Advances in ATLAS@Home towards a major ATLAS computing resource

David Cameron, University of Oslo
Wenjing Wu, IHEP & CAS
Alexander Bogdanchikov, Budker Institute of Nuclear Physics
Riccardo Maria Bianchi, University of Pittsburgh
On behalf of the ATLAS collaboration

Why volunteer computing for ATLAS?

- “Free” computing resources
- Outreach, letting people contribute and feel part of HEP experiments
- ATLAS@Home started in 2013, in 2017 joined forces with LHC@Home
- More info on LHC@Home and BOINC in [LHC@Home talk]

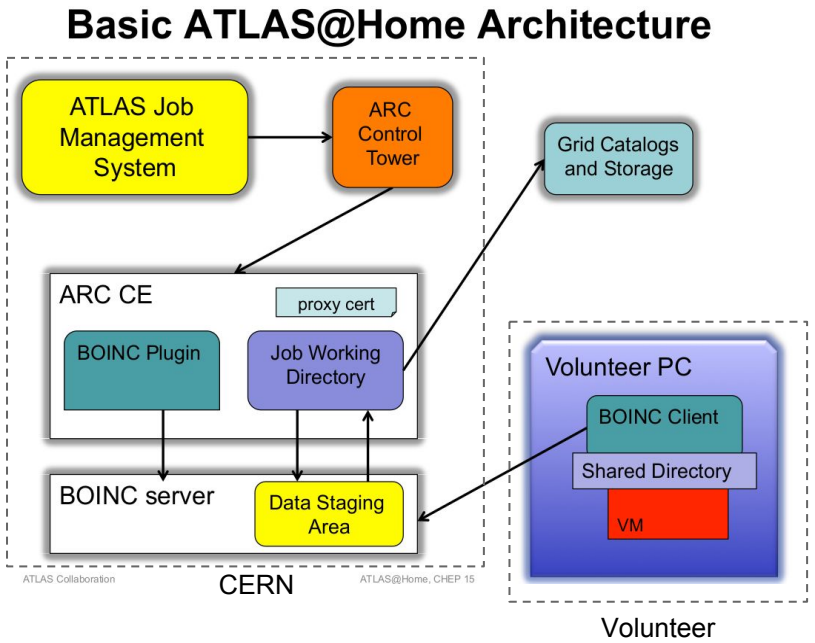

Expanding beyond traditional volunteers

- Traditional volunteer resource of home PCs is static or declining
- Entry barrier to ATLAS@Home for non-experts is rather high compared to other projects
- At the same time ATLAS does not exploit fully Grid CPUs
  - Software has inefficiencies
  - Initialisation and finalisation steps in multi-core jobs
  - Draining and ramping up around downtimes
  - Scheduling inefficiencies (draining a node to run a multi-core job)
- This led to investigating using ATLAS@Home as a backfilling platform
  - Independent from batch system and grid jobs
  - Not affected by grid services downtime
  - By design BOINC does backfilling

Removing the need for VirtualBox

- ATLAS@Home was built around using VirtualBox to provide a uniform execution environment (75% of volunteers use Windows)
- Not necessary for Linux machines
  - Admins not keen on installing VirtualBox on worker nodes
  - Not easily possible to run VM inside another VM (on cloud infrastructures)
- -> Creation of native Linux version
  - Runs natively on SLC6/Centos7
  - Uses a singularity container on other Linux OS
    - Image is standard SLC6 image used on many ATLAS sites
- Install BOINC client on worker nodes and configure to run in background
  - E.g. with high nice value
  - Doesn't affect batch system scheduling since it doesn't know about BOINC processes

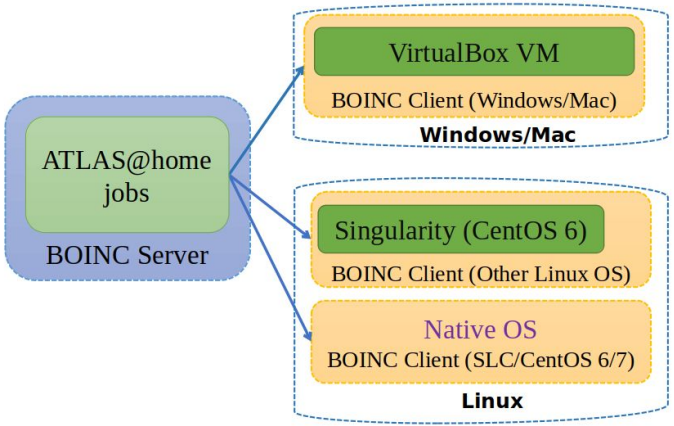

Usage of different apps and versions

CPU days per day for each LHC@Home app in May 2018

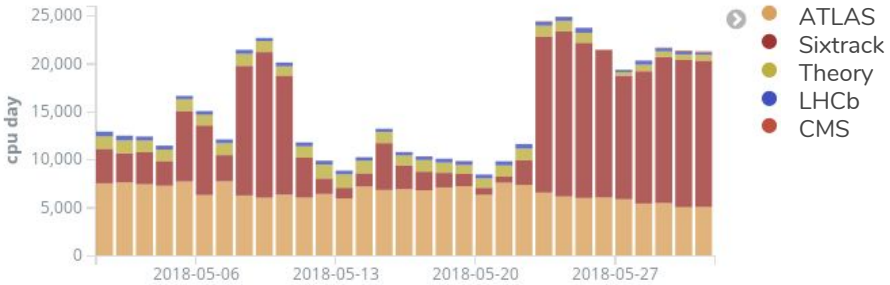

CPU days per day for each version of ATLAS@Home in May 2018

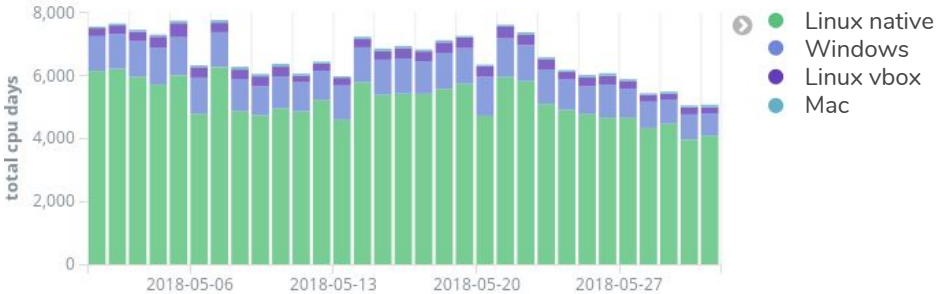

- ATLAS and sixtrack are roughly equivalent in terms of CPU time consumed
- The majority of CPU for ATLAS@Home now comes from the native app

Note: these plots do not include ~3k CPU days/day for ATLAS from LHC@Home development server

Current volunteers

Top 6 (~85%) are ATLAS/CERN-related resources:

1. **Agile Boincers**: Machines being commissioned/decommissioned in CERN computing centre
2. **TRIUMF**: Canadian T1 site ~5k cores in backfill
3. **BEIJING**: Chinese T2 site ~500 cores in backfill
4. **Rod Walker**: German T2 site (LRZ) ~200 cores in backfill
5. **WLCG Performance Cluster**: used for ATLAS software performance testing, ATLAS@Home in background
6. **MPI für Physik**: Office desktops in MPI Munich

More information on how Grid sites are used in "**Backfilling the Grid with Containerized BOINC in the ATLAS computing", Wed 11:45, Track 7**

CPU consumption per volunteer, May 2018

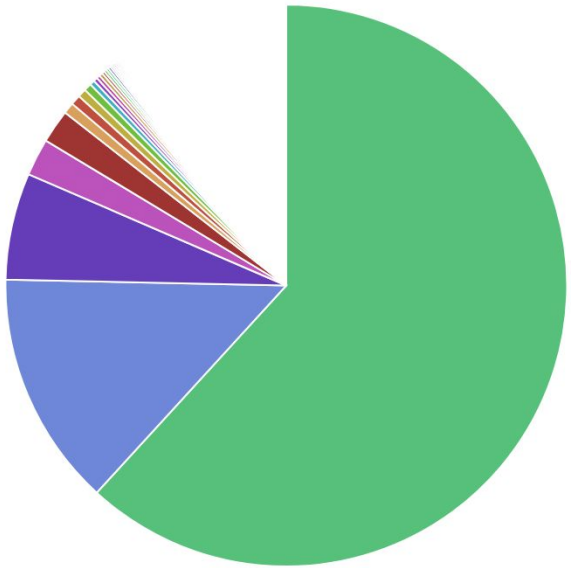

Impact of ATLAS@Home

Events processed per week by ATLAS@Home June 2017 - May 2018

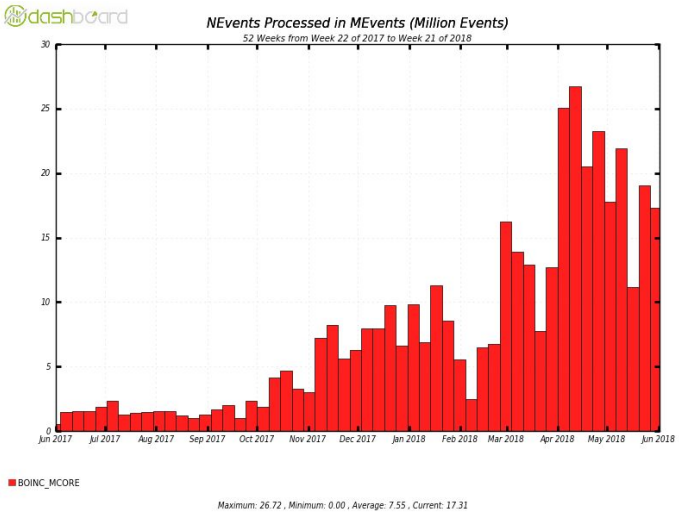

CPU consumption per day of ATLAS@Home jobs in April/May 2018 - equivalent of 10,000 continuous running cores

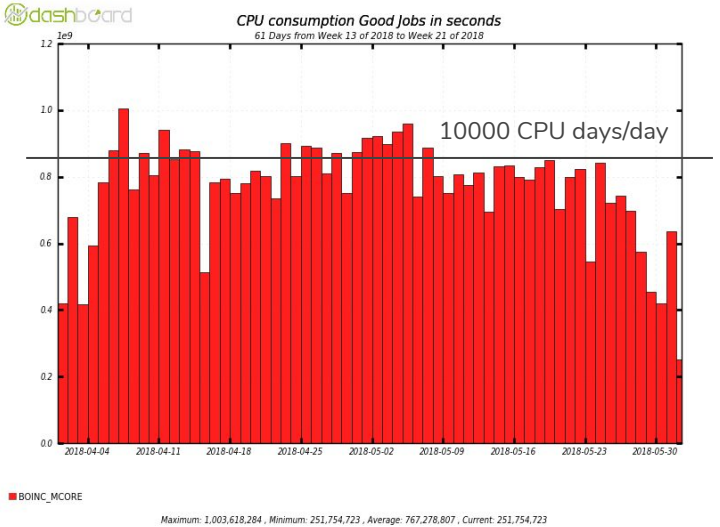

Impact of ATLAS@Home

ATLAS simulation events processed in May 2018

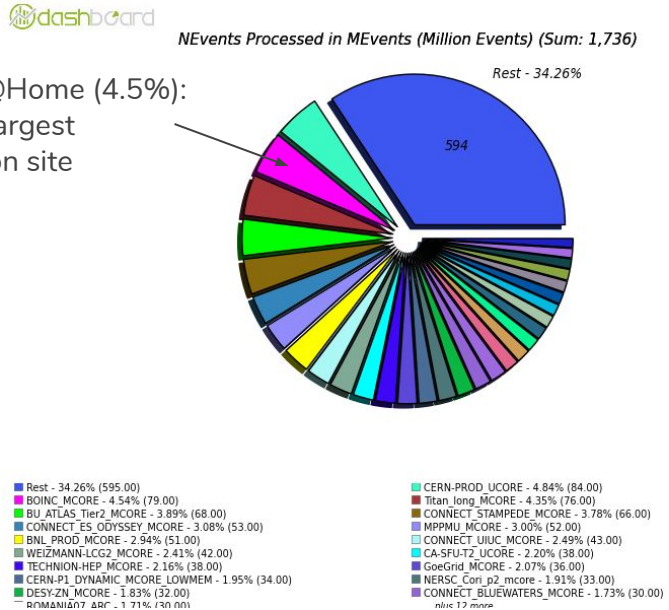

ATLAS@Home (4.5%): second largest simulation site

Average efficiency of ATLAS@Home jobs in May 2018

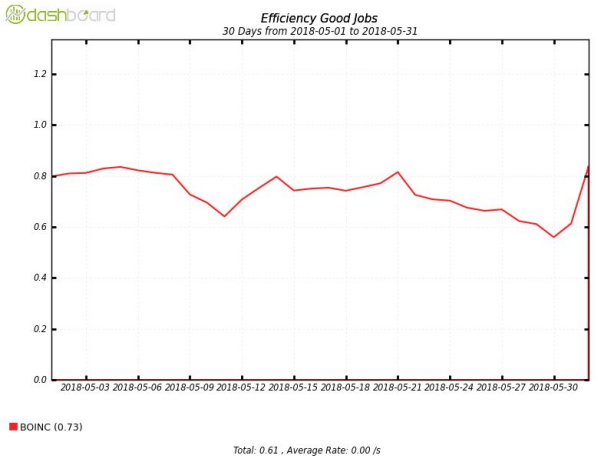

Average efficiency of ATLAS@Home jobs is **73%** compared to **84%** for all ATLAS simulation

Ongoing research and development

Dynamic data staging with event-level granularity
using the ATLAS event service

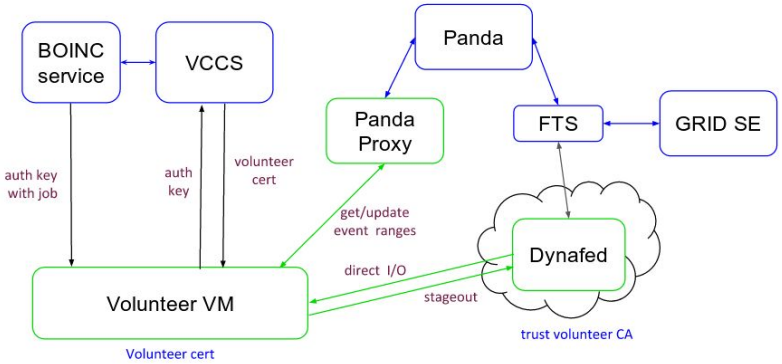

Thank you Mario Rossi!

ATLAS Qualification task E. Rye (UiO):
“Enhanced graphical interface and workloads for ATLAS@Home”

Extending the graphical interface to show events in real time
Exploring new workloads outside MC simulation

Publications/projects

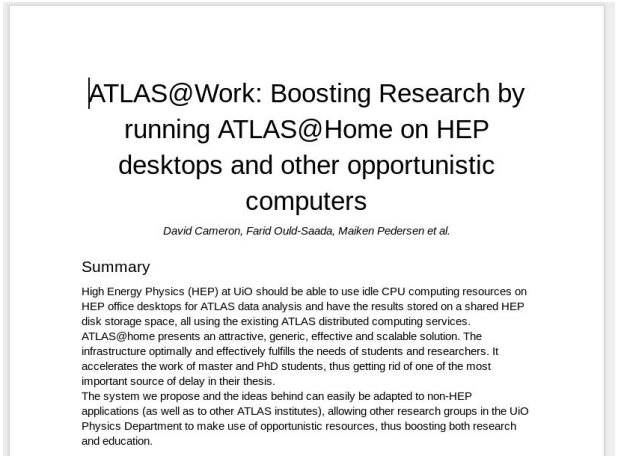

ATLAS@Work: Boosting Research by running ATLAS@Home on HEP desktops and other opportunistic computers

David Cameron, Farid Ould-Saada, Maiken Pedersen et al.

Summary

High Energy Physics (HEP) at UiO should be able to use idle CPU computing resources on HEP office desktops for ATLAS data analysis and have the results stored on a shared HEP disk storage space, all using the existing ATLAS distributed computing services. ATLAS@home presents an attractive, generic, effective and scalable solution. The infrastructure optimally and effectively fulfills the needs of students and researchers. It accelerates the work of master and PhD students, thus getting rid of one of the most important source of delay in their thesis.
The system we propose and the ideas behind can easily be adapted to non-HEP applications (as well as to other ATLAS institutes), allowing other research groups in the UiO Physics Department to make use of opportunistic resources, thus boosting both research and education.

Local University of Oslo project ATLAS@Work: implement private analysis cluster at UiO using ATLAS@Home on office desktops

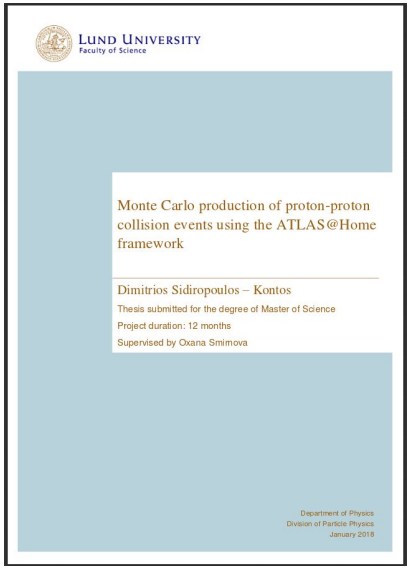

LUND UNIVERSITY
Faculty of Science

Monte Carlo production of proton-proton collision events using the ATLAS@Home framework

Dimitrios Sidiropoulos – Kontos
Thesis submitted for the degree of Master of Science
Project duration: 12 months
Supervised by Oxana Smirnova

Department of Physics
Division of Particle Physics
January 2018

https://lup.lub.lu.se/student-papers/search/publication/8932453

Masters thesis of D. Sidiropoulos Kontos (Lund University): validation of ATLAS@Home as a platform for event generation

Conclusions

- ATLAS@Home has expanded far beyond its traditional “@home” base
- However it is still a useful tool to connect and involve the general public in HEP
- The move towards backfilling Grid sites provides ATLAS extra resources at no extra cost

- Join us! http://lhcathome.web.cern.ch/projects/atlas

Acknowledgements

- Thanks to
  - CERN IT for hosting the BOINC infrastructure and help integrating to LHC@Home
  - All our volunteers not just for resources provided over the years but for support in helping out others with problems
  - Grid site admins willing to risk running ATLAS@Home alongside their Grid jobs