



BOINC Pentathlon 2018 - Observations

Patrick Schöfer

24 Jul 2018

About me

- BOINC enthusiast since 2004
- Co-Administrator of SETI.Germany
- Member of the Pentathlon Organizing Group

- final-year PhD student (stellar astrophysics)

What is the BOINC Pentathlon?

- 14-day multi-project Volunteer Computing team competition
- 5 'disciplines' with different durations at different BOINC projects
- Teams may choose to concentrate on single disciplines, but need a coordinated effort in all disciplines to win the overall crown
- Short-term boost and new volunteers for the projects
- More information at https://www.seti-germany.de/boinc_pentathlon/

This year's Pentathlon projects



How well did the projects handle the load?

- we were pretty lucky with the projects this year, no major problems
 - Universe@Home validator was overloaded in the final few hours
 - Asteroids@home ran dry for a few hours
- great increase in computing power:
 - PrimeGrid, Rosetta@home, Universe@Home were part of the 2016 Pentathlon, total output of Pentathlon teams ~ **tripled**

Problems I – Plan Classes

- Asteroids@home uses different plan classes for different levels of optimization (SSE2, SSE3, AVX)
- server needs some time to figure out which one is the best, but enthusiasts are too impatient 😊
- more advanced users use anonymous platform to force the best-suited application
- no big deal with fixed-credit projects like Asteroids@home, but trouble for projects using CreditNew

Problems II – Credit Systems

- fortunately, three of this year's Pentathlon projects grant fixed credit
- Rosetta@home: average-based system tailored to their needs, but short-term competitions are too short to average out *luck*
- NumberFields@home: CreditNew(ish?) + not completely predictable runtimes may lead to unfairness, e.g. HT vs. non-HT CPUs

Problems III – Pending Credits

- WUs with a minimum quorum > 1 may be validated too late
- some participants claimed that some clients are just created to download lots of tasks and never return them
- some participants started to run only WUs that were already returned by another host

Suggestions

- Plan classes: some projects already have better solutions in place (merged applications or wrappers which choose the application based on CPU generation or internal benchmarks)
- Credit systems: long discussion 😊
- Pending credits:
 - do not show the status of other tasks to the user until the user returns his results (not a true solution, but avoids cherry-picking and makes validation more cheater-proof)
 - PrimeGrid Challenge System