

ARTS Project Proposal

MacOSX@LHC

Submitted by the Laboratoire de l'Accélérateur Linéaire

Sébastien Binet, PhD, ATLAS collaboration

Charles Loomis, PhD, EGEE-III

Guy Barrant, PhD, LHCb collaboration

10 novembre 2009

1. Summary

Proposal

Scientists carrying out High Energy Physics (HEP) experiments are increasingly using MacOS X as their daily personal computing platform. Despite that, MacOS X is not currently used for mainstream scientific data processing. As a result, the software suites of the main HEP experiments do not run on the physicist's laptop or desktop, preventing those physicists from using their personal computer for program development, data analysis preparation, and application tuning. We propose to take advantage of MacOS X principal strengths (Unix underpinning, advanced features such as LLVM, GCD support, OpenCL support, and graphics capabilities) to make it a supported platform in two of the main HEP experiments.

Expertise at LAL

The Linear Accelerator Laboratory (LAL) is deeply involved, on one hand, in the development of the large software infrastructures of two of the next generation experiments at CERN, and on the other hand, in the deployment of the worldwide grid of computers which will handle the data management and analysis of these experiments. It has thus acquired a recognized expertise in the fields of multithreaded programming, grid middleware deployment, and particle physics detector data visualization, among others.

Deliverables

This project proposes to:

- Port to the MacOS X platform the part of the grid middleware that will allow the Macintosh to act as a worker node in the European Computing Grid EGEE.
- Make the ATLAS and LHCb software frameworks run on the Macintosh.
- Use the technologies offered by Snow Leopard to develop the next generation of these software frameworks in order to take advantage of the potential computing power of the modern machines, namely the many cores and GPUs, that they cannot leverage in their present implementation.
- Adapt the event display and analysis programs to run on the image wall (WILD project) developed by the LRI in order to yield to the HEP community a new way of analyzing its data.

People behind this proposal

Sébastien Binet, after defending a PhD in High Energy Physics in 2006 on jet calibration studies in Atlas, spent two years at LBL, Berkeley, CA, working on performance monitoring aspects of the Atlas software framework, and parallelization of that framework using forked sub-processes. He's now working at LAL on reconstruction integration issues, performance monitoring and parallelization aspects of the Atlas software framework.

Charles Loomis is a research engineer at LAL who has been working with grid and distributed computing technologies since 2001. He has held critical posts in each of the EGEE series of projects (European DataGrid, EGEE, EGEE-II, and EGEE-III) giving him extensive experience with the grid middleware, grid deployment, and use of the grid by scientists throughout Europe.

Guy Barrand has developed a comprehensive data visualization and analysis framework, OpenScientist, which is used in several HEP experiments, which runs natively on the three main platforms (Linux, MacOS X and Windows). He is now interested in adapting OpenScientist to take advantage of the innovative features of the WILD project for scientific analysis.
ARTS@LAL

2. The main components of this project

Enabling Grids for E-science (EGEE)

Europe through its e-infrastructures programs has invested significantly in computing platforms for its scientists and researchers. The flagship distributed computing resource is the pan-European grid infrastructure provided by the Enabling Grids for E-science (<http://www.eu-egee.org/>) series of projects. This grid federates cluster-based resources from more than 260 institutes in 55 countries and serves more than 14000 users.

The Linear Accelerator Laboratory (LAL) of IN2P3 has contributed to the EGEE projects at all levels: production of tools, integration of grid middleware components, provision of resources, and use of the infrastructure. This broad experience allows us to see problems related with the infrastructure and the means to address them. One issue relates to the narrow support of operating systems by the grid middleware, which impedes the adoption of the grid by the more diverse user communities that EGEE is now attracting. A second issue is the scheduling model used to maximize the utilization of resources. The scheduling model, conceived for mono-core machines, is suboptimal for today's multi-core reality. Incorporating Mac OS X platforms into the production grid infrastructure provides a means to address both of these concerns.

Scheduling Model and Multi-threaded Applications

The EGEE grid infrastructure is a federation of cluster-based resources. At the cluster level, most sites run a standard batch scheduler like LSF or Torque to handle tasks from the grid users. The scheduling model used by most sites allocates a single job per processor in a machine, assuming that most jobs treated on the grid are mono-threaded. With the ubiquity of multi-processor/multi-core machines, this scheduling model had been trivially expanded to schedule one job per core.

On the infrastructure we now regularly see memory intensive jobs that require 4 GB or more of RAM. With these jobs, the "one job per core" scheduling model is inefficient because of memory bandwidth and memory capacity limitations of the machines. Because of this, grid users are having to rethink their single thread application model and seriously looking into writing true multi-threaded applications despite the additional complexity.

This is where having Mac OS X platforms as part of the EGEE grid infrastructure could greatly benefit the grid users. The recent introduction of Grand Central Dispatch (GCD) would allow the application developers to take advantage of multiple cores easily, without each developer needing to become proficient in the complexities of concurrent programming. Perhaps more importantly they would also be able to take advantage of the Mac's excellent debugging tools to ensure that applications rewritten for multi-core processors function correctly.

Multiplatform Support

One of the current drawbacks of the EGEE grid infrastructure is the homogenous platform support including only RedHat Enterprise Linux (RHEL) and its close derivatives. To date, this has served the existing user community well; new user communities, however, view this as an impediment as they do not have a tradition of running on RHEL platforms. Given that Mac OS X has a true Unix system underneath and that it uses the gcc compilers that are pervasive in the academic and grid communities, it is an easy first step for the diversification of the platform support on the grid infrastructure. Much of the work for expanding the grid middleware to support Mac OS X has already been done. However providing dedicated resources for consistent building of the software stack on Mac OS X and having a set of users relying on the availability of the platform via the grid will ensure that the last problems are tackled quickly and that the platform remains supported in future releases of the grid middleware.

The ATLAS collaboration

ATLAS is a particle physics experiment at the Large Hadron Collider (LHC) at CERN, consisting of more than 2900 physicists from around the world. The LHC is expected to start its operation at the end of this year (2009). In order to cope with the petabytes of data it will collect and analyze, ATLAS is using the EGEE grid infrastructure described above. Although the MacOS X platform is the preferred individual computing device by a growing fraction of the ATLAS physicists, the software suite of the experiment does not run on the platform. Porting (or more exactly finishing porting) this suite to the MacOS X environment would allow the ATLAS physicist to run her favorite data analysis directly on her laptop or desktop, and use any Macintosh-based grid nodes optimally.

LAL has been deeply involved in the development of the ATLAS software for many years, and especially in the collaboration-wide effort to adapt it to multi-threading and multi-processing paradigms. This qualifies it as one of the best places to conduct the porting of this software to the MacOS X platform, taking advantage of the last technologies Apple offered in Snow Leopard, among which LLVM, GCD and OpenCL.

It is worth noting that these advances could also benefit other LHC experiments, which share some code with ATLAS, and become a de facto standard in high performance computing within the entire community.

The LHCb collaboration

LHCb is another LHC experiment in which LAL is involved. One significant contribution of LAL to the software infrastructure of LHCb is its graphics event display, PANORAMIX. We plan to develop a version of this display for the image wall WILD developed at LRI, in order to explore new ways of analyzing data in HEP experiments by taking into account the innovative man-machine interaction this device is offering. The ARTS grant would be a considerable help for this project, allowing to install at LAL a reduced version of the WILD installation for development and debugging purposes. One unresolved challenge presently not solved is how to make a reasonable "reduced" version of the input sensors of WILD.

3. Conclusions

In short, including Mac OS X platforms on the grid infrastructure allows us to ensure true multi-platform support and gives the user community a valuable means of improving their applications through parallelization. Making MacOS X a supported platform within the ATLAS and LHCb collaborations both at the workstation and computing node levels will establish a new level of quality and efficiency for their software.

Appendix 1

In order to achieve the goals presented in this proposal, we intend to use the grant to set up at LAL a small but significant infrastructure: 6 MacPro with 4 GPUs in each, and a 4x3 monitor "image wall" to develop the programs that will be deployed on the WILD project (32x30 inch monitors) at LRI.



Node Configuration

Two 2.26GHz Quad-Core Intel Xeon
16Go (8x2Go)
1To 7200-rpm Serial ATA 3Go/s
4x NVIDIA GeForce GT 120 512MB



24 inches 1 920 x 1 200 pixels