

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Final Report of the Simulation Optimization Task Force

Permalink

<https://escholarship.org/uc/item/67t404rc>

Authors

Rimoldi, A.
ATLAS Collaboration

Publication Date

2009-04-09

Final Report of the Simulation Optimization Task Force

A. Rimoldi¹, T. Carli², A. Dell'Acqua², D. Froidevaux², F. Gianotti²,
C. Guyot^{2,3}, I. Hinchliffe⁴, K. Jakobs⁵, Z. Marshall^{6,7}, A. Nisatis,
D. Quarrie⁴, G. Unal², C. Young⁹


This work was supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.



ATLAS NOTE



December 15, 2008

Final Report of the Simulation Optimization Task Force

A. Rimoldi¹, T. Carli², A. Dell'Acqua², D. Froidevaux², F. Gianotti²,
C. Guyot^{2,3}, I. Hinchliffe⁴, K. Jakobs⁵, Z. Marshall^{6,7}, A. Nisati⁸,
D. Quarrie⁴, G. Unal², C. Young⁹

¹*INFN Pavia and Università di Pavia, Dipartimento di Fisica Nucleare e Teorica,
Via Bassi 6, IT-27100 Pavia, Italy*

²*CERN, CH - 1211 Geneva 23, Switzerland*

³*CEA, DSM/DAPNIA, Centre d'Etudes de Saclay, FR - 91191 Gif-sur-Yvette, France*

⁴*Lawrence Berkeley National Laboratory and University of California, Physics Division,
MS50B-6227, 1 Cyclotron Road, Berkeley, CA 94720, United States of America*

⁵*Physikalisches Institut, Universitaet Freiburg, Hermann-Herder Str. 3,
D - 79104 Freiburg i.Br., Germany*

⁶*Columbia University, Nevis Laboratory, 136 So. Broadway, Irvington,
NY 10533, United States of America*

⁷*California Institute of Technology, Physics Department, Pasadena, CA 91125, United States of America*

⁸*INFN Roma I and Università La Sapienza, Dipartimento di Fisica, Piazzale A. Moro 2,
IT-00185 Roma, Italy*

⁹*Stanford Linear Accelerator Center, 2575 Sand Hill Road, Menlo Park, CA 94025, United States of America*

Abstract

This is the final report of the ATLAS Simulation Optimization Task Force, established in June of 2007. This note justifies the selected Geant4 version, physics list, and range cuts to be used by the default ATLAS simulation for initial data taking and beyond. The current status of several projects, including detector description, simulation validation, studies of additional Geant4 parameters, and cavern background, are reported.



1 Introduction

The Simulation Optimization Task Force was set up by the TOB in June of 2007 to optimize the GEANT4-based simulation [1] for ATLAS in terms of physics content and technical performance, to recommend a baseline GEANT4 release and set of parameters, and to address long-term validation and strategy issues. The task force has met every two to four weeks since, collecting input from and helping to coordinate the efforts of the relevant detector, performance, and physics groups. The first report of the task force was produced earlier this year [2].

Since the previous report, several new studies have been undertaken, and several issues have been resolved. A new version of GEANT4 has been adopted for Athena release 14.4.0 [3]. Optimizations of disk space consumption and truth storage are in their final stages.

2 Geant4 Version

Until November 2008, simulation production was done with GEANT4 version 8.3 with official patch #2 and three private patches: backported support for boundary represented solids, a fix for a bug in G4Tubs that was causing aborted events during production, and a fix for a bug in G4PropagatorInField that was causing rare crashes during production. After these private patches were introduced, the crash rate for simulation jobs during production went essentially to zero. No GEANT4-related bugs have been reported for several months.

Owing to the delay of the LHC, the decision was made to migrate to GEANT4 version 9 somewhat earlier than was originally anticipated. For Athena release 14.4.0 (November 2008), ATLAS has adopted GEANT4 version 9.1 with official patch #3 and three private patches: a fix to G4UrbanMscModel which caused very long events during simulation with LUCID and fixes to G4ExtrudedSolid and G4TriangularFacet which may have been causing unfaithful construction of some solids.

The new version of GEANT4 has been validated in a computing sense and has not crashed in running during private and limited automatic testing. Table 1 shows ratios of simulation time per event for various processes in several recent GEANT4 releases to that of the production version of GEANT4 8.3. The errors associated with this type of benchmarking are discussed further in Section 8, but are generally $\sim 5\%$. It appears that the newest releases all provide a 10% performance improvement relative to the current production releases. It is not yet possible to ascertain the stability of the new GEANT4 production release. This early adoption, however, allows a normal validation cycle prior to the beginning of major production in 2009.

Process	G4.9.1.p2	G4.9.1.p3	G4.9.2.beta	G4.9.1.p3.a1
J2	0.94	1.02	0.90	0.92
$H(130) \rightarrow ZZ \rightarrow 4\ell$	0.95	1.07	0.93	0.93
Minbias	0.90	0.93	0.85	0.89
SU3 SUSY	0.93	1.12	0.92	1.03
$Z \rightarrow \mu^+ \mu^-$	0.92	0.99	0.88	0.93
$Z \rightarrow e^+ e^-$	0.92	1.04	0.89	0.94
$Z \rightarrow \tau^+ \tau^-$	0.96	1.05	0.95	0.95

Table 1: The ratio of simulation time for various recent GEANT4 releases to that of the production version of GEANT4 8.3. A $\sim 10\%$ improvement can be seen in most of the releases. Errors on the ratios are $\sim 5\%$. For each release, “pN” indicates official patch N and “a1” indicates the patches to G4UrbanMscModel, G4ExtrudedSolid, and G4TriangularFacet described in the text.

The beta release of GEANT4 9.2 was tested successfully in Athena, but a final release is not due until mid-December 2008. There is little time for validating the release before the beginning of production in 2009, so the move to GEANT4 9.2 will be delayed until there is a gap in the production schedule or there is a convincing physics argument that can be made for the migration. At the moment, however, the new physics lists available in GEANT4 9.2 do not show any significant improvement in agreement with the test beam data relative to those in 9.1.

3 Core Simulation Software Updates

A simulation software infrastructure update is underway. A migration nightly is in use, with about 30 packages already modified, and the migration is targeted for Athena release 15. There are several longstanding issues being addressed, including the elimination of a great deal of obsolete code and the migration to new header files in several core packages.

Several optimizations of simulation code have also been undertaken. Most importantly, these include the removal of string comparisons throughout the core and detector packages. Simple changes to using pointers for particle types have already been implemented for Athena release 14.4 and 14.5. Modifications to check process type using an integer rather than a string have also been made in preparation for Athena release 14.5. Since these comparisons are executed at every step, their removal is expected to benefit the simulation time noticeably. It is also critical to understand what can be achieved with these “free” optimizations before more lossy optimizations are considered.

During each simulation job, the detector description is first built in GeoModel format. It is then translated into GEANT4 format, and all subsequent simulation relies only upon the GEANT4 description of the geometry. It is possible, therefore, to release the GeoModel geometry from memory once it has been translated. This is expected to save approximately 100 MB of memory once it is complete. The memory release is on hold, pending some modifications to GeoModel itself that may be available by Athena release 15.

4 File Size and Storage

Output hit file sizes increase 75% after the transition from the QGSP_EMV physics list to the QGSP_BERT physics list early in 2008. In QGSP_BERT, approximately three times as many thermal neutrons are produced, and each propagates approximately three times further than it did during simulation with the QGSP_EMV physics list. To reduce the computing impact of these neutrons, a neutron time cut was added to the simulation that removed all neutrons 150 ns after the primary interaction. The cut was chosen so as to not affect calorimeter response. The neutron cut has been approved by the physics validation group and is now a part of the standard simulation. Several observables were tested and remained unchanged. Output files produced with the QGSP_BERT physics list and the neutron cut enabled are approximately the same size as those produced with the QGSP_EMV physics list.

In the meantime, the hit storage in the liquid argon calorimetry and tile calorimeter has been improved. The time binning in the liquid argon calorimetry was optimized to reduce disk space without affecting the reconstruction. As a result of these improvements, the removal of the neutron cut now only results in a 30% increase in disk consumption. These optimizations may protect us from later modifications to physics models in GEANT4.

The storage of hits in the silicon detectors was also optimized in Athena release 14.2. Rather than storing individual hits in the silicon, hits are chained together along the path of the minimally ionizing particle traversing the silicon. In this way, a great deal of duplicate information can be removed from the subsequent hits, and after some additional optimization of information storage total disk space requirements are reduced by as much as a factor of two.

After these optimizations, the breakdown of output hit files is shown in Table 2 for the average of 50 $t\bar{t}$ events. Calibration hits are optionally and are generally only included in production of Monte Carlo samples to be used for calorimeter calibration. The hits of the transition radiation detector by far dominate the file size. The optimization of this container has been undertaken and should be completed soon. Because disk space is currently the limiting factor in Monte Carlo production, optimization of containers and reduction of file size will be an ongoing effort.

Collection Name	Size [kb/event]	Percentage of File
Silicon pixel tracker	82	4%
Silicon strip tracker	356	16%
Transition radiation tracker	921	46%
Electromagnetic Barrel Calorimeter	89	4%
Electromagnetic Endcap Calorimeter	104	5%
Hadronic Barrel Calorimeter	29	1%
Hadronic Endcap Calorimeter	22	1%
Forward Calorimeter	42	2%
Calorimeter calibration hits	243	12%
Muon system (all collections)	3	<1%
Truth (all collections)	134	7%
Total	1987	100%

Table 2: Hit collection size, in kb per event, by subdetector. The average was taken of 50 simulated $t\bar{t}$ events. Calorimeter calibration hits are hits in the dead material of the calorimeters stored for studying Monte Carlo-based calorimeter calibration schemes.

5 Truth Optimization

In order to study fake rates in very low transverse momentum tracks in the inner detector, a revision of the truth storage schemes was undertaken. The modifications were undertaken when a few seemingly good tracks at moderate energies (500 MeV to 1 GeV) were found in $t\bar{t}$ events at a rate of about 0.5 Hz. These tracks had no association with truth and were, therefore, classified as fakes. The original requests of the Monte Carlo Truth Task Force [4] had been correctly implemented. Several modifications to the strategies were made as follows:

- Compton scattering interactions are stored in the same way as ionization or bremsstrahlung.
- Rather than cutting on total particle energy, cuts are made on particle transverse momentum.
- Hadronic interactions are checked for high transverse momentum secondaries.
- Cuts are lowered for all interactions to 100 MeV of p_T .

These cuts are sufficient for most tracking performance studies, however for studies of low energy (~ 100 MeV) tracks in minimum bias events the cuts may need to be further reduced. All cuts have been made configurable, so such a modification can be made only for limited production if desirable. The effect of these modifications on file sizes from the hit file through to the AOD was examined and was found to be negligible. The modifications were validated using Athena release 14.2.24.1 and will shortly become the default for production.

6 Steppers and Parameters

A stepper in GEANT4 is the code for calculating the next position of a charged particle in a magnetic field. Such steppers numerically solve the ordinary differential equation for motion, limiting the errors on the calculation based on parameters provided by the user. Currently the stepping parameters used by ATLAS are the GEANT4 defaults in the calorimetry and muon systems and approximately ten times tighter parameters in the inner detector. No step is allowed by GEANT4 to cross more than one volume, and the error on the step is in fact a bias that accumulates as the particle crosses more volume boundaries.

As a part of the core software migration and the migration to GEANT4 9.1, a new “stepper dispatcher” has been implemented. The dispatcher allows selection of the stepper parameters according to the properties of the step. Such a dispatcher allows very relaxed stepping of electrons through the calorimeter while at the same time tracking very carefully high energy muons through the inner detector and calorimeter. Thus the accuracy of muon tracks can be improved without harming the overall speed of the simulation. The dispatcher can change the type of stepper and the error parameters used by the stepper. It also has a parameter to control the distance over which the magnetic field in ATLAS can be considered uniform, a good approximation for a step that is only a few tens of microns long. More time is spent accessing the magnetic field map during simulation than in any other method (up to 20% of the total simulation time). By using a lower-order stepper, the number of values of the magnetic field requested for each step can be reduced from four to two, and by considering the magnetic field uniform over the length of the step only one value of the magnetic field is required.

Representatives from each subdetector group are now studying the effect of these stepping parameters on charged particles propagating up to and through their detector, so that a first set of parameters can be tested for computing performance improvement. It is expected to relax the default parameters and only tighten them on the request of the subdetector, rather than having a tight default which can be loosened by request.

7 Detector Description

The Simulation Optimization Task Force has continued its oversight of detector description and conditions modifications in simulation and digitization. One critical problem of late has been the dissemination of information regarding modifications to the detector geometry and the contents of new geometry versions. To address the problem, each subdetector group was asked to report at the biweekly meetings on any recent or anticipated updates to their geometry. This also provided a forum for discussions of necessary validation samples and time scales for any geometry versions that were to be used for production. For the time being, magnetic field maps are coupled directly to geometry tags, and a new geometry tag is required any time the magnetic field map must be updated. Therefore, the magnetic field group was also asked to regularly report any changes to the maps, and the validation of the newest (e.g. vertically offset) maps was discussed. The magnetic field map is expected to be decoupled from geometry tags with Athena release 15.

Almost every subdetector had at least minor updates in the last few months, but the most noteworthy changes were made to the muon geometry. Cutouts were added for many chambers that until now had unphysical overlaps with support structures and dead materials. Several muon chambers were flipped to reproduce their as-installed position. The most significant undertaking was a package for translating XML geometry descriptions into GeoModel geometries, prepared for the muon system’s dead material. The translation package was validated by declashing (i.e. removing volume overlaps) and studying the support structures of the end cap wheels. Now its use is being extended to the barrel region, where the feet and toroid support structure will benefit greatly from a more realistic description.

With the migration to GEANT4 release 9.1, it is now possible to add parallel geometry navigation

and scoring volumes to the simulation. Scoring volumes can be placed around the muon chambers, and neutrons and photons traversing them can be recorded and removed during simulation of a typical event. These particles comprise the “cavern background” which until now has been taken from a standalone GEANT3/GCALOR simulation. By adding scoring volumes to the standard ATLAS geometry it is possible to regenerate cavern background events with consistent simulation versions, test new physics lists and options for their ability to describe neutron propagation, and generate considerably higher statistics for the cavern background samples than are currently available. The implementation of the scoring volumes is not expected for the immediate future, since their use will require considerable testing and validation.

8 Computing Benchmarking

One of the longstanding issues in the task force was the proper benchmarking of the computing performance of ATLAS simulation in each Athena release. Although the current samples are sufficiently diverse to provide a realistic picture of the overall resource requirements per event, several difficulties have been encountered in making fair comparisons between subsequent tests. Attempting to time jobs in the CERN batch system, variations with an RMS of 5-10% prohibit any strong statements about changes in computing time per event.

It has recently been observed that private running on a quiet benchmarking machine is stable to better than 1% from one run to the next provided *taskset* is used for the Athena jobs. *Taskset* locks a process to a single core of a computer, whereas previously Athena jobs would occasionally jump from one core to another, losing any cached information. Now that stability in a single setup has been achieved, efforts are underway to understand the differences between the different setups so that eventually stable timing results can be produced even in the CERN batch system. Until then, either benchmark jobs must be run privately, which itself is an extremely time-consuming process, or there will be a 5-10% error on all reported simulation times per event.

9 Status of the Mandate

As this is the final report of the Simulation Optimization Task Force, we will review here the mandate of the group and the progress made with respect to the intended path. The overall goal of the task force, to “optimize the G4-based simulation in terms of physics content and technical performance, recommend baseline releases and parameters, address long-term validation and strategy issues,” we believe has been accomplished with great success.

9.1 Phase I

- The baseline GEANT4 release was chosen to be version 8, and patches have continued on that version up to the present. Range cuts, physics lists, and particle cuts have all been defined by the group and validated.
- Urgent technical issues were addressed in a systematic way. The near-zero crash rate of simulation in production is a direct reflection of this effort. Geometry overlaps were addressed with the help of subdetector experts.
- An open dialogue was continued with the subdetector experts to ensure that their simulation needs were being addressed. The description of the dead material in the muon system presented the most pressing challenge during the early work of the group.

- Optimization studies were pursued in several directions. Disk space consumption optimization became the most urgent challenge when it became clear that it was the limiting factor for Monte Carlo production. The disk consumption was kept stable or reduced slightly, and several avenues for future improvement were identified.
- A consistent approach to maintaining the test beam simulation was agreed upon. Barrel and end-cap combined test beam studies will be important to the understanding of various aspects of the calibration of the calorimeters. The barrel CTB also includes the only useful test beam data for incident particle energies between 2 and 9 GeV. Comparisons of Monte Carlo to CTB data were crucial to the optimization of the physics performance of the simulation, and the comparisons guided many of the task force's parameter choices. Support of the CTB reconstruction outside of the calorimetry was dropped after release 12.0.95. However, the simulation for the barrel and end-cap CTB should continue to be maintained, even if the inner detector and muon system data will not be reconstructed in future analyses. Also, the standalone test beams of the individual ATLAS calorimeters will be maintained for the sake of continued comparisons with new physics lists and cuts in GEANT4, for example, for as long as possible.
- Although collisions did not play a role in physics validation of the Monte Carlo, some cosmic ray data has been examined [5] and detector conditions (e.g. dead channels and detector misalignments) from real data taking runs were examined.

9.2 Phase II

- The baseline GEANT4 release for future production has been identified as version 9.1, and parameters for the production have been selected. All are consistent with those chosen during Phase I.
- Optimization studies have continued with the stepper dispatcher and code optimization most recently, and a clear path has been defined for future work.
- Monte Carlo truth was reviewed and the strategies were revised. Cavern background was examined, and an approach for regenerating the events when the manpower becomes available has been defined. The status of pile-up and data overlay was closely monitored, although no intervention of any kind was necessary. Problems with production of beam gas and beam halo events were quickly resolved, and at the present time their production is straightforward.
- Changes in the geometry to be used for simulation were constantly monitored and discussed.
- The faster simulations (including shower libraries) have been validated, and large scale samples will be provided to the physics groups for closer examination prior to large-scale production. The Simulation Strategy Task Force helped assure the completeness of the benchmarking samples being used [6].
- The group gave a report and met during almost every Trigger and Physics / Collaboration Week and Software Week.

9.3 Unfinished Items

In our view, the key unfinished items we leave for future simulation groups to address are (in no particular order):

- Validation of GEANT4 version 9

- Completion and validation of the muon system’s barrel dead material description
- TRT hit container size reduction
- Validation of scoring volumes for GEANT4-based cavern background simulation
- Completion of the simulation core software migration
- Development of an accurate benchmarking scheme
- GeoModel memory release
- Validation of the new stepper dispatcher
- Completion of the simulation of noisy and dead channels

9.4 Recommendations for the Future

We had the benefit of very active validation and production coordinators, but in the original composition of the task force there was no obvious connection to these two groups. In order to get rapid feedback on high-priority issues and to help prioritize efforts, future simulation groups should have direct contact with the physics validation group. The mandate would have been impossible were it not for this support.

A stronger link could have been maintained with the computing performance group. This group, which deals with computing optimization from a technical point of view, has dealt mainly with the reconstruction. Although the simulation has benefited from some of their implemented improvements, it would have been useful to address some simulation-specific issues, or at least determine what exactly the most important issues to the computing performance of the simulation are.

This group had some limited connection with the fast simulations. The Fast G4 simulation (shower libraries approach) relies on GEANT4 for the bulk of the simulation. It was, therefore, more naturally a part of this groups work than was Atlfast-I or Atlfast-II. There must be strong interaction, however, between the strategic approach to simulation of all samples and the process of optimization of the GEANT4 simulation. If more groups rely on GEANT4 for simulation, there is considerable pressure to improve the speed of the simulation. On the other hand, if most groups can rely on some flavor of fast simulation and only use GEANT4 simulation for small statistics tests or detailed validation, greater allowances can be made in favor of physics performance (as was the case with the selection of the QGSP_BERT physics list). Simulation optimization can only be effective when it is done with an understanding of the balance struck between fast and GEANT4 simulation¹⁾.

10 Summary and Conclusions

In the 18 months since its inception, the ATLAS Simulation Optimization Task Force has collected input from subdetector groups and performance groups on a variety of subjects. It has selected a version of GEANT4, a physics list, range cuts, neutron time cuts, and neutrino cuts for the experiment’s default GEANT4 simulation production. It has helped to study the computing validation of the simulation. The various subdetectors have also given reports on their progress in finalizing their geometry and dead material descriptions. The pile-up and digitization groups have given frequent status reports and have continued to make rapid progress towards realistic and robust conditions.

¹⁾A more detailed discussion of the performance of full and fast simulation can be found in [7].

Several studies are ongoing and will be continued past the term of the task force. Stepper dispatcher studies are ongoing with the core software migration proposed for Athena release 15. The recent migration to GEANT4 9.1 for Athena release 14.4 will allow additional studies of parallel geometries and the regeneration of cavern background events with the current GEANT4 simulation and detector description.

Although the Simulation Optimization Task Force is ending its term, having successfully completed its mandate, there remain several open and pending studies that should be continued by the subdetector, performance, and computing groups. The optimization of the Monte Carlo is an ongoing process that will be critical in the coming years.

11 Acknowledgements

The authors would like to thank all those who contributed to the group's success, including: Makoto Asai, Ketevi Assamagan, Sasha Baranov, Andrew Beddall, Joe Boudreau, John Chapman, Dave Charlton, Laurent Chevalier, Marina Cobal, Lidia Dell'Asta, Andrea Di Simone, Rob Duxfield, Grant Gorfine, Andreas Hoecker, Andrei Kiryunin, Alexander Khodinov, Oliver Kortner, Thorsten Kuhl, Guilherme Lima, Bill Lockman, Pawel Majewski, Frank Paige, Troels Petersen, Antonello Sbrizzi, Margar Simonyan, Sanya Solodkov, Vakho Tsulaia, Wouter Verkerke, Iacopo Vivarelli, and Pat Ward.

References

- [1] J. Allison et al., IEEE Transactions on Nuclear Science **53** (2006) 270–278.
- [2] A. Rimoldi et al., ATL-SOFT-PUB-2008-002 (2008).
- [3] Athena Core Software, see <http://atlas-computing.web.cern.ch/atlas-computing/packages/athenaCore/athenaCore.php>.
- [4] F. Akesson et al., ATL-SOFT-INT-2007-002 (2007).
- [5] M. Cooke et al., ATL-LARG-PUB-2007-013 (2007).
- [6] M. Cobal et al., ATL-COM-SOFT-2008-001 (2008).
- [7] K. Assamagan et al., (2009), To be submitted to JINST.