Research article

Javier Barranco, Yunhai Cai, David Cameron, Matthew Crouch, Riccardo De Maria, Laurence Field, Massimo Giovannozzi*, Pascal Hermes, Nils Høimyr, Dobrin Kaltchev, Nikos Karastathis, Cinzia Luzzi, Ewen Maclean, Eric McIntosh, Alessio Mereghetti, James Molson, Yuri Nosochkov, Tatiana Pieloni, Ivan D. Reid, Lenny Rivkin, Ben Segal, Kyrre Sjobak, Peter Skands, Claudia Tambasco, Frederik Van der Veken, and Igor Zacharov

LHC@Home: a BOINC-based volunteer computing infrastructure for physics studies at CERN

Abstract: The LHC@Home BOINC project has provided computing capacity for numerical simulations to researchers at CERN since 2004, and has since 2011 been expanded with a wider range of applications. The traditional CERN accelerator physics simulation code SixTrack enjoys continuing volunteers support, and thanks to virtualisation a number of applications from the LHC experiment collaborations and particle theory groups have joined the consolidated LHC@Home BOINC project. This paper addresses the challenges related to traditional and virtualized applications in the BOINC environment, and how volunteer computing has been integrated into the overall computing strategy of the laboratory through the consolidated LHC@Home service. Thanks to the computing power provided by volunteers joining LHC@Home, numerous accelerator beam physics studies have been carried out, yielding an improved understanding of charged particle dynamics in the CERN Large Hadron Collider (LHC) and its future upgrades. The main results are highlighted in this paper.

Keywords: volunteer computing, high energy physics, particle accelerators, beam dynamics

1 Introduction

This paper addresses the use of volunteer computing at CERN, and its integration with Grid infrastructure and applications in High Energy Physics (HEP). The motivation for bringing LHC computing under the Berkeley Open Infrastructure for Network Computing (BOINC) [1] is that available computing resources at CERN and in the HEP community are not sufficient to cover the needs for numerical simulation capacity. Today, active BOINC projects together harness about 7.5 Petaflops of computing power, covering a wide range of physical application, and also particle physics communities can benefit from these resources of donated simulation capacity.

The structure of the paper is the following: in Section 2 an overview of the LHC@Home BOINC project is given, while the detail and specificities of the various applications running under LHC@Home are given in Section 3, with separate sections, from 3.1 to 3.5, to cover the various applications. A detailed analysis of the SixTrack case is provided in Section 4, covering the current studies (see Section 4.1) the performance analysis (see Section 4.2) and an outlook on future applications (see Section 4.3). Finally, conclusions are drawn in Section 5.

David Cameron: University of Oslo, Postboks 1048 Blindern N-0316 Oslo, Norway

Javier Barranco, Tatiana Pieloni, Lenny Rivkin, Claudia Tambasco, Igor Zacharov: Particle Accelerator Physics Laboratory, École polytechnique fédérale de Lausanne (EPFL) 1015 Lausanne, Switzerland

Peter Skands: School of Physics & Astronomy, Monash University, VIC-3800, Australia

Yunhai Cai, Yuri Nosochkov: SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Dobrin Kaltchev: TRIUMF, 4004 Wesbrook Mall, Vancouver, B.C., Canada V6T 2A3

Ivan D. Reid: Brunel University, London, Uxbridge UB8 3PH, UK
Matthew Crouch, Riccardo De Maria, Laurence Field, Pascal Hermes, Nils Høimyr, Nikos Karastathis, Cinzia Luzzi, Ewen Maclean, Eric McIntosh, Alessio Mereghetti, Ben Segal, Kyrre Sjobak, Frederik Van der Veken: CERN, 1211 Geneva, Switzerland
*Corresponding Author: Massimo Giovannozzi: CERN, 1211 Geneva, Switzerland
James Molson: LAL, Univ. Paris-Sud, CNRS/IN2P3, Université Paris-Saclay, Orsay, France

2 LHC@Home project

In 2002, as part of the ongoing search for ever better price-performance ratio computing, as CERN had moved from mainframes to workstations and then PCs, an article on the use of PlayStations suggested the use of even lower cost alternatives. Neither the PlayStation 2 nor 3, however, provided IEEE 754 compliant double precision floating-point arithmetic which was, and is, considered essential for most CERN applications. Instead, an informal project, Compact Physics ScreenSaver (CPSS) [2, 3], was established to attempt to use the several thousand Windows desktop PCs at CERN during nights and weekends when otherwise idle. It was then proposed to use the BOINC infrastructure to extend the potential usage worldwide. Thus, volunteer computing has been used successfully at CERN since 2004 with the LHC@Home project, and has provided additional computing power for CPU-intensive applications with small data sets, as well as an outreach channel for CERN activities. LHC@Home started off with the accelerator code SixTrack [4, 5], which had been successively ported from mainframe to supercomputer to emulator farms and PCs, and later on a gas detector simulation program [6]. However, as applications running under BOINC had to be compiled for each and every possible client operating system, only the SixTrack application was ported to Windows, Linux and later MacOSX clients. Note that most HEP codes, such as the analysis frameworks of the LHC experiments, run almost exclusively under the Linux operating system and are therefore run in virtual machines as described below.

2.1 Virtualisation with BOINC

Thanks to developments started at CERN, and later brought into the BOINC distribution, such Linux programs can now run on a Virtual Machine (VM) distributed to the volunteer computers via BOINC and running on volunteer PCs within the Oracle VirtualBox hypervisor. This use of virtualisation under BOINC was pioneered by the Test4Theory LHC@Home project during 2008-2011 [7–9, 11]. This development has allowed the LHC experiment collaborations to run their simulations also under BOINC, in the CernVM virtual machine [24].

The CernVM project provides virtual images tailored for the LHC experiments' software and these images can run seamlessly inside the virtualisation layer provided by BOINC. CernVM is a compact Linux virtual machine based on Scientific Linux 6 combined with a kernel adapted for virtualisation. CernVM does not include a physical disk, and the current Micro-CernVM only contains core libraries, compilers and handlers for an http file system [25]. In this way, the core image size is of only ∼20MB, and hence suitable to download for volunteer computing applications. Application-specific software is downloaded via a networked http file system, CernVMFs.

Several experimental groups have been running pilot BOINC projects for their collaborators to contribute simulations via BOINC and virtualisation. Following the experience with Test4Theory, ATLAS@Home and other pilot projects, with a view to include volunteer computing into the production computing infrastructure for HEP [12], a major effort has been undertaken to consolidate the original LHC@Home and host additional applications utilising virtualisation.

It is worth mentioning that use of Docker containers as a lighter alternative to virtual machines has been tested as a proof of concept for the ATLAS application use case in 2016, but this currently requires more work for the current CERN applications, although used elsewhere [10].

2.2 LHC@Home consolidation

Adding more applications to a BOINC project is straightforward. However, to make multiple applications appeal to volunteers and users from different communities, application-specific credit was deployed. The credit for the applications running in a VM environment is based on the CPU consumption of the VM on the volunteer host and is gathered via the vboxwrapper application. SixTrask obtains traditional BOINC credit based on the CPU consumed by the Sixtrack application. In the recent versions of the BOINC library, this credit is calculated in a uniform way to level the ground between the different BOINC projects.

These steps pave the way to a consolidated CERN infrastructure, which implied tackling the task of porting the applications from the old to the new infrastructure. The accounts and BOINC credit of volunteers who had been contributing to the pilot projects Test4Theory/vLHC@home and ATLAS@Home were migrated to the consolidated LHC@Home project by means of a set of SQL scripts, as all the information is stored in the database. The volunteer's email address was used as unique key for the data, as the user ID differs in each project depending on when the volunteer joined the BOINC project.

On the consolidated LHC@Home, users have a choice of applications that is enabled via LHC@Home project preferences. The SixTrack application, which does not require VirtualBox, is enabled by default for volunteers.

Once registered, volunteers can enable e.g. ATLAS, CMS or Theory simulations via the LHC@Home project preferences.

In terms of computing power provided by the volunteers to LHC@home, the average is about 1×10^5 simulation tasks. For SixTrack, peaks of 3.5×10^5 simultaneously running tasks on 2.4×10^4 hosts have been observed during SixTrack simulation campaigns, but note that every SixTrack task is run twice to eliminate random host errors and minimise the impact of a failing host. This can be compared against the average of 2.5×10^5 running tasks on 1.4×10^5 processor cores in the CERN computer centre, that is fully loaded with tasks of analysis and reconstruction of collisions recorded by LHC experiments, and has limited spare capacity for beam dynamics simulations. The applications of the LHC experiments that require virtualisation support on volunteer computers have operated with a sustained load of about 7000 tasks for ATLAS, 6000 for Theory, 3500 for LHCb, and 1000 for CMS.

3 Applications

3.1 SixTrack

SixTrack is an open source program for the simulation of charged particle trajectories in circular accelerators; it has been running under LHC@Home since 2004. Some 1.5×10^5 users with more than 3×10^5 PCs have been active LHC@Home volunteers since its launch. This has provided significant computing power for accelerator physics studies, for which there was no equivalent capacity available in the regular CERN computing clusters. Volunteers contributing to SixTrack have delivered a sustained processing capacity of more than 45 TeraFlops. Figure 1 shows the time evolution of the volunteers, active tasks, and cumulative number of workunits (WU) since Feb. 2017. Note that each WU is submitted at least twice for ensuring numerical stability of the results. Note that the number of volunteers underestimates the actual CPU capacity available, as each volunteer could provide several machines and each machine might be multi-core.

The SixTrack code is mainly Fortran-based, vectorized to take advantage of vector instructions, pipelining, and hardware features such as SSE and AVX. It was ported for use with BOINC to Windows, MacOSX and Linux by incorporating calls to the BOINC application programming interface (API) library and re-compiling and re-linking the source code to produce executables for each client platform. Since 2004, the application code has undergone sev-

eral updates to adapt to new BOINC versions as well as to improvements to SixTrack itself (see [13] for a recent account on the code state). The principal functional changes for consistent and reliable operation are outlined in [14], but subsequent improvements now allow the use of several Fortran compilers, at any Fortran standard compliant level of optimisation, providing identical results, i.e. 0 Units difference in the Last Place (ULPs), on any IEEE 754 compliant hardware (E. McIntosh, in preparation). In order to achieve this, Fortran expressions, which could be evaluated in a different order as allowed by the standard, were parenthesised (H. Renshall, personal communication). SixTrack can be built in many different configurations, e.g. for dynamic aperture (see Sections 4.1 and 4.3) or collimation studies, and with or without support for checkpoint/restarting, compressed input/output, correct and consistent rounding of mathematical functions [15], BOINC, and more. Furthermore, it can run natively on most major platforms (Linux, MacOSX, Windows including XP, FreeBSD, NetBSD, OpenBSD, and GNU Hurd on x86 and x86_64, as well as Linux on armv6, armv7, 64bit armv8 (including Android systems), and PPC64, as long as a UNIX-like build environment is available; on Windows this is provided by MSYS2 [16]. The present CMake-based build system can compile the sources [17] and tests the reproducibility of the results using GNU, Intel, or NAG Fortran compilers. Consistency down to 0 ULP is automatically verified between the versions, platforms, and compilers using a CTest-based test suite, which includes automatic building reports and test coverage published on CDash [18].

3.2 Test4Theory

Since 2011, Monte-Carlo (MC) computer simulations of both ongoing and historical collider experiments have been performed in a CernVM virtual machine sent to volunteers using BOINC [7] (see also Section 3.3 for more detail on CernVM). Such so-called event-generator programs (see [8] for an introduction and review) are used extensively in HEP, as explicit numerical models of the (often highly complicated) particle dynamics and to provide theoretical reference calculations for the experimental measurements. Via the BOINC project Test4Theory (later renamed vLHC@home), which pioneered the use of virtual-machine technology for volunteer cloud applications, more than 3 trillion events have been simulated with different simulation programs. The generated events are compared against a large (and ever growing) library of particle-physics measurements, via the Rivet analysis
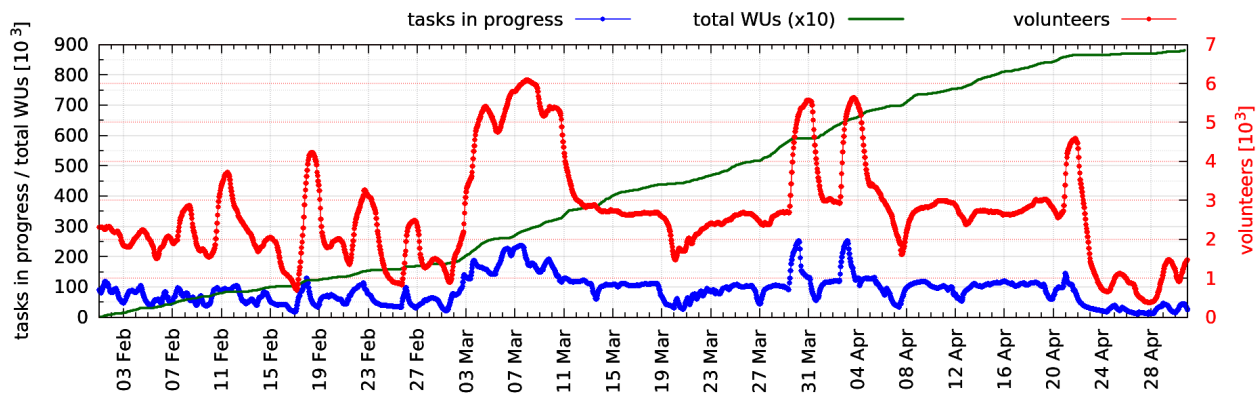
Figure 1: Time evolution of the cumulative number of WUs, volunteers, and tasks sent to BOINC from Feb. 2017.

preservation tool [19]. The results are stored as histograms and reference plots, in the on-line MCPlots database [11], which is available to the global particle-physics community. It is used by both the authors of the simulations and by their users, as a validation tool, and to guide further efforts to improve the physics models and to optimise their parameters (see e.g. [20]).

The upper part of Fig. 2 shows a time slice from the summer of 2012, of the number of new users per day signing up for the Test4Theory project. On July 4th that year, CERN announced the discovery of the Higgs boson, prompting hundreds of new users to join the project. The lower part shows one of the many thousands of plots that are available at the MCPlots site [11]. Several state-of-the-art models for particle collisions (coloured lines) are compared against an archived measurement performed in 1996 by the ALEPH experiment (black squares) [21], of the probability distribution for observing N charged particles (N_{ch} on the x axis) in electron-positron collisions at the LEP collider. (The lower pane shows the ratio of theory divided by data.) One clearly sees that the average of about 20 charged particles per collision is well reproduced by all the models, while their predictions differ in the tails of the distribution, where the uncertainty on the measurement (yellow band) was large.

3.3 A Toroidal LHC ApparatuS (ATLAS)

ATLAS@Home started in 2014 as an independent project where volunteers run Geant4 [22] MC simulation of particles passing through the ATLAS detector [23]. These simulations are well-suited to volunteer computing for several reasons: they involve less data transfer compared against other workloads; in ATLAS they are the largest consumer of CPU resources and hence there is always a reliable

source of work; many simulation campaigns run over several months, so a fast turnaround is not expected.

ATLAS relies on virtualisation to allow its simulation software to run on non-Linux hosts. ATLAS software is provided to the VM through the CernVM File System (CVMFS) [26], a remote read-only filesystem using aggressive local caching which is mounted inside the image. To avoid downloading the software every time the VM is started, the CVMFS cache inside the image is pre-filled with the required software, by running an example job, saving a snapshot of the image, and using that snapshot as the final image to distribute to volunteers.

One critical requirement when starting the project was that no sensitive ATLAS credentials should be distributed to volunteers. The solution was to use the model deployed in NorduGrid [27] and other environments such as High Performance Computing (HPC) centres which have restricted access to the outside world from the job worker nodes. The architecture of this model is shown in Fig. 3.

The Advanced Resource Connector (ARC) Computing Element (ARC CE) [28] takes care of data staging before and after the job runs, and the ARC Control Tower (aCT) [29] provides the link with the ATLAS workload management system, PanDA [30]. Jobs which are assigned to ATLAS@Home by PanDA are picked up by the aCT, and sent to an ARC CE connected to the BOINC server. ARC CE copies the required input files from Grid storage to a staging area inside the BOINC server. ARC CE supports many batch systems and a new plugin for a BOINC "batch system" was written to allow injection of jobs as work units in the BOINC server. Instead of calling batch system commands, this plugin uses the create_work command to inject jobs into the BOINC server and queries the BOINC database to find out when jobs have completed. The BOINC client on the volunteer's PC only has access to the BOINC server data staging area and no access to Grid storage or
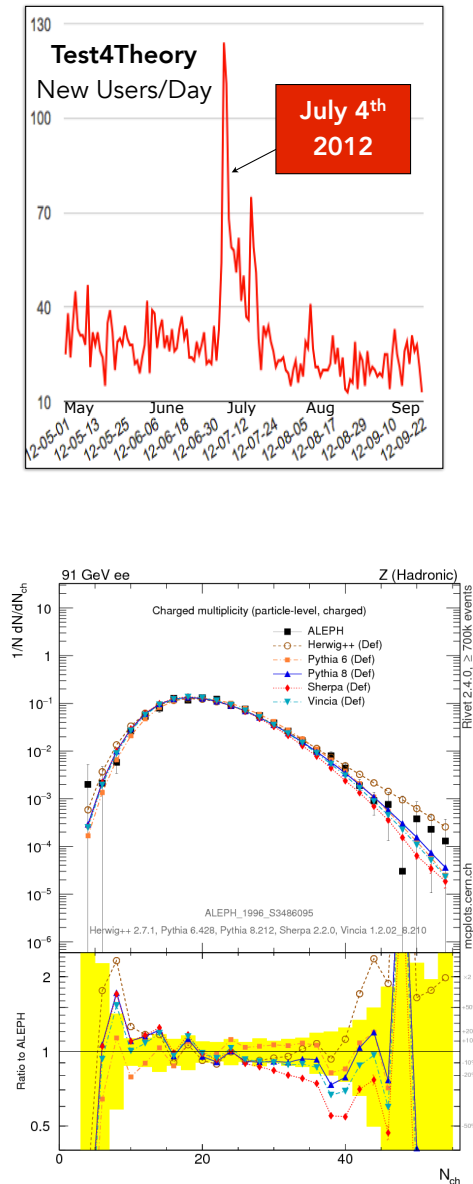
Figure 2: New users per day on Test4Theory during 2012 (upper) and comparison of modern event generators to a legacy measurement (lower, from the MCPlots web site [11]).

Grid credentials and so there is no chance of accidental or deliberate tampering with ATLAS data. Because ARC CE and aCT are services which are part of the regular ATLAS computing Grid, ATLAS@Home looks from the outside like a regular Grid site, which means no special treatment is needed when it comes to defining tasks, monitoring, accounting etc.

ATLAS@Home is one of the most demanding volunteer computing applications in part due to its high memory usage. A job using a single core can require a virtual machine with up to 2.5 GB of memory, and for many machines this means that it is not possible to fill all cores

with ATLAS@Home tasks. However, ATLAS software can run on several cores inside a single virtual machine and can take advantage of sharing memory between processes running on each core. These multi-core jobs provide a significant memory saving, with an 8-core job typically using 5-6 GB of memory in total. Previously, BOINC only allowed a fixed memory limit per WU no matter how many cores were used. The ATLAS@Home jobs' memory requirements are dependent on the number of cores and so the project team implemented in BOINC a way of dynamically determining the memory required based on the number of cores. Two new parameters were added to the plan class, which describes the characteristics of the virtual machine. A base memory and memory per core can be specified and the memory of the virtual machine is calculated as base memory + (memory per core × number of cores). This feature was passed upstream and is now part of the standard BOINC software.

At the time of writing, ATLAS volunteers have simulated almost 170 million ATLAS events (one event typically takes around 5 minutes of CPU time to simulate) and the combined resources add up to around 2% of overall ATLAS computing resources.

3.4 Compact Muon Solenoid (CMS)

CMS [31] is one of two general-purpose detectors at the LHC project, alongside ATLAS. Development began on a CMS@Home project in 2015 using a modified CMS Remote Analysis Builder v3 (CRAB3) [32] server VM submitting jobs running CMS standard software (CMSSW) [33] to a dedicated HTCondor [34] server VM rather than the normal submission to the Worldwide LHC Computing Grid (WLCG) [35]. The VMs were run at Rutherford Appleton Laboratory (RAL), UK.

Care was taken to match the type of jobs being run to the limitations of the volunteer environment. Of particular concern was the amount of data to be transferred, since many users still have ADSL connections which may have upload speeds as low as 1 Mbps. This obviously ruled out analysis of CMS data, but still allowed the generation of MC simulations of collision events. The MC job parameters were adjusted to give average run-times of about one hour, and output files of the order of 50 MB. The BOINC server distributed tasks which ran in the volunteers' VMs and executed MC jobs retrieved from the HTCondor server. Job output files were returned to a dedicated Data Bridge service [36] from where they could then be transferred to the normal CMS computing infrastructure. After a job completed, if the task had run for less than 12 hours it fetched
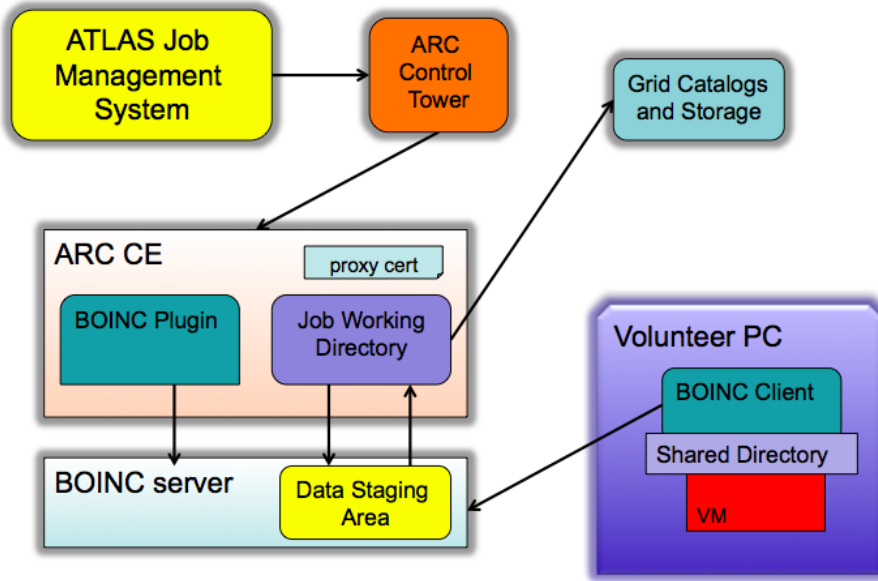
Figure 3: Architecture of ATLAS@Home.

another job to process, otherwise it terminated. Tasks were scheduled by BOINC according to the volunteers' preferences, taking into account other projects they may also have been running.

As a comparison with standard Grid jobs, batches of 2×10^3 jobs consisting of 25 events producing top-antitop ($t\bar{t}$, or ttbar) pairs were submitted to both CMS@Home and the Grid. The number of result files received over time from submission is shown in Fig. 4.
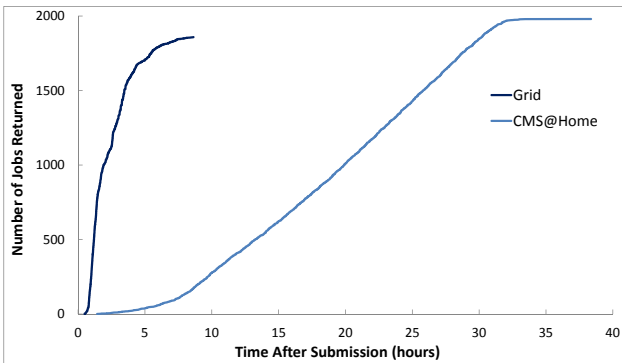

Figure 4: The distribution of result files received for 2×10^3 25-event ttbar simulation jobs, as a function of time from submission: dark curve – results from the Grid; light curve – results from CMS@Home volunteers.

Since the Grid has a large number of fast hosts, the first results started arriving after just 30 minutes, with 90% (1800) of the expected results received in about 6 hours. Unexpectedly, 7.1% (142) of the result files were never received. Meanwhile, CMS@Home results began arriving after ∼80 minutes, but due to the small number of available volunteer hosts (∼100) only a limited number could run at any one time. Thus the graph of return times (Fig. 4) has a costant slope for much of its duration as results returned at a constant rate. 90% of the results were received in 29.5 hours; in total 99% (1980) arrived in 38 hours.

As a test of a scientifically valuable process, the project turned to the simulation of the production of Λ_b^0 in LHC collisions, and its decay to a proton, a muon, and a neutrino. This is of interest as a background in measurements of a B_s decaying to two muons, since the proton may be misidentified as a muon. Because the Λ_b^0 is more massive (5.62 GeV/c^2) than B_s (5.37 GeV/c^2), the reconstructed mass of the $p + \mu$ overlaps the B_s mass spectrum, since the undetectable ν carries away a variable amount of energy. However, the production ratio is small, around 3×10^{-5}, so many proton-proton collisions need to be simulated to provide a significant number of desired events. Jobs simulating 2×10^5 collisions were used (median run-time 2h20m, result files ∼16 MB). In the last half of 2016, as the project developed and was incorporated into the larger LHC@Home, the number of simultaneous jobs increased and altogether several tens of billions of collisions were simulated, returning more than 2 million filtered events.

The project has now turned to the use of the workflow management system (WMAgent) [37] for job submission. WMAgent gives the ability to specify a destination site within the CMS infrastructure to which results are automatically replicated using the transport software PhEDEx [38]. Thus fully end-to-end running of CMS MC production jobs has been demonstrated and the project will be able to contribute a significant computing resource

to the CMS Collaboration. At the time of writing, volunteers are providing around 800 job slots to production, a figure that is expected to rise in the future.

3.5 Large Hadron Collider beauty experiment (LHCb)

The LHCb [39] experiment detector has been designed to filter out from the different particles generated by LHC those containing beauty and anti-beauty quarks (*B*-mesons) and the products of their decay. Unlike the other LHC experiments that surround the entire collision point with layers of sub-detectors, the LHCb detector extends along the beam pipe, with its sub-detectors piled behind each other. This is because the *B*-mesons do not travel in all directions, but rather stay close to the line of the beam pipe. Considering the growing need of computing power, the LHCb computing group has created a first prototype of the Beauty@Home project in 2013 to profit from volunteer computing resources.

The project uses the CernVM Virtual Software Appliance [40], the BOINC framework, and the Distributed Infrastructure with Remote Agent Control (DIRAC) system for distributed computing [41, 42]. At the beginning, the project was used only by users belonging to the LHCb Virtual Organisation. This because the architecture did not provide a secure technique to authenticate volunteers, but a trusted host certificate was contained in the machine dispatched to the volunteer.

The original problem was that pilot jobs needed to contact central DIRAC services such as the job matching or the job status update. They also needed to perform data management operations, such as the upload of the output files and the deployment of real credentials (proxy or server certificate), on untrusted machines, which was representing a big security threat. The necessity of having a secure authorization and authentication process to open the project to the outside world triggered the development of a DIRAC gateway service called Workload Management System Secure Gateway (WMSSecureGW). The service had the aim to interface untrusted volunteers machines to the DIRAC System authorizing BOINC users to execute LHCb jobs.

The WMSSecureGW service runs on a trusted machine, which has a valid certificate and accepts a dummy Grid certificate signed by a dummy certification authority (CA). The service receives all calls coming from the job and directed to different DIRAC services and it dispatches them as appropriate. Before the real storage upload is performed, the output data produced by the volunteer machines are uploaded on the gateway machine where a check has to be performed to avoid storing wrong data on LHCb storage resources. The architecture of the WMSSecureGW service is shown in Fig. 5.

Through this service, the Beauty@Home has been integrated in the LHCb Grid infrastructure and the BOINC volunteers are running LHCb simulation jobs as all others Grid resources.

Currently, almost 3.5×10^3 simulation jobs are performed per day by volunteer computing resources, hoping that this number will grow in the near future, thanks to the increasing contribution of volunteers.

4 A closer look at SixTrack

4.1 CERN Large Hadron Collider and its High-Luminosity upgrade

Modern particle colliders are all based on superconducting magnets to generate high-magnetic field and hence high-energy beams. This class of magnets comes with intrinsic field errors that generate non-linear effects in the charged particle dynamics. Non-linearities are potentially harmful for particle's motion as they could drift away from the central trajectory, eventually hitting the beam pipe. This would induce beam losses or, even worse, a transition from the super- to normal-conducting state. Both events would entail an overall loss of accelerator performance. The only means to determine whether a charged particle will be eventually lost is via numerical simulations. The aim of these simulations is to determine the so-called dynamic aperture (DA), i.e. the region in phase space where the particle's motion is stable for a given number of turns.

Each simulation requires generating a set of initial conditions to be tracked through the accelerator structure for $10^5 - 10^6$ turns, which, in the case of the CERN Large Hadron Collider (LHC) corresponds to only $\sim 9 - 90$ s out of a cycle of several hours. The DA depends on several physical parameters and scan over these quantities is essential to better understand the beam behaviour. Moreover, magnetic field errors are treated statistically and the DA computations are repeated for several realisations of these errors, typically 60, to ensure enough statistical relevance of the results. Overall, this implies that a typical study is made of $\approx 1 - 3 \times 10^6$ WUs each performing tracking over $10^5 - 10^6$ turns. This makes LHC@Home the ideal system for DA simulations that, otherwise, would not be possible to perform on standard computing resources.
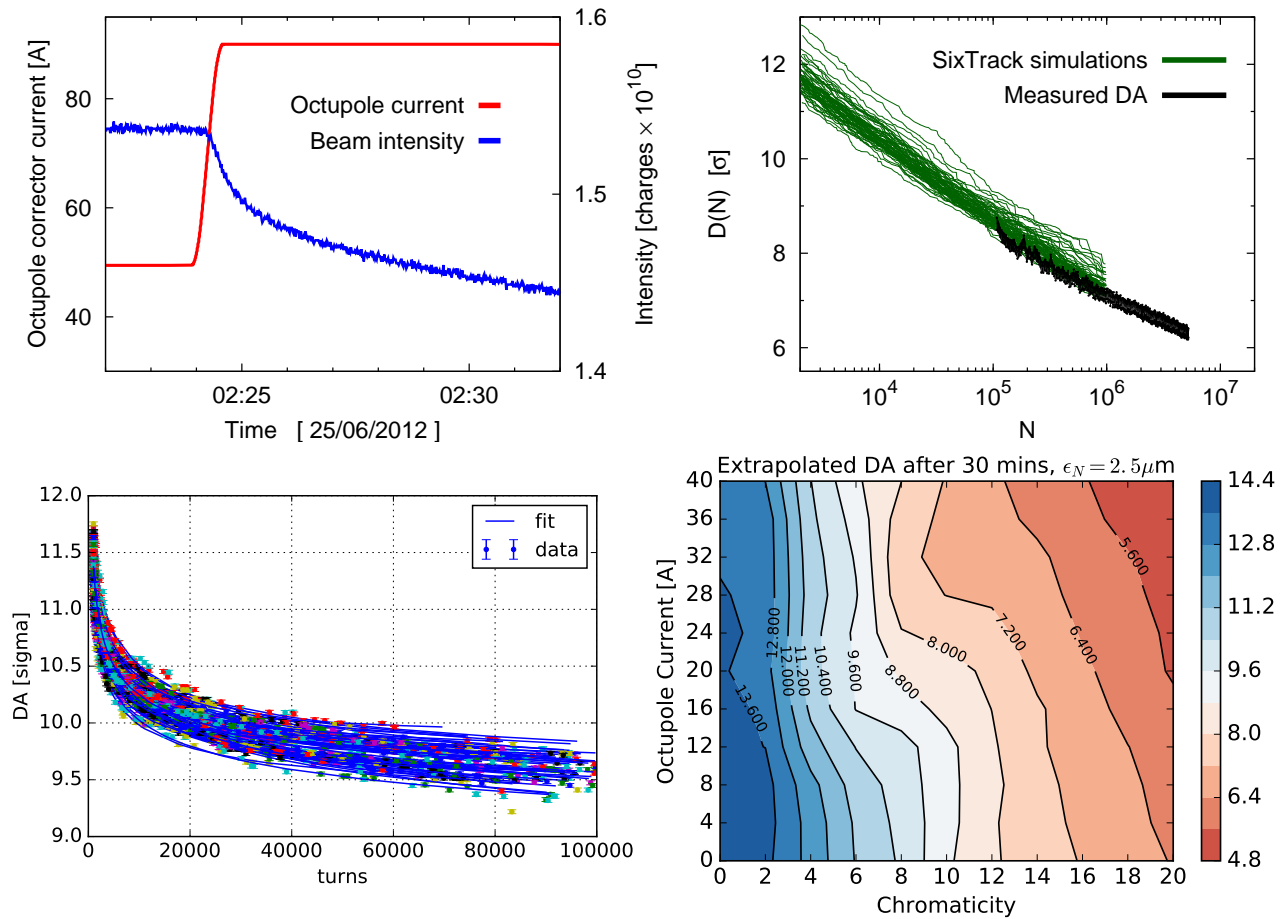
Figure 5: The whole gateway architecture, including the WMSSecureGW service and all services necessary to interface volunteers to the DIRAC framework.

The limited number of turns that can be explored requires special techniques to extrapolate the particle behaviour to more relevant time scales [43] and dedicated measurement campaigns have been carried out to benchmark numerical simulations in the LHC without (E. Maclean, M. Giovannozzi, R. Appleby, submitted for publication) and with [44, 45] beam-beam effects. Examples of these studies are shown in Fig. 6, where the upper row shows comparison of measured and simulated DA, while in the lower row a typical scan of the extrapolated DA vs key parameters is shown.

For the LHC high-luminosity upgrade (HL-LHC) [46], beam simulations are essential for a reliable estimate of the collider's performance, also to guide the design of the new hardware. In Fig. 7 (left) the DA is shown as a function of the phase advance (horizontal and vertical) between the collision points in ATLAS and CMS, while (right) the DA as a function of transverse tunes including beam-beam interaction between bunches of 2.2×10^{11} protons is depicted (see also [47]). Note that these studies are essential to select the parameters' values providing the maximum DA, hence optimising the accelerator's design.

4.2 SixTrack performance with BOINC

The processing time of SixTrack studies submitted to BOINC over the month of September 2017, those that were not yet purged from the result database, have been analysed. Data extraction has been performed in order to select all results that completed with error-free processing.

Figure 6: Upper left: measured beam intensity evolution during an experimental session. Upper right: comparison between simulated and measured DA of the LHC at injection. Lower left: DA evolution with the number of turns from SixTrack simulations compared against fits of the data for the individual seeds. Lower right: Extrapolated DA of LHC at 30 minutes after injection as a function of different chromaticities and octupole settings.

For the sake of the time analysis, SixTrack tasks in BOINC can be divided in two categories, based on the total number of turns used in the beam dynamic simulations, which translate directly into CPU time. A sample of $95,773$ runs with 10^5 turns and $115,245$ runs with 10^6 turns have been prepared, the latter set represents studies including beam-beam effects (note that these effects are particularly expensive in terms of CPU-time). In Fig. 8 the actual processing time on volunteers' computers is shown.

The distribution of the SixTrack computing time on volunteers' computer is determined by the properties of the initial conditions. If they are located in an unstable region of phase space, their amplitude increases quickly thus reaching the limit and hence the simulation stops. On the other hand, if the the initial conditions belong to a stable region of phase space they survive until the maximum number of turns is achieved (either 10^5 or 10^6). The first case is represented by the large peak at a short time scale,

i.e., a few minutes for both 10^5 and 10^6 turns (see the inset of the right plot in Fig. 8). The second case is represented by cluster around 8 – 10 hours (10^6 turns) or 40 minutes (10^5 turns) of processing time. Note that the second peak in the distribution of the execution time is approximately shifted by a factor of ten for the case of 10^6 turn with respect to that of 10^5 turns. The simultaneous presence of stable and unstable initial conditions makes the distribution of SixTrack calculation time bi-modal.

The spread of the peak in calculation time for stable initial conditions is due to the presence of volunteers' computers with "slow" and "fast" processors. The BOINC WU processing time includes the delay due to the WU queueing in the system, calculation time on volunteers' computer, and the time needed to register the result. The registration time is important, because we use the two-way redundancy, in which each WU is run at least twice on different volunteers' machines. The results are compared and if they do not match, then the WU is submitted again un-
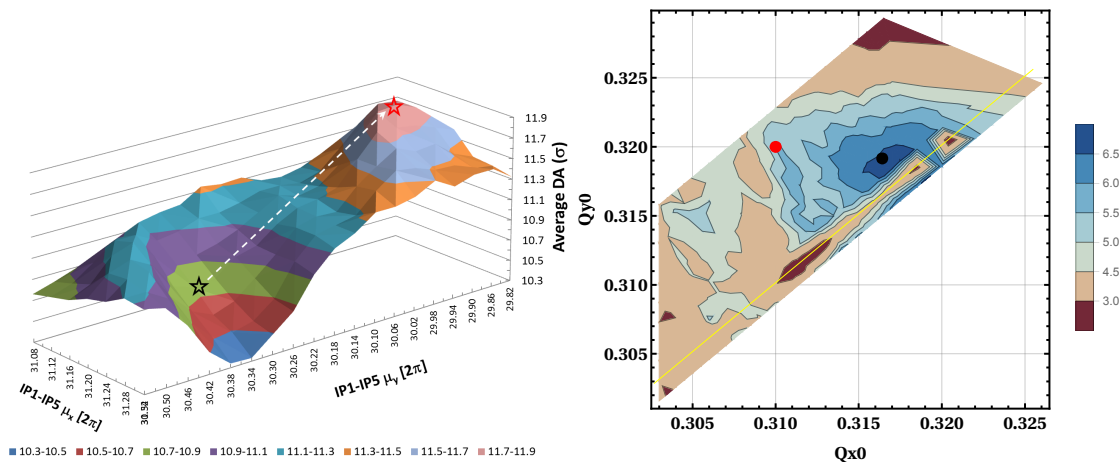
Figure 7: Left: DA averaged over the 60 realisation of the magnetic field errors as a function of the phase advance (horizontal and vertical) between the collision points in ATLAS and CMS. Right: DA as a function of transverse tunes including beam-beam interaction between bunches of 2.2×10^{11} protons.

til a matching pair is found or a maximum number of trials is reached. The current system configuration considers 5 maximum trials. For instance, a 7-day sample statistics for a randomly-selected processing week is reported in Table 1.

The BOINC server transitions its state from UNSENT to IN_PROGRESS to OVER and the last is the final server state when the outcome is known. All OVER results are marked either with the SUCCESS label or with different categories of errors. In particular, the SUCCESS is set only if the reply from the processing is received and there is no client error. If the result has state SUCCESS it is submitted to the validation with the state transitions from INIT to VALID or INVALID or INCONCLUSIVE or other errors detected during the processing. The INCONCLUSIVE state requires BOINC server to resubmit the WU for processing to obtain additional candidates to enter the validation and reach the consensus.

About 10% of computations need to be re-submitted to recover from a computational error or other failures. The analysis of these errors has not been completed yet. Nevertheless, the BOINC submission process makes the recovery transparent for the user.

It is remarkable that there are multiple self-consistent and successful results that differ when compared bit-for-bit. Note that SixTrack is a program that provides 0-ULP identical output for any variation of the compiler and OS it is running on. It has been found previously that the failure to validate self-consistent results is correlated to the usage of computers over-clocked by the volunteers. Non-valid results are eliminated from the physics analysis when using the two-way redundancy as described. It is worth noting that the rejection rate due to the invalids is only 0.37%.

This parameter is being scrutinised for consistency with different studies, which may result in relaxing the need for redundancy in the future.

The WU total processing times and the queue waiting time are plotted in Fig. 9.

The queue waiting time is computed from difference between the sent_time and the create_time for each WU. The total time is computed as the elapsed time between the creation time of the WU and the registration of the result after validation. This is shown as blue lines in the plots in Fig. 9, separately for the runs with 10^5 (upper panel) and 10^6 (lower panel) turns. The total time includes also the execution time on the volunteers' machines shown in Fig. 8. The random process of the results arrival has the distinctive feature that the majority of results is sent for processing within the first 30 minutes since submission. This concerns 60% of the WUs for the 10^5 turns run and 90% of the WUs for the 10^6 turns run. Then the distribution features a long tail of WUs proceeding to execution in small batches over an extended period of time.

The percentage of the WUs submitted immediately may depend on the occupancy of the volunteers' machines and on the state of the queued tasks on the BOINC server. As resources become available, the WUs are sent out for processing at the rate of $\approx 1 - 10$ tasks/minute. This may be estimated from the plots in Fig. 9 given the binning is 10 minutes (10^5 turns run) and 30 minutes (10^6 turns run) in the two histograms. Nevertheless, the long tail of the distribution is also generated by the two-fold redundancy, since the result can only be validated if two or more results are available for the comparison.
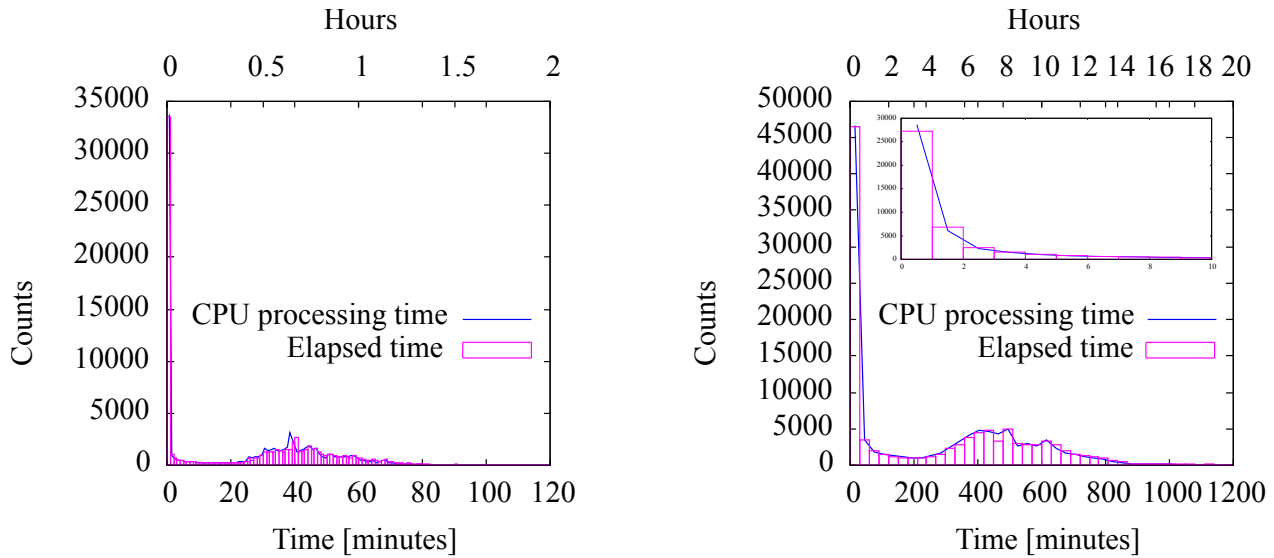
Figure 8: Distribution of the processing time of SixTrack runs on volunteers' computers. The bin width is 1 minute for the 10^5 turn runs (left panel) and 30 minutes for the 10^6 turn runs (right panel). The inset shows the distribution of the first 10 minutes using a 1 minute bin.

Table 1: Statistical distribution of different category computational results for a 7-day time interval.

Category	Size	Fraction [%]	Comments
Total number of completed results (`OVER` state)	639, 225	100	There are in addition 77, 600 results in progress (for a total of 716, 825 results)
Computational success (`SUCCESS` state)	576, 179	90.14	Results delivered without errors to the BOINC server
Computational errors	47, 269	7.39	Automatic WU re-submission
Other (no reply, abort, etc.)	15, 777	2.47	
Valid results	528, 490	100	There are in addition 45, 406 `INIT` results that need the pairing result for validation
Invalid results	1, 966	0.37	Results that do not pass the two-way redundancy
Others (no consensus, etc.)	317		

This first statistical study of the SixTrack execution profile shows the need for tuning the WU submission parameters to reduce the long tail in the server queue. It may suggest also a different strategy for treating the results. In fact, it may be more efficient to cut off the long tail in the total processing time by aborting all the WUs that have not completed before the cut-off time. After this clean-up, the same WUs can be resubmitted as a new batch that will be processed faster.

4.3 Future challenges for SixTrack

The CERN Future Circular Collider (FCC) [48], a 100 TeV centre-of-mass energy collider, is one of the options for future large-scale particle physics experiments [49]. Design studies involving world-wide collaborative efforts are in full swing [50]. FCC is a true challenge, both in terms of accelerator physics as well as from the computational standpoint and the huge capacity offered by volunteer computing is an added value. In fact, while the LHC lattice is made
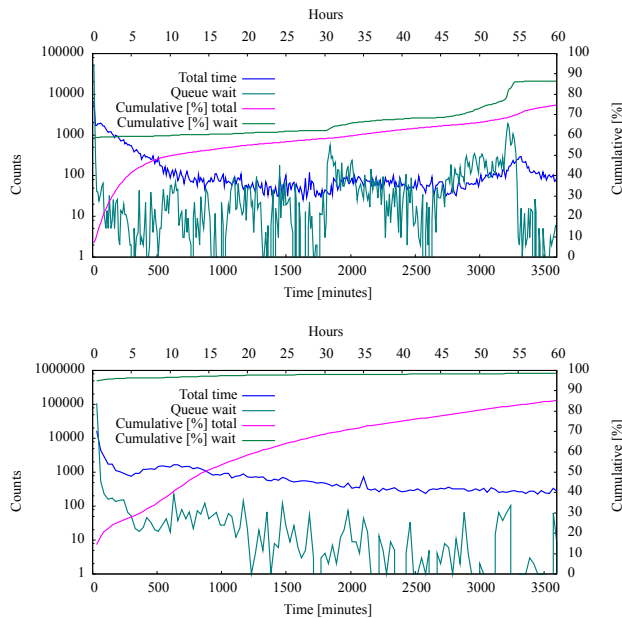
Figure 9: Distribution of the total processing time from the WU submission to the result acquisition and the waiting time component shown for the 10^5 turns run (upper panel) and 10^6 turns run (lower panel). The cumulative distribution is plotted against the right linear axis, showing the fraction of acquired results. Note the log-scale for the histogram entries on the left axis.
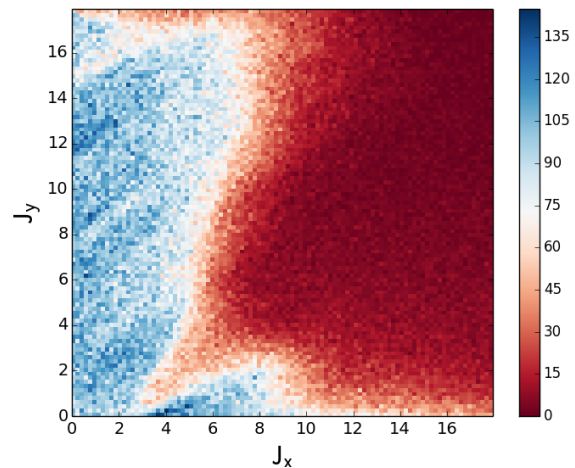

Figure 10: Particle distribution in action space as computed from numerical simulations performed with SixTrack. The configuration refers to the LHC ring at injection energy with Landau octupoles powered with a current of 35 A. The colour scale indicates the number of particles per bin.

of 2.3×10^4 elements, the FCC hadron collider is built out of at least more than 10×10^4 elements. Furthermore, the longer straight sections in which the experiments are hosted will increase the number of beam-beam interactions and these interactions are particularly expensive in terms of CPU-power [51]. Therefore, for FCC DA studies, an increase in CPU power of about a factor 3 – 5 is expected. As a consequence, a single case of dynamic aperture studies might turn to be out of reach for typical batch systems, while about 6 – 7 days would be needed on LHC@Home.

The study of the evolution of distributions of initial conditions to mimic a real beam is yet another challenge ahead of us. This could address questions concerning, e.g., collective instabilities in presence of beam-beam effects [52–54], with particular emphasis on the aspects of Landau damping and its loss turning the beam unstable [55].

The Landau damping studies [54] performed for the LHC, have shown the huge amount of computational resources needed to describe accurately the beam dynamics in the presence of beam-beam effects and the magnetic non-linearities. The reason for such an increase in computational needs is mainly due to the large number of macro particles, of the order of 10^4 – 10^5, needed to describe the actual beam distribution in real space.

An example of tracked particle distribution in the presence of strong non-linearities introduced by octupole magnets at injection energy in the LHC is shown in Fig. 10. A non-uniform change of the particle distributions in action space is visible with direct implications to Landau damping. The changes in the distribution comes from particles losses and clustering around resonances. For these studies, assuming a computing time of 500 µs per turn and 10^6 turns, corresponding to only $\sim$ 90 s of orbit revolutions in the LHC, several days of CPU-time would be needed on a typical batch system. On the other hand, about 1 – 2 days would be needed on LHC@Home system, assuming an average of 4×10^3 CPUs available. Once more, LHC@Home proves to be an essential tool for very detailed beam dynamics simulations.

Finally, the simulation of beam losses induced by the interaction between the beam and the jaws of collimators used to clean the beam halo [56] is also another domain of beam dynamics where the computing needs are beyond the capabilities of standard facilities and volunteer computing represents an ideal solution.

5 Conclusions and Outlook

This paper provided an overview of the activities carried out at CERN in terms of setting up an efficient infrastructure for making use of the huge CPU capacity offered by volunteer computing. Indeed, volunteer computing with BOINC has proven to bring in significant resources for simulations for accelerator physics and HEP community.

Thus expanding the number of volunteers taking part to LHC@Home is our long-term goal.

Thanks to virtualisation, the scope of applications that may run under BOINC has been widened. As use of virtualisation with volunteer computing implies more complexity and overhead for the volunteers, potential simplification, e.g., with container technology, should be investigated further.

The volunteer computing strategy at CERN is to integrate the volunteer computing tool chain with the HTCondor batch system used for computing on batch, cloud, and Grid computing resources. This approach will make it easier for scientists to submit work to different resources allowing the IT team to direct workloads to the appropriate ones. In this respect, further attention is needed to evolve the BOINC middleware and to improve the integration with HTCondor. Development at CERN on the HTCondor-BOINC remote submission gateway is on-going, and will be brought into the BOINC code base. It should be mentioned that also other BOINC projects have posed questions on BOINC software governance and the inclusion of contributed code. Hence an effort to evolve the BOINC community software with contributions from major BOINC projects and stakeholders is required to ensure a long-term future for BOINC and the current volunteer computing community [57].

More specifically for the case of SixTrack, the computational problem in accelerator physics is largely one of throughput and the number of processors available is more important than the per processor performance. Therefore, by providing support for ARM processors with Android (tablets and smart phones) and for Raspberry Pi, an even larger number of processors can be made available for at least the SixTrack application. We are also working on the porting of the SixTrack application to use GPU resources. In fact, since most computers used by volunteers have graphics processors, usage of GPUs might generate an estimated five- to ten-fold increase of the throughput of SixTrack jobs running on same number of volunteers' computers.

It is worth stressing that SixTrack is undergoing major development efforts to open up new domains of accelerator physics, needed for a better understanding of current and future circular particle colliders. LHC@Home is the ideal environment to exploit at best the new code capabilities in view of massive numerical simulations.

References

[1] https://boinc.berkeley.edu

[2] McIntosh E. and Wagner A., CERN Modular Physics Screensaver or Using Spare CPU Cycles of CERN's Desktop PCs, In: A. Aimar, J. Harvey, N. Knoors (Eds.), Proceedings of 14th International Conference on Computing in High-Energy and Nuclear Physics (27 September 1 October 2004, Interlaken, Switzerland), CERN Geneva 2005, 1055-1058

[3] http://cpss.web.cern.ch/cpss/

[4] Schmidt F., SixTrack Version 4.5.17, Single Particle Tracking Code Treating Transverse Motion with Synchrotron Oscillations in a Symplectic Manner, User's Reference Manual, CERN/SL/94-56 (AP)

[5] http://sixtrack.web.cern.ch/sixtrack-ng/

[6] http://garfield.web.cern.ch/garfield/

[7] Lombraña González D., Grey F., Blomer J., Buncic P., Harutyunyan A., Marquina M., et al., Virtual machines & volunteer computing: Experience from LHC@Home: Test4Theory project, PoS ISGC2012 036 (2012)

[8] Buckley A., Butterworth J., Gieseke S., Grellscheid D., Hoche S., Hoeth H., et al., General-purpose event generators for LHC physics, Phys. Rept. 504, 145 (2011), http://www.montecarlonet.org

[9] Høimyr N., Blomer J., Buncic P., Giovannozzi M., Gonzalez. A., Harutyunyan A., et al., BOINC service for volunteer cloud computing, J. Phys.: Conf. Ser. 396 032057 (2012)

[10] Boinc2docker an approach to run BOINC applications with Docker: https://github.com/marius311/boinc2docker

[11] Karneyeu A., Mijovic L., Prestel S. and Skands P.Z., MCPLOTS: a particle physics resource based on volunteer computing, Eur. Phys. J. C 74, 2714 (2014), http://mcplots.cern.ch

[12] Høimyr N., Marquina, M., Asp T., Jones P.,Gonzalez, A., Field L., Towards a Production Volunteer Computing Infrastructure for HEP, J. Phys.: Conf. Ser. 664 022023 (2015)

[13] Sjobak K., De Maria R., McIntosh E., Mereghetti A., Barranco J., Fitterer M., et al., New features of the 2017 SixTrack release, In: G. Arduini, M. Lindroos, J. Pranke, V. RW Schaa, M. Seidel (Eds.), Proceedings 8th International Particle Accelerator Conference (14-19 May 2017, Copenhagen, Denmark), JaCOW, 2017, 3815-3817

[14] McIntosh E., Schmidt F. and de Dinechin F., Massive Tracking on Heterogeneous Platforms, In Proceedings of 9th International Computational Accelerator Physics Conference (2-6 Oct 2006, Chamonix, France), 2006, 13-16

[15] Daramy C., Defour D., de Dinechin F., Muller J.-M., CR-LIBM: a correctly rounded elementary function library, In: F. T. Luk (Ed.) Proceedings of Advanced Signal Processing Algorithms, Architectures, and Implementations XIII, (5-7 August 2003, San Diego, California, USA), SPIE 2003, 458-464 doi:10.1117/12.505591

[16] MSYS2 software distribution homepage, http://www.msys2.org/

[17] SixTrack source repository http://github.com/SixTrack/SixTrack

[18] https://abp-cdash.web.cern.ch/abp-cdash/index.php?project=SixTrack

[19] Buckley A., Butterworth J., Grellscheid D., Hoeth H., Lonnblad L., Monk J., et al., Rivet user manual, Comput. Phys. Commun. 184, 2803 (2013)

[20] Skands P.Z., Carrazza S. and Rojo J., Tuning PYTHIA 8.1: the Monash 2013 Tune, Eur. Phys. J. C 74, no. 8, 3024 (2014)

[21] Dissertori G., Hörtnagl A., Kuhn D., Marie L.K., Rudolph G., Betteridge A.P., et al. [ALEPH Collaboration], Studies of quantum chromodynamics with the ALEPH detector, Phys. Rept. 294, 1 (1998)

[22] Agostinelli S., Allison S., Amako K., Apostolakis J., Araujo H., Arce P., et al., Geant4 - A Simulation Toolkit, Nucl. Instrum. & Methods A 506 250-303 (2003)

[23] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, J. Inst. 3 S08003 (2008)

[24] Buncic P., Aguado Sanchez C., Blomer J., Franco L., A. Harutyunian A., Mato P., et al., CernVM - a virtual software appliance for LHC applications, J. Phys.: Conf. Ser. 219 042003 (2010)

[25] Blomer J., Berzano D., Buncic P., Charalampidis I., Ganis G., Lestaris G., et al., Micro-CernVM: slashing the cost of building and deploying virtual machines, J. Phys.: Conf. Ser. 513 032007 (2014)

[26] Aguado Sanchez C., Blomer J., Buncic P., Franco L., Klemer S. and Mato P., CVMFS - a file system for the CernVM virtual appliance, J. Phys.: Conf. Ser. 52 042003 (2010)

[27] NorduGrid Collaboration, http://www.nordugrid.org

[28] Ellert M., Grønager M., A. Konstantinov A., Kónyae B., Lindemann J., Livenson I., et al., Advanced Resource Connector middleware for lightweight computational Grids, Future Gener. Comput. Syst. 23 pp. 219-240 (2007)

[29] Filipcic A. for the ATLAS Collaboration, arcControlTower: the System for Atlas Production and Analysis on ARC, J. Phys.: Conf. Ser. 331 072013 (2011)

[30] Maeno T. for the ATLAS Collaboration, PanDA: distributed production and distributed analysis system for ATLAS, J. Phys.: Conf. Ser. 119 062036 (2008)

[31] CMS collaboration, The CMS experiment at the CERN LHC, J. Inst. 3, S08004 (2008)

[32] Mascheroni M., Balcas J., Belforte S., Bockelman B.P., Hernández J.M., Ciangottini D., et al., CMS distributed data analysis with CRAB3, J. Phys.: Conf. Ser. 664, 062038 (2015)

[33] Jones C.D., Paterno M., Kowalkowski J., Sexton-Kennedy L. and Tanenbaum W., The New CMS Event Data Model and Framework, In: (Ed.), Proceedings of 15th International Conference on Computing in High-Energy and Nuclear Physics (13-17 February

2006, Mumbai, India), McMillan Mumbai 2006, 248-251

[34] Bockelman B., Cartwright T., Frey J., Fajardo E.M., Lin B., Selmeci M., et al., Commissioning the HTCondor-CE for the Open Science Grid, J. Phys.: Conf. Ser. 664, 062003 (2015)

[35] http://wlcg-public.web.cern.ch/

[36] https://cds.cern.ch/record/2002390

[37] Fajardo E., Gutsche O., Foulkes S., Linacre J., Spinoso V., Lahiff A., et al., A new era for central processing and production in CMS, J. Phys.: Conf. Ser. 396, 042018 (2012)

[38] Sanchez-Hernandez A., Egeland R., Huang C-H., Ratnikova N., Maginie N. and Wildish T., From toolkit to framework - the past and future evolution of PhEDEx, J. Phys.: Conf. Ser. 396, 032118 (2012)

[39] LHCb Collaboration, The LHCb Detector at the LHC, J. Inst. 3, S08005 (2008)

[40] Buncic P., Aguado Sanchez C., Blomer J., Franco L., Harutyunian A., Mato P., et al., CernVM - a virtual appliance for LHC applications, J. Phys.: Conf. Ser. 219 042003 (2010)

[41] Tsaregorodtsev A., Bargiotti M., Brook N., Ramo A.C., Castellani G., Charpentier P., et al., DIRAC: a community Grid solution, J. Phys.: Conf. Ser. 119 062048 (2008)

[42] http://diracgrid.org

[43] Giovannozzi M., Proposed scaling law for intensity evolution in hadron storage rings based on dynamic aperture variation with time, Phys. Rev. ST Accel. Beams 15, 024001 (2012)

[44] Crouch M., Appleby R., Barranco García J., Buffat X., Giovannozzi M., Maclean E., et al., Dynamic aperture studies of long-range beam-beam interactions at the LHC, In: G. Arduini, M. Lindroos, J. Pranke, V. RW Schaa, M. Seidel (Eds.), Proceedings 8th International Particle Accelerator Conference (14-19 May 2017, Copenhagen, Denmark), JaCOW, 2017, 3840-3842

[45] Crouch M., Luminosity Performance Limitations due to the Beam-Beam Interaction in the Large Hadron Collider, The Manchester University PhD thesis, in press

[46] Apollinari G., Béjar Alonso I., Brüning O., Lamont M., Rossi L. (eds.), et al., High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report, CERN, Geneva, 2015, https://cds.cern.ch/record/2116337/

[47] Pieloni T., Banfi D., Barranco J., Dynamic Aperture Studies for HL-LHC with beam-beam effects, CERN-ACC-NOTE-2017-0035 (2017)

[48] FCC design studies, https://fcc.web.cern.ch/

[49] Kramer M., The update of the European strategy for particle physics, Phys. Scr. 2013 014019 (2013)

[50] EuroCirCol, EU Horizon 2020 design study project, http://www.eurocircol.eu/

[51] Barranco J., Pieloni T., Buffat X., Furuseth S.V., Beam-Beam Studies for the FCC-hh, In: G. Arduini, M. Lindroos, J. Pranke, V. RW Schaa, M. Seidel (Eds.), Proceedings 8th International Particle Accelerator Conference (14-19 May 2017, Copenhagen, Denmark), JaCOW, 2017, 2109-2111

[52] Buffat X., Herr W., Mounet N., Pieloni T., Stability Diagrams of colliding beams, Phys. Rev. ST Accel. Beams 17 111002 (2014)

[53] Tambasco C., Buffat X., Barranco J., Pieloni T., Impact of incoherent effects on the Landau Stability Diagram at the LHC, In: G. Arduini, M. Lindroos, J. Pranke, V. RW Schaa, M. Seidel (Eds.), Proceedings 8th International Particle Accelerator Conference (14-19 May 2017, Copenhagen, Denmark), JaCOW, 2017, 2125-2127

[54] Tambasco C., Beam Transfer Function measurements and transverse beam stability studies for the Large Hadron Collider and its High Luminosity upgrade, École Polytechnique Fédérale de

Lausanne PhD Thesis, 2017

[55] Berg J.S., Ruggero F., Landau Damping with two-dimensional betatron tune spread, CERN-SL-96-071-AP (1996)

[56] Mereghetti A., Bruce R., Cerutti F., De Maria R., Ferrari A., Fiascaris M., et al., SixTrack for Cleaning Studies: 2017 Updates, In: G. Arduini, M. Lindroos, J. Pranke, V. RW Schaa, M. Seidel (Eds.), Proceedings 8th International Particle Accelerator Conference (14-19 May 2017, Copenhagen, Denmark), JaCOW, 2017, 3811-3813

[57] BOINC workshop 2017, https://indico.cern.ch/event/648533/overview