Are Low-Power SoCs Feasible for Heterogenous HPC Workloads?

Max Plauth and Andreas Polze

Operating Systems and Middleware Group
Hasso Plattner Institute for Software Systems Engineering
University of Potsdam, Germany
{firstname.lastname}@hpi.uni-potsdam.de

Abstract. Energy efficiency has become a crucial aspect in the domain of High Performance Computing since running costs for electricity often exceed the initial acquisition costs. In consequence, low-power System-on-a-Chip designs are drawing much attention from the HPC community. Driven by the demand for high performance and long battery life in mobile consumer devices, all building blocks of SoCs are undergoing drastic improvements. In addition to the end-user availability of SoCs based on the ARMv8-A instruction set architecture, heterogenous aspects ranging from the big.LITTLE paradigm to compute-capable GPUs are gaining popularity. Focusing on the heterogenous nature of SoCs, we investigate both performance and energy consumption of todays state-of-the-art SoCs for heterogenous workloads using the Rodinia benchmark suite. Based on the results, we anticipate the potential of forthcoming SoC designs in the HPC domain.

1 Introduction

In the domain of High Performance Computing (HPC), energy consumption has always been an important factor in the design of new HPC setups. Recently however, the aspect of energy efficiency has acquired an additional facet for the HPC sector, as increasing energy costs have reached a level where the running costs exceed the initial acquisition costs. This issue is reflected by the emergence of numerous conferences[1], workshops[2] and projects[3] that are solely dedicated to the goal of increasing energy efficiency in the HPC domain. Animated by the motive of energy efficiency, the search for new approaches has engaged the interest of the HPC community in low-power System-on-a-Chip (SoC) designs [1, 5, 19–21].

Originating from the embedded and mobile sectors, the prime design goals of SoC designs are energy efficiency and reduced manufacturing costs [8]. However, the tremendous demand for compute intensive applications such as multimedia

[1] Energy-Aware High Performance Computing, http://www.ena-hpc.org
[2] UnConventional High Performance Computing, http://uchpc.lrr.in.tum.de
[3] Green500 List, http://www.green500.org

capabilities of high-end smartphones has driven major advances in the compute performance of modern SoCs as well. In addition to steady improvements of the chip design, SoCs have a history of employing special purpose hardware accelerators to meet high computational demands at a minimum level of energy consumption [23].

Focusing on the heterogenous property of SoCs, this paper evaluates the feasibility of state-of-the-art SoCs for heterogenous HPC workloads using the Rodinia benchmark suite [6]. Furthermore, the memory subsystem of all test platforms is evaluated using the STREAM [16, 17] benchmark and TinyMemBench [22]. In addition to mere *Time-to-Computation* (TtC) performance measurements, we also provide energy readings and the corresponding *Energy-to-Computation* (EtC) ratings. This evaluation targets recent trends such as heterogenous multiprocessing introduced by the *big.LITTLE* CPU paradigm as well as general purpose compute capabilities of SoC-grade GPUs. Furthermore, this work extends the scope of the evaluation to include SoCs based on the 64 bit *ARMv8-A* architecture, since previous evaluations of SoC hardware for HPC use cases have mostly dealt with hardware based on the 32 bit *ARMv7-A* architecture. For the measurements, we employed a variety of Single Board Computers (SBCs) to cover the range of the aforementioned characteristics.

To provide a representative for the current state of the art in x86_64 based systems, all tests were also performed on a rack-scale HPE Moonshot system. Even though technically the employed cartridges do not comprise SoC hardware, the growing density of rack-scale server systems indicates increasing opportunities for the use of SoCs in future rack-scale systems.

We provide the following contributions:

- We investigate the heterogenous properties of state-of-the-art SoCs, elaborating on both the compute capabilities of SoC-grade GPUs and the heterogenous multiprocessing feature of *big.LITTLE* CPUs.
- We analyze improvements of *ARMv8-A* compared to *ARMV7-A* SoCs.
- Based on the narrowing gap between ARM and x86_64 based SoCs, we anticipate the potential of forthcoming ARM designs in the HPC domain.

This work is structured as follows: In Section 2, we are going to present an overview of related publications that deal with the evaluation of low-power hardware for HPC workloads. Based on the research gaps, the Section 3 presents the choice of target platforms that will be evaluated in the further course of this work. In Section 4, all details of our benchmarking procedure are explicated. Subsequently, Section 5 presents the detailed results retrieved from our benchmark procedure and provides a thorough discussion. Finally, Section 6 concludes this work with final thoughts.

2 Related Work

In 2002, one of the first attempts at using low-power hardware in HPC scenarios was GreenDestiny, which accommodates a cluster of 240 Transmeta TM5600 667-MHz CPUs in a single rack and achieved 13.5 MFLOPS/Watt roughly an improvement of a factor two compared to competing systems built from COTS hardware at the time [24]. In a similar approach presented in 2005, MegaProto is built from 512 Transmeta TM8820 CPUs and achieved roughly 100 MFLOPS/Watt [18]. Unfortunately, even though the specifications of both approaches looked promising, neither one found its way into larger installations.

In an economic study conducted by HP Labs in 2011, the authors prognosticated that with further improvements of the production node, employing SoCs in server and HPC systems would yield major economic benefits in the future [15]. In 2013, the Mont-Blanc project [19] demonstrated the practical feasibility of employing ARM-based SoCs for HPC workloads using the Tibidabo cluster [20]. Tibidabo was the first ARM-based HPC cluster and was built from 128 NVIDIA Tegra 2 dual-core ARM Cortex-A9 processors. The setup achieved roughly 120 MFLOPS/Watt, which was competitive with CPUs from Intel (Xeon X5660) and AMD (Opteron 6128) at the time. The biggest issue of the hardware was the lack of Gigabit Ethernet and the unstable implementation of PCI Express. On the software side, linux distributions lacked the support for hardware floating point units and there was no software support for employing GPU-based accelerators for general purpose computations using CUDA or OpenCL. However, the authors reached a positive verdict and stated that the Cortex-A15 might reach competitive performance levels and that switching to the *ARMv8-A* architecture would yield drastic performance improvements.

In 2015, two publications have surfaced that investigate the performance characteristics of *ARMv8-A*-based SoCs. The publication by Rusitoru [21] takes on a very abstract approach, as the architectural properties of the *ARMv8-A* architecture are evaluated for HPC workloads using simulated hardware. This work explores the sensitivities of HPC workloads for in-order and out-of-order cpu configurations. A more practical approach has been applied in the work of Abdurachmanov et al. [1], which evaluates the performance of the Applied Micro X-Gene SoC for high throughput scientific computing tasks and compares the performance to an Intel Xeon CPU and a Xeon Phi accelerator card. While the X-Gene fails to keep up with the performance of the Intel-based platforms, it yields a much better performance per joule ratio. The authors conclude that ARM-based SoCs like the X-Gene are very promising and that with improving maturity of compilers and further hardware iterations, ARM-based SoCs might become valid competitors to x86_64-based CPUs.

Focusing on energy consumption versus performance tradeoffs, a very recent publication by Calore et al. (2015) [5] investigated the interactions between HPC workloads and energy consumption using Tegra K1 SoCs, which are based on the ARM Cortex-A15. Using a time-accurate measurement setup, the authors observed both compute performance and energy consumption for a benchmark

that has been implemented for both the CPU and the GPU portion of the SoC. As the measurements were performed at different clock speeds and energy states, the evaluation yielded *Time-to-Computation* (TtC) as well as *Energy-to-Computation* (EtC) figures for many configurations. The authors came to the conclusion that due to significant background energy dissipation, the best overall energy efficiency can be reached when all resources are driven at the fastest clock-speed in order to complete computations faster and then cut power.

Finally, the performance behavior of the heterogenous multiprocessing feature of *big.LITTLE* CPUs has been investigated by Butko et al. (2016) [4]. In their work, the authors used the OpenMP-based CPU implementations from the Rodinia benchmark suite [6] to evaluate the performance of the Samsung Exynos 5422 SoC, which comprises four ARM Cortex-A7 cores and four ARM Cortex-A15 cores. Contrary to the naïve assumption that running all cores simultaneously would yield certain performance improvements, the authors demonstrated that the best performance could be achieved using the Cortex-A15 cores only. In most benchmarks, using the Cortex-A7 cores as well resulted in severely decreased performance, which suggests that current parallel programming paradigms are unable to deal with the heterogenous CPUs where cores come with varying performance characteristics.

3 Hardware targets

In this work, we are using the latest generation of Single Board Computers (SBCs) in order to cover three major points of interest:

- Heterogeneous multiprocessing in *big.LITTLE* CPUs.
- Performance and compute capabilities of SoC-grade GPUs.
- Architectural improvements of *ARMv8-A* compared to *ARMV7-A*.

We decided to use SBCs due to their wide availability even though high-end ARMv8-A SoCs are available in server-scale hardware like the HPE ProLiant m400 Server Cartridge [12], which is based on the Applied Micro X-Gene SoC. The large demand from the hobbyist community has created a large spectrum ranging from low-cost, low-power hardware to more performant products. From the upper range of the performance spectrum, we are employing the Odroid-XU4 [3, 10] to cover both the aspects of heterogenous multiprocessing in *big.LITTLE* CPUs as well as the compute capabilities of SoC-grade GPUs. However, the XU4 still uses *ARMv7-A*-based CPU designs (ARM Cortex-A7 and A15) and thus does not represent recent advances in ARM processor design. For the investigation of advances introduced by *ARMv8-A*, we are using the Odroid-C2 [11] and the Raspberry Pi 3 [7], which are both based on the ARM Cortex-A53. It should be noted that the A53 is a low-power design, which is intended to be used as the energy efficient counterpart to the more powerful A57 in *ARMv8-A*-based *big.LITTLE* CPUs. The detailed specifications of all SBCs utilized in this work are denoted in Table 1.

Table 1: Detailed specifications of the employed Single Board Computers (SBCs).

	Raspberry Pi 3 [7]	Odroid C-2 [11]	Odroid XU-4 [3, 10]
SoC	Broadcom BCM2837	Amlogic S905	Samsung Exynos 5422
CPU	4×ARM Cortex-A53 1.2GHz, in-order	4×ARM Cortex-A53 2.0GHz, in-order	ARM big.LITTLE octa core 4×A7, 1.5GHz, in-order 4×A15, 2.0GHz, out-of-order
Arch	ARMv8-A (64 bit)	ARMv8-A (64 bit)	ARMv7-A (32 bit)
L1$ (I/D)	32KB/32KB	32KB/32KB	32KB/32KB
L2$	512KB	512KB	512KB (A7), 2MB (A15)
Memory	1GB LPDDR2 900 MHz	2GB DDR3 32 bit / 912Mhz	2GB LPDDR3 32 bit / 933MHz, PoP
GPU	BCM VideoCore IV	ARM Mali-450	ARM Mali-T628 MP6
Compute	no	no	OpenCL 1.1
OS	Ubuntu MATE 15.10	Ubuntu MATE 16.04	Ubuntu MATE 15.10
Kernel	4.1.18-v7+ (armv7l)	3.14.29-29 (aarch64)	3.10.96-78 (armv7l, HMP)
Compiler	GCC v5.2.1	GCC v5.3.1	GCC v5.2.1

To evaluate whether energy efficient SoCs might become a serious threat to the predominance of x86_64 in the field of High Performance Computing and to the server market in general, we extended our measurements to include the HPE ProLiant m710p Server Cartridge (detailed specifications denoted in Table 2). Systems belonging to the class of so-called Rack-Scale Computers [14] are particularly interesting, as energy-efficient but performant SoCs are the prime target for a system architecture that aims at highly increased levels of hardware density. Even though the Intel Xeon E3-1284L v4 employed in the m710p is no SoC, it represents the latest advances of performance-optimized x86_64 hardware targeting compact and energy efficient form factors.

Table 2: Detailed specifications of the x86_64 based reference system.

	HPE ProLiant m710p Server Cartridge [13]
CPU	Intel Xeon E3-1284L v4, 4C/8T, 2.90GHz, out-of-order
Arch	x86_64
L1$ (I/D)	32KB/32KB (per core)
L2$	256KB (per core)
L3$	6MB (shared)
L4$	128MB eDRAM
Memory	32GB, 4x8GB PC3L-12800 (DDR3-1600) SODIMM
GPU	Iris Pro Graphics P6300 BroadWell GT3
Compute	OpenCL 1.2
OS	Ubuntu 16.04 LTS
Kernel	Linux 4.4.0-21 (x86_64)
Compiler	GCC v5.3.1

4 Benchmark procedure

Corresponding to the focus on heterogenous hardware, we used a selection of tests from the Rodinia benchmark suite [6], which is specialized on heterogenous computing environments. The Rodinia suite comprises more than 20 benchmarks from various application domains, each implemented in OpenMP, OpenCL and CUDA. To provide an additional level of categorization, each benchmark is classified according to the *Berkley Dwarves* [2].

With the goal of prohibiting any bias, we ignored the specific hardware characteristics of each target platform and did not apply any modifications to the benchmark implementations, except for minor adaptations of the makefiles. While the OpenMP-based implementations for CPUs worked flawlessly on all platforms, the OpenCL-based implementations had certain issues on the ARM Mali-T628 MP6 GPU of the Odroid-XU4. Implemented with high-end workstation and server-grade GPUs in mind, several benchmarks failed to run properly on the XU4 without profound alterations of the implementation. As this forced us to narrow down the choice of benchmarks, we decided to pick one benchmark per *Berkley Dwarf* covered by the Rodinia benchmark suite:

- *Structured Grid*: Leukocyte tracking
- *Unstructured Grid*: Computational Fluid Dynamics
- *Dense Linear Algebra*: k-Nearest Neighbors
- *Graph Traversal*: Breadth-First Search

Except for the leukocyte tracking benchmark, all workloads are memory-bound [6] and thus sensitive to memory performance. Hence, memory bandwidth was measured using the STREAM benchmark [16, 17], whereas TinyMemBench [22] was used to obtain memory latency measurements.

All performance measurements were performed on a clean, freshly rebooted system with no other active users or background tasks running. In order to retrieve a sufficiently meaningful dataset, each benchmark was executed 10 times. Furthermore, each benchmark was preceded by a warm-up run in order to eliminate any confounding factors. The power consumption values denoted in Table 3 were measured using a power consumption meter switched in between the SBC power supply and the wall socket. For the m710p cartridge, the readings were retrieved from the management interface of the cartridge chassis.

Table 3: Power Consumption in Watts for the states *Off*, *Idle* and *Load*.

	RPI 3	C2	XU4 (A7)	XU4 (A15)	XU4 (GPU)	m710p (CPU)	m710p (GPU)
Off	0.50	1.00	0.70	0.70	0.70	9.85	9.85
Idle	1.70	2.30	3.80	3.80	3.80	20.65	20.65
Load	2.70	4.10	5.10	11.70	6.60	79.45	67.93

5 Results & Discussion

In this section, we provide *Time-to-Computation* (TtC) and *Energy-to-Computation* (EtC) results for application domains representing four major *Berkley Dwarves*: *Structured Grid* (see Figure 1), *Unstructured Grid* (see Figure 2), *Dense Linear Algebra* (see Figure 3) and *Graph Traversal* (see Figure 4). Except for the *Structured Grid* benchmark, all presented benchmarks are memory-bound [6].

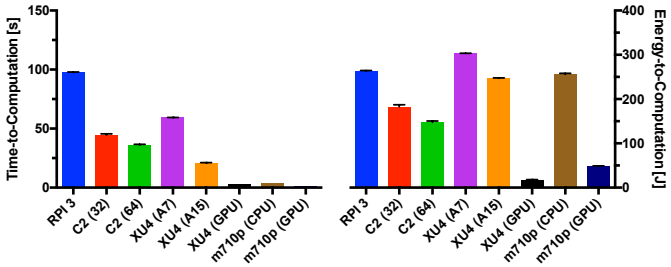

Fig. 1: Results of the *Structured Grid* benchmark (Leukocyte Tracking).

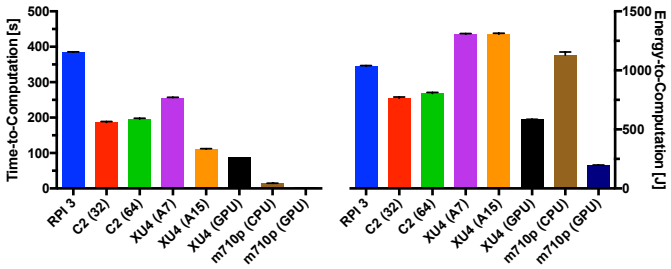

Fig. 2: Results of the *Unstructured Grid* benchmark (Computational Fluid Dynamics).

5.1 Heterogenous Properties

Analyzing the heterogeneous multiprocessing properties of the Odroid-XU4 (XU4), we were able to reproduce the findings of Butko et al. (2016) [4], that using both the A7 and A15 cores simultaneously results in decreased performance levels (data not shown). As this result has been expected, we were more interested in conditions where only the A7 or the A15 cores are used exclusively in order to evaluate whether the decreased compute performance of A7 cores might yield better *EtC* performance.

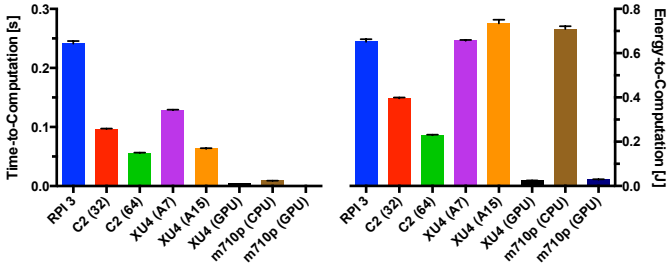

Fig. 3: Results of the *Dense Linear Algebra* benchmark (k-Nearest Neighbors).

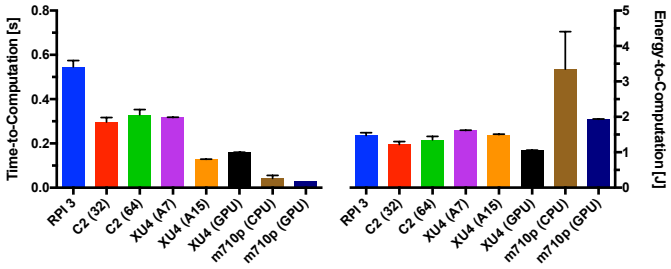

Fig. 4: Results of the *Graph Traversal* benchmark (Breadth-First Search).

Except for the *Dense Linear Algebra* benchmark presented in Figure 3, the A15 cores delivered much better *TtC* values, even mitigating the higher power consumption (see Table 3) in terms of *EtC* performance. These observations agree with Calore et al. (2015) [5], who recommended that running at high performance configurations and then turning of the system yields the best energy efficiency.

Comparing the OpenCL compute performance of the ARM Mali-T628 MP6 GPU with the CPU performance of the A15, the GPU provides significantly faster *TtC* performance in all disciplines except for *Graph Traversal* (see Figure 4). With almost half the power consumption however (see Table 3), the GPU exceeds CPU performance in the *EtC* category. While these results are already impressive enough, it should be noted that without further modifications to the benchmark code, only four out of six shader cores can be utilized and various optimization schemes remain untapped, including the use of vector types and zero-copy memory transfers. As a result thereof, the GPU employed in the XU4 offers a lot of potential for further performance improvements.

5.2 Improvements of ARMv8-A

The Raspberry Pi 3 (RPI 3) and the Odroid-C2 (C2) represent two different SoC designs that are based on the A53 CPU. Both in terms of *TtC* and *EtC* performance, the RPI 3 represents a low-key configuration, whereas the C2 demonstrates the strengths of the A53 design. In addition, the availability of a Linux kernel supporting the *aarch64* state makes the C2 the prime target for analyzing the improvements introduced by the *ARMv8-A* architecture.

Foremost, all benchmarks demonstrate that regardless of running in 32 or 64 bit mode, the A53 provides decent improvements compared to the A7 both with respect to *TtC* and *EtC* performance. Providing decent performance improvements at reduced energy uptake levels, the A53 delivers *EtC* performance improvements of 105% (*Structured Grid*, Figure 1), 71% (*Unstructured Grid*, Figure 2) and 186% (*Dense Linear Algebra*, Figure 3).

Focusing on the 64 bit mode of operation, both the *Structured Grid* benchmark (see Figure 1) and the *Dense Linear Algebra* benchmark (see Figure 3) demonstrate performance improvements of 24% and 73% compared to the 32 bit mode of operation, respectively. In contrast to that, the 64 bit mode causes a mild slowdown for the *Unstructured Grid* benchmark (see Figure 2) and the *Graph Traversal* benchmark (see Figure 4) with 5% and 10% respectively. While the beneficial benchmarks seem to profit from the increased register count and width of the *aarch64* mode, the slight performance drop in the remaining disciplines might be caused by the amenable state of platform specific optimizations of current *aarch64* compilers [1]. Both with improvements on the compiler side and manual optimizations on the benchmark side, it seems probable that decent performance improvements could be achieved for all conditions.

5.3 Competitiveness with x86_64

In terms of *TtC* performance, the x86_64 CPU evaluated in our benchmarks is far ahead of the ARM-based CPUs. However, this outcome should not come as a surprise, as comparing SBCs to high end server cartridges is close to a comparison between apples and oranges. Most notably, SBCs employ a much more basic configuration level of the memory subsystem, which is demonstrated by the memory subsystem performance discussed in Section 3. While the investigated SBCs employ single channel memory controllers and lower clock frequencies of the memory chips, the x86_64 system features a multi-channel memory controller and much higher memory frequencies. While this limitation is inherent to the class of SBCs, ARM-based SoCs targeted at server applications such as the Applied Micro X-Gene have demonstrated that memory subsystems with competitive performance can be implemented [9].

Taking on the position of *EtC*performance, the overall picture changes drastically as most ARM-based SoCs delivered better results in this discipline. With todays available technology, this win comes at the price of much higher computation times. However, if we consider the heterogenous aspects of SoCs and compare SoC-grade GPUs with x86_64 CPUs, the XU4 SBC manages to get

close to the CPU performance of the m710p cartridge in the *Structured Grid* and *Dense Linear Algebra* benchmarks (see Figures 1 and 3, respectively). Of course, the Iris Pro GPU on the x86_64 side also manages to provide performance boosts both in terms of *TtC* and *EtC*, but the SoC-grade GPU of the XU4 still manages to keep up in the latter discipline.

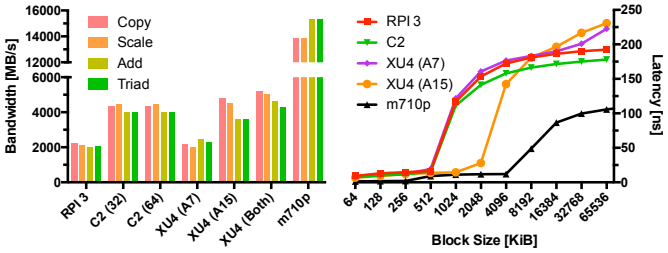

Fig. 5: Memory bandwidth as reported by the STREAM benchmark [16,17] (left). Dual read memory latency relative to L1$ as retrieved by TinyMemBench [22] (right).

6 Conclusion

Focusing on the heterogenous capabilities and architectural improvements of SoCs, this paper evaluated the feasibility of state-of-the-art SoCs for heterogenous HPC workloads using the Rodinia benchmark suite [6]. Furthermore, the memory subsystem of all test platforms was evaluated using the STREAM [16,17] benchmark and TinyMemBench [22]. Next to *Time-to-Computation* measurements, we also provided *Energy-to-Computation* figures to address the pressing subject of energy efficiency.

In our evaluation, we incorporated recent trends such as heterogenous multiprocessing introduced by the *big.LITTLE* CPU paradigm as well as general purpose compute capabilities of SoC-grade GPUs. Furthermore, this work investigated the improvements introduced by the *ARMv8-A* architecture, since previous evaluations of SoC hardware for HPC use cases have mostly dealt with hardware based on the 32 bit *ARMv7-A* architecture.

Considering that already lower-end in-order execution SoCs based on the Cortex-A53 indicated decent performance improvements compared to previous generations, promising improvements in the field of ARM-based CPUs are on their way. Similarly, surprising performance levels both in terms of *Time-to-Computation* and *Energy-to-Computation* were observed for SoC-grade GPUs. In an attempt to anticipate the near future based on our findings, the trend towards heterogeneous designs will be the key to superior performance. Hence, the potential of utilizing energy efficient hardware in compute intensive scenarios seems to be growing steadily.

Acknowledgement

This paper has received funding from the European Union's Horizon 2020 research and innovation programme 2014-2018 under grant agreement No. 644866.

Disclaimer

This paper reflects only the authors' views and the European Commission is not responsible for any use that may be made of the information it contains.

References

1. Abdurachmanov, D., Bockelman, B., Elmer, P., Eulisse, G., Knight, R., Muzaffar, S.: Heterogeneous High Throughput Scientific Computing with APM X-Gene and Intel Xeon Phi. Journal of Physics: Conference Series 608(1), 012033 (May 2015)
2. Asanovic, K., Bodik, R., Catanzaro, B.C., Gebis, J.J., Husbands, P., Keutzer, K., Patterson, D.A., Plishker, W.L., Shalf, J., Williams, S.W., Yelick, K.A.: The landscape of parallel computing research: A view from berkeley. Tech. Rep. UCB/EECS-2006-183, EECS Department, University of California, Berkeley (Dec 2006)
3. Aufranc, J.L.: ARM Cortex A15/A17 SoCs Comparison. http://www.cnx-software.com/2014/05/21/comparison-nvidia-tegra-k1-samsung-exynos-5422-rockchip-rk3288-allwinner-a80/ (May 2014)
4. Butko, A., Bessad, L., Novo, D., Bruguier, F., Gamatié, A., Sassatelli, G., Torres, L., Robert, M.: OpenMP Scheduling on ARM big.LITTLE Architecture. In: Proceedings of the Ninth International Workshop on Programmability and Architectures for Heterogeneous Multicores (MULTIPROG). Prague, Czech Republic (Jan 2016)
5. Calore, E., Schifano, S.F., Tripiccione, R.: Energy-Performance Tradeoffs for HPC Applications on Low Power Processors. In: Hunold, S., Costan, A., Giménez, D., Iosup, A., Ricci, L., Gómez Requena, E.M., Scarano, V., Varbanescu, L.A., Scott, L.S., Lankes, S., Weidendorfer, J., Alexander, M. (eds.) Euro-Par 2015: Parallel Processing Workshops. pp. 737–748. Springer International Publishing, Cham, Switzerland (2015)
6. Che, S., Boyer, M., Meng, J., Tarjan, D., Sheaffer, J.W., Lee, S.H., Skadron, K.: Rodinia: A Benchmark Suite for Heterogeneous Computing. In: Proceedings of the 2009 IEEE International Symposium on Workload Characterization (IISWC). pp. 44–54. IEEE (Oct 2009)
7. element14 Community: Raspberry Pi 3 Specifications. https://www.element14.com/community/community/raspberry-pi?ICID=rpimain-pi3doc-techspecs (2016)
8. Flautner, K., Flynn, D., Roberts, D., Patel, D.I.: IEM926: An Energy Efficient SoC with Dynamic Voltage Scaling p. 30324 (Feb 2004)
9. Gelas, J.D.: X-Gene 1, Atom C2000 and Xeon E3: Exploring the Scale-Out Server World. http://www.anandtech.com/show/8357/exploring-the-low-end-and-micro-server-platforms (Mar 2015)
10. Hardkernel co., Ltd.: ODROID-XU4 Hardware Information. http://odroid.com/dokuwiki/doku.php?id=en:xu4_hardware (2015)

11. Hardkernel co., Ltd.: ODROID-C2 Hardware Information. http://odroid.com/dokuwiki/doku.php?id=en:c2_hardware (2016)
12. Hewlett Packard Enterprise: HPE ProLiant m400 Server Cartridge QuickSpecs. https://www.hpe.com/h20195/v2/GetDocument.aspx?docname=c04384048 (2015)
13. Hewlett Packard Enterprise: HPE ProLiant m710p Server Cartridge QuickSpecs. https://www.hpe.com/h20195/v2/GetDocument.aspx?docname=c04760473 (2015)
14. Intel Corporation: Intel Rack Scale Architecture Overview. http://presentations.interop.com/events/las-vegas/2013/free-sessions—keynote-presentations/download/463 (Sep 2013)
15. Li, S., Lim, K., Faraboschi, P., Chang, J., Ranganathan, P., Jouppi, N.P.: System-level integrated server architectures for scale-out datacenters. In: Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture - MICRO-44 '11. p. 260. ACM Press, New York, New York, USA (dec 2011)
16. McCalpin, J.D.: Stream: Sustainable memory bandwidth in high performance computers. Tech. rep., University of Virginia, Charlottesville, Virginia (1991-2007), `http://www.cs.virginia.edu/stream/`, a continually updated technical report. http://www.cs.virginia.edu/stream/
17. McCalpin, J.D.: Memory bandwidth and machine balance in current high performance computers. IEEE Computer Society Technical Committee on Computer Architecture (TCCA) Newsletter pp. 19–25 (Dec 1995)
18. Nakashima, H., Nakamura, H., Sato, M., Boku, T., Matsuoka, S., Takahashi, D., Hotta, Y.: MegaProto: 1 TFlops/10kW Rack Is Feasible Even with Only Commodity Technology. In: ACM/IEEE SC 2005 Conference (SC'05). pp. 28–28. IEEE
19. Rajovic, N., Carpenter, P.M., Gelado, I., Puzovic, N., Ramirez, A., Valero, M.: Supercomputing with commodity CPUs: Are mobile SoCs ready for HPC? In: Proceedings of the 2013 International Conference for High Performance Computing, Networking, Storage and Analysis (SC). pp. 1–12. ACM Press, New York, NY, USA (Nov 2013)
20. Rajovic, N., Rico, A., Puzovic, N., Adeniyi-Jones, C., Ramirez, A.: Tibidabo: Making the case for an ARM-based HPC system. Future Generation Computer Systems 36, 322–334 (Jul 2014)
21. Rusitoru, R.: ARMv8 Micro-architectural Design Space Exploration for High Performance Computing Using Fractional Factorial. In: Proceedings of the 6th International Workshop on Performance Modeling, Benchmarking, and Simulation of High Performance Computing Systems. pp. 8:1–8:10. PMBS '15, ACM, New York, NY, USA (Nov 2015)
22. Siamashka, S.: TinyMemBench, `https://github.com/ssvb/tinymembench`
23. Silven, O., Jyrkkä, K.: Observations on Power-Efficiency Trends in Mobile Communication Devices. EURASIP Journal on Embedded Systems 2007, 1–10 (2007)
24. Warren, M., Weigle, E.: High-Density Computing: A 240-Processor Beowulf in One Cubic Meter. In: ACM/IEEE SC 2002 Conference (SC'02). pp. 61–61. IEEE (2002)