Affordable and power efficient computing for high energy physics: CPU and FFT benchmarks of ARM processors

Mitchell A Cox, Robert Reed and Bruce Mellado

School of Physics, University of the Witwatersrand. 1 Jan Smuts Avenue, Braamfontein, Johannesburg, South Africa, 2000

E-mail: mitchell.cox@students.wits.ac.za

Abstract. Projects such as the Large Hadron Collider at CERN generate enormous amounts of raw data which presents a serious computing challenge. After planned upgrades in 2022, the data output from the ATLAS Tile Calorimeter will increase by 200 times to over 40 Tb/s. ARM System on Chips are common in mobile devices due to their low cost, low energy consumption and high performance and may be an affordable alternative to standard x86 based servers where massive parallelism is required. High Performance Linpack and CoreMark benchmark applications are used to test ARM Cortex-A7, A9 and A15 System on Chips CPU performance while their power consumption is measured. In addition to synthetic benchmarking, the FFTW library is used to test the single precision Fast Fourier Transform (FFT) performance of the ARM processors and the results obtained are converted to theoretical data throughputs for a range of FFT lengths. These results can be used to assist with specifying ARM rather than x86-based compute farms for upgrades and upcoming scientific projects.

1. Introduction

Projects such as the Large Hadron Collider (LHC) generate enormous amounts of raw data which presents a serious computing challenge. After planned upgrades in 2022, the data output from the ATLAS Tile Calorimeter will increase by 200 times to over 41 Tb/s (Terabits/s) [1]. It is not feasible to store this data for offline computation.

A paradigm shift is necessary to deal with these future workloads and the cost, energy efficiency, processing performance and I/O throughput of the computing system to achieve this task are vitally important to the success of future big science projects. Current x86-based microprocessors such as those commonly found in personal computers and servers are biased towards processing performance and not I/O throughput and are therefore less-suitable for high data throughput applications otherwise known as Data Stream Computing [2].

ARM System on Chips (SoCs) are found in almost all mobile devices due to their low energy consumption, high performance and low cost [3]. The modern ARM Cortex-A range of processors have 32- and 64-bit cores and clock speeds of up to 2.5 GHz, making them potential alternatives to common x86 CPUs for scientific computing. This paper presents benchmarks of three common ARM Cortex CPUs, namely the Cortex-A7, A9 and A15 with more information on these platforms in Table 1. The benchmark results are useful for specifying an ARM-based system in new scientific projects such as ATLAS read-out and trigger system upgrades.

A brief discussion of the ATLAS Triggering and Data Acquisition System (TDAQ) and where ARM processors may potentially be used is presented in Section 2. CPU benchmark results are given in Section 3. Fast Fourier Transform (FFT) benchmarks are in Section 4. Section 5 concludes with a brief discussion of future work.

Table 1: Specifications and other details of the ARM platforms used.

	Cortex-A7	**Cortex-A9**	**Cortex-A15**
Platform	Cubieboard A20	Wandboard Quad	ODROID-XU+E
SoC	Allwinner A20	Freescale i.MX6Q	Samsung 5410
Cores	2	4	4 (+ 4 Cortex-A7)
Max. CPU Clock (MHz)	1008	996	1600
L2 Cache (kB)	256	1024	2048
Floating Point Unit	VFPv4 + NEONv2	VFPv3 + NEON	VFPv4 + NEONv2
RAM (MB)	1024	2048	2048
RAM Type	432 MHz 32 bit	528 MHz 64 bit	800 MHz 64 bit
2014 Retail (USD)	65	129	169
Linux Kernel	3.4.61	3.10.17	3.4.5
GCC	4.7.1	4.7.3	4.7.3

2. ATLAS Triggering and Data Acquisition System

The ATLAS experiment is composed of several sub-detectors, each of which has separate data processing requirements. The massive amount of raw data is reduced by a process called triggering. In the Tile Calorimeter, there are currently three main levels of triggering, shown in Figure 1. The read-out system is based on FPGAs (Field Programmable Gate Arrays) and DSPs (Digital Signal Processors) to form the level one trigger which serves to reduce the data rate (event rate) from 40 MHz to about 100 kHz. Each ATLAS event consists of about 1.5 MB data, of which a portion is made up from TileCal data. Some components of the read-out subsystems will be replaced by the Super Read Out Driver (superROD or sROD) in the 2022 upgrade [1].

An ARM-based Processing Unit (PU) is in development at the University of the Witwatersrand, Johannesburg, to complement the sROD with higher level processing tasks on the raw data, before it has been reduced by the triggering system.

The level two and three triggers (LVL2 and Event Filter) are implemented with compute clusters. Data from the level one trigger system is fed to the Read Out System (ROS) which passes the data to the Event Builder and LVL2 trigger at a rate of about 120 GB/s. [4]

2.1. Level Two Trigger, Event Builder and Event Filter

The LVL2 filter only works on a portion of the data to determine whether it is interesting - if the portion is interesting then all of the associated data is let through to the Event Builder. The LVL2 cluster is built from approximately 500 machines, each of which has two 2.5 GHz quad-core Intel Harpertown CPUs, two 1 Gb/s Ethernet interfaces and 16 GB RAM [5]. This CPU achieves about 59 000 CoreMarks.

The Event Builder consists of about 100 rack-mounted servers, each of which has two 2.6 GHz AMD Opteron 252 CPUs with 1 Gb/s Ethernet. According to the CoreMark online database, the AMD Opteron 254 CPU - which is a dual core variant of the Opteron 252 used in the Event Builder - achieves about 13700 CoreMarks [6].

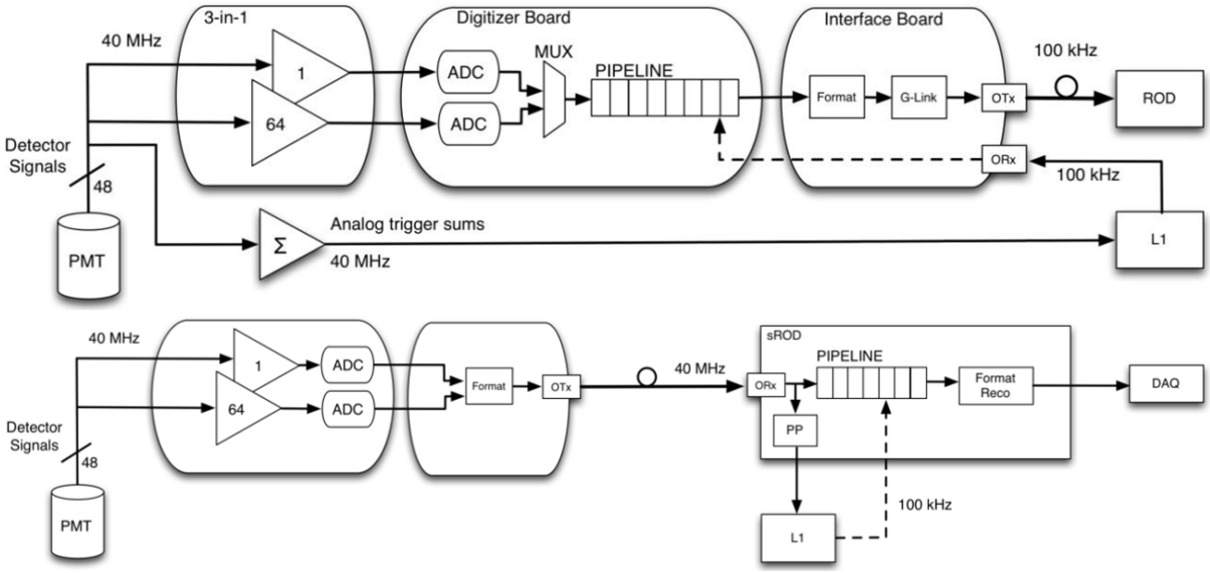

Figure 1: ATLAS TileCal Trigger and Data Acquisition System flow diagram before (above) and after the 2022 upgrade (below) [1].

The Event Filter is a much larger cluster of about 1900 machines which is the last level of triggering. The same machines as for the LVL2 filter are used. Data is output from the Event Filter at about 300 MB/s for storage [5].

3. General CPU Benchmarks

High Performance Linpack (HPL) is an industry standard benchmark application for measuring floating point performance [7]. HPL is the primary benchmark used when rating supercomputers on the Green500 and Top500 lists [8]. Both single and double precision floating point were tested.

CoreMark is another industry standard benchmark for measuring general CPU performance with a balance of integer, floating point and common algorithms [6]. ARM has recommended that CoreMark be used to replace the older Dhrystone benchmark for integer MIPS performance [9].

All three CPUs were manually forced to 1 GHz for a fair comparison. The peak power consumption of the test platforms were measured. The HPL and CoreMark results as well as the power consumption is presented in Table 2.

A cluster of eight Wandboards has been built at The University of the Witwatersrand in order to test scientific algorithms on ARM. In total, 32 1 GHz Cortex-A9 cores are available, with 16 GB RAM and interconnected with 1 Gb/s Ethernet. A photo of the cluster mounted in a rack is shown in Figure 2. No in depth testing is available at present.

4. Fast Fourier Transform Benchmarks

FFTs have numerous uses in science, mathematics and engineering. FFTs are computationally intensive and are an $O(n \log(n))$ algorithm which stresses CPU and memory subsystems.

FFTW is an open source, high performance FFT library which has a benchmark facility which reports the time, t, (in microseconds) and the estimated MFLOPS of a run [10]. The length, N, and type of FFT is specified for each run. One dimensional, single-precision complex FFTs (8 bytes per point) were tested. The CPUs were manually set to their maximum frequencies.

Figure 3 a) shows the MFLOPS results for a wide range of FFT lengths. The maximum results of multi-core and summed multi-process runs are reported. This methodology was used

Table 2: General CPU benchmarking results and power consumption.

	Cortex-A7	**Cortex-A9**	**Cortex-A15**
CPU Clock (MHz)	1008	996	1000
CPU Cores	2	4	4
HPL (Single Precision GFLOPS)	1.76	5.12	10.56
HPL (Double Precision GFLOPS)	0.70	2.40	6.04
CoreMark	4858	11327	14994
Peak Power (W)	2.85	5.03	7.48
Double Precision GFLOPS/Watt	0.25	0.48	0.81
Single Precision GFLOPS/Watt	0.62	1.02	1.41

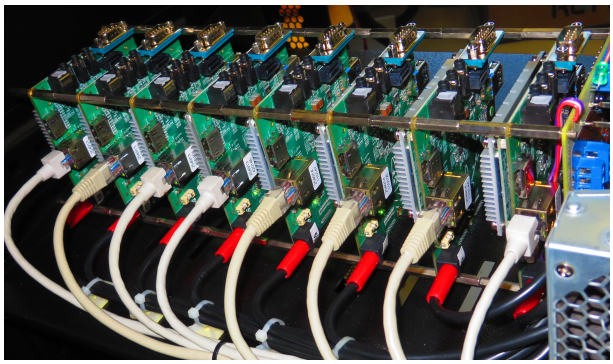

Figure 2: Photo of the 32 core Cortex-A9 cluster at The University of the Witwatersrand, Johannesburg.

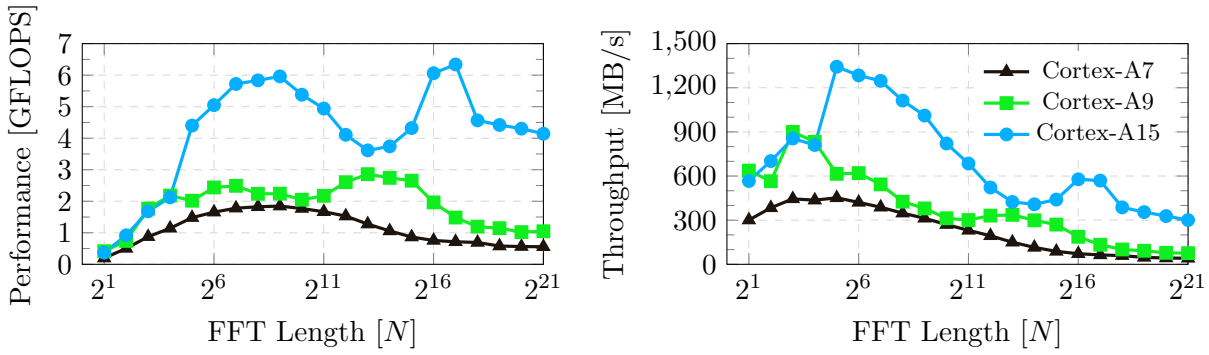

Figure 3: a) The maximum score of multi-core and multi-process runs to always utilise all processors. b) Theoretical maximum calculated FFT throughput.

to ensure 100% utilisation of the CPU for the tests. Figure 3 b) shows the calculated theoretical FFT throughput based on the FFT size in Bytes and the run time: Throughput = $(8N)/t$ [MB/s].

5. Discussion, Conclusions and Future Work

The processing performance of the Cortex-A7, A9 and A15 CPUs is comparable to an older x86 CPU, but the power efficiency is excellent. The Cortex-A15 achieves 0.81 GFLOPS/W with double precision floating point operations. The ARM processors that have been tested are optimised for single precision floating point and so are likely to be used in applications where double precision floating point is not necessary. The power efficiency is above 1 GFLOPS/W for the Cortex-A9 and A15 and is 0.62 GFLOPS/W for the Cortex-A7.

The compute performance of the Cortex-A15 is significantly higher when it is run at its maximum clock speed of 1.6 GHz but this would be true for any processor at a higher clock speed. The CoreMark results confirm that the Cortex-A15 is significantly higher performance than the Cortex-A7 and A9.

It should be noted that the Cortex-A9 and A15 SoCs that were tested could be used as a substitute based on CoreMark results and specifications for the AMD CPU used in the Event Builder section of the ATLAS TDAQ system. If ARM SoCs were used, power consumption would decrease by an order of magnitude with the ARM SoC consuming approximately 5 W and the current AMD-based system consuming over 68 W [11].

The FFT performance also indicates that the Cortex-A15 is superior to the Cortex-A9 and A7. The theoretical FFT throughputs are over 300 MB/s for most FFT lengths with the Cortex-A15 sustaining over 300 MB/s up to the largest FFT tested at 2097152 points.

Based on the results presented in this paper, specifically the low cost, high performance and good power efficiency, it is clear that ARM SoCs should be considered for future upgrades and new computing systems in big science projects.

Acknowledgements

The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. We would also like to acknowledge the School of Physics, the Faculty of Science and the Research Office at the University of the Witwatersrand, Johannesburg.

References

[1] Carrió F *et al.* 2014 *Journal of Instrumentation* **9** C02019–C02019 ISSN 1748-0221 URL http://stacks.iop.org/1748-0221/9/i=02/a=C02019

[2] Cox M A, Reed R, Wrigley T, Harmsen G and Mellado B 2014 *In Review*

[3] Krazit T 2006 ARMed for the living room URL http://news.cnet.com/ARMed-for-the-living-room/2100-1006_3-6056729.html

[4] Beck H P *et al.* 2008 *IEEE Transactions on Nuclear Science* **55** 176–181

[5] Winklmeier F 2009 The ATLAS High Level Trigger infrastructure, performance and future developments *2009 16th IEEE-NPSS Real Time Conference* (IEEE) pp 183–188 ISBN 978-1-4244-4454-0 URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5321918

[6] EEMBC 2014 CoreMark Online Database URL http://www.eembc.org/coremark/

[7] Petitet A, Whaley R C, Dongarra J and Cleary A 2008 HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers URL http://www.netlib.org/benchmark/hpl/

[8] Feng W c and Cameron K 2007 *Computer* **40** 50–55 ISSN 0018-9162 URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4404810

[9] ARM 2011 Application Note 273: Dhrystone Benchmarking for ARM Cortex Processors URL http://infocenter.arm.com/help/topic/com.arm.doc.dai0273a/DAI0273A_dhrystone_benchmarking.pdf

[10] Frigo M and Johnson S 2005 *Proceedings of the IEEE* **93** 216–231 URL http://www.fftw.org

[11] CPU World 2014 AMD Opteron 252 Specifications URL http://www.cpu-world.com/CPUs/K8/AMD-Opteron252-OSP252FAA5BL.html