.

184 4. Asynchronous Callbacks

185 ATLAS uses the Incident Service to execute callback functions at certain well-defined times
 186 following the well-established observer patterns. Clients register interest in certain “Incidents”
 187 with the Service, such as BeginEvent, FileOpen, or EndMetaData. When components fire these
 188 Incidents, execution flow is passed to the IncidentSvc, which triggers the appropriate callback
 189 function in the registered observers (see Fig. 3). There are many issues with this design in
 190 AthenaMT, where there can be multiple instances of any Algorithm, executing simultaneously
 191 in different events. If a cloned Algorithm is an Incident observer, should all instances execute
 192 the callback? What if an instance is currently executing in a different thread? Fixing the design
 193 in a generic way looked to be an impossible task.

194 Instead, we did a study of exactly how Incidents were being fired and used, and discovered that
 195 the vast majority were fired outside the event execution loop (*i.e.* before or after all Algorithms
 196 are executed for one event), and being used to signal discrete state changes, such as BeginEvent.
 197 We realized that we could significantly limit the scope of the IncidentSvc without losing any
 198 functionality. Incidents instead became schedulable (see Fig. 4), where the IncidentSvc would
 199 add special IncidentAlgs at the beginning or end of the event processing loop, which would
 200 interact with event context aware Services to perform the same function as the old Incident
 201 callback functions. Clients would then interact with these Services, passing them the current
 202 event to extract the relevant information.

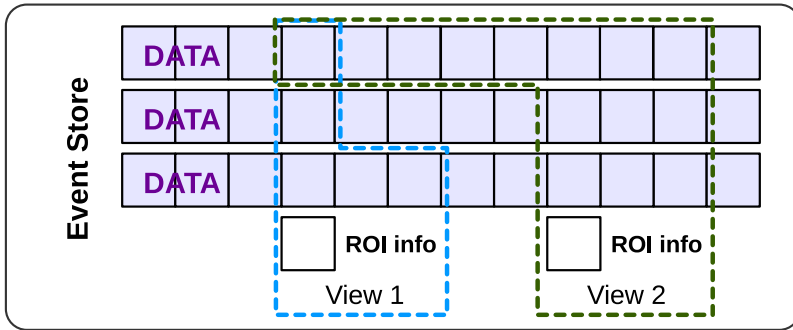


Figure 5. Structure of EventViews and ROIs in the EventStore

203 5. Event Views

204 One of the major requirements put forth by the *Future Framework Requirements Task Force*,
 205 was that the online and offline software be unified, and use the same code base. The High Level
 206 Trigger (HLT), used by the online, functions in a significantly different manner than offline
 207 reconstruction and analysis. In order to maximise performance and perform quick rejection, the
 208 HLT operates independently on geometrical Regions of Interest (ROI). Multiple filtering and
 209 hypothesis Algorithms are applied to each ROI in chains, with the ability to abort a chain if an
 210 unsatisfactory conclusion is reached. Each of these Algorithms sees only the data relevant to
 211 the ROI that it is operating upon, unlike the offline analysis Algorithms which operate on the
 212 entire detector. But since the same Algorithms are to be shared by the online and offline, the
 213 framework must be able to present the data to the Algorithms in a ROI or detector agnostic
 214 manner.

215 We were able to make use of the fact that Algorithms access Event data via smart
 216 DataHandles, and having the framework modify the data presented to the Algorithm internally
 217 within the DataHandle itself. This was done by implementing an EventView class, that can be
 218 used interchangeably with the whole Event Store. It has the same interface as the Event Store,
 219 but adds a Data Object that describes the corresponding ROI. In the case of online processing,
 220 each Event View is populated with the data corresponding to a single ROI, and the DataHandle
 221 of a particular Algorithm is updated to point to that View before execution by the framework.

222 By this mechanism, we are not only able to use the same algorithmic code for online and
 223 offline analysis, but it also offers us the possibility of increasing concurrency in the offline by
 224 doing sub-Algorithmic parallelism in various ROIs.

225 6. Re-Entrant Algorithms

226 One of the features of AthenaMT which frees Algorithm authors from having to implement
 227 thread safe code, is that any single instance of an Algorithm is guaranteed to process an event in
 228 its entirety in only a single thread at a time. Concurrent event processing is achieved by cloning
 229 Algorithms, *i.e.* multiple instances of the same Algorithm are created by the framework, so if
 230 the same Algorithm is scheduled simultaneously in different Events, each Event gets its own
 231 copy of the Algorithm. This frees Algorithm authors from having to worry about most forms of
 232 thread safety. Thread hostile behavior, however, such as the use of global statics, must still be
 233 avoided.

234 The downside of this cloning mechanism is that memory usage will increase as each Algorithm
 235 is copied. While the user can set the number of copies of each Algorithm during job configuration,
 236 and limit the cloning, this comes at the expense of limiting possible concurrency. There is always
 237 a trade off between performance and memory usage.

238 Because of this, we decided to add a new type of Algorithm - a fully re-entrant Algorithm.
 239 For these Algorithms, the framework will only create one instance, and the same instance may

240 be executed by the scheduler concurrently in multiple events. This eliminates the memory
241 overhead of multiply cloned Algorithms. The re-entrant Algorithms must be made fully thread
242 safe, and also stateless. In order to facilitate this design, the signature of the Algorithm's
243 event process method had been made *const*, and the EventContext, a class which is used to
244 provide information about the currently executing event, is explicitly passed into the method.
245 Furthermore, the name of the base class of the Algorithm has been changed, to ensure that it is
246 not used accidentally.

247 Having re-entrant Algorithms available gives us greater flexibility in designing Algorithms.
248 While it is unlikely that the majority of existing Algorithms will be migrated to a re-entrant
249 design, it is recommended to users that they first attempt to write new Algorithms in the re-
250 entrant fashion. Older Algorithms will be made re-entrant on a case by case and as-needed
251 basis, in all likelihood by experts who understand the complexity of re-entrant designs.

252 7. Conclusions

253 ATLAS has begun the migration of core framework elements to function in AthenaMT. While in
254 some cases this is a relatively straight forward task, it often requires significant design changes
255 beyond mere thread safety. Some redesigns were able to preserve the full functionality of the
256 serial versions, but in other cases it became apparent that the serial version was completely
257 incompatible with a concurrent environment. In these cases we examined the real use of the
258 code, and limited the re-design to reproduce the actual use cases.

259 We have made design choices that minimized alterations to client code, due to its enormous
260 volume. By leveraging on existing features of the framework, such as DataHandles, data
261 dependency hierarchies of Algorithms, and the structure of the Scheduler itself, we have been able
262 to vastly simplify the task. We anticipate on-schedule finalization of design, and implementation
263 of essential core Services by the end of 2016, with full support of multi-threaded concurrency
264 by the end of 2017. We already have production level ATLAS Geant4 simulations running in
265 multi-threaded environment on the Knights Landing platform[8].

266 Changes to Algorithmic client code that use these framework elements are also underway. We
267 believe that, for the most part, there is a relatively straight forward recipe for this migration, but
268 it will nevertheless require a significant amount of manpower to effectuate. The broad migration
269 of ATLAS's Algorithm client code to function in AthenaMT will take place in 2017.

270 References

- 271 [1] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider," *JINST* **3**, S08003
272 (2008).
- 273 [2] G. Barrand, I. Belyaev, P. Binko, M. Cattaneo, R. Chytracek, G. Corti, M. Frank and G. Gracia *et al.*, "GAUDI
274 - A software architecture and framework for building HEP data processing applications," *Comput. Phys.*
275 *Commun.* **140**, 45 (2001).
- 276 [3] P. Calafiura, W. Lavrijsen, C. Leggett, M. Marino and D. Quarrie, "The Athena control framework in
277 production, new developments and lessons learned," *CHEP 2004 Conf. Proc.* **C04-09-27** pp 456-458 (2005)
- 278 [4] Binet S *et al.*, 2012 Multicore in production: Advantages and limits of the multiprocess approach in the
279 ATLAS experiment *J. Phys.: Conf. Series* **368** 012018 (ACAT2011 proceedings)
- 280 [5] P. Calafiura, W. Lampl, C. Leggett, D. Malon, G. Stewart and B. Wynne, "Development of a Next Generation
281 Concurrent Framework for the ATLAS Experiment," *J. Phys. Conf. Ser.* **664**, no. 7, 072031 (2015).
282 doi:10.1088/1742-6596/664/7/072031
- 283 [6] G. A. Stewart *et al.*, "Multi-threaded software framework development for the ATLAS experiment," *J. Phys.*
284 *Conf. Ser.* **762**, no. 1, 012024 (2016). doi:10.1088/1742-6596/762/1/012024
- 285 [7] M. Clemencic, B. Hegner, P. Mato and D. Piparo, "Introducing concurrency in the Gaudi data processing
286 framework," *J. Phys. Conf. Ser.* **513**, 022013 (2014).
- 287 [8] S. Farrell, P. Calafiura, C. Leggett, V. Tsulaia, A. Dotti *et al.*, "Multi-threaded ATLAS Simulation on Intel
288 Knights Landing Processors," *Proceedings of the CHEP 2016 conference J. Phys.: Conf. Ser.*