

УДК: 004.27

The development of an ARM system on chip based processing unit for data stream computing

Mitchell A. Cox^a, Robert Reed, Bruce Mellado

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

E-mail: ^amitchell.cox@cern.ch

Received October 2, 2014

Modern big science projects are becoming highly data intensive to the point where offline processing of stored data is infeasible. High data throughput computing, or Data Stream Computing, for future projects is required to deal with terabytes of data per second which cannot be stored in long-term storage elements. Conventional data-centres based on typical server-grade hardware are expensive and are biased towards processing power. The overall I/O bandwidth can be increased with massive parallelism, usually at the expense of excessive processing power and high energy consumption. An ARM System on Chip (SoC) based processing unit may address the issue of system I/O and CPU balance, affordability and energy efficiency since ARM SoCs are mass produced and designed to be energy efficient for use in mobile devices. Such a processing unit is currently in development, with a design goal of 20 Gb/s I/O throughput and significant processing power. The I/O capabilities of consumer ARM System on Chips are discussed along with to-date performance and I/O throughput tests.

Keywords: high data throughput computing, big data, arm system on chips

Разработка системы ARM на базе блока обработки данных для вычислений потока данных, реализованного на основе ИС

М. А. Кокс, Р. Рид, Б. Мелладо

Университет Витватерсранда, Южная Африка, 2000, Йоханнесбург, 1 Ян Смут Авеню

Современные масштабные научные проекты становятся все более информационно ёмкими, и обработка хранимых данных в режиме offline является невозможной. Требуется высокая пропускная способность при вычислениях или Вычисления Потокa Данных, чтобы иметь возможность обрабатывать терабайты данных в секунду; такие данные не могут быть элементами длительного хранения. Общепринятые дата-центры, основанные на стандартном аппаратном обеспечении, являются дорогими и настроены на вычислительную мощность. Общая пропускная способность может быть увеличена с помощью массивного параллелизма, чаще всего за счет повышенной вычислительной мощности и потребления энергии. Система ARM на основе ИС (SoC) может решить проблему системы ввода/вывода и соотношение CPU, доступность и эффективность использования энергии, так как ARM SoC являются элементами массового производства и разработаны на основе эффективного использования энергии в мобильных устройствах. На данный момент такой элемент обработки находится в разработке и нацелен на пропускную способность ввода/вывода в 20 Гб/с и значительную вычислительную мощность. Рассмотрены возможности ввода/вывода потребления системы ARM на основе ИС вместе с вычислением производительности и тестами на пропускную способность ввода/вывода.

Ключевые слова: высокая вычислительная пропускная способность, большие данные, система на ARM чипе

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.

Citation: *Computer Research and Modeling*, 2015, vol. 7, no. 3, pp. 505–509.

1. Introduction

Projects such as the Large Hadron Collider (LHC) at CERN and the Square Kilometer Array (SKA) in South Africa generate enormous amounts of raw data which presents a serious computing challenge.

A simple plot, shown in Fig. 1, of the increase in CPU processing power in MIPS (Million Instructions Per Second) and hard drive read-write speed in MB/s (MegaByte/s) over many years clearly demonstrates the fact that hard drive I/O rates are insufficient, and will not become sufficient in the near future, to store the entirety of raw data from modern scientific experiments such as the SKA and LHC [Dursi, 2012].

The increase in Ethernet throughput, however, is at a similar rate to the increase in CPU processing power. Based on Amdahl's Laws it has been recommended that approximately one compute instruction per bit of data is required for a balanced system and this relationship is clear when comparing CPU and Ethernet in Fig. 1 [Szalay et al., 2010]. It appears upon inspection that CPU performance and Ethernet throughput are well balanced but in reality high-end Ethernet is not commonly available except on very high-end systems. For example, 1 Gb/s Ethernet from 2002 is suitably balanced with a 2002 performance CPU. It is imbalanced when it is coupled with a modern CPU with an order of magnitude higher performance, however this is a very common situation since high-end CPUs are more prevalent than cutting-edge Ethernet.

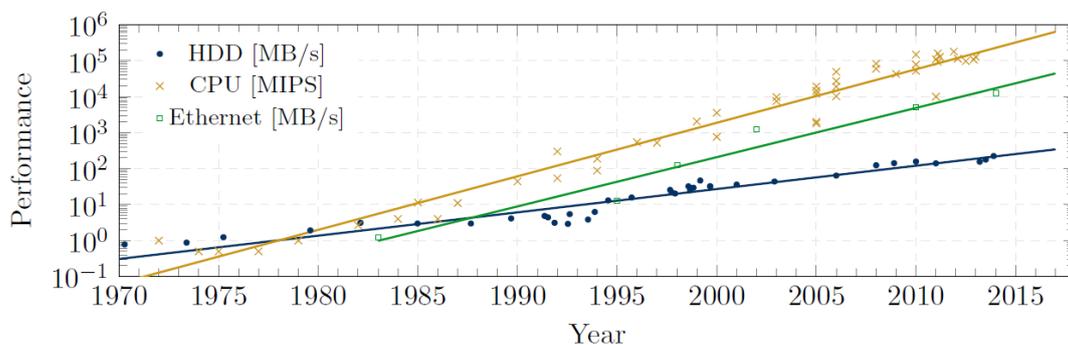


Fig. 1. Hard drive (HDD) and Ethernet 802.3 data throughput and CPU performance on a log scale since 1970 [Dursi, 2012; Wikipedia IEEE 802.3, 2014]

A specialised triggering and data acquisition system is currently employed by the LHC to reduce the amount of data produced to a manageable quantity for offline storage. This solution is not always suitable and so a paradigm shift is necessary to deal with future workloads and new projects. The cost, energy efficiency, processing performance and I/O throughput of the computing system to achieve this task is 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 typically biased towards processing performance and not I/O throughput and are therefore less-suitable for cost-effective high data throughput applications due to the necessity for massive parallelism.

High Volume throughput Computing (HVC) provides a suitable paradigm for data stream computing applications [Zhan J et al., 2012]. HVC is a datacenter based computing paradigm where the focus is on loosely-coupled throughput-oriented workloads in terms of either requests (service type applications), processed data (big data applications) or the maximum number of simultaneous subscribers (interactive real-time applications). The definition does not include data-intensive MPI workloads since these are suitably covered by High Performance Computing (HPC).

One of the first steps to the development of an effective HVC system is a high data throughput Processing Unit (PU). This PU should be well balanced in terms of CPU performance and I/O throughput and latency to maximise energy efficiency and cost.

ARM System on Chips (SoCs) are found in almost all mobile devices due to their low energy consumption, high performance and low cost and are the basis for the PU under develop-

ment [Rajovic et al., 2013]. Section 2 provides a brief overview of the specifications and performance for the SoC that was used for the PCIe testing. Two of these SoCs were connected via their PCI-Express interface and tested. This test setup is described and preliminary results are given in Section 3. Section 4 concludes.

2. ARM System on Chips

ARM System on Chips (SoCs) are low cost, energy efficient and high performance which has led to their extensive use in mobile devices. Several ARM platforms have been tested by the group at the University of the Witwatersrand, Johannesburg but only the specifications and test results for the Freescale i.MX6 quad-core ARM SoC is presented in Tab. 1 [Reed et al., 2014].

3. PCI-Express pair testing

PCI-Express throughput tests have been performed on a pair of Freescale i.MX6 quad-core ARM Cortex-A9 SoCs clocked at 1 GHz, located on Wandboard development boards [Wandboardorg 2012 Wandboard]. The results are presented in Tab. 2 and a photo of the custom test setup designed by the author is in Fig. 2. Three tests were run to ascertain the maximum data throughput that can be obtained from the i.MX6 SoC: a simple CPU based memcopy command and two Direct Memory Access (DMA) transfers, initiated by the Endpoint (EP) or slave and the Root Complex (RC) which is the host.

Table 1: CPU benchmark results and specifications of the Freescale i.MX6Q Cortex-A9 SoC [Reed et al., 2014]

Core Revision	r2p2
Clock (MHz)	996
Cores	4
Feature Size (nm)	40
SP GFLOPS	5.12
DP GFLOPS	2.40
Load Power (W)	5.03
Idle Power (W)	2.02
Calc. Power (w)	3.01
DP GFLOPS/W	0.80
Ethernet (Mb/s)	470
PCIe (Gb/s)	5

Unfortunately the i.MX6Q SoC does not have a DMA unit on the PCIe controller and so the Image Processing Unit DMA unit was used instead. This is a workaround provided by the manufacturer.

Table 2: PCI-Express throughput results of a i.MX6 (Wandboard) pair

	CPU memcopy	DMA (EP)	DMA (RC)
Read (MB/s)	94.8 ± 1.1%	174.1 ± 0.3%	236.4 ± 0.2%
Write (MB/s)	283.3 ± 0.3%	352.2 ± 0.3%	357.9 ± 0.4%

The theoretical maximum throughput for the PCI-Express Gen 2 x1 link that was used is 500 MB/s. The best result is using DMA initiated by the RC but it is only 72% of the theoretical maximum. The RC-mode drivers are more optimized than the EP-mode drivers due to limited manufacturer support for EP-mode. The read results are lower than write because of overheads to

initiate the read. The PU architecture will take these differences into account and use a data push rather than a pull based approach.

4. Discussion, conclusions and future work

Data stream computing, or more formally High Volume throughput Computing (HVC), is required for projects such as the LHC and SKA which produce enormous amounts of raw data. A general purpose ARM System on Chip based processing unit is being developed at the University of the Witwatersrand, Johannesburg which hopes to enable affordable and energy efficient HVC.

PCI-Express is superior to Ethernet in energy efficiency, I/O throughput and latency. Typical commodity ARM SoCs do not support Ethernet faster than 1 Gb/s however PCI-Express may be used for higher data throughput communications. Unfortunately, PCI-Express is not suitable for longer distance communications but the solution to this may be found in a PCI-Express to Ethernet bridge.

Initial throughput measurements presented for a pair of Freescale i.MX6 quad-core Cortex-A9 SoCs are 72 % of the theoretical maximum 500 MB/s for the available x1 link. Six of these SoCs would therefore be connected in parallel to provide 20 Gb/s throughput at a power consumption of less than 50 W. As a proof of concept the final Cortex-A9 prototype aims to provide 20 Gb/s aggregated throughput.



Fig. 2. PCI-Express test setup for a pair of i.MX6 SoCs (Wandboards)

The next stage of research by the author will be to test a small PCIe cluster of Cortex- A9 SoCs. The use of multiple energy efficient commodity ARM SoCs interconnected via PCI- Express and a single higher-end SoC for external communications via multiple 10 Gb/s Ethernet connections is theoretically well suited as a HVC Processing Unit.

Future big science experiments may be jeopardised by prohibitive data processing costs but the research presented in this paper, as well as future research and development of a HVC processing unit, may lead to a possible solution to this problem with its high data throughput, energy efficient and affordable computing capabilities.

References

- Dursi J.* Parallel I/O doesn't have to be so hard: The ADIOS Library Tech. rep. SciNet. 2012. URL: <http://wiki.scinethpc.ca/wiki/images/8/8c/Adios-techtalk-may2012.pdf>
- Rajovic N. et al.* Journal of Computational Science 4 439-443 ISSN 18777503. 2013. URL: <http://www.sciencedirect.com/science/article/pii/S1877750313000148>
- Reed R. et al.* A CPU Benchmarking Characterization of ARM Based Processors // The 6th International Conference: Distributed Computing and Grid-technologies in Science and Education (Dubna, Russia). 2014.

-
- Szalay A. S. et al.* ACM SIGOPS Operating Systems Review 44 71 ISSN 01635980. 2010. URL: <http://dl.acm.org/citation.cfm?id=1740390.1740407>
- Wandboardorg 2012 Wandboard - Freescale i.MX6 ARM Cortex-A9 Opensource Community Development Board accessed: 18 February 2014 URL: <http://www.wandboard.org/>
- Wikipedia. IEEE 802.3 — Wikipedia, The Free Encyclopedia accessed:16 September 2014 URL: http://en.wikipedia.org/wiki/IEEE_802.3
- Zhan J. et al.* High Volume Throughput Computing:Identifying and Characterizing Throughput. Oriented Workloads in Data Centers 2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum (IEEE). 2012. PP 1712-1721 ISBN 978-1-4673-0974-5 URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6270846>