







The fight against COVID-19: Running Folding@Home simulations on ATLAS resources

Alessandra Forti¹ , Ivan Glushkov² , Lukas Heinrich³ , Mario Lassnig³ , David South⁴ , and Rodney Walker⁵ 

¹School of Physics and Astronomy, University of Manchester, Oxford Rd, Manchester, M139PL, UK

²University of Texas at Arlington, 701 South Nedderman Drive, Arlington, TX 76019 USA

³CERN, Esplanade des Particules 1, 1211 Meyrin, Switzerland

⁴Deutsches Elektronen-Synchrotron (DESY), Notkestr. 85, 22607 Hamburg, Germany

⁵Ludwig-Maximilians-Universität, Am Coulombwall 1, 85748 Garching bei München, Germany

Abstract. Following the outbreak of the COVID-19 pandemic, the ATLAS experiment considered how it could most efficiently contribute using its distributed computing resources. After considering many suggestions, examining several potential projects and following the advice of the CERN COVID Task Force, it was decided to engage in the Folding@Home initiative, which provides payloads that perform protein folding simulations. This paper describes how ATLAS made a significant contribution to this project over the summer of 2020.

1 Introduction

As the outbreak of a novel coronavirus grew into a pandemic in early 2020, the question began to be asked as to what extent the significant resources of the Worldwide LHC Computing Grid WLCG [1] could contribute, in particular to support molecular biology workloads investigating the molecular structure of SARS-CoV-2. The WLCG and wider LHC computing community have been active in volunteer computing since 2004 through the development of the LHC@Home project [2], in which volunteers donate their computational resources to support the LHC physics program. A dedicated COVID task force [3] was established at CERN in March 2020, with relevant contacts such as the World Health Organisation [4] and the European Bioinformatics Institute [5]. Of the many COVID-based volunteer computing projects [6–10] gaining visibility at the time, Folding@Home [6] was identified as the most useful and scientifically sound. At the same time, capitalising on prior developments to streamline software deployment via containers and to integrate heterogeneous resources such as GPUs, ATLAS [11] had already begun to look at Folding@Home, and, in the end, became a significant donor of computational resources to the Folding@Home project.

2 Protein folding as a computing challenge

Viruses are parasites that cannot replicate on their own. They use proteins both to encode their genetic material and to attach themselves to a cell host to use the cell mechanism to replicate the virus. Proteins are macro-molecules with key functions in a large number of biological

processes. Their function is determined by their shape. A core problem in the study of proteins is understanding their spatial configuration (their “folding”) and the dynamics that these large molecules follow through time as they assume their shape. Usually proteins fold into a minimum energy configuration and finding this configuration is computationally very expensive. The movement of different parts of the protein and their folding can be simulated by Markov State Models [12]. The aim is to find a shape in some of the COVID–19 proteins which can be potentially targeted by drugs.

3 The Folding@Home project

Folding@Home is a distributed project of volunteer computing to study protein dynamics, with the objective to help medical and pharmaceutical professionals to develop effective drugs against a variety of diseases. The system was designed originally to run the applications with low priority on the volunteers personal computers but has since then grown to run also on larger resources. The Folding@Home software has a client–server architecture with different components, the most important of which is the FAHClient which downloads the work, checkpoints and uploads the results to the server.

The volunteers are called donors and are organised in teams. There is a cumulative credit system for the work done which results in a ranking of donor/teams, the more credits the higher the ranking. The work is split into Work Units (WU) which are organised into projects. This is very similar to ATLAS organisation of jobs (WU) organised into tasks (projects). A project represents a particular study and donors can chose the projects they want to donate their processor time to.

Credits are given based on the speed and the number of WUs processed. Credits are calculated by simply benchmarking each WU on a reference CPU and using the time it takes for the same WU to run on another processor to assign the credits: the shorter the time the more the credits. This is different from WLCG use of HEPSPEC06 (HS06) [13] benchmarks to calculate the work done by each job on different processors. The Folding@Home system of credit calculation flattens CPUs and GPUs on the same reference CPU by only looking at the walltime. It is a practical decision considering the donors system but not applicable in ATLAS where the resources are pledged.

4 Containerisation of the clients

The standard distribution mechanism for software run on grid resources uses the read-only filesystem CVMFS [14], and standardised workflows exist to make HEP software such as ATLAS Reconstruction and Analysis releases available on it. Recently, an alternative method for software distribution based on Linux container images has been enabled, allowing analysis users and production operators to easily specify full root–filesystems as the execution environment of the workload. The latter method has the advantage of the ability to fully specify the required dependencies within the container image. Thus the workload requires no matching or testing to work on the existing operating system of the compute elements within the grid resources. It also is the primary method to distribute software targeted to run on GPU–enabled queues within ATLAS [15]. Containers are thus suitable for new, non–HEP workloads such as Machine Learning or, as in this case, molecular biology simulation.

Container images were prepared by adding the FAHClient to NVIDIA’s Ubuntu–based images (`nvidia/opencl:runtime-ubuntu18.04`) [16], which allowed access to both CPU and GPU resources. The images were then distributed via the CVMFS filesystem using its container image distribution repository (`unpacked.cern.ch`).

5 Containerised Folding@Home as a workflow

In ATLAS, production workflows are based around so called transforms [17], and the adoption of any new workflow, external or otherwise, requires the development of a dedicated transform for that case. Therefore, the easiest way to introduce a new activity such as Folding@Home was to run it as a user job, which allows for some degree of flexibility. As described above, the Folding@Home application is containerised and the first ATLAS jobs were submitted on a handful of GPUs, using the already established ability to use stand alone containers. More GPUs were later included in the ATLAS effort, as described in section 7.

Turning to CPUs, many more of which are available to ATLAS, the question was how to coordinate the WLCG sites to contribute their resources in such a way as to not affect the standard, day-to-day computing operations of the experiment. This was done by requesting each site, via the associated funding agency representatives, for an opt in of 5% of their resources. This opt-in, which was later raised to 10%, was met with an overwhelmingly positive response. This also avoided the situation where sites would individually redirect resources towards the Folding@Home effort in their own chaotic way, whilst also placing all of the required effort on the shoulders of the central ATLAS distributed computing team.

From the site perspective, ATLAS just runs another activity on their resources whose usage is accounted in HS06 hours, keeping everything consistent. To achieve this, the COVID workflow was further integrated as if it were a physics group activity and the following aspects were included:

- **Authorisation:** A way to authenticate users who submit official ATLAS COVID workflow. This was solved by adding a covid VOMS group with a production role: `/atlas/covid/Role=Production`
- **Resource Allocation:** A way to control the amount of resources allocated for COVID simulations. Once the VOMS group was created we added a dedicated global share, COVID, at the same level as production and analysis.
- **Storage of logs:** A dedicated scope (`group.covid`) created in the data management system, Rucio [18].
- **Monitoring:** The possibility to identify COVID jobs in the monitoring happened automatically with the setting up of the share. The ATLAS has grouping and selecting capabilities based on global share.

All of this allowed ATLAS CPU resources to be grouped under one single Folding@Home donor, `ATLAS_CPU`, with the contributions of the individual sites visible in the ATLAS monitoring, under a well defined COVID activity, thus maximising the visibility both within Folding@Home and ATLAS.

6 Contributions to Folding@Home from CPUs

The ramp up of running COVID jobs on ATLAS CPU resources is shown in Figure 1. The first COVID jobs were submitted to the Tier-0 resource at CERN, which served as a proof of concept. Shortly afterwards, 20,000 job slots (around 22% of the total) were allocated from the High Level Trigger (HLT) farm, which is normally employed [19] for Monte Carlo (MC) simulations during years where there are no collisions from the LHC, such as 2020. A gradual increase of the contributing sites can also be seen, as well as an increase in the total running jobs from May, when the global share of both the HLT farm and the sites were increased to around 30,000 job slots each. An approximate 50:50 ratio of the HLT and the WLCG sites can be seen across the whole period, showing the effectiveness of the global

shares, despite fluctuations in site availability. The main drops in running jobs on the HLT farm occur when it is needed by the Trigger and Data Acquisition (TDAQ) system, so called "TDAQ Weeks".




Figure 1. Running Folding@Home COVID jobs on ATLAS CPU resources from the start of April to the end of June 2020. Each colour represents a different site. The use of global shares to regulate the amount of resources dedicated to COVID can also be seen.

ATLAS ran COVID simulation jobs for Folding@Home for 6 months from April to October 2020, using up to 60,000 job slots, equally divided between the HLT farm and the WLCG sites. Figure 2 shows all ATLAS running jobs for the same period, where the different activities of the experiment can be seen, including all steps of MC simulations, a significant data reprocessing campaign (in yellow), user and group analysis, and the contribution from COVID jobs (in red). More than 60 sites participated in the ATLAS COVID effort, contributing to the *ATLAS_CPU* donor as part of the *CERN & LHC Computing* Folding@Home team alongside CERN and the other LHC experiments as well as many donors, ranging from Laboratories to individual users. In the end, ATLAS completed almost 6 million Folding@Home Work Units, using about 220 kHS06years, which is about 4% of the total computing output of the year, and is equivalent to simulating around 2.1 billion events (3.25 kHS06s/event). At the time of writing, the *CERN & LHC Computing* team sits in position 20 of all time Folding@Home contributors, where the *ATLAS_CPU* donor is in first place in that team [20].

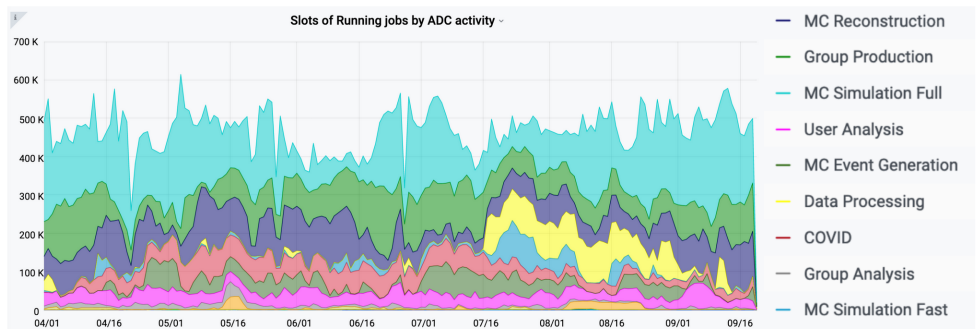


Figure 2. Running job slots on ATLAS CPU resources during 2020, mainly consisting of WLCG sites and the HLT farm, as well as several opportunistic HPC and Cloud resources. The different activities of the experiment are indicated in the legend.

7 Contributions to Folding@Home from GPUs

Whilst Folding@Home payloads exist for CPUs [21], there are also those which benefit from the architecture and high number of cores that only GPUs can provide. With this in mind, ATLAS also made an effort to run Folding@Home on the GPU resources that had been made available to the experiment for Machine Learning workflows. Whilst the job submission mechanism was the same, and the GPUs jobs similarly appeared under COVID activities, the ATLAS monitoring does not account for GPU processing time correctly because there is no benchmark that can be used like HS06. For this reason GPU resources were set up as individual donors in the *CERN & LHC Computing Folding@Home* team, even though the job submission was still centralised.

ATLAS successfully ran Folding@Home jobs on six different sites with GPUs, including two in the UK, two in the USA, one in Italy and one in Germany. Up to 32 simultaneous GPU jobs were running at one time, as can be seen in Figure 3, although it should be noted due to the differences in the GPU installations at each site the jobs in the histogram and not necessarily equal. Nevertheless, it is worth pointing out that four of the six GPU sites employed appear in the top in 50 in the *CERN & LHC Computing Folding@Home* team, as individual donors. The case of the GPUs from DESY, Germany is particularly unique, as it involved introducing a hitherto separate and multi-VO resource into the central ATLAS distributed computing infrastructure, where Folding@Home then also gained from the use of CMS, University of Hamburg and DESY-IT GPU resources.

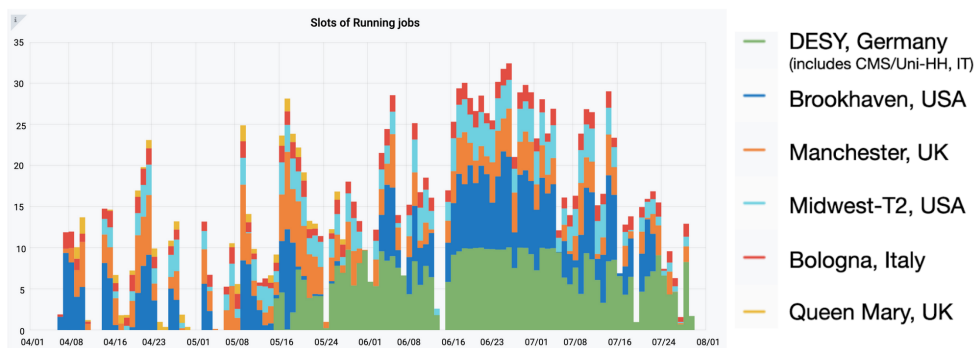


Figure 3. Running Folding@Home COVID jobs centrally submitted by ATLAS on a variety sites with GPUs during 2020. The different sites, which do not feature identical GPU resources, indicated in the legend.

8 Rucio as a data management model for Folding@Home

Rucio is the scientific data management system built by the ATLAS Experiment, and is now in use by a diverse set of communities, from radio astronomy with the Square Kilometre Array to neutrino studies with DUNE. Rucio is horizontally scalable and well-suited to heterogeneous data management across continents and diverse administrative entities, and makes it easy for scientist to interact with their data.

During initial discussions with the Folding@Home team, it became clear that in the future their data requirements both in volume and throughput will grow beyond manually manageable means, and that they are looking for an open solution that is both reliable and efficient.

We agreed that a Rucio service will be set up for Folding@Home, such that their scientists will have a demonstrator and can evaluate if Rucio meets their needs. As a specific use-case, the Rucio instance was setup at CERN with the goal to deliver science data products from the Folding@Home servers to the Hartree High-Performance Computing Centre in the UK for end-user analysis.

This proved surprisingly difficult due to the heterogeneous work servers from Folding@Home using different distributions of Linux, and the absence of a central authentication and authorisation entity, which is the typical case in the HEP world. The central deployment of Rucio itself was straightforward using the new Kubernetes-based installation mechanism. However, since there are dependencies which require deployment at the Folding@Home servers, a significant portion of work had to be invested to make them easily deployable in this environment, ranging from packaging to fixing bugs related to assumptions on the environment. One example includes the GFAL library, which at the necessary version required either dependencies unavailable on the Debian distribution, or was not easily compiled. XrootD was deployed as the storage frontend, however it was quickly discovered that several assumptions on the security infrastructure were made related to certificate expiration, which were not compatible with the newly instituted Folding@Home X.509 environment. While this was eventually fixed in XrootD, it required custom compilation of the packages.

Since the Hartree HPC is within a secure network area, transfers had to be tunneled through RAL's CEPH storage cluster. A custom service is running at Hartree HPC which continuously queries Rucio for the datasets available on the RAL storage, and securely moves data from RAL to Hartree.

Eventually, a working setup has been demonstrated, and is now being moved into production for end-user analysis. The Fermi National Accelerator Laboratory (FNAL) has also expressed interest to host the Rucio installation for Folding@Home in production for the long-term, as the CERN installation has only been procured for demonstration purposes. This will be accomplished through a move of the Rucio database from CERN to FNAL, as FNAL is using the same Kubernetes-based installation mechanism.

9 Conclusions

Through the use of the worldwide distributed computing infrastructure available to them, the LHC experiments were able to quickly react to the global coronavirus pandemic by donating a significant amount of resources to the external volunteer computing project Folding@Home. For ATLAS, a sustained rate of over 60,000 job slots consisting of a mixture of CPU and GPU workloads was achieved, placing the ATLAS experiment within the leading daily individual contributors of the project at large.

The recent infrastructure work within the WLCG context such as container image support for analysis jobs, image distribution via CVMFS and GPU-enabled queues were crucial in terms of deployment. In addition, the coordinated steps taken to fully integrate the Folding@Home workflow into the ATLAS distributed computing environment were key to the success. Not only has the experience gained from employing such resources and containerised workflows been of great benefit to ATLAS, but future adaptation and integration of further external workflows can draw on our experience with Folding@Home. Furthermore, scientific data generated by these and other volunteer's workloads is now managed using the Rucio data management software.

References

- [1] J. Shiers, The Worldwide LHC Computing Grid (Worldwide LCG), *Comput. Phys. Commun.* **177** (2007) 219–223, doi:10.1016/j.cpc.2007.02.021
- [2] J. Barranco *et al.*, LHC@Home: A BOINC-based volunteer computing infrastructure for physics studies at CERN, *Open Eng.* **7** (2017) no.1, 378-392, doi:10.1515/eng-2017-0042
- [3] *CERN against COVID-19*, <https://against-covid-19.web.cern.ch> (2021), accessed: 2021-05-12
- [4] *WHO: Coronavirus disease (COVID-19) pandemic*, <https://www.who.int/emergencies/diseases/novel-coronavirus-2019> (2021), accessed: 2021-05-12
- [5] *Coronavirus data resources at EMBL-EBI*, <https://www.ebi.ac.uk/about/news/announcements/coronavirus-data> (2021), accessed: 2021-05-12
- [6] *Folding@Home*, <https://foldingathome.org> (2021), accessed: 2021-05-12
- [7] *Rosetta@Home*, <https://boinc.bakerlab.org> (2021), accessed: 2021-05-12
- [8] *Gene*, <https://www.ncbi.nlm.nih.gov/gene> (2021), accessed: 2021-05-12
- [9] *BOINC@TACC*, <https://boinc.tacc.utexas.edu> (2021), accessed: 2021-05-12
- [10] *Open Pandemics: COVID-19*, <https://www.worldcommunitygrid.org/research/opn1/overview.do> (2021), accessed: 2021-05-12
- [11] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, *JINST* **3** (2008) S08003, doi:10.1088/1748-0221/3/08/S08003
- [12] V. S. Pande *et al.*, Everything you wanted to know about Markov State Models but were afraid to ask, *Methods* (San Diego, Calif.) vol. 52,1 (2010) 99-105, doi:10.1016/j.ymeth.2010.06.002
- [13] M. Michelotto *et al.*, A comparison of HEP code with SPEC1 benchmarks on multi-core worker nodes, *J. Phys. Conf. Ser.* **219** (2010) 052009, doi:10.1088/1742-6596/219/5/052009
- [14] J. Blomer *et al.*, Status and future perspectives of CernVM-FS, *J. Phys. Conf. Ser.* **396** (2012) 052013, doi:10.1088/1742-6596/396/5/052013
- [15] A. C. Forti, L. Heinrich and M. Guth, Hardware accelerated ATLAS workloads on the WLCG Grid, *J. Phys. Conf. Ser.* **1525** (2020) no.1, 012059, doi:10.1088/1742-6596/1525/1/012059
- [16] L. Heinrich, M. Feickert and M. Lassnig. (2021, January 29). *lukasheinrich/folding-at-home-docker*. Zenodo, doi:0.5281/zenodo.4479723
- [17] G. A. Stewart *et al.*, ATLAS Job Transforms: A data driven workflow engine, *J. Phys. Conf. Ser.* **513** (2014) 032094, doi:10.1088/1742-6596/513/3/032094
- [18] M. Barisits *et al.*, Rucio: Scientific data management, *Comput. Softw. Big Sci.* **3** (2019) 11, doi:10.1007/s41781-019-0026-3
- [19] S. Ballestrero *et al.*, Evolution and experience with the ATLAS simulation at Point 1 Project, *J. Phys. Conf. Ser.* **898** (2017) no.8, 082012, doi:10.1088/1742-6596/898/8/082012
- [20] *Folding@Home Team: CERN & LHC Computing*, <https://stats.foldingathome.org/team/38188> (2021), accessed: 2021-05-12
- [21] *New COVID-19 small molecule screening simulations are running on full Folding@Home*, <https://foldingathome.org/2020/03/30/new-covid-19-small-molecule-screening-simulations-are-running-on-full-foldinghome> (2021), accessed: 2021-05-12