

Mainstream vs. Emerging HPC: Metrics, Trade-offs and Lessons Learned

Milan Radulovic^{1,2}, Kazi Asifuzzaman^{1,2}, Darko Zivanovic^{1,2}, Nikola Rajovic¹,
Guillaume Colin de Verdière³, Dirk Pleiter⁴, Manolis Marazakis⁵,
Nikolaos Kallimanis⁵, Paul Carpenter¹, Petar Radojković¹, and
Eduard Ayguadé^{1,2}

¹ Barcelona Supercomputing Center (BSC), Barcelona, Spain
`first.last@bsc.es`, `eduard@ac.upc.edu`

² Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

³ CEA, DAM, DIF, F-91297 Arpajon, France
`guillaume.colin-de-verdiere@cea.fr`

⁴ Jülich Supercomputer Centre, Forschungszentrum Jülich, Jülich, Germany
`d.pleiter@fz-juelich.de`

⁵ Foundation For Research & Technology - Hellas (FORTH)
`{maraz, nkallima}@ics.forth.gr`

Abstract. Various servers with different characteristics and architectures are hitting the market, and their evaluation and comparison in terms of HPC features is complex and multi-dimensional. In this paper, we share our experience of evaluating a diverse set of HPC systems, consisting of three mainstream and five emerging architectures. We evaluate the performance and power efficiency using prominent HPC benchmarks, High-Performance Linpack (HPL) and High Performance Conjugate Gradients (HPCG), and expand our analysis using publicly available specialized kernel benchmarks, targeting specific system components. In addition to a large body of quantitative results, we emphasize five usually overlooked aspects of the HPC platforms evaluation, and share our conclusions and lessons learnt. Overall, we believe that this paper will improve the evaluation and comparison of HPC platforms, making a first step towards a more reliable and uniform methodology.

1 Introduction

Each year we see a greater variety of HPC systems in the market. In addition to mainstream x86 architectures, emerging architectures based on POWER, ARM, SPARC and others are steadily appearing and catching attention [1]. Instead of hosting a single type of platform, supercomputing centers already provide a diverse set of systems. Making the right choice of architecture is critical, but evaluating and comparing HPC systems is hard.

Our study evaluates and compares three generations of mainstream x86 architectures: Intel Nehalem, Sandy Bridge and Haswell, and five emerging architectures: Intel Knights Landing (KNL), IBM Power8, Cavium ThunderX, and Applied Micro (APM) X-Gene 1 and X-Gene 2. In addition to presenting a large body of quantitative results, we emphasize six usually overlooked aspects of HPC platform evaluation, and share our conclusions and lessons learnt.

First, we show that a platform’s **performance and energy-efficiency depend significantly (n -fold) on the characteristics of the target applications**. For example, the ThunderX platform has 50% better energy efficiency

than Haswell when running memory-bound HPCG, but Haswell shows $3.6\times$ better efficiency for the compute-intensive HPL. We strongly advocate that any comparison between the platforms should start with the performance and energy-efficiency of HPL and HPCG as the boundaries of the compute-intensive and memory-intensive HPC applications. However, none of the previous studies that compare emerging and mainstream HPC platforms includes these results.

Second, our performance-per-watt results show that emerging **ARM-based platforms generally show much lower energy efficiency** than the KNL and mainstream Intel platforms. This finding opposes the conclusions of previous studies that report or estimate high energy-efficiency of the ARM platforms (reasons for this are discussed in detail in Section 4).

Third, it is important to understand whether HPC systems should be based on strong or weak CPU cores. **ARM-based platforms typically use weak cores** that are slower than x86, KNL and POWER8 servers in terms of floating-point performance and memory bandwidth. The ARM approach can still be used to build massively parallel systems that reach the target performance by using a larger number of cores. However this scale-out approach may cause significant performance and energy-consumption penalties [2], and it is important to confirm that it is indeed a viable alternative.

Forth, we detect a significant range in the **main memory access latency**, with a factor of three difference between the fastest and slowest platforms under study (90 ns–285 ns). Since memory latency has a direct performance impact for many applications [3], it should be minimized in HPC servers, and any increment above about 100 ns should be analysed and justified.

Fifth, we also analyse the **Byte/FLOP ratio** and detect a huge difference of up to $21\times$ among the platforms under study. The Byte/FLOP ratio is one of the most important design decisions, and we hope that our results will resurface a discussion on its desired range.

Sixth, our measurements show significant, **up to 70% differences between theoretical and sustained performance**, especially for emerging platforms. Therefore, we note the importance of measuring performance using specialized kernel benchmarks rather than relying on theoretical numbers from datasheets, even for the first order evaluation of the system. Also, hopefully these results will motivate further development of the emerging HPC compilers and scientific libraries.

In summary, given the substantial investment of time and money to deploy an HPC system, it is important to carefully evaluate and compare the available mainstream and emerging architectures. The conclusions of such an analysis depend significantly on the applied methodology, and the previous studies report the findings based on different experimental set-up, statistics of interest and benchmarks. Overall, we believe that this paper will improve the evaluation and comparison of HPC platforms, making first steps towards more reliable and uniform methodology.

2 Experimental environment

In this section, we explain efforts in evaluation of HPC systems, together with workloads and experimental platforms we used in our analysis.

2.1 HPC benchmarks

HPC benchmarks are important for bounding the sustainable performance of different components in a system.

High-performance Linpack (HPL) [4] has been the only metric for ranking HPC systems for a long time. It measures the sustained floating-point rate (GFLOPs/s) for solving a dense system of linear equations using double-precision floating-point arithmetic.

High-performance Conjugate Gradients (HPCG) [5] is based on an iterative sparse-matrix conjugate gradient kernel. The performance of HPCG largely depends on the available memory bandwidth [6].

HPC Challenge (HPCC) [7] is a benchmark suite that is designed to approximately bound computations of high and low spatial and temporal locality. We used DGEMM and STREAM benchmarks from HPCC suite. **DGEMM** is a floating-point intensive benchmark that represents the corresponding Level 3 Basic Linear Algebra Subprograms (BLAS) routine. The benchmark calculates the product of dense double precision matrices: $C \leftarrow \alpha A \times B + \beta$. It is used for measuring the sustainable FLOP performance, at the *per-core* or *per-node* level. The **STREAM** benchmark performs operations on arrays that are several times larger than the last level cache, effectively measuring the system’s sustained memory bandwidth [8]. It comprises four kernels: Copy, Add, Scale and Triad. In our analysis, we report the results of the Triad operation, since it is the most similar to kernels used in HPC applications.

LMbench suite [9] contains several benchmarks which measure performance of different hardware and software components in a system. We used the memory read latency benchmark in order to measure access latencies of different levels in memory hierarchy. The benchmark reads the input dataset in a random order to mitigate the impact of the data prefetching. By varying the input load size, we measure access latency to all memory hierarchy levels.⁶

2.2 HPC platforms

For the last decade, the dominant HPC architectures have been Intel architectures such as Nehalem, Sandy Bridge and Haswell. Apart from these, many-core systems, of which Intel’s KNL is an example, are becoming popular, while other vendors are also emerging architectures that are promising for HPC. For our study, we included mainstream HPC architectures which have been predominantly used in HPC systems so far, as well as emerging architectures which have been recently introduced to the market and are set to be used in future HPC systems. The architectures under study with their most important features are summarized in Table 1.

Comparing different HPC architectures under study is challenging. Architectures developed by different vendors essentially have different Instruction Set Architectures (ISAs) and therefore different system software such as compilers and scientific libraries. For each platform, we identified system software that provided

⁶ The measured latency comprises not only the latency of the hardware components (caches, memory controller, main memory), but also the latency of the system software such as virtual-to-physical memory translation.

Table 1: Summary of the most important features of platforms under study

Platforms	Mainstream architectures			Emerging architectures				
	Nehalem X5560	Sandy Bridge E5-2670	Haswell E5-2698v3	Knights Landing	Powers8	ThunderX	X-Gene 2	X-Gene 1
Manuf.	Intel	Intel	Intel	Intel	IBM	Cavium	APM	APM
Arch.	Nehalem	Sandy Bridge	Haswell	MIC	POWER8	ARMv8-A	ARMv8-A	ARMv8-A
Released	2009	2012	2014	2016	2014	2014	2015	2013
Sockets	2	2	2	1	2	2	1	1
Cores per Socket	4	8	16	68	10	48	8	8
CPU freq. [GHz]	2.8	2.6	2.3	1.4	3.49	1.8	2.4	2.4
Out-of-order	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
DP Flops, per cycle, per core	4	8	16	32	8	2	2	2
L1i	32kB	32kB	32kB	32kB	32kB	48kB	32kB	32kB
L1d	32kB	32kB	32kB	32kB	64kB	32kB	32kB	32kB
L2	256kB	512kB	256kB	1MB	512kB	16MB	256kB	256kB
L3	8MB	20MB	40MB	/	80MB	/	8MB	8MB
Memory conf. per socket	3 ch. DDR3 1333	4 ch. DDR3 1600	4 ch. DDR4 2133	8 ch. MCDRAM ^a + 6 ch. DDR4 2400	4 ch. DMI 28.8GBps	4 ch. DDR3 1600	4 ch. DDR3 1600	4 ch. DDR3 1600
Memory capacity per node	24GB	32GB	128GB	16GB (MCDRAM) + 192GB (DDR4)	256GB	128GB	128GB	64GB


^a KNL system has been set to *flat mode*, therefore both memories, MCDRAM and DDR4, are exposed as separate NUMA nodes, and the user can choose in which memory the workload executes.

the best performance. It has been used as is, and has not been tuned for each of the platforms. Hence, our conclusions should not be understood as a comparison between different hardware (CPUs and memory only), but a comparison of the platforms (systems) that also include the corresponding system software.

To our knowledge, there are no studies which analyze this many platforms, three mainstream and five emerging ones. Unlike some of the previous studies which performed first-order evaluation of the emerging platforms by using their developer kits, all the platforms under study are fully-fledged production servers that could be used in an HPC system. We argue that it is important to compare fully-fledged servers since their performance features and power consumption differ significantly from the corresponding developers kits.

2.3 Power measurements

For all platforms under study we measure the power consumption at the server level, which may comprise multiple sockets, as detailed in Table 1. The power measurements are performed with on-board or external power meters, that account for the overall server consumption including CPUs, memory, power supply, and so on. We also used power measurements to calculate the power efficiency of the platforms under study for HPL and HPCG benchmarks.



(a) HPL and HPCG performance (b) HPL and HPCG performance-per-watt

Fig. 1: Except for KNL, emerging plat. are generally behind the mainstream ones.

3 Results

In this section, we present results from the evaluation of multiple HPC architectures with different benchmarks. We start with the most prominent HPC benchmarks, HPL and HPCG, and later expand the analysis to include other benchmarks, which together give a more complete picture of a system’s performance.

3.1 HPL and HPCG benchmarks

This section gives insights on the performance and power efficiency of platforms under study, while executing the HPL and HPCG benchmarks.⁷ Figures 1a and 1b show the HPL and HPCG performance and performance-per-watt of platforms under study. These measurements are obtained with multi-threaded version of the benchmarks that use all the available cores.

Figure 1a lists the performance measurements. The KNL platform shows by far the best performance for HPL, followed by the Haswell and POWER8 servers, which reach 42% and 18% of the KNL’s HPL performance. Emerging ARM platforms show significantly lower performance. ThunderX, X-Gen 2 and X-Gen 1 deliver 5%, 1.3% and 0.1% of the KNL HPL scores, respectively. The results also show the notable improvement in the HPL performance over the various generations of mainstream platforms, from Nehalem to Sandy Bridge and Haswell.

The HPCG results show a slightly different trend. KNL using MCDRAM is still the highest-ranked platform, followed by Haswell, POWER8, Sandy Bridge and ThunderX. However, the gap in performance between these platforms is much lower for HPCG than for HPL.

Figure 1b lists the performance-per-watt results for the studied platforms. These measurements are important because one of the main drivers for research on the feasibility of the HPC on emerging ARM-based platforms is improved energy efficiency over mainstream HPC servers. Mainstream platforms show increasing power efficiency for HPL, with KNL as the best. POWER8, ThunderX and X-Gen 2 show significantly lower energy efficiencies, at just 7.6%, 8.3% and 2.9% of the KNL’s performance-per-watt. HPCG power efficiency increases from Nehalem to Sandy Bridge and then stagnates for Haswell and KNL using DDR4. On the other hand, KNL using MCDRAM achieves the highest power efficiency. Emerging platforms show a much lower power efficiency, except for the ThunderX platform, which is the second best, only 31% lower than KNL using MCDRAM.

The results show that, regarding power efficiency, it is very important to identify the target application. When targeting floating-point intensive applications (such as HPL), using low-power/low-performance cores seems not to be the best approach for overall energy efficiency. However, when targeting applications with lower processing requirements (and higher stress to other resources such as

⁷ For X-Gen 1 platform, we could not obtain power measurements.




Fig. 2: Sustainable FLOPS and memory bandwidth: The position of each platform in the chart shows the per-node sustained GFLOPs/s (x -axis) and memory bandwidth (y -axis). The marker size is proportional to the per-core performance.

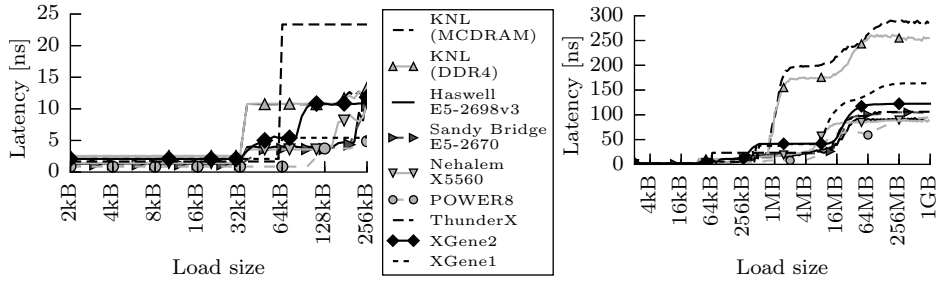
main memory) the ThunderX approach may deliver the energy efficiency, which significantly exceeds the x86 and POWER8 platforms.

3.2 Sustainable FLOP performance and memory bandwidth

In addition to the HPL and HPCG measurements, we also plot the raw measurements, sustained FLOPS and memory bandwidths of the studied platforms, at two levels of granularity, *per-node* and *per-core*. For node-level measurements, we execute the multi-threaded implementations of DGEMM and STREAM, using all the cores in the system. For core-level measurements, we execute the DGEMM benchmark on a single CPU core, and compute the per-core ratio of the sustainable bandwidth as the node STREAM result divided by the number of cores per node. This would correspond to a fair share of the sustainable memory bandwidth delivered to each application process assuming that the system is fully utilized.

The results are summarized in Figure 2. The position of each platform in the chart shows the per-node sustained GFLOPs/s (x -axis) and memory bandwidth (y -axis), while the size of the marker is proportional to the per-core performance. We also use different hatch patterns to distinguish between mainstream, ARM-based, POWER8 and KNL platforms. The results correlate with the HPL and HPCG analysis. Again, we see that the relative difference in the FLOPs performance (x -axis) is much higher than the memory bandwidth differences (y -axis). The Figure 2 also indicates that POWER8, Haswell and KNL platforms with DDR x memory interfaces reached plateau in terms of the sustained memory bandwidth, while KNL with MCDRAM provided a huge leap forward. Finally, the figure clearly shows that the emerging ARM platforms have lower performance on the server level (they are clustered in the left bottom corner of the chart) and their per-core capabilities are n -fold below the capabilities of the x86, POWER8 and KNL platforms (the size of the marker is much smaller).

This result launches an important discussion: should one use powerful cores similar to Haswell, POWER8 or KNL as the building blocks for large-scale HPC systems or use weaker cores, such as X-Gene or ThunderX? Reaching the target performance with weaker cores would require building of massively parallel systems, using higher number of cores, sockets and servers. This *scale-out* approach could indeed show good HPL and HPCG performance since these benchmarks have very good scalability and their performance is close to being proportional to the total system FLOPS or system memory bandwidth, respectively [10]. However, scale-out approach with production HPC applications may cause significant performance penalties. A recent study from Zivanovic et al. [2] analyzes the scale-out overhead of production HPC applications [11] running on a large HPC



(a) L1 and L2 cache access latency (b) L1, L2, L3 and main memory access latency

Fig. 3: Cache and main memory latency can vary significantly among the studied platforms. KNL memory access latency exceeds $3\times$ the latency on other platforms.

system [12]. Their results show that, even if the applications are substantially optimized, increasing the number of application processes to solve a fixed problem leads to significant increase in both energy and node-hours.

We therefore believe that before claiming that HPC systems based on weak cores are a viable alternative, it is essential to perform a profound analysis of the trade-offs between horizontal and vertical scaling in HPC.

3.3 Caches and main memory access latency

For decades, the memory system has imposed a fundamental limitation on system performance. This is recognized by the HPC community: HPL scores are frequently complemented by HPCG performance; sustained memory bandwidth is one of the main HPC performance metrics [7] [13], and high-bandwidth memory solutions caused a lot of interest by the HPC users. However, although the community invests significant effort to understand the memory bandwidth, the cache and main memory latencies are usually overlooked. This is surprising because the memory latency has a direct performance impact, and the memory wall itself was defined in terms of latency, not bandwidth [3].

In this section, we compare the access time of the caches and main memory for platforms under study. The results are plotted in Figure 3. The x -axis of the figures shows the input dataset size. In Figure 3a the load size ranges from 2 KB to 256 KB, which focuses on the L1 and L2 caches. In Figure 3b dataset size reaches up to 1 GB, covering all levels of caches and main memory. Even for the L1 and L2 caches we detect a significant difference in the latencies. At the L1 cache (2–32 KB load) the latency varies from 1.25 ns (Sandy Bridge, Haswell, POWER8) to 2.5 ns for KNL. In the L2 cache (128 KB load), the difference is even more significant, from 3.6 ns (Nehalem, Sandy Bridge, Haswell, POWER8) to 23 ns (ThunderX). The main memory latency (256 MB load) ranges from 90 ns (Nehalem, Haswell and POWER8) and 105 ns (Sandy Bridge, ThunderX) up to 250 ns and 285 ns for KNL using DDR4 and MCDRAM, respectively.

Overall, our measurements show that the cache and main memory latencies can vary significantly among platforms. Mainstream platforms and POWER8 perform well on all memory hierarchy levels. Emerging ARM platforms, ThunderX and X-Gene, have somewhat higher latency. KNL has significantly higher latency especially for the datasets that exceed 1 MB. Since these latencies have a direct performance impact, especially for the workloads with a high rate of dependent memory accesses, they are an important parameter to consider. KNL platform is




Fig. 4: Ratio between sustained memory bandwidth and FLOPS of the platforms under study can differ up to $21\times$.

especially interesting since it incorporates high-bandwidth MCDRAM, based on 3D-integration. While KNL delivers memory bandwidth far superior to any other platform under study, it comes with the cost of the memory access latency that exceeds $3\times$ the latency on mainstream platforms. Finally, it is also important to notice that most of the KNL memory access penalty does not come from the memory device itself. DDR3 and DDR4 modules timing parameters are standardized by JEDEC [14], and the variation between them (in nanoseconds) is negligible. Still, KNL DDR4 access is around 150 ns slower than other platforms. Therefore, the KNL memory access penalty originates mainly from handling the memory request between the last-level cache and the memory device, i.e. from the memory queues and memory controller. It is interesting to see whether future architectures will succeed in incorporating 3D-stacked memory without a significant latency overhead.

3.4 Byte/FLOP ratio

Using the node-level measurements of FLOPs and memory bandwidth, in Figure 4 we show the ratio between sustained memory bandwidth and FLOPS. This ratio presents the amount of data (in bytes) which has to be transferred from main memory, in order to perform one floating-point operation. Platforms with a low Byte/FLOP ratio are well suited for compute-intensive applications such as HPL. In these platforms, for real applications memory bandwidth may easily become a performance bottleneck. The platforms with a high Byte/FLOP ratio perform well with applications that put a high pressure on memory bandwidth, such as HPCG. In this case, floating-point processing power may limit the performance.

We detect a huge difference in the Byte/FLOP ratio among the platforms under study. The measured Byte/FLOP ratio ranges from 0.05 (KNL-DDR4) to 1.07 (X-Gen 1), a difference of more than $21\times$. For mainstream HPC systems (Nehalem, Sandy Bridge and Haswell), the Byte/FLOP ratio is significantly below 1, and it has the tendency of decreasing [15], which does not serve well for memory-bound HPC workloads. As seen in Section 3.2, current DDR x technology cannot keep up with aggressive FLOPs performance increases, so further progress in memory bandwidth relies on high-bandwidth memory solutions based on 3D-integration. In this respect, the KNL platform has a much higher Byte/FLOP ratio using MCDRAM than DDR4. Emerging systems, on the other hand, show a promising ratio, which is higher than mainstream platforms. This is mostly because the sustainable memory bandwidth is currently comparable between mainstream and emerging platforms, while the FLOPs performance of emerging systems is significantly below the mainstream ones. If this ratio keeps up with future developments of emerging platforms, we could see systems that cope better with memory-bandwidth intensive HPC workloads. Since HPC system performance strongly depends on the Byte/FLOP ratio, we advocate for this ratio to be constrained more precisely for HPC systems.




Fig. 5: Sustained FLOPS and memory bandwidth show significant difference to theoretical maximums, especially for emerging platforms.

3.5 Theoretical vs. sustained FLOPs/s and memory bandwidth

As the final step of our analysis, we compare the maximum theoretical FLOPS performance and memory bandwidth from platform datasheets with the sustained values measured using DGEMM and STREAM. This comparison is important because sometimes theoretical numbers are used to compare platforms or estimate large-scale system performance before they are built. Our results, however, show that the differences between theoretical and measured numbers may be significant.

The results are displayed in Figure 5. Mainstream HPC systems based on Sandy Bridge and Haswell deliver sustained FLOPS performance and memory bandwidth close to theoretical maximums. Some emerging architectures, however, reach moderate FLOPS and memory bandwidth utilization even when running the DGEMM and STREAM benchmarks. For example, X-Gen 1 and KNL reach only 48% and 56% of the maximum theoretical FLOPS, while X-Gen 2 and POWER8 achieve similar rates for memory bandwidth. An explanation could be that the overall system cannot fully utilize SIMD floating-point execution units or data-transfer mechanisms. By the overall system we include both hardware and system software, including the pipeline, out-of-order (OoO) engine, caches, compilers and scientific libraries. The HPC system software for emerging platforms is still under development; for example, the first math libraries for ARM-based servers were released two years ago [16]. Similar studies confirm that system software stack on emerging platforms is relatively immature, which limits the achievable performance [17–19]. Finally, ThunderX shows very low FLOPS and memory utilization of 23% and 27%, respectively. In this case, additional problem is the simplicity of the in-order core and poor performance of inter-socket communication.

This analysis has two outcomes. Firstly, we would avoid using maximum theoretical performance even for first-order provisioning or an early evaluation of the HPC system, especially for the emerging platforms. Secondly, for some of the platforms under study, the results also show notable room for performance improvement, which will hopefully motivate further development of compilers and scientific libraries for emerging HPC platforms.

4 Related work

In addition to mainstream x86 architectures, emerging architectures based on POWER, ARM, and SPARC are steadily appearing and catching the attention of the HPC community. Although making the right choice of architecture is critical for the HPC infrastructure providers, only few studies evaluate and compare available emerging and mainstream HPC platforms.

The study of Rajovic et al. [20] is the first to analyze the suitability of mobile ARM processors for HPC. The study compares the performance and energy efficiency of development boards using mobile ARM 32-bit SoCs against a laptop

with a Intel Sandy Bridge CPU. ⁸ Based on these measurements, the authors conclude that the performance and energy efficiency of mobile ARM platforms is competitive to the mainstream x86 HPC servers.

Abdurachmanov et al. [18] compare an X-Gene 1 development board with a dual-socket Intel Sandy Bridge server and Intel Xeon Phi PCIe add-on card. The authors compare only the CPU power consumption using on board sensors for X-Gene 1 development kit and Xeon Phi card, and RAPL interface [21] on the Sandy Bridge CPU. The study analyzes performance and energy efficiency of a single benchmark, ParFullCMS, and it concludes that the Sandy Bridge and Xeon Phi CPUs have similar performance that is $2.5\times$ higher than X-Gene 1. Performance-per-watt results position Sandy Bridge as the most efficient platform, followed by X-Gene 1 (approximately 10% lower efficiency) and Xeon Phi (more than 35% lower efficiency w.r.t. Sandy Bridge).

Early evaluation of emerging platforms using developer kits is valuable and needed. However, we argue that energy-efficiency analysis requires measurements on the fully-fledged production servers, as performed in our study.

Rajovic et al. [22] also deploy a prototype cluster with nodes based on mobile ARM 32-bit SoCs and compare it with a production HPC Sandy Bridge cluster. The study also estimates the performance of the potential successor mobile SoCs with advanced ARM cores and embedded GPUs. The authors conclude that emerging ARM-based systems would offer performance equivalent to mainstream x86 systems, while saving 40% energy, and achieving higher integration density. However, these conclusions are based on two non-trivial HPC application requirements. First, the HPC applications would have to fully utilize the GPUs embedded into emerging SoCs, which is not the case for most current production HPC codes. For the applications that can fully utilize the GPUs, the CPU+GPU emerging systems should be compared with similar (CPU+GPU) mainstream platforms, not with respect to the CPU-only systems. Second, the application should have perfect parallel efficiency and load balancing when scaling-out from strong x86 cores to a approximately $4\times$ larger number of weaker ARM cores. However, scale-out of production HPC applications typically leads to significant performance penalties [2]. Finally, the authors do not consider the performance and energy impact of RAS features (RAS: Reliability, availability and serviceability), such as memory ECC, available in the contemporary HPC systems, and not available on the emerging system under study.

Laurenzano et al. [17] compare the performance, power and energy consumption, and bottlenecks of Sandy Bridge, Atom, Haswell and X-Gene 1 servers. This analysis is based on system measurements with a large number of benchmarks and statistical modeling. The authors conclude that on average, for all the benchmarks under study, the X-Gene 1 and Atom servers have comparable performance, which is significantly below the Haswell and Sandy Bridge systems. When focusing on the workloads with an intensive use of floating-point and vector units, the study *estimates* similar performance on Sandy Bridge and X-Gene 1 systems. Our

⁸ In order to reduce the non-essential power consumption the authors switch off the laptop's screen.

measurements, however, show that DGEMM performance on Sandy Bridge CPUs exceed X-Gene 1 and X-Gene 2 by $9\times$ and $6\times$, respectively (see Section 3.2). Regarding the energy efficiency, Laurenzano et al. measure similar results for the X-Gene 1 and Sandy Bridge, somewhat below the Atom and Haswell servers. For all the platforms under study, the authors perform server-level measurements, but then extract the *power resulting from executing the application* as a subtraction between the server power executing the application and the idle server power. Our position is that using this metric to quantify and compare energy efficiency is misleading and unfavorable for servers with higher energy proportionality, in which power consumption is highly correlated to server performance.

The conclusions of the studies that evaluate and compare emerging and mainstream HPC platforms depend significantly on the methodology and benchmarks used. Still, the related work shows there is no unified approach for this analysis, and that the conclusions are sometimes based on a methodology and assumptions open to discussion. In addition to a large body of quantitative results, this paper emphasizes usually-overlooked and important aspects of the HPC platforms evaluation. We believe this will improve the evaluation and comparison of HPC platforms, making a first step towards a uniform and more reliable methodology.

5 Conclusions

In our study, we perform an extensive analysis of HPC architectures, three mainstream and five emerging ones. To the best of our knowledge this is the first study to include so many platforms. In addition to presenting a large body of quantitative results, we highlight six important features in HPC systems evaluation that require higher attention by the community.

First, we show a platform’s performance and energy-efficiency depend significantly (n -fold) on the characteristics of the target applications. We strongly advocate that any comparison among platforms should start with measurements using HPL and HPCG, which form the boundaries of compute-intensive and memory-intensive HPC applications.

Second, contrary to the conclusions of previous studies, our measurements on fully-fledged HPC servers show that emerging ARM platforms generally show much lower energy efficiency than the KNL and mainstream Intel platforms.

Third, we (re-)open a discussion on whether HPC systems should be based on strong or weak CPU cores. Using weak cores to deploy a scale-out approach may cause significant performance and energy-consumption penalties due to the imperfect scalability of production HPC applications. It is therefore important to confirm that this approach is indeed a viable alternative.

Forth, our results show a huge range of memory access latencies, from 90 ns to 285 ns for the studied platforms. While KNL with MCDRAM has the highest mem. bandwidth, it also has the highest mem. latency, due to complex memory controller and its handling of memory requests. Since memory latency has a direct performance impact any increment above about 100 ns should be analysed and justified.

Fifth, we detect that the Byte/FLOP ratio can differ by a factor of up to $21\times$ between platforms. While mainstream platforms show a decreasing tendency, emerging platforms trend upwards in this metric. We propose for a community

to properly define this ratio for HPC applications, since it has a direct impact on system performance.

Sixth, our results show that sustainable performance on the emerging platforms can deviate more than 70% from theoretical performance. Therefore, we strongly suggest not rely on theoretical performance, even in a first-order system provisioning. These results will hopefully motivate further development of the compilers and scientific libraries for emerging HPC platforms.

Overall, our study provides a significant body of useful information for HPC practitioners and infrastructure providers. Even more important, we believe it will considerably improve the future evaluations and comparisons of HPC platforms, making a first step towards a more reliable and uniform methodology.

References

1. TOP500 List: Supercomputing Centers Have Become Showcases for Competing HPC Technologies. <http://www.top500.org/news/supercomputing-centers-have-become-showcases-for-competing-hpc-technologies/> (June 2017)
2. Zivanovic, D., Radulovic, M., Llort, G., et al.: Large-Memory Nodes for Energy Efficient High-Performance Computing. In: Proc. of the International Symposium on Memory Systems. (2016)
3. Wulf, W.A., McKee, S.A.: Hitting the memory wall: Implications of the obvious. ACM SIGARCH Computer Architecture News **23**(1) (March 1995) 20–24
4. Petitet, A., Whaley, R.C., Dongarra, J., Cleary, A.: HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers. <http://www.netlib.org/benchmark/hpl/> (September 2008)
5. Dongarra, J., Heroux, M., Luszczek, P.: The HPCG Benchmark. <http://www.hpcg-benchmark.org> (2016)
6. Marjanović, V., Gracia, J., Glass, C.W. In: Performance Modeling of the HPCG Benchmark. Springer International Publishing (2014) 172–192
7. Luszczek, P.R., Bailey, D.H., Dongarra, J.J., et al.: The HPC Challenge (HPCC) Benchmark Suite. In: Proc. of the ACM/IEEE Conference on Supercomputing. (2006)
8. McCalpin, J.D.: STREAM: Sustainable Memory Bandwidth in High Performance Computers. Technical report, University of Virginia (1991-2007)
9. McVoy, L., Staelin, C.: Lmbench: Portable Tools for Performance Analysis. In: Proc. of the Annual Conference on USENIX Annual Technical Conference. (1996)
10. Zivanovic, D., Pavlovic, M., Radulovic, M., et al.: Main Memory in HPC: Do We Need More or Could We Live with Less? ACM Transactions on Architecture and Code Optimization **14**(1) (March 2017) 3:1–3:26
11. Partnership for Advanced Computing in Europe (PRACE): Unified european applications benchmark suite. www.prace-ri.eu/ueabs/ (2013)
12. Barcelona Supercomputing Center: MareNostrum III System Architecture. <http://www.bsc.es/marenostrum-support-services/mn3> (2013)
13. Williams, S.W., Waterman, A., Patterson, D.A.: Roofline: An Insightful Visual Performance Model for Floating-Point Programs and Multicore Architectures. Technical Report UCB/EECS-2008-134, EECS Department, University of California, Berkeley (October 2008)
14. Joint Electron Device Engineering Council. <http://www.jedec.org>
15. McCalpin, J.D.: SC16 Invited Talk: Memory Bandwidth and System Balance in HPC Systems
16. ARM Ltd: ARM Accelerates Mathematical Computation on 64-bit ARM-based HPC Systems. <http://www.arm.com/about/newsroom/arm-accelerates-mathematical-computation-on-64-bit-arm-based-hpc-systems.php> (November 2015)
17. Laurenzano, M.A., Tiwari, A., Cauble-Chantrenne, A., et al.: Characterization and bottleneck analysis of a 64-bit ARMv8 platform. In: IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS). (April 2016) 36–45
18. Abdurachmanov, D., Bockelman, B., Elmer, P., et al.: Heterogeneous High Throughput Scientific Computing with APM X-Gene and Intel Xeon Phi. CoRR (2014)
19. Reguly, I.Z., Keita, A.K., Giles, M.B.: Benchmarking the IBM Power8 Processor. In: Proc. of the International Conference on Computer Science and Software Engineering. (2015)
20. Rajovic, N., Carpenter, P.M., Gelado, I., et al.: Supercomputing with Commodity CPUs: Are Mobile SoCs Ready for HPC? In: Proc. of the International Conference on High Performance Computing, Networking, Storage and Analysis. (2013) 40:1–40:12
21. Intel Corporation: Intel[®] 64 and IA-32 Architectures Software Developer’s Manual. Technical report (July 2017)
22. Rajovic, N., Rico, A., Mantovani, F., et al.: The Mont-blanc Prototype: An Alternative Approach for HPC Systems. In: Proc. of the International Conference for High Performance Computing, Networking, Storage and Analysis. (2016) 38:1–38:12