GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

# ARM porting and benchmark studies for ATLAS and CMS

**David Abdurachmanov, Joshua Wyatt Smith and Graeme Stewart**
**on behalf of ATLAS and CMS**

**CERN openlab Technical Workshop 2016**

**9 December 2016**

BMBF-Forschungsschwerpunkt
ATLAS-EXPERIMENT
Physik bei höchsten Energien mit dem ATLAS-Experiment am LHC
FSP 103
ATLAS

Compact Muon Solenoid
experiment at CERN's LHC
CMS

Bundesministerium für Bildung und Forschung

# Motivation

GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

- **Explore new hardware and software platforms that, in the future, may be more suited to its bulk production workloads, i.e. Simulation**
  - **Performance benchmarks**
  - **Are results consistent? (i.e. validation)**
- **Improves overall code quality**
- **More efficient computing - less energy / computation**
- **Geopolitics plays a role - server farms might be different architectures for various reasons (Russia, Asia, Africa etc.)**
- **Business model of ARM is very flexible -**
  - **Competition, freedom, flexibility**

- **What will our computing model be in Run 3??**

**Disclaimer: I am on the ATLAS experiment (and not a computer scientist)**

# Platforms

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

| Name | Processor | Cores | RAM | Cache | Fabrication (Release) | OS |
|---|---|---|---|---|---|---|
| HP Moonshot | X-Gene, 2.4 GHz | 8 Armv8 | 64 GiB DDR3 (1600 MHz) | 32 KiB L1/core, 256 KiB L2/core pair, 8 MiB L3 | 40 nm (2014) | Ubuntu 14.04 |
| Aarch64_Proto | -, 2.1 GHz | 32 Cortex-A57 | 128 GiB DDR3 (1866 MHz) | 32 KiB L1, 1 MiB L2 | 16 nm (-) | Ubuntu 14.04 |
| Intel Atom | Intel Atom Processor C2750, 2.4GHz | 8 | 32 GiB DDR3 (1600 MHz) | 24 KiB L1d, 32 KiB L1i, 1 MiB L2 | 22 nm (2013) | Fedora 21 |
| Intel | Intel Xeon CPU E5-4650, 2.70 GHz | 32 | 512 GiB DDR3 (1600 MHz) | 32 KiB L1(d)(i)/core, 256 KiB L2/core, 20 MiB L3 | 32 nm (2012) | Scientific Linux CERN 6 |

| | Cavium ThunderX | AMD A1100 | APM X-Gene 1 | Intel Xeon D-1540 | Intel Xeon E5-2699 v3 | Intel Xeon E5-2699 v4 |
|---|---|---|---|---|---|---|
| **Cores / Threads** | 48C/48T | 8C/8T | 8C/8T | 8C/16C | 18C/36T | 22C/44T |
| **Speed (GHz)** | 2.5 | 2.0 | 2.4 | 2.0 (2.6) | 2.3 (3.6) | 2.2 (3.6) |
| **Cache (L3)** | 16MB | 8MB | 8MB | 12MB | 45MB | 55MB |
| **SMT** | 1 | 1 | 1 | 2 | 2 | 2 |
| **Year** | Q4/15 | Q1/2014 | Q3/13 | Q1/15 | Q3/14 | Q1/16 |

| | Intel Xeon E5-2698 v3 | Intel Xeon E5-2683 v4 | Intel Xeon E5-2680 v4 | Intel Xeon D-1581 | IBM POWER8 | IBM POWER8 |
|---|---|---|---|---|---|---|
| **Cores / Threads** | 16C/32T | 16C/32T | 14C/28T | 16C/32T | 8C/64T | 10C/80T |
| **Speed** | 2.3 (3.6) | 2.1 (3.0) | 2.4 (3.3) | 1.8 (2.4) | 3.857 | 3.491 |
| **Cache (L3)** | 40MB | 40MB | 35MB | 24MB | 64MB | 80MB |
| **SMT** | 2 | 2 | 2 | 2 | 8 | 8 |
| **Year** | Q3/14 | Q1/16 | Q1/16 | Q1/16 | Q3/13 | Q3/13 |

# Software stacks

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

ATLAS EXPERIMENT

- ATLAS Codebase (Athena) ~6.5 million lines of code
  - ~2400 packages
- AthSimulation is a subset of Athena
  - ~350 packages

- Validate results
- Benchmark Performance

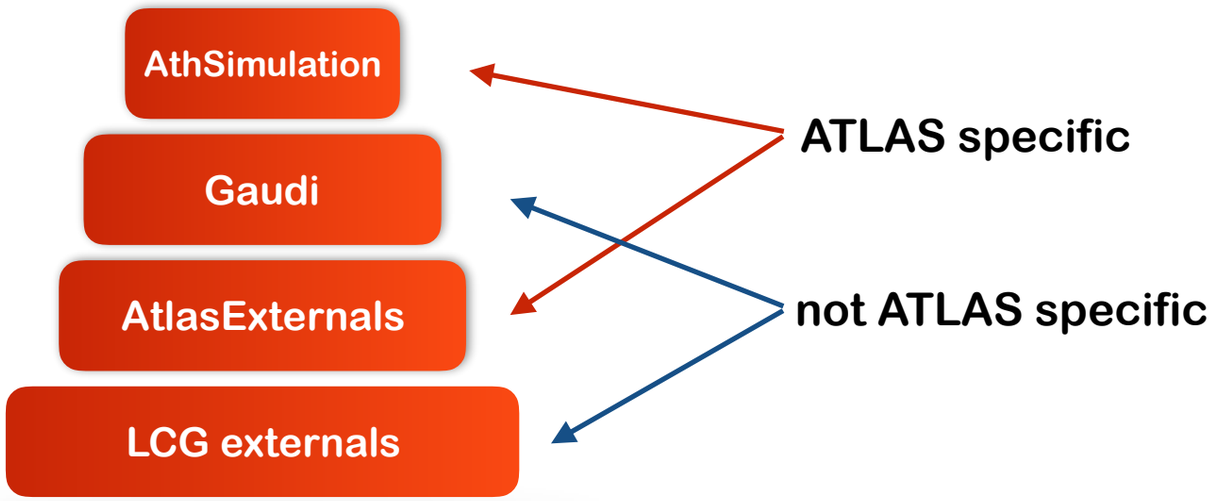


http://lcgsoft.web.cern.ch/lcgsoft/
(for full list of packages)

CMS

- CMSSW ~6 million lines of code
- Open-source and available in GitHub
- Full software stack
- "CMSSW is like Software Collection package or Linux Container without actually being any of them" - David A.

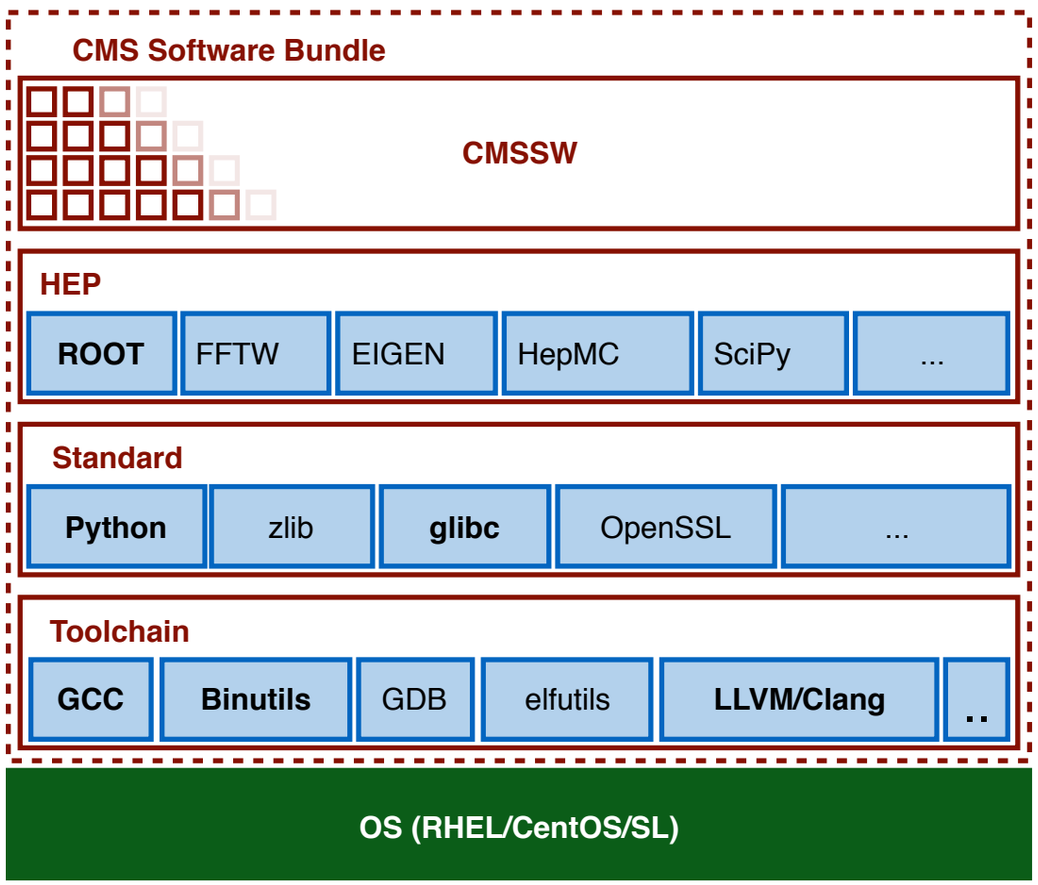


CVMFS

# Validation



- **Simulation is a MC process**
  - **numerical identity isn't expected**
  - **But different trends are clear warning signs!**

## Reconstructed hits in ATLAS SCT detectors

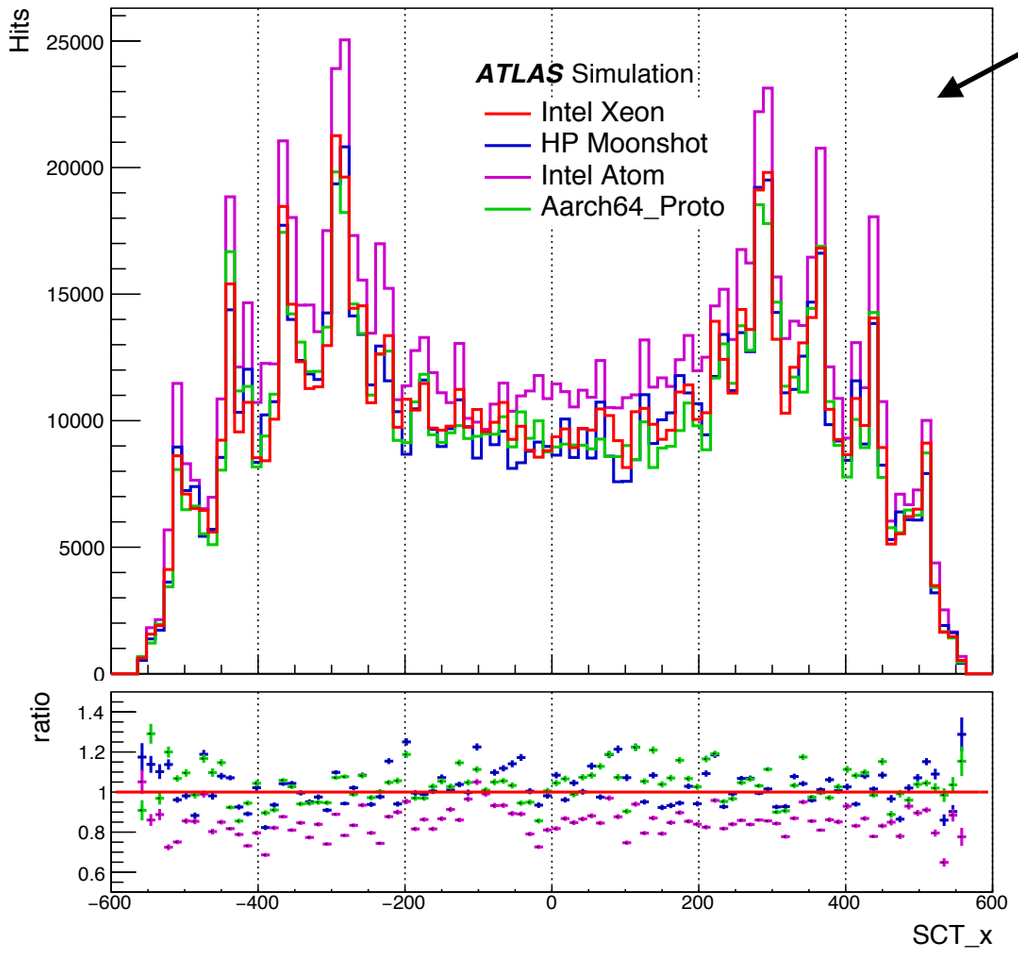


**Intel Atom shows clear trend of more events**

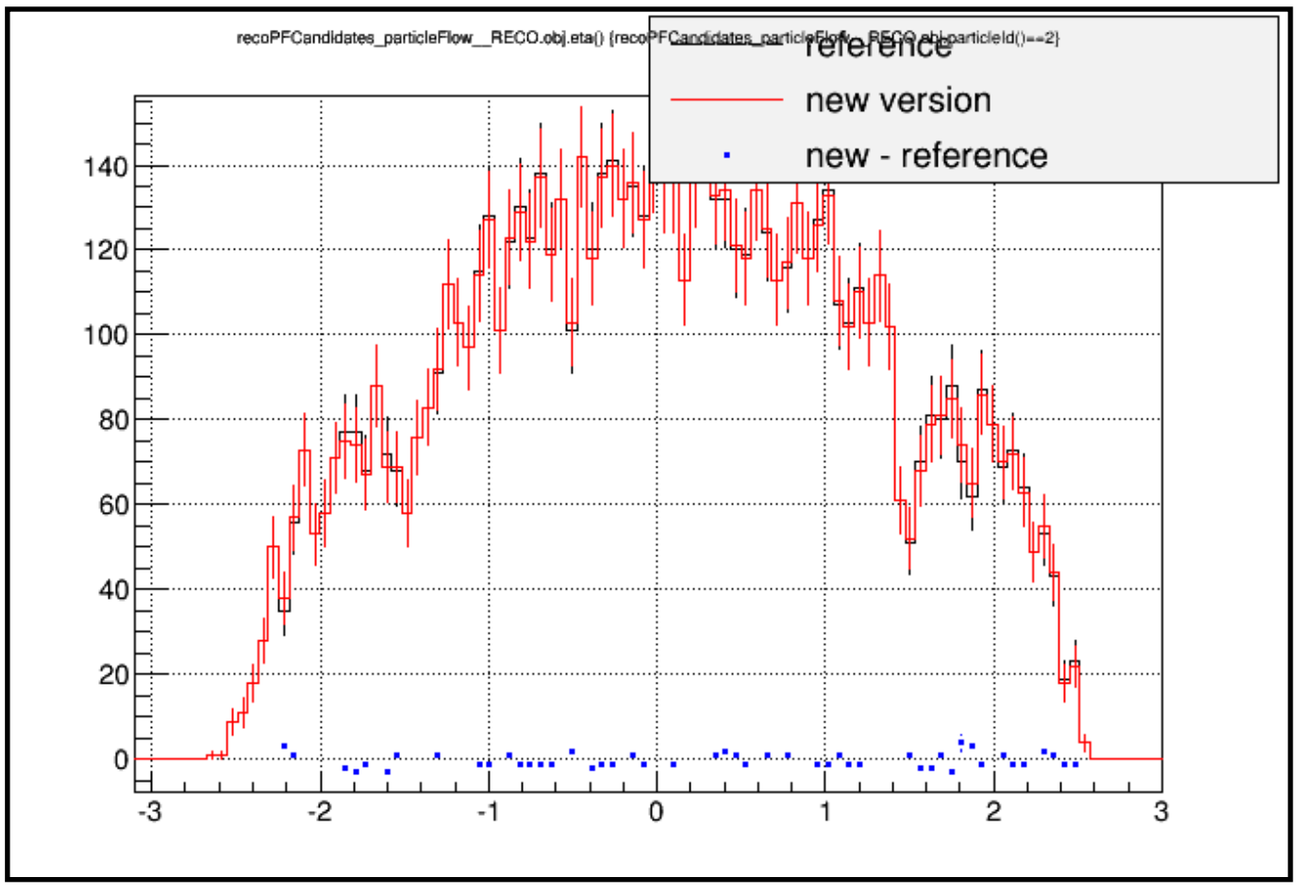

# Validation

GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

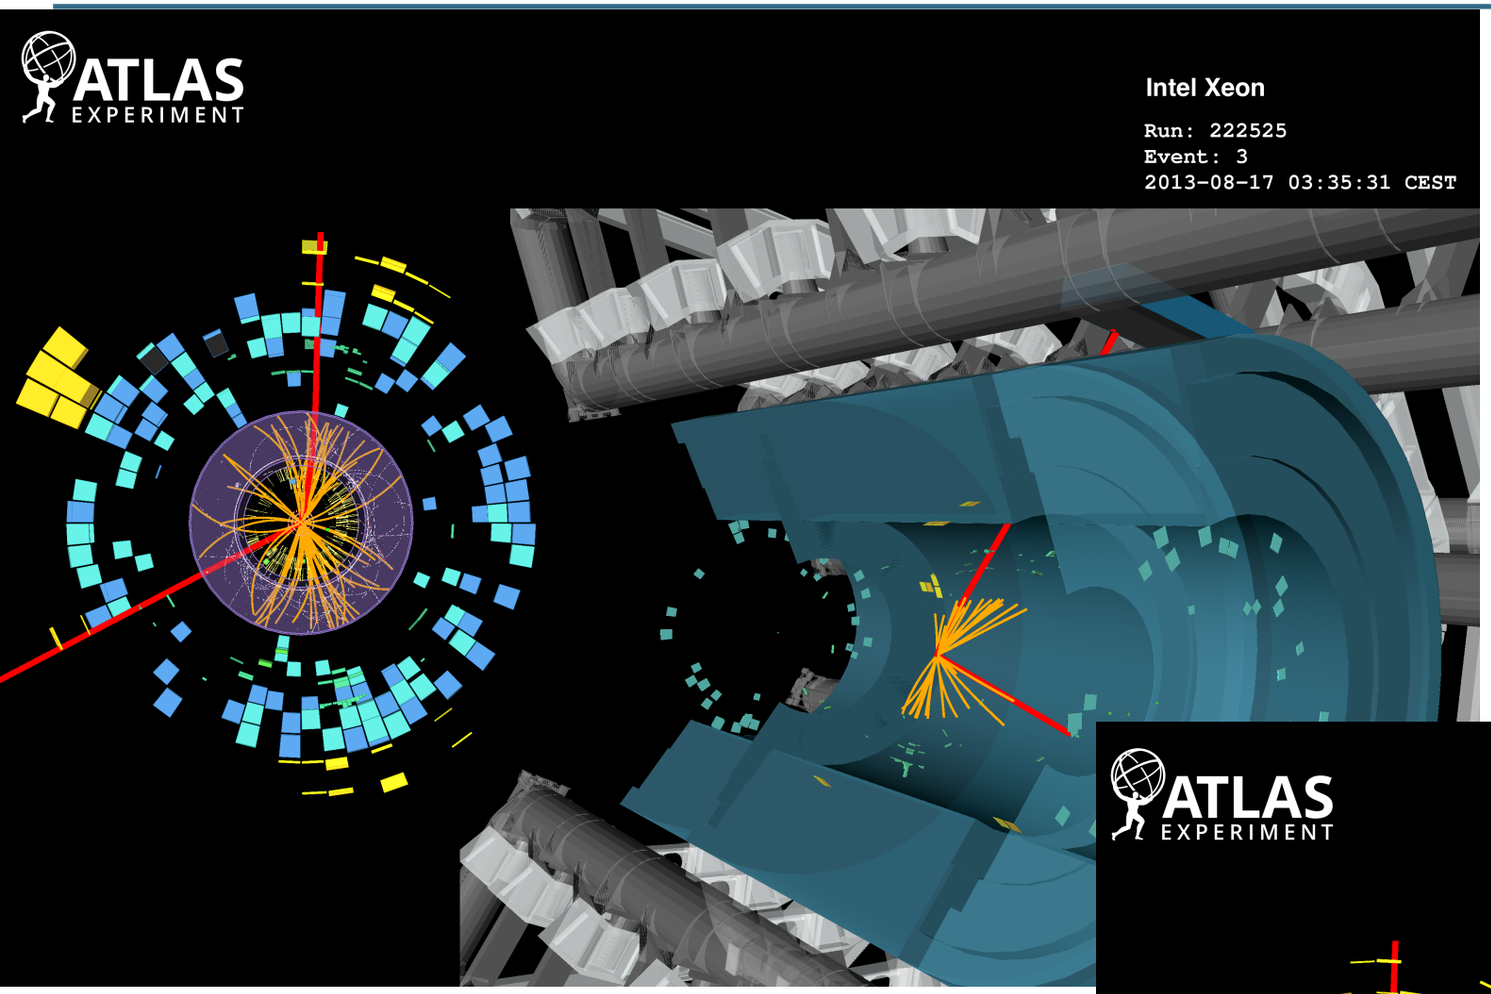


## Intel Xeon

**Same, single $t\bar{t}$ event**

**Spot the differences…**

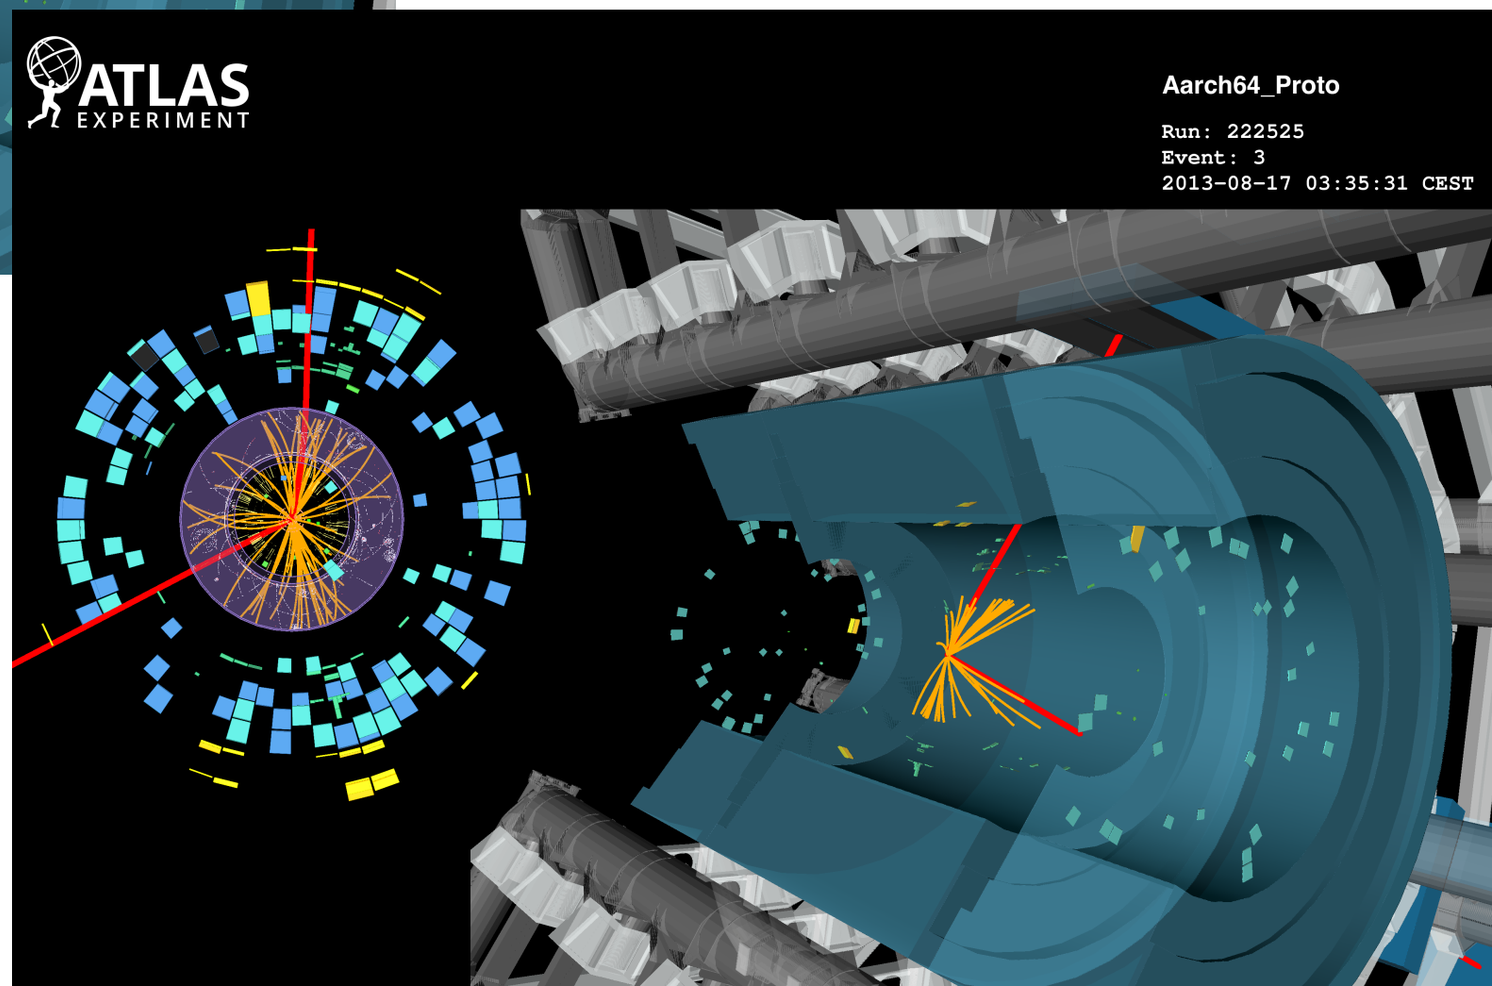


## Aarch64_Proto

Joshua Wyatt Smith

# Power Measurements - Single core

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

**Total power**

**Events per unit power**

**Total - idling power**

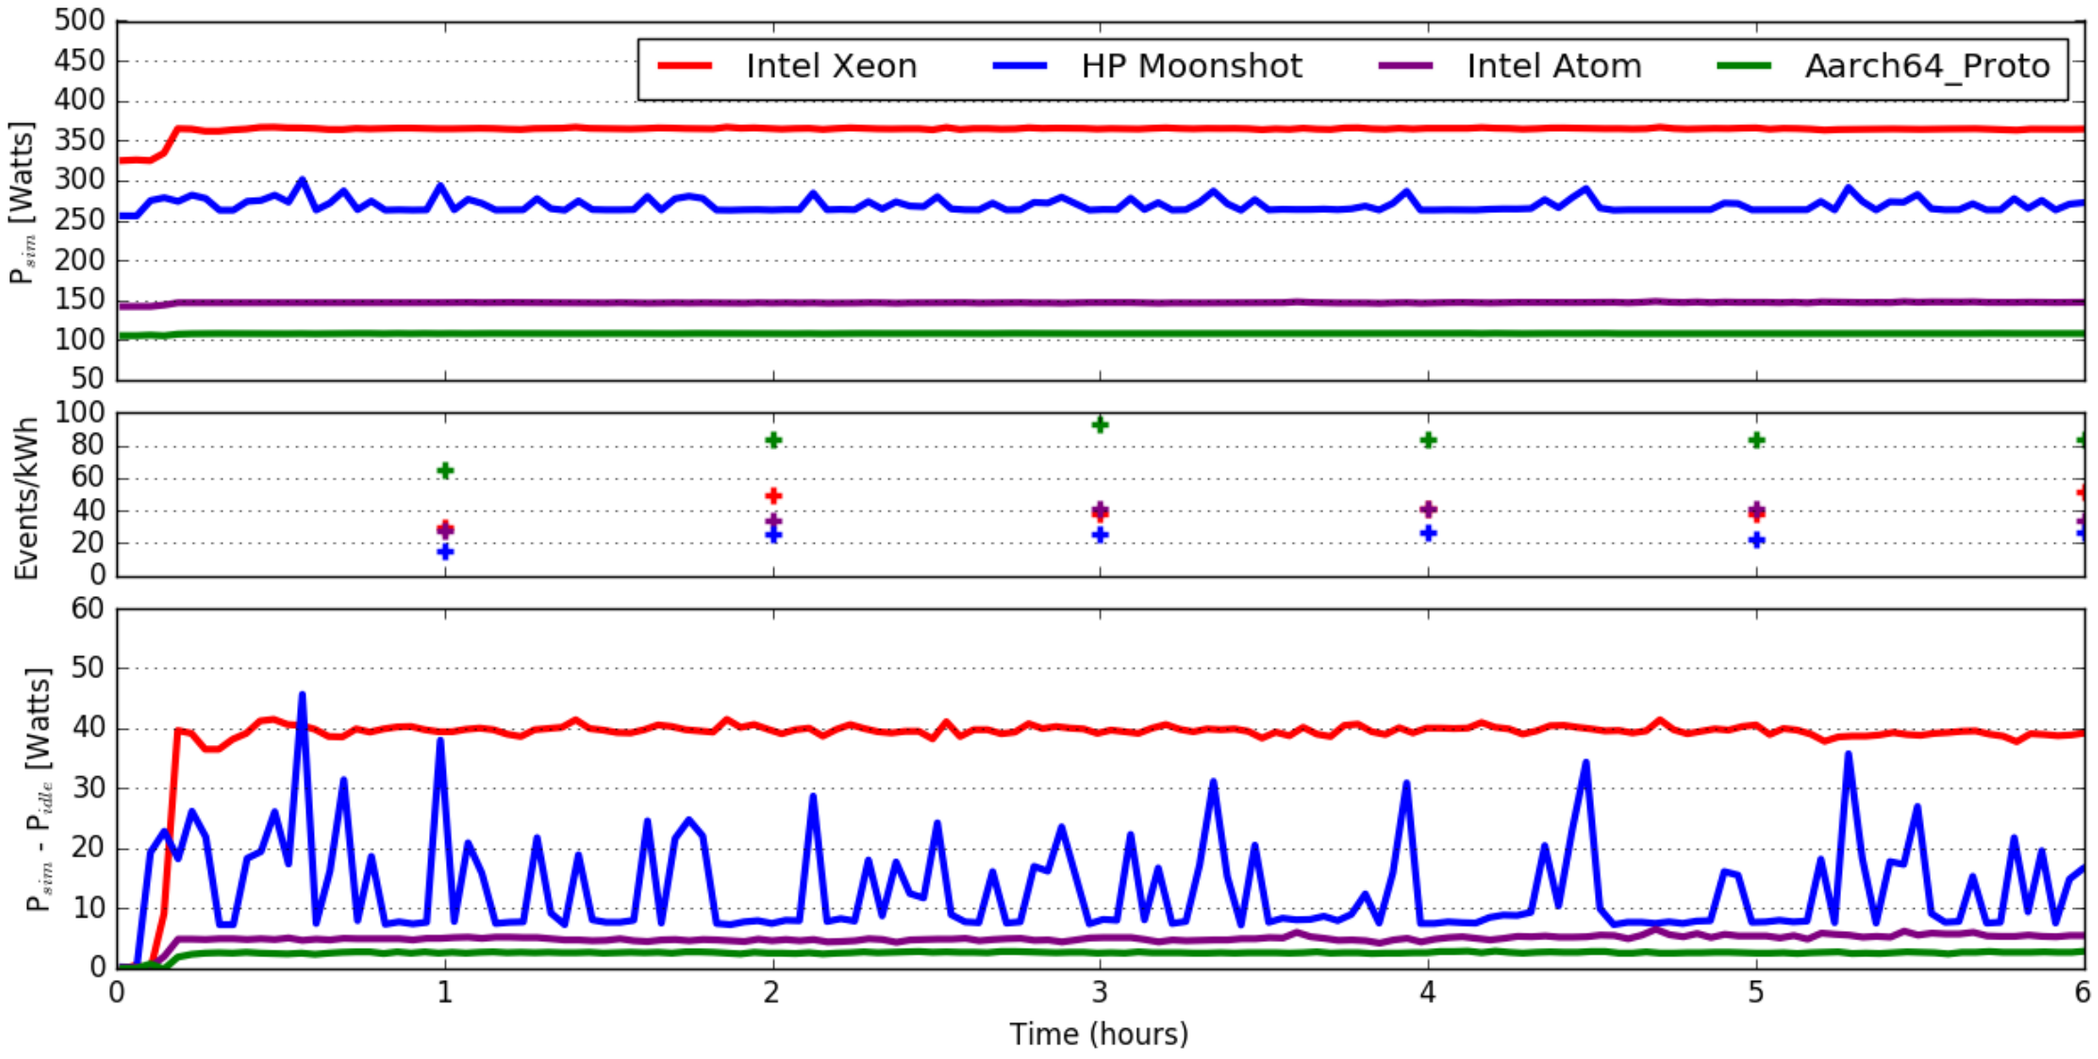


| Name | Time (Hours) |
|---|---|
| HP Moonshot | 15.10 |
| Aarch64_Proto | 10.46 |
| Intel Atom | 18.03 |
| Intel | 6.33 |

# CMSSW RECO: Raw performance – single socket

GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

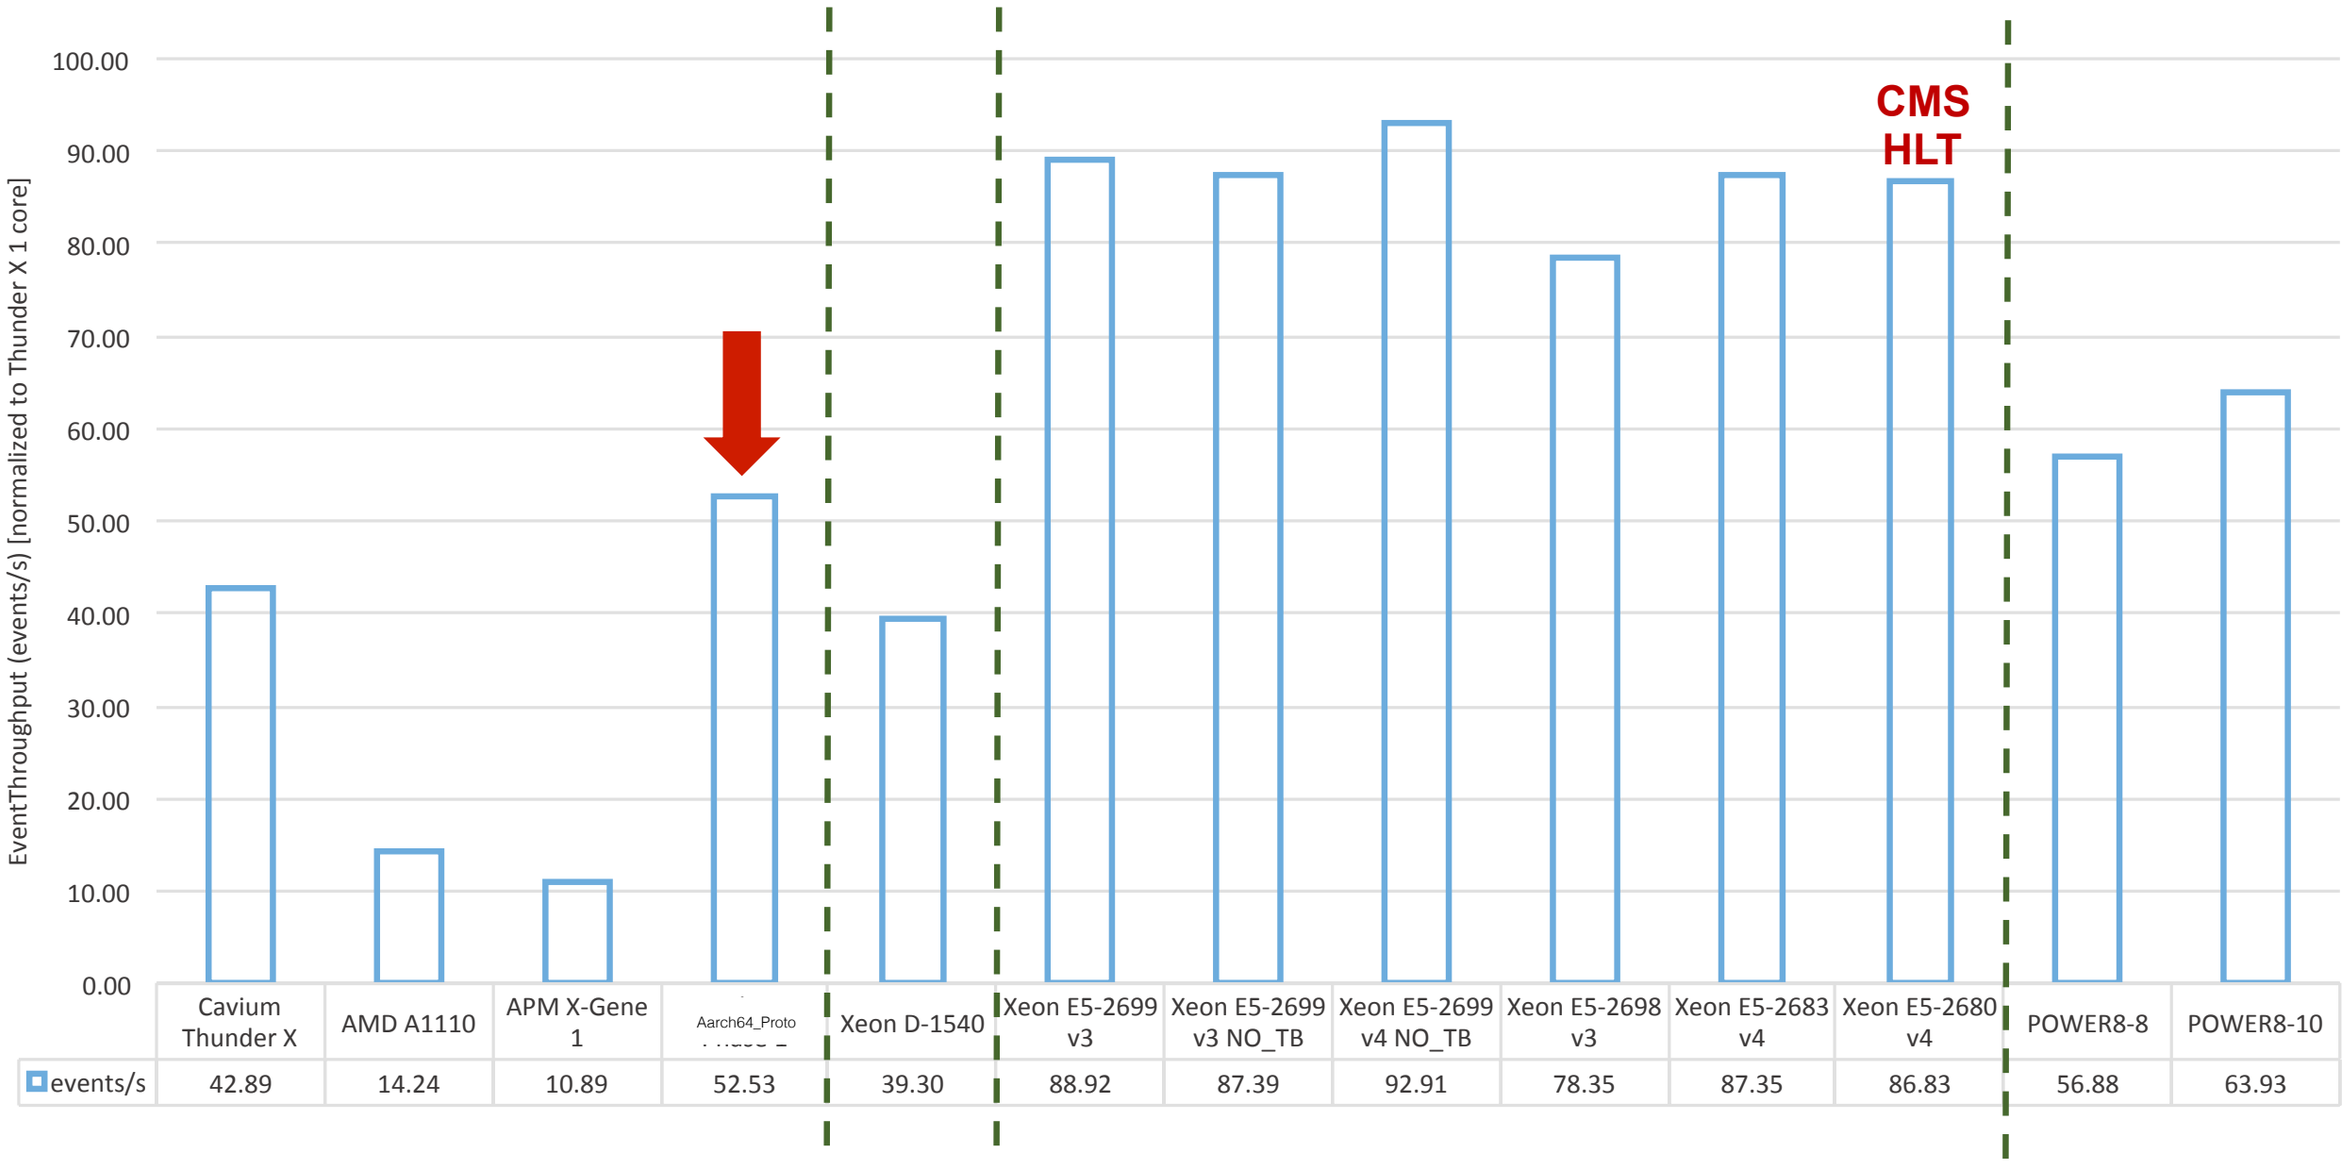


**NOTE:** All numbers normalized to a single core Thunder X performance (lowest perf. per core)
Multi-threaded (**4 threads**) jobs were used to fill full socket due to memory constraints.

* **HLT (High Level Trigger) marks Broadwell machines that CMS acquired in 2016**

# Power Measurements - Multiple cores

GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

**Simulated 8 $t\bar{t}$ events while loading an increasing number of cores**

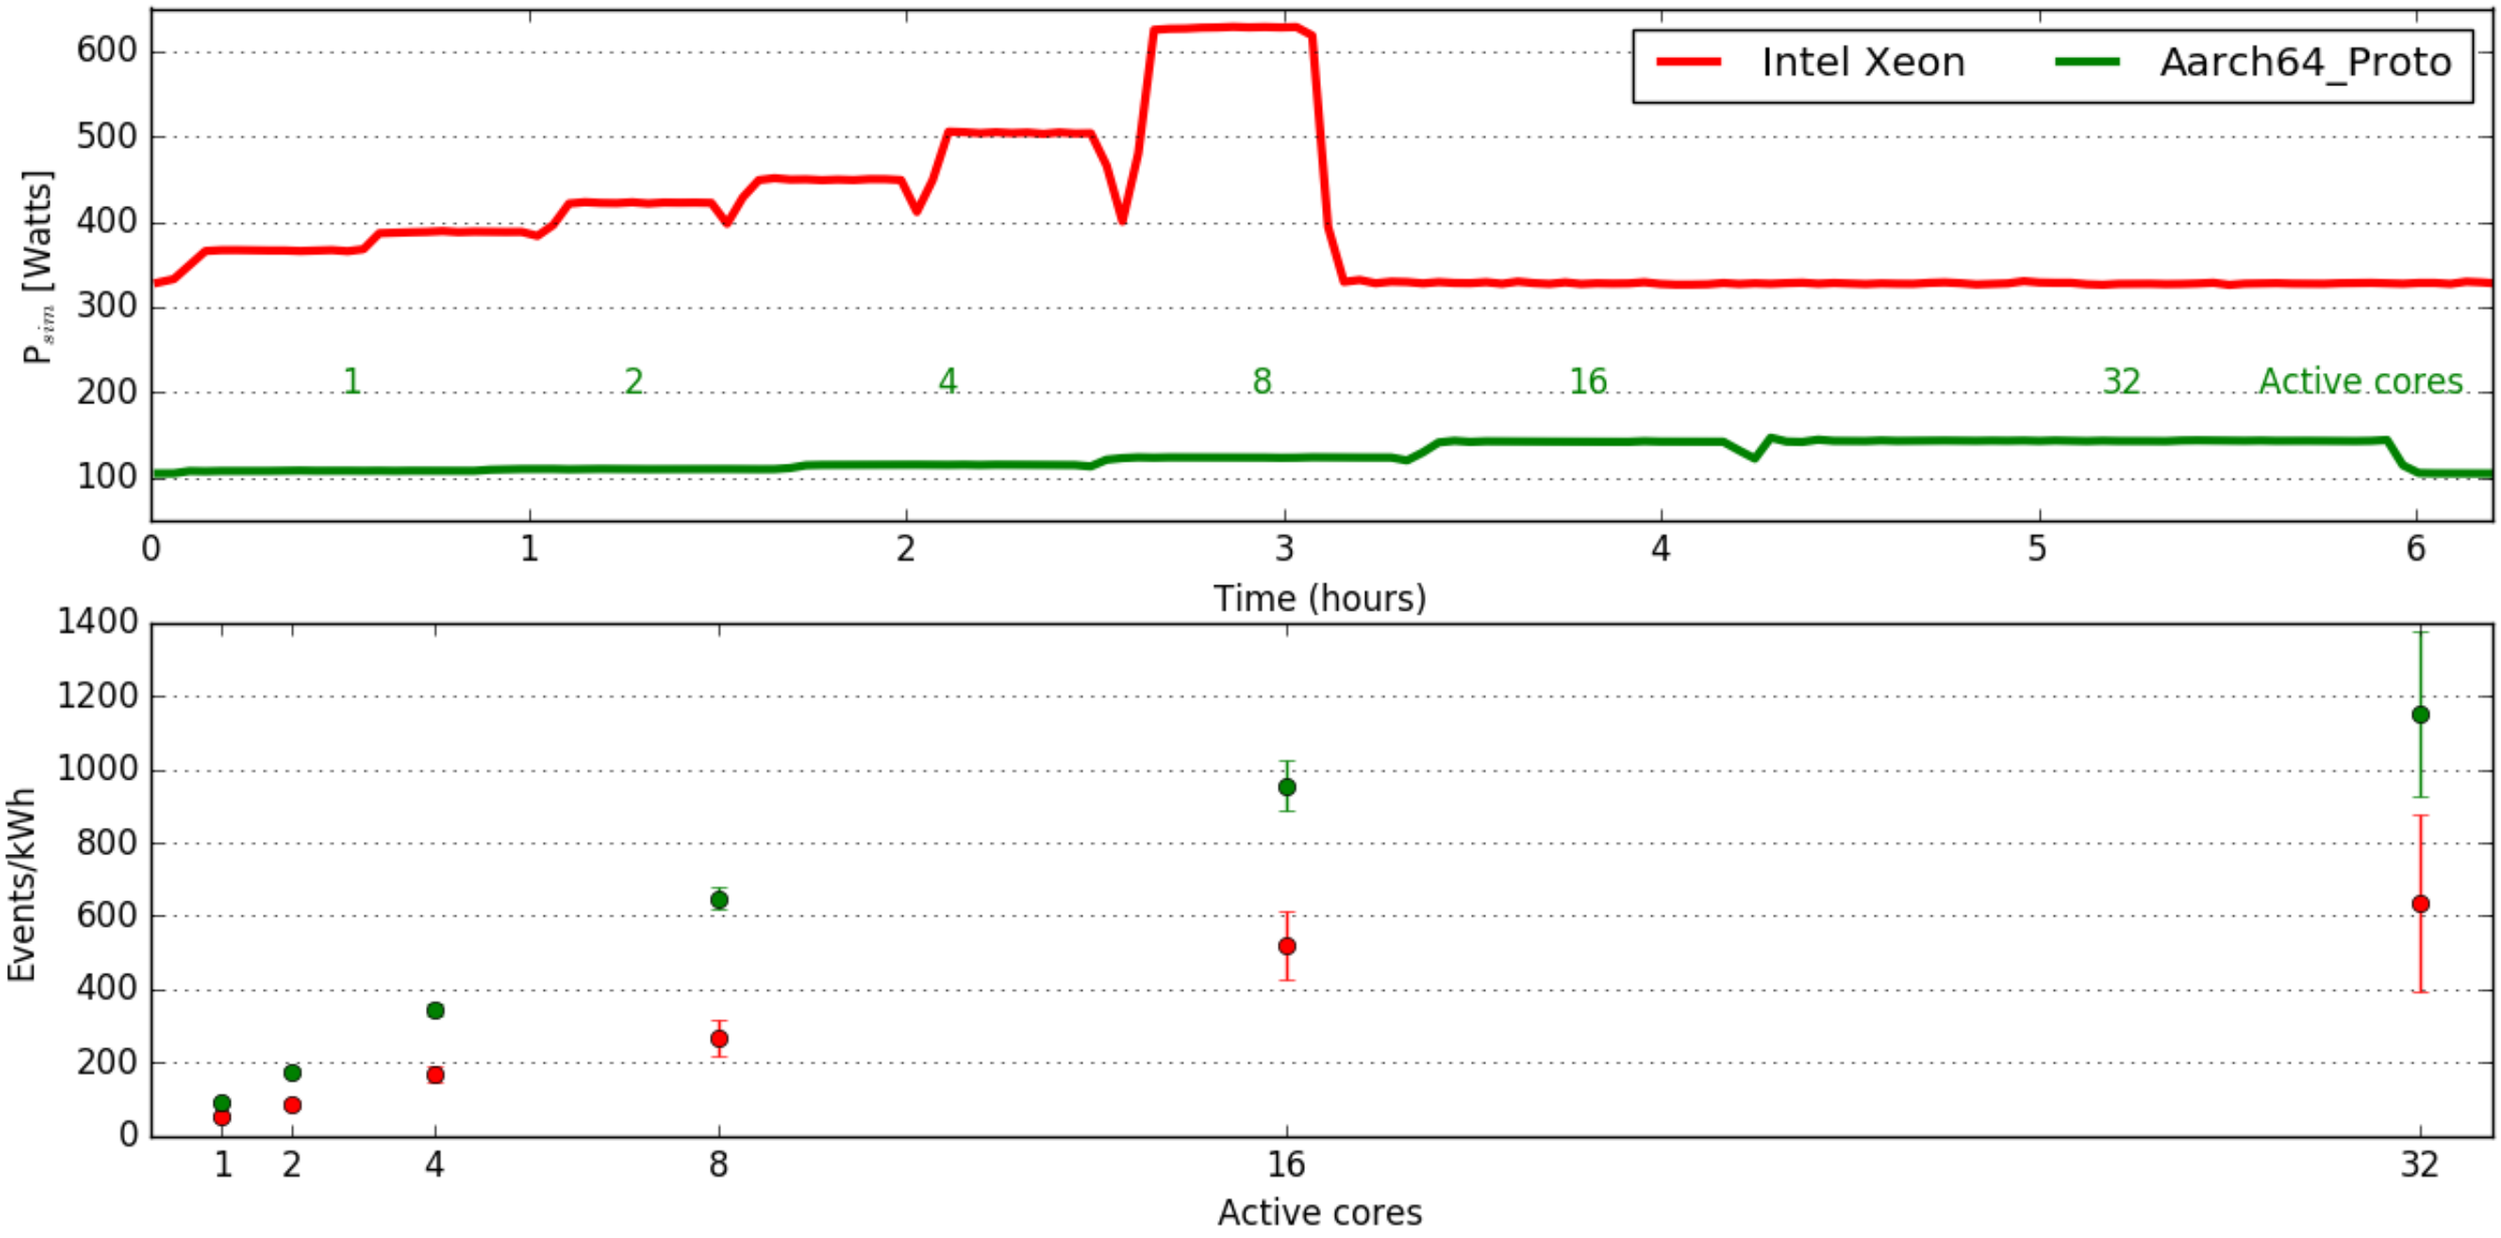

# ParFullCMS: Scalability

GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

- **Geant4 simulation with full CMS detector geometry**

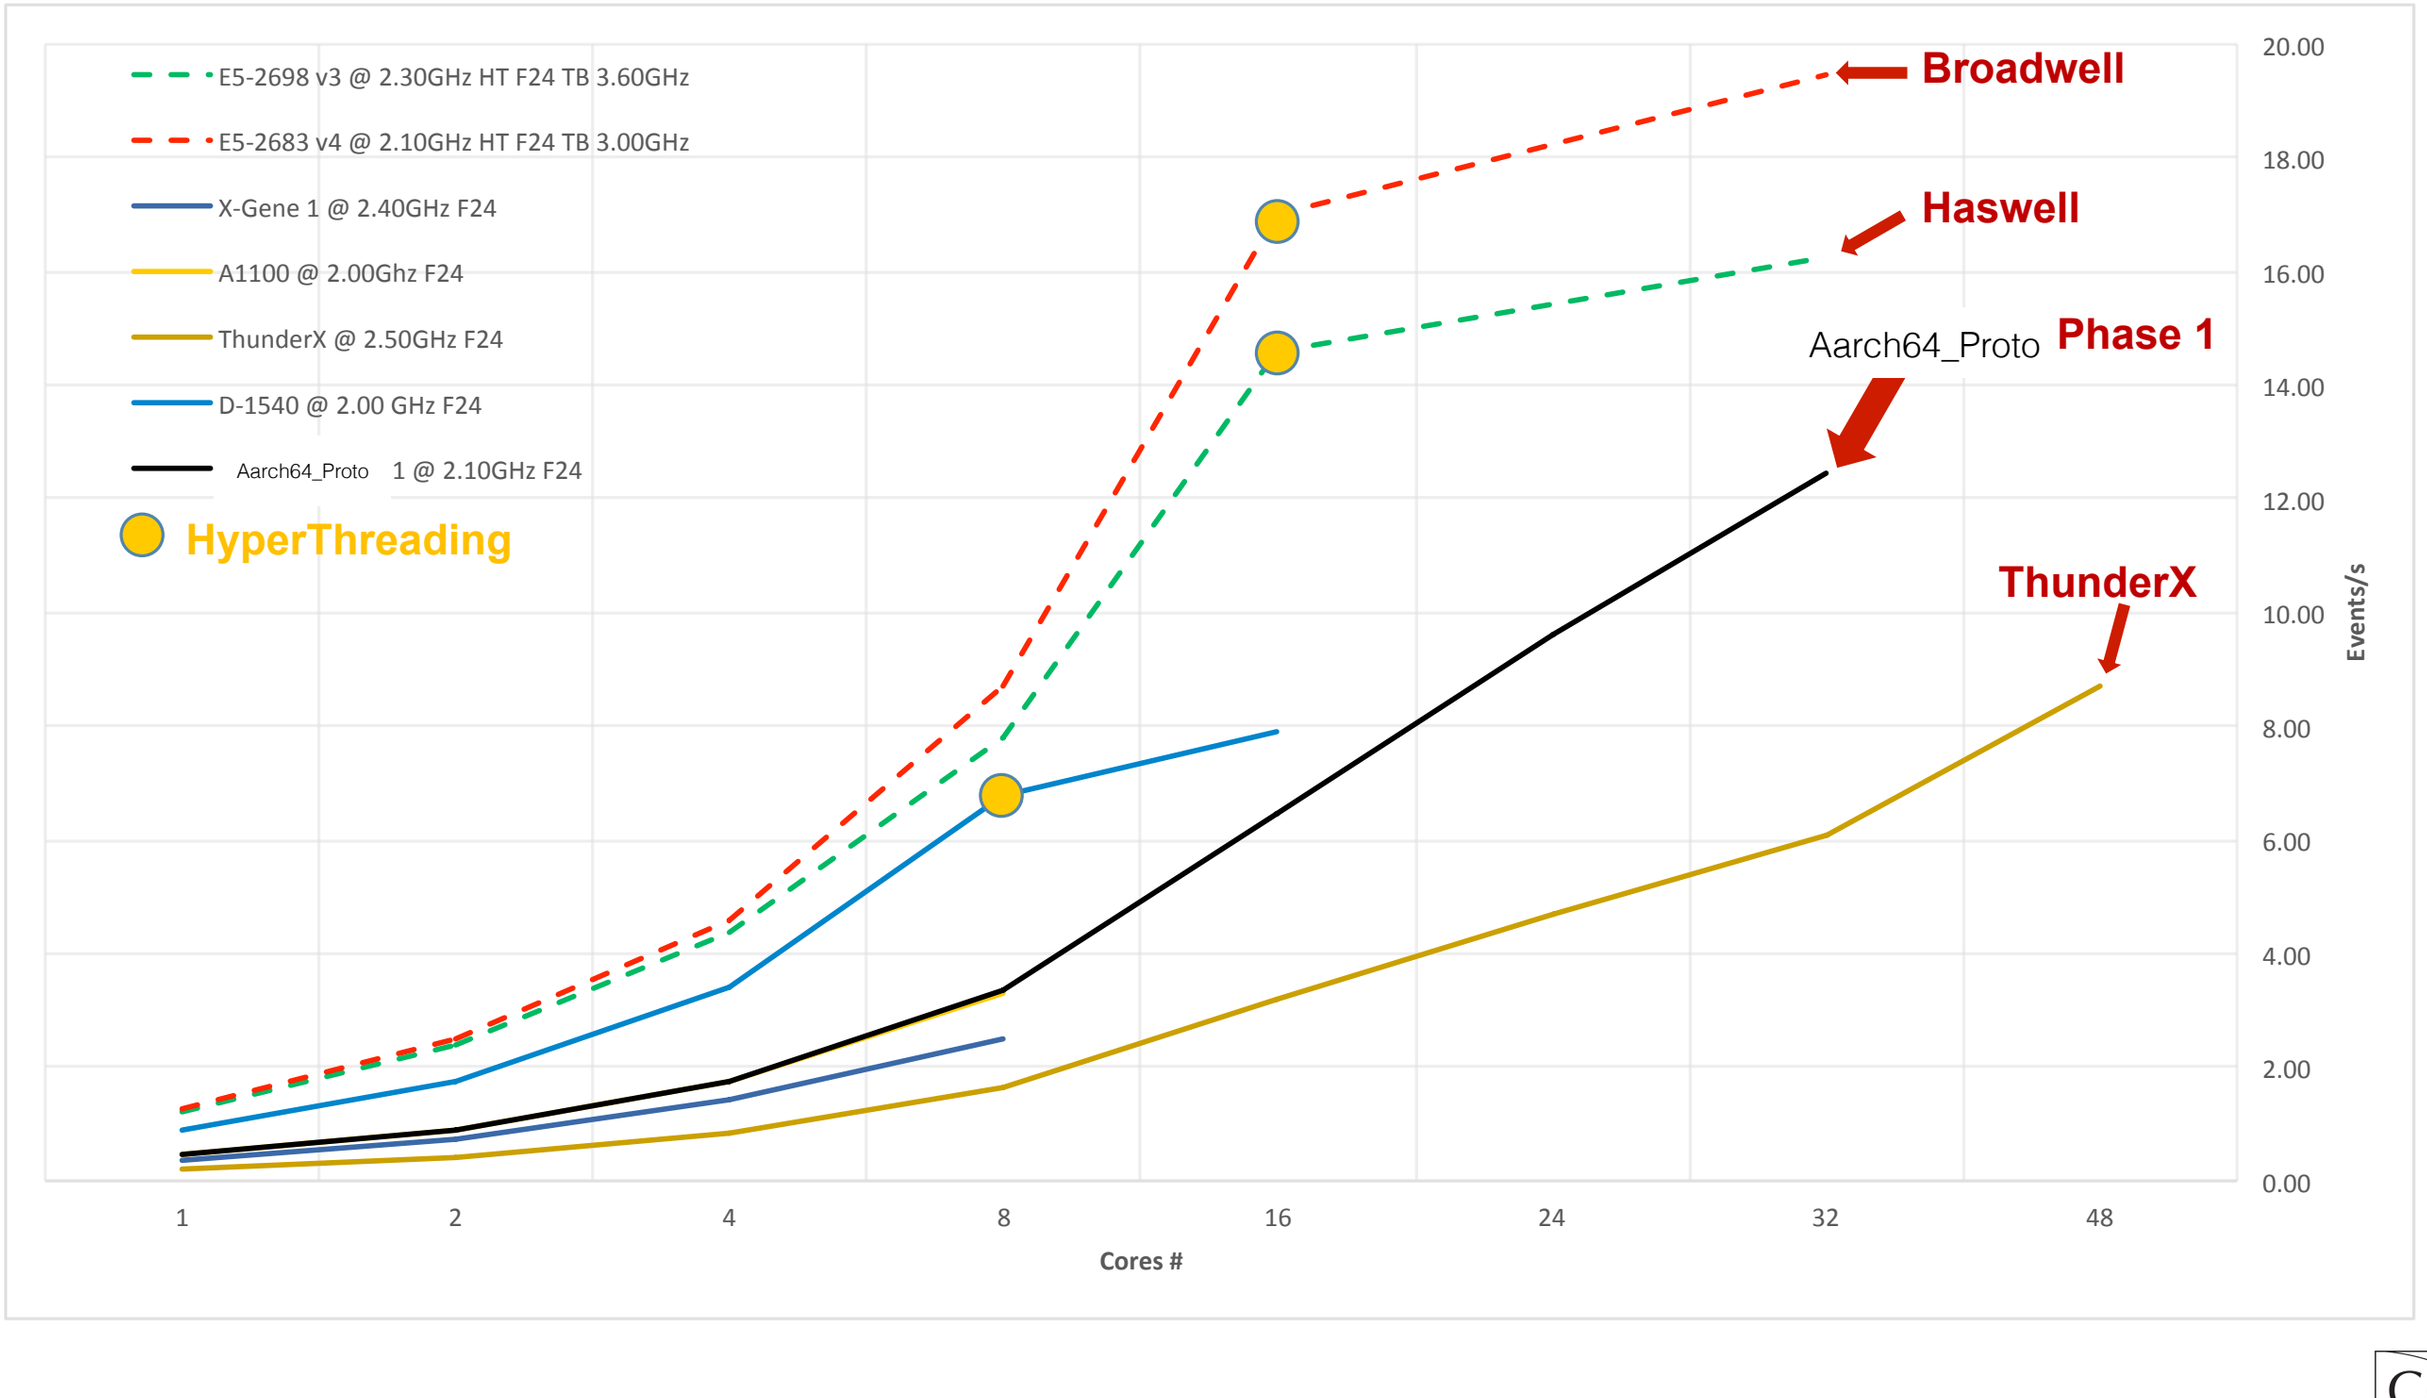

# Conclusions

- ATLAS has performed initial validation from Simulation to Reconstruction
  - More needs to be done!
- Next goal is to slot AthSimulation ARMv8 build into nightlies
  - Easily deploy on server farms
- Maybe proceed to full codebase
- CMS has been deploying Aarch64 builds to the grid (via CVMFS) since 2014
- In 2015 prepared T3_US_Princeton_ARM as demonstrator site
- CMS has made contributions to WLCG infrastructure packages, RHEL Fedora/CentOS, toolchains, etc.)

- ATLAS and CMS are committed to pushing alternative architectures
- I'll leave interpretation of benchmark results up to you :)