BOINC service for volunteer cloud computing

N. Høimyr[1], J. Blomer[1], P. Buncic[1], M. Giovannozzi[1], A. Gonzalez[1], A. Harutyunyan[1], P. L. Jones[1], A. Karneyeu[1], M. A. Marquina[1], E. Mcintosh[1], B. Segal[1], P. Skands[1], F. Grey[2], D. Lombraña González[2], I. Zacharov[3]

[1] CERN, CH-1211 Geneva, Switzerland
[2] Citizen Cyberscience Centre, CERN, CH-1211 Switzerland
[3] Particle Accelerator Physics Laboratory, EPFL, CH-1015 Lausanne, Switzerland

E-mail: Nils.Hoimyr@cern.ch

Abstract. Since a couple of years, a team at CERN and partners from the Citizen Cyberscience Centre (CCC) have been working on a project that enables general physics simulation programs to run in a virtual machine on volunteer PCs around the world. The project uses the Berkeley Open Infrastructure for Network Computing (BOINC) framework. Based on CERNVM and the job management framework Co-Pilot, this project was made available for public beta-testing in August 2011 with Monte Carlo simulations of LHC physics under the name "LHC@home 2.0" and the BOINC project: "Test4Theory". At the same time, CERN's efforts on Volunteer Computing for LHC machine studies have been intensified; this project has previously been known as LHC@home, and has been running the "Sixtrack" beam dynamics application for the LHC accelerator, using a classic BOINC framework without virtual machines. CERN-IT has set up a BOINC server cluster, and has provided and supported the BOINC infrastructure for both projects. CERN intends to evolve the setup into a generic BOINC application service that will allow scientists and engineers at CERN to profit from volunteer computing. This paper describes the experience with the 2 different approaches to volunteer computing as well as the status and outlook of a general BOINC service.

1. Introduction

Volunteer computing provides processing power up to the petaflop scale to a wide range of scientific and technical projects. SETI@home, Folding@home, the World Community Grid, Climatepredicitons.net and also LHC@home [1] are well known examples, The first of them, SETI, attracted so much volunteer interest and CPU power that one of its lead developers, David Anderson, split its middleware infrastructure into what became the Berkeley Open Infrastructure for Network Computing (BOINC) [2,3]. BOINC enables institutes and individual researchers to harness computing power from volunteer PCs around the world.

In addition to access to computing power that otherwise would not be available, BOINC and volunteer computing also provide a way for outreach and publicity for scientific projects.

The BOINC model is based on a server hosting a given project, and sends out jobs on its behalf to all volunteer client nodes which have attached to it. The project server receives results from the submitted jobs, validates them, and deals with issues of client down-time or unreliability. BOINC also comprises a system of project message boards to allow communication between volunteers and project administrators. Volunteers are awarded BOINC credit to encourage their collaboration. It is neither material nor financial, but nevertheless represents a major incentive to some volunteer communities, who compete actively for it.

Until recently, each BOINC project required considerable effort to set up, e.g., to port its computing application to the wide variety of volunteer platforms. These are typically Windows-based, but the number of Linux and Mac OS X systems is increasing. The next step is developing suitable job submission scripts to manage the flow of work provided by the project scientists to the volunteers, via the project's BOINC server. A key development in this sense has been adapting BOINC to support Virtualization [4, 5, 6], which has been successfully pioneered at CERN. The effort to broaden the use of volunteer computing for Particle Physics, and earlier experience with the Sixtrack [7] accelerator beam dynamics simulation code, has led CERN-IT Department to re-establish a BOINC-based computing service. The Sixtrack application server that was hosted since 2007 at Queen Mary University of London was repatriated ahead of a major CERN outreach event featuring LHC@home, held at the Frankfurt book fair in October 2011.

2. Classic BOINC volunteer computing

Volunteers can connect and contribute computing cycles to a classic BOINC project by simply installing the client software and then connect to one or more of the BOINC project web sites where the user wishes to contribute CPU cycles. The BOINC client will then automatically download an executable program and start to download "tasks" with input jobs for the executable program. A pre-requisite for contributing to a project, is that there is an executable available matching the operating system and processor architecture of the volunteers PC. Once a job is finished, the resulting output will be uploaded back to the BOINC project server.

2.1. Sixtrack (LHC@home classic)

An example of a classic BOINC project is Sixtrack, the starting point for the LHC@home system. In 2003, the first tests started of an in-house screensaver called the Compact Physics Screen Saver (CPSS) [8], running Sixtrack on desktop computers at CERN. CPSS proved to be a useful first step with this kind of distributed computing application. On the occasion of CERN's 50th anniversary in 2004, it was relocated as the BOINC-based project LHC@home/Sixtrack.

Some 60,000 users with about 100,000 PCs have been active LHC@home volunteers since 2004. This provided huge computing power for special accelerator physics studies [9, 10], for which there was no equivalent available capacity at the CERN batch clusters. The simulations carried out with Sixtrack under BOINC have allowed for better understanding of the improvement to the LHC machine performance resulting from the sorting algorithm applied to the main cryodipoles [11].

The Sixtrack code is FORTRAN-based and was ported for use with BOINC on Windows and Linux by incorporating calls to the BOINC application programming interfaces (API) library and recompiling and re-linking the source code to produce executables for each client platform. Since 2004, the application code has undergone several updates to adapt to new BOINC versions, as well as to improvements to Sixtrack itself [12, 13].

3. Virtualisation support with BOINC

Following the success of LHC@home, the question was raised whether CERN could benefit from BOINC to assist with its LHC physics research programme as well as for accelerator design. Volunteer computers are one of the most heterogeneous resources for running any type of scientific experiments because of different operating systems, hardware architectures, versions of libraries, applications, operating systems, APIs, etc.

In order to ease access to such heterogeneous resources, BOINC provides an API and a set of wrappers. However, researchers have to either port their existing applications to the supported BOINC programming languages (C/C++, FORTRAN and Java), start from scratch a new version, or provide a statically linked version that could be used with one of the provided wrappers.

While Microsoft Windows is the most popular volunteer platform OS, the vast majority of LHC experiments' code runs only on Linux systems. Furthermore, the LHC experiments have very large code repositories of libraries, applications, modules, etc. which are updated frequently. Hence, porting all the source code to Windows machines in order to run these applications under BOINC would be time-consuming and could lead to errors with the simulations.

A solution to porting issues is to use virtualization technology [14]. In a guest Virtual Machine (VM), it is possible to load an OS and all of its software applications. The VM runs under a hypervisor package that is installed on the local computer, e.g. ORACLE's "Virtual Box" [15]

In 2006, CERN-IT began work on inclusion of VMs into BOINC, primarily as master or summer student projects. Results were promising, but the large size of the resulting VM images containing typical LHC physics code was a major problem.

In 2008, a CERN R&D project called CernVM [4] was launched by the SFT team in the CERN Physics Department, offering a general solution to the problem of virtual machine image management for physics computing at the LHC experiments [5]. Essentially, CernVM is a minimum virtual machine appliance – less than 250 MB in size – that then downloads by itself all the required libraries for LHC applications, and also provides a secure gateway to the experiments' own job management systems [6].

It was quickly realized that this technology could be incorporated into BOINC using a "CernVM Wrapper" - a modified version [16] of the official BOINC wrapper. It is in charge of uncompressing the VM after it has been downloaded to the client, launching it and controlling it. For example, it will pause or resume the VM when the BOINC task has to be paused or resumed to follow the work flow of BOINC. This solution preserves the standard BOINC project model.

3.1. Test4Theory (LHC@home 2.0)

After a joint effort by the CernVM team and CERN-IT, the first BOINC project using virtual machines was born in 2011: the LHC@home Test4Theory project. A particularly interesting feature of this system is that the BOINC-CernVM volunteer machines appear as cloud nodes to the application scientists [17]. Based on the described technology, a new web-based resource for validation of Monte Carlo (MC) models of high-energy collider physics has been established at CERN [18]. The aim of the MCPLOTS [18] website is to provide a simple online repository of plots, comparing the most widely used MC event-generator tools (currently ALPGEN, HERWIG++, PYTHIA 6 & 8, SHERPA, and VINCIA [19]) to a large database of experimental results, encoded in the RIVET analysis tool [20]. The repository is continually updated, and the computing power is supplied by volunteers via Test4Theory.

CernVM has allowed a very clean factorization between the physics software and the architecture of the volunteer machines. Another important aspect is a very high CPU-to-bandwidth ratio. The typical job will run for hours with zero load on the volunteer's network bandwidth (the events being analysed are simulated locally on the volunteer computer itself), and only a very small number of bytes is required to send job I/O back and forth at the start and end of jobs. After downloading the initial install package each volunteer machine only needs to be sent a single ASCII parameter card at the start of each simulation run, and the output sent back to the MCPLOTS server is likewise a single ASCII histogram.

At present, the volunteers connected to the Test4Theory project have generated a total of more than 300 billion collision events, providing a large amount of simulated MC statistics to compare to the experimentally measured reference distributions. The data comes from a large range of collider experiments, including four at LHC (ALICE, ATLAS, CMS, LHCb), three at LEP (ALEPH, DELPHI, OPAL), two at SPS (UA1, UA5), two at the Tevatron (CDF, D0), and one at the RHIC (STAR), with both more generators and new data continually being added.

4. BOINC server infrastructure

The BOINC cluster set-up currently consists of 5 virtual server nodes: one hosts the classic applications, currently Sixtrack, and another hosts the VM-based project Test4Theory. A third node acts as a web-server redirector to the 2 projects and the main LHC@home portal. The two remaining nodes are reserved for tests and upgrades. The cluster structure is flexible and will be extended with additional nodes as necessary. (E.g. additional server nodes are now added to host the Co-Pilot [6] infrastructure, hosted up to now on R&D nodes maintained by the SFT team.)
The application servers are running with a MySQL [21] back end for the BOINC application and forums. The database is currently located on the local machine hosting the project and backed up with database exports and regular system backups.

4.1. BOINC application support model at CERN

The current support model for BOINC is to host projects operated by dedicated teams and to provide the underlying application and server support to these. Volunteers, who wish to connect and contribute via BOINC, can attach to one of the projects by selecting the corresponding URLs from the BOINC client, namely:

- http://lhcathome.cern.ch/sixtrack
- http://lhcathome.cern.ch/test4theory

As the two utilize respectively a different underlying technology (classic BOINC vs Virtualisation), each is served through separate URLs as the requirements on the volunteer PCs are different. (The aliases “lhcathomeclassic” and “lhcathome2” have been defined with this in mind.)

This project-based approach relies on the corresponding teams to carry out outreach work to attract and retain volunteers. If new projects are joining, they need to be ready to make an effort communicating the goals of their simulations and the research done by their experiments.

4.2. Virtualisation versus classic BOINC

From the volunteer user's point of view, contributing to a VM-based project such as Test4Theory requires an additional step to get started; the user must download the Virtual Box hypervisor software in addition to the BOINC client. Thus, for number crunchers who already contribute to several BOINC projects like SETI@home, this additional step may be perceived as a hurdle, and they may choose another project instead.

Currently the LHC@home application Sixtrack offers such an alternative for users who wish to contribute their computing power to CERN. However, the classic BOINC distribution method requires porting of the application software to different client operating systems. (Currently only Windows and Linux are supported for Sixtrack.)

As discussed earlier, the software used by the High Energy Physics community is mainly available for Linux only; it takes a lot of effort to port experiment software to Windows and MacOS for use with native BOINC clients. Thus at CERN we believe that the future of volunteer computing for High Energy Physics lies with Virtualisation and the Volunteer cloud computing approach, as pioneered by Test4Theory. Integrating the download of the hypervisor with that of the BOINC client would be one step to simplify the use of virtualisation technology with BOINC for the volunteers, and make the underlying technology transparent.

5. Outlook and further plans

Currently the focus is on consolidating the BOINC infrastructure and to bring the Co-Pilot servers under the CERN-IT infrastructure like done earlier for the CernVmFS [22] storage. In order to distribute the database load from the BOINC project servers and improve database backups the plan is to relocate the data from the BOINC servers to a Database-on-Demand service [23].

Furthermore, a web-proxy load balancing setup is planned with a common web server frontend that directs transparently user connections to one of the project backend servers. Once the consolidated setup is available, the CERN BOINC service would be in a position to welcome new customers, e.g. from the LHC experiments if there should be demand for use of Volunteer cloud as a computing resource. As pointed out by Haratunyan et.al [6], the Co-Pilot job management framework is flexible and should allow for submission of jobs to the LHC computing grid, external cloud providers like Amazon or volunteer clouds to use available capacity. Thus the BOINC-based volunteer cloud computing service at CERN could be considered as yet another computing resource that can be tapped by applications with small data sets.

If several new projects at CERN would like to benefit from the BOINC service and computing resource, it would make sense to combine the efforts in terms of web-portal and infrastructure for outreach. A combined portal with a unique URL for the use by BOINC clients, would allow for sending of computing simulations from a range of different experiments to the same group of volunteers. The outreach efforts via forums could also be combined, instead of being project-specific as today. A single portal with a high level of activity is more attractive to volunteers than multiple projects within the same organisation. When in the future, all the CERN BOINC projects will rely on a single technology, and virtualisation support is transparently integrated with the BOINC client, sharing the same project URL for all projects and distribute tasks for volunteers from different CERN projects would provide a solution that scales better, both in terms of infrastructure and for volunteers providing CPU power. Rotation of jobs from different sources, such as accelerator studies with Sixtrack and other beam dynamics codes, Theory studies or the LHC experiments Alice, ATLAS and CMS would also ensure that the volunteers do not run out of tasks.

6. Conclusion

Volunteer computing relying on the BOINC framework and the Sixtrack application has been successfully running beam dynamics simulations for the LHC for several years. However, the hurdles adapting physics applications to BOINC and different operating system environments has up to now discouraged adaptation of other CERN simulation packages to BOINC. Thanks to the efforts of teams at CERN's PH and IT Departments, as well as CCC collaborators, this hurdle has been overcome by adapting BOINC to support Virtualization utilising the popular CernVM machine image.

Based on these developments, the Test4Theory project is running particle physics Monte Carlo simulations using the computing power offered by volunteer computing. The virtual machine approach to Volunteer Cloud computing potentially opens up for a wider range of physics applications, and the aim is to evolve the current project-based BOINC service towards a general service hosting Volunteer Cloud computing applications at CERN.

7. Acknowledgements

We would like to thank the authors of the software packages and IT infrastructure teams who provide the foundation for the setup of the BOINC and LHC@home project servers.

Last but not least, we would like to express our gratitude to all volunteers that contribute computing power to LHC@home and other BOINC projects.

References:

[1] LHC@home portal at CERN. URL: http://lhcathome.cern.ch

[2] Anderson, D *Grid Computing, 2004. Proceedings. Fifth IEEE/ACM International Workshop on , vol., no., pp. 4- 10*

[3] BOINC URL: http://boinc.berkeley.edu/

[4] Buncic P e a 2010 *Journal of Physics: Conference Series, 219, 042003*

[5] Harutyunyan A e a 2010 *Journal of Physics: Conference Series, 219, 072036*

[6] Harutyunyan A e a 2011 *Journal of Physics: Conference Series, 331*

[7] Schmidt F 1994 *CERN-SL-94-56*

[8] McIntosh E and Wagner A 2005 *CERN-2005-002, 1055.*

[9] Herr W, Kaltchev DI, McIntosh E and Schmidt F 2006 *LHC-Project-Report-927*

[10] Giovannozzi M and McIntosh E 2006 *LHC-Project-Report-925*

[11] Fartoukh S and Giovannozzi M 2012 *Nucl. Instrum. & Methods A 671*

[12] McIntosh E, Schmidt F and de Dinechin F "Massive Tracking On Heterogeneous Platforms" 2006 *9th International Computational Accelerator Physics Conference.*

[13] Giovannozzi M e a "LHC@home" 2012 *International Particle Accelerator Conference*.

[14] Smith J E and Nair R 2005 Virtual Machines *S. Francisco: Morgan Kaufmann, Elsevier*

[15] Oracle VM Virtual Box URL: https://www.virtualbox.org/

[16] Lombraña-González D 2012 "CernVM wrapper", Technical report, Citizen Cyberscience Centre. URL: https://github.com/citizen-cyberscience-centre/cernvmwrapper

[17] Segal B e a 2010 *Proceedings of the XIII International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT10), Jaipur*

[18] Skands P e a "MC production statistics", URL: http://mcplots.cern.ch

[19] Buckley e a 2011 *Phys.Rept. 504, 145*

[20] Buckley e a "Rivet user manual", arXiv:1003.0694 [hep-ph].

[21] MySQL URL: http://www.mysql.com/

[22] Blomer J e a "Distributing LHC application software and conditions databases using the CernVM file system" 2011 *J. Phys.: Conf. Ser. 331 042003*

[23] DB on Demand URL: http://cern.ch/dbondemand