ATLAS software stack on ARM64

Joshua Wyatt Smith[a], Graeme A Stewart[b], Rolf Seuster[c] and Arnulf Quadt[a]
on behalf of the ATLAS Collaboration
[a]University of Goettingen, [b]University of Glasgow, [c]University of Victoria
joshua.wyatt.smith@cern.ch

Introduction

The ATLAS experiment explores new hardware and software platforms that, in the future, may be more suited to its data intensive workloads. One such alternative hardware platform is the ARM architecture, which is designed to be extremely power efficient and is found in most smartphones and tablets. Using the more recent 64-bit ARM servers, one can explore how various types of high energy physics related tasks perform, specifically simulation. Using a "real-world" benchmark, we can compare power usage between ARM and traditional Intel servers.

AthSimulation

The ATLAS codebase (Athena) is extraordinarily powerful and complex, consisting of around 2400 packages. Due to its size, this makes porting to alternative architectures more difficult. Thus, we chose a subset of the packages that make up a project called AthSimulation. AthSimulation is capable of carrying out all intensive simulations needed for the experiment. At around 350 packages, the porting process becomes significantly easier and faster. The figure shows the subset of packages ported to ARM.

The benchmark is Geant4 intensive and simulates the passage of particles through the ATLAS detector. For the purpose of this benchmark, the simulation of 100 $t\bar{t}$ events was chosen when running the benchmark on a single core, and 8 events when increasing the active cores.

The projects ported to ARM. AthSimulation is a subset of Athena, in itself consisting of ~350 packages.

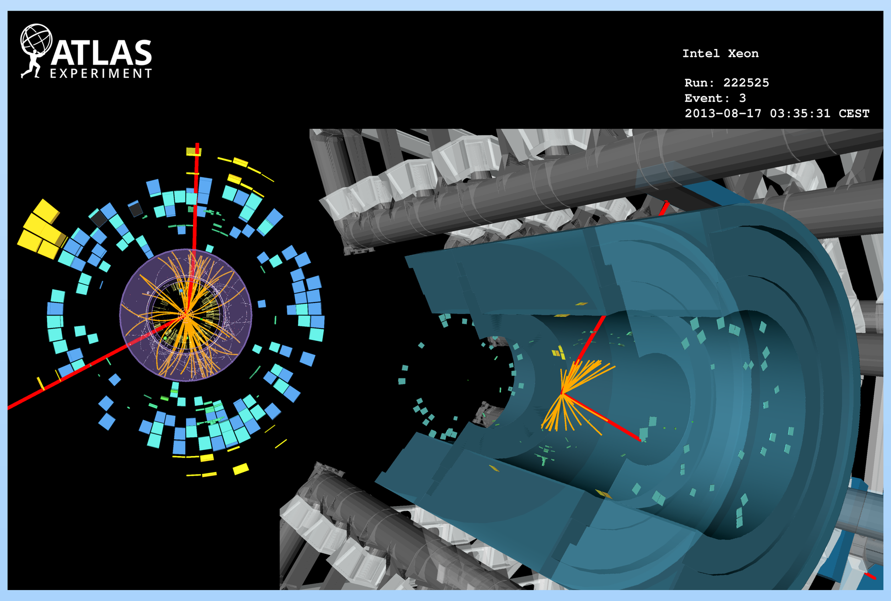

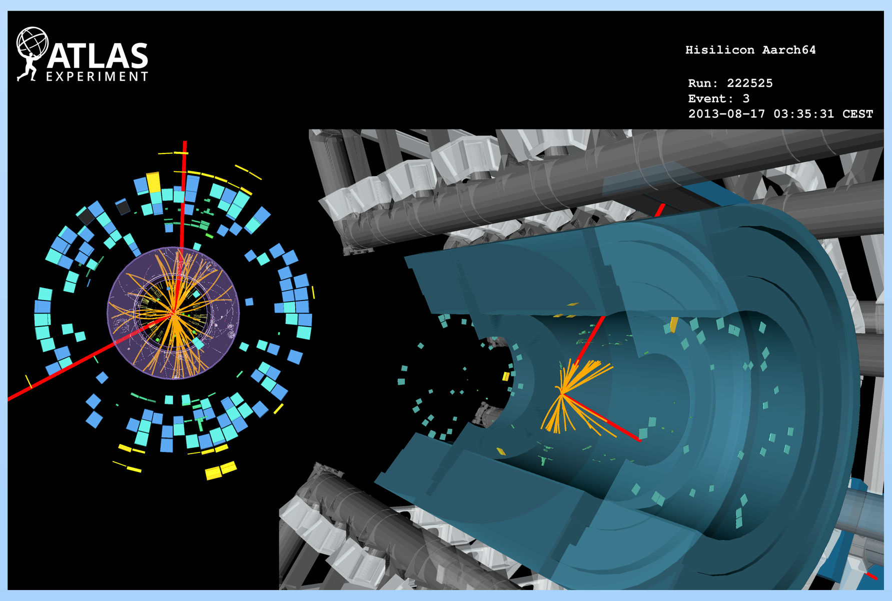

Hardware

ARM (Advanced RISC (Reduced Instruction Set Computing) Machine) is an architecture whereby less instructions are used on the CPU. This reduces the transistors required and thus reduces overall power consumption. Techlab at CERN maintains two different types of ARM servers (HP Moonshot and Hisilicon) as well as two types of Intel servers (Intel Atom and Intel Xeon). It is important to note the Intel Atom has been discontinued but still provides interesting results.

The different servers with their configurations.

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
Hisilicon	Hisilicon PV660, 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

Validation

Our strategy was to first verify that the results from the different architectures made sense. Thus, initial validation was carried out. The benchmark was run on each server, with the output "HITS" files put through another reconstruction phase on a traditional x86 server running Athena. Simulation is a monte-carlo process and so numerical identity is not expected. Results are shown in the figures below.

When compared to the Intel Xeon, the ARM servers give reasonable behavior. Intel Atom clearly registers more hits.

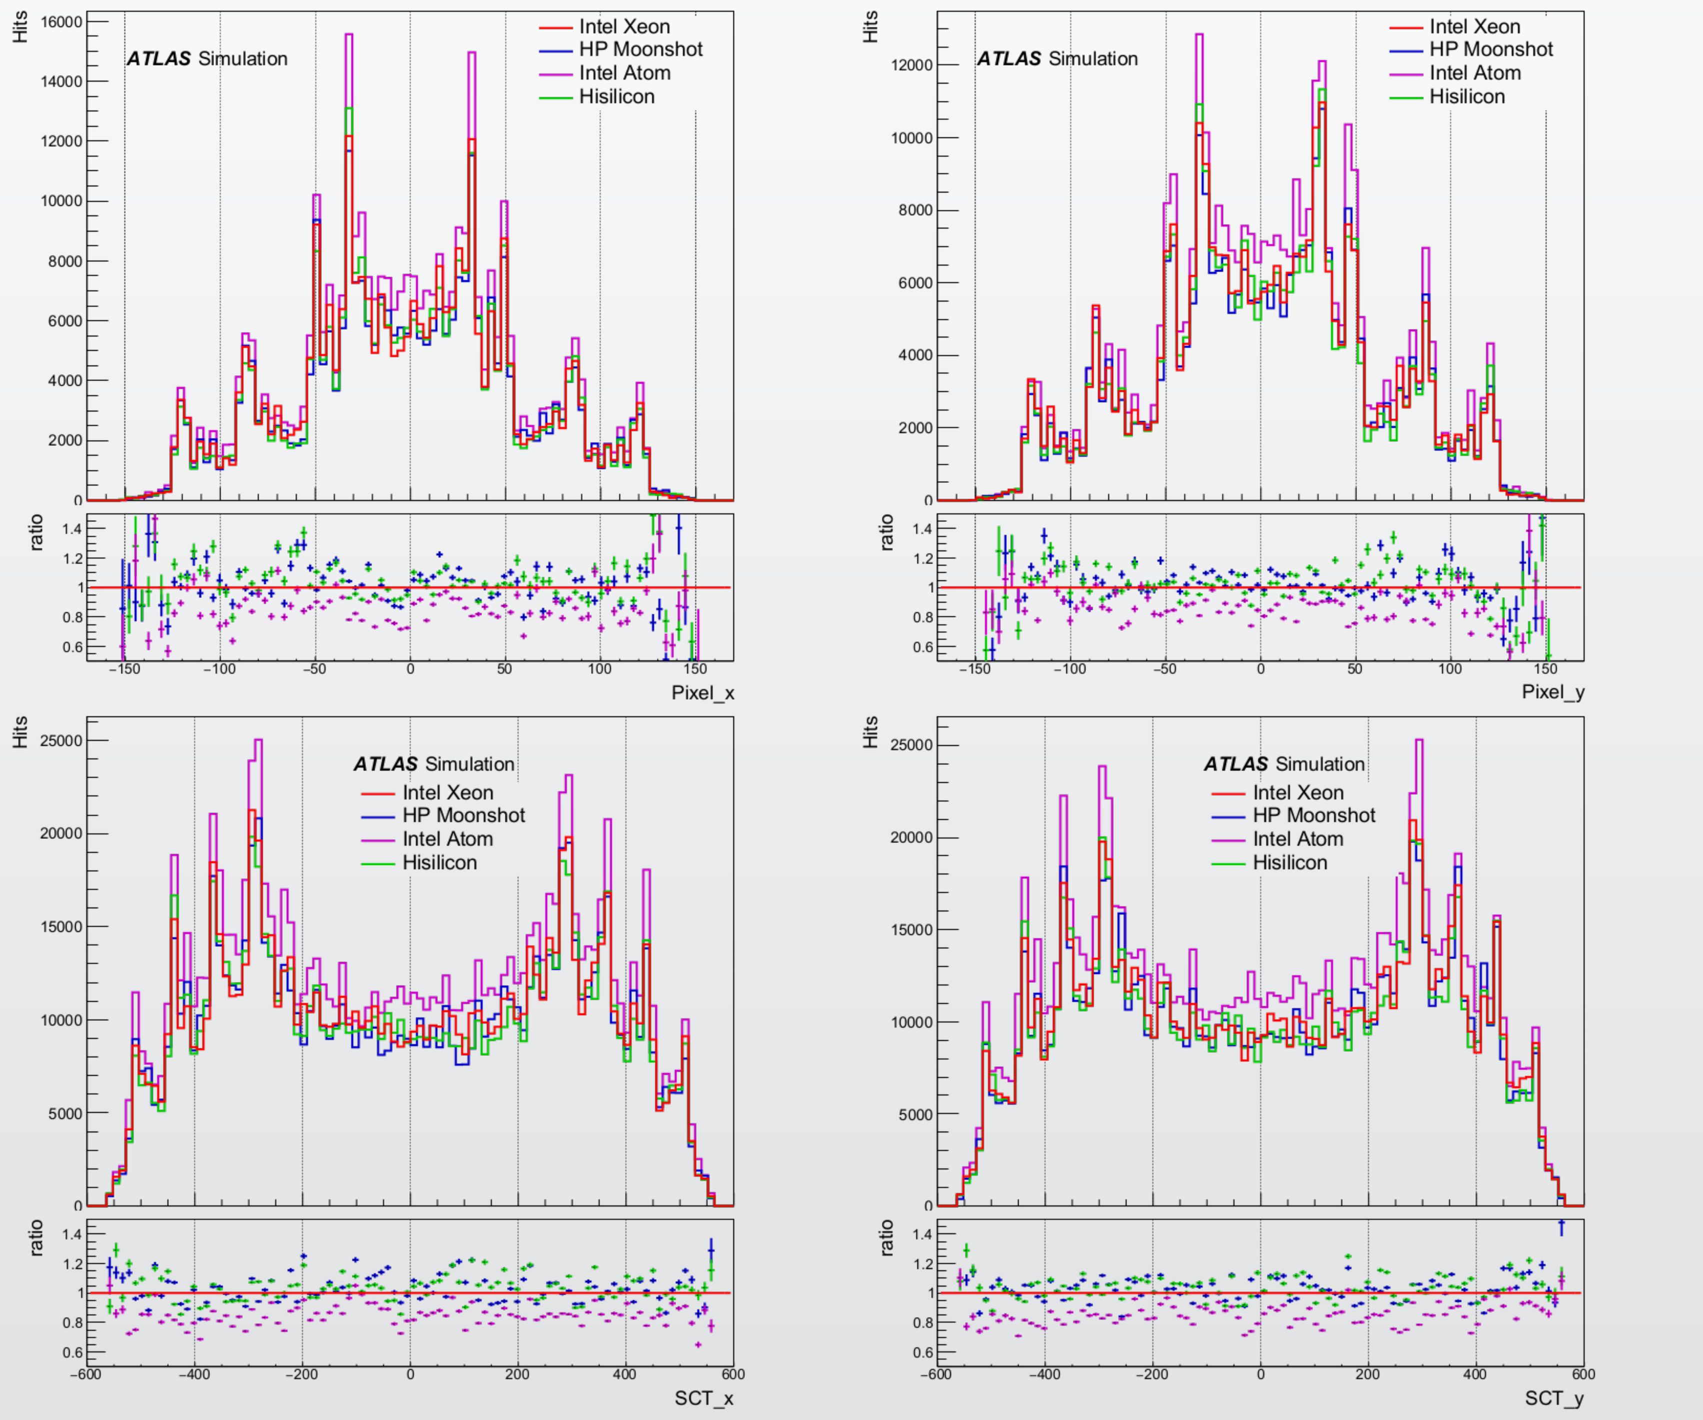

Results showing the hits in the pixel and silicon microstrip (SCT) detectors on ATLAS. The ratio between the three servers and the "Intel Xeon" is shown in the ratio plots. The "Intel Xeon" is taken as the "accepted" distribution.

Power Measurements

Power measurements were taken every 10 seconds. The below graph shows the results of the benchmark for a period of 6 hours on a single core. The table shows the total times for each server.

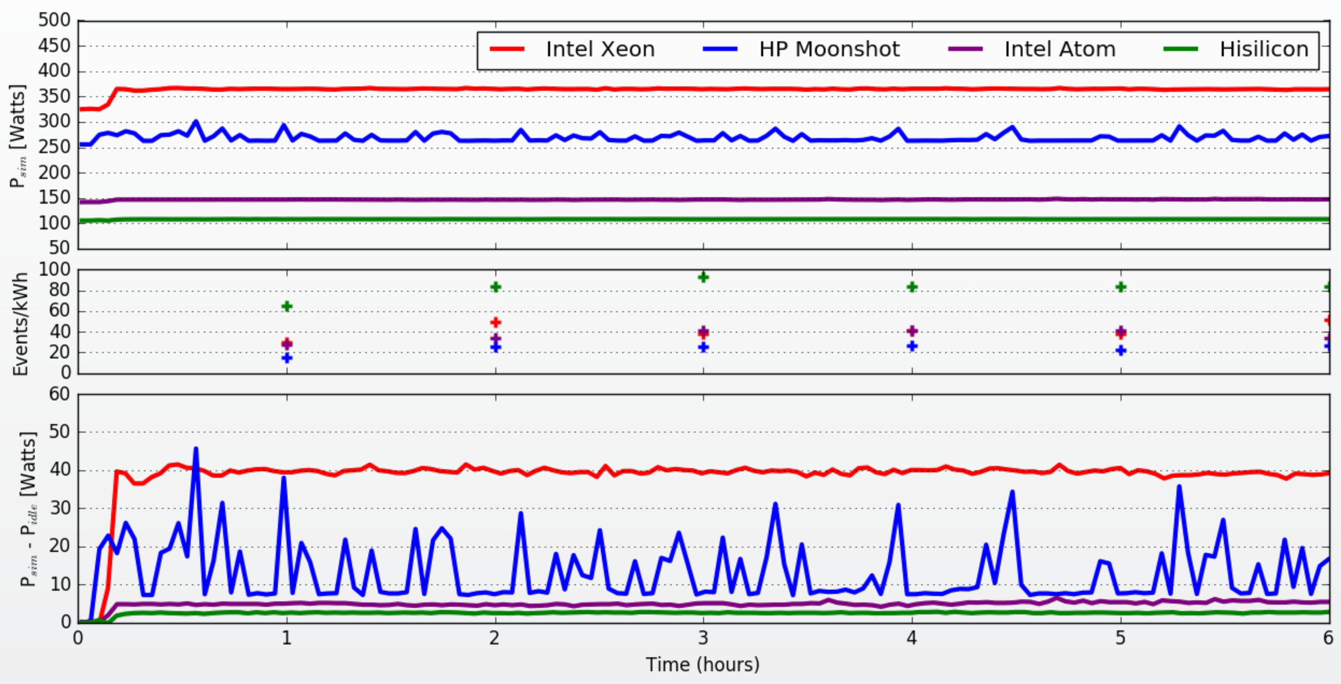

Total time of benchmark for each server.

Name	Time (Hours)
HP Moonshot	15.10
Hisilicon	10.46
Intel Atom	18.03
Intel	6.33

Power measurements for each of the servers. Top: Total power for 6 hours. Middle: Events/kWh calculated for each hour. Bottom: Idling power of each server subtracted from total power when running simulation.

To take I/O into account, the benchmark was repeated on the Hisilicon and Intel servers with multiple jobs. The below plot shows the results of the same benchmark while the server is running an increasing number of identical jobs. Due to a technical error, test times were read from the top graph in calculating kWh. The uncertainty in reading these measurements is incorporated into the error bars of the lower graph.

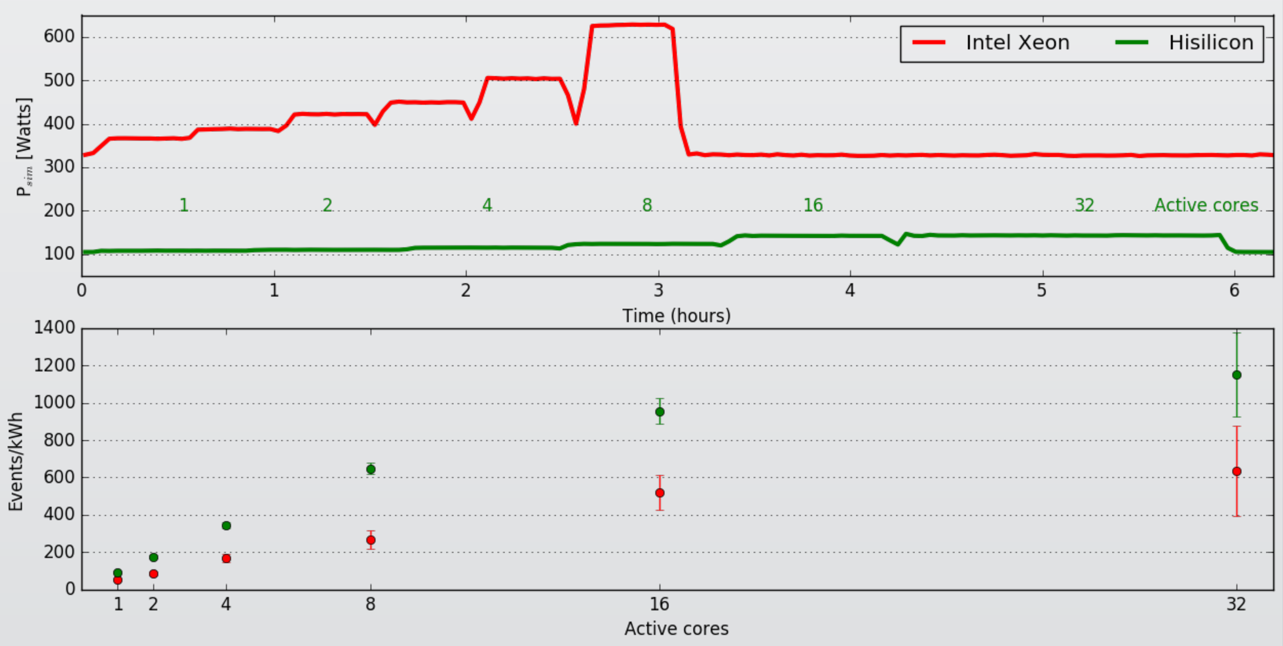

Power measurements for identical benchmark running on increasing number of cores. Top: total time and power for all tests. Bottom: Events/kWh calculated for each test.

Results

- Validation tests show that while the ARM servers are comparable to the Intel Xeon, the Intel Atom has a clear trend of more hits. This can originate from many factors including how the architecture treats floating numbers to how random numbers are generated. Further validation over more events and event by event reconstruction needs to be done for the ARM servers.
- Power measurements on a single core show that the Hisilicon server performs the best on a event/watt basis. Even though the time to process 100 events is slower the overall power consumption is much better.
- When loading the servers with an increasing amount of jobs, the ARM server clearly uses significantly less power. At 32 active cores the benchmark test time increases due to bottlenecks developing in the ARM sever. Taking total processing time and power consumption into the equation, the ARM server performs between ~1.7 and ~2.4 times more efficiently than the Intel server while under load.

Conclusions

Initial validation as well as power measurements are presented. Four servers, 2 ARM and 2 Intel, are compared alongside each other in a "real-world" ATLAS simulation benchmark that was ported to the ARM architecture. The results show that ARM servers have improved dramatically over the past few years. Their 64-bit architecture is now competitive with the traditional Intel machines. In terms of performance per watt, these results show the ARM servers perform more effectively than the standard Intel server.

Acknowledgments

The authors would like to thank Attila Krasznahorkay for useful discussions and tips when porting the code. Thank you to Zachary Marshall for providing the knowledge required to run the AthSimulation benchmark. Finally, thanks to CERN openlab and Techlab, specifically Aritz Brosa Iartza for helping with power measurements and for maintaining the hardware.

ATLAS EXPERIMENT
Tech Lab
GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN