



SC 19 snapshot

Marc Paterno
5 Dec, 2019

Tutorial on Floating-Point software reproducibility

- Variety of tools presented, see <http://fpanalysistools.org>.
- GPU-support tools starting to be available. CUDA behaves as if all floating-point exceptions are masked, which means one must take great care.
 - Floating-point (IEEE-754) exceptions are not C++ exceptions!
- None of the tools are suitable for use on experiment-scale code bases. Important routines would need to be extracted and analyzed.
 - Another reason to want modularity!

Interactive HPC computing workshop

- Data Science/AI applications dominate most of discussion.
- Keynote recognizes also *data exploration* and *workflow experimentation*.
- Almost everyone is using some Jupyter-related thing (-HUB, -Lab, or just Jupyter).
- Everyone is talking about edge services.
- Every speaker was speaking about his or her own system; little seems shared.
- Speaker from NERSC noted LHC use – ATLAS named, nobody else.

Intel OneAPI tutorial

- More information about Aurora was released
- Major news: release of OneAPI
 - An “Open industry specification”: others are welcome to implement it. But who will?
 - Extension of C++ (Data Parallel C++) for direct programming of cpus, GPUs, tensor cores, FPGAs.
- No expectation that one algorithm will be efficient on multiple platforms.
 - You need to write different algorithms for each type of hardware.
 - Contrast this with competitors (e.g. Kokkos, OpenMP offloading).

Other bits

- OpenMP for programming GPUs is being pushed by both IBM and AMD.
- AMD is also pushing ROCm and HIP.
- Met with one of my proposed co-PIs for LDRD!