

Perspectives in High Performance Computing

Michael Resch

Abstract High Performance Computing is undergoing a major change in the coming years. This paper discusses the perspectives that HPC has in the coming years and how these perspectives are going to change the way we operate and use HPC systems. The paper discusses the technologies that HPC will use, discusses the current trends based on the TOP500 list and argues that a further improvement in performance will basically be driven by software rather than by hardware.

Keywords High performance computing · Trends · Hardware · Software

1 Introduction

High Performance Computing (HPC) architectures are a hotly debated issue as the designers of such systems are increasingly facing new challenges. Looking at current developments traditional approaches seem to be running out of steam. A few years ago, HPC centers were concerned about the lack of variety of architectures and suspected that a monocultural world would take hold of the HPC market. In fact, a monopoly of architectures can already be seen today with many vendors having left the market. In recent discussions on architectures at the International Supercomputing Conference at Dresden [1] it became clear that this monopoly is triggering a new development of architectures. Most of them are not yet mature enough such that it is unclear which of them will reach a level of maturity that would allow their usage in everyday production. This has made it increasingly difficult to make investment decisions when it comes to designing large-scale systems. On the other hand, HPC centers are still concerned that they might be running out of options when it comes to procuring for next-generation systems.

M. Resch (✉)

High Performance Computing Center Stuttgart (HLRS), University of Stuttgart,
Nobelstrasse 19, 70569 Stuttgart, Germany
e-mail: resch@hlrs.de

The situation in the field of HPC is actually more complicated than it was 10 years ago. IBM recently decided to sell its manufacturing facilities. The step that was widely assumed to be a starting point of an exit strategy from the hardware business. At the same time, IBM made the design for the Power architecture publicly available—giving the market another visible signal for a retreat from the HPC market. These were two steps that followed the decision of IBM to sell its x86 activities to Lenovo—a Chinese vendor that already took over when IBM dropped their laptop business. Furthermore, the market receives signals that the highly successful BlueGene line of IBM may not see a follow-on product. With IBM giving mixed signals for HPC, and with the disappearance of vendors like SUN, the HPC market is left with few options—having experienced a continuous decline in number of stable vendors over the last years.

Technically HPC is facing the end of a development that used to be called Moore's law [2]. Processor clock frequencies—which carried the main load of speeding up hardware—cannot be further increased since 2004. Multi-core processors have become standard. So-called accelerators provide solutions that push the number of cores on a single chip to extremes but leave the users with adapting their codes to a new architecture and a new programming model.

In this paper, we will investigate a number of questions that come up in this context. We will explore the messages that the history of the TOP500 [3] list provides. In most recent editions, we have seen interesting developments that will be important for centers and users alike. We will furthermore look into technically new trends that may help to overcome some of the limitations that we face with massively parallel systems. Finally, we will try to explore and evaluate new technologies that might be available to the market in the near future.

2 A Little Bit of History

HPC has for a while been dominated by a development that was described by Gordon Moore in 1965 [2]. Basically, Moore figured that the number of components on a chip of the same size had doubled consistently over a certain period of time. From this, he concluded—by making an economical argument rather than diving too deep into technical details—that a similar development could be expected in the near future. Originally, he assumed a doubling of components every 12 months and later on modified this to a doubling every 18 months. As a corollary from this, it was assumed for a long time that clock frequencies of processors could be doubled every 18 months. The basic assumption was that reducing the feature size would reduce the distance for a signal to travel and hence increase the clock frequency.

For several decades, Moore's expectation proved to be right. Clock frequencies actually increased and basically followed the expected path. The development started to slow down for high-end processors in the mid-1990s. Clock frequencies were at about 0.5 GHz and higher. At the same time so-called standard processors (at the time provided by Intel and AMD) rapidly caught up with HPC systems—driven by

increased clock frequencies and by a market that was eager to absorb whatever new processor became available for the consumer market.

The slow down in clock rate increase was foreseeable and parallelism was early on investigated as a concept to overcome the problem. With the introduction of parallel processing the focus shifted from single processor speed to number of cores available in a single system. Early adopters of the concept failed but provided the necessary groundwork for our current technology in HPC. By 2004—when clock speeds started to stall at a value around 2–3 GHz—parallelism took over as the leading principal in HPC [4].

3 Technology Trends

To make up for the lack of acceleration based on increased clock frequencies parallel computing was pushed to the extreme over the last decade. Parallelism is not a new paradigm. It was exploited over time in a variety of architectures. Actually, even a standard technology like pipelining is in a sense some sort of parallelism exploited in the architecture. However, at the processor level parallelism arrived relatively late and the level of parallelism employed in high performance computing systems remained relatively moderate for a while. The number of processors used was hovering around 512–1024—with the Thinking Machine approach being the notable exception.

The currently fastest system in the world is based on about 3 million compute cores bundled in a single system [3]. This development is based on the fact, that the principal idea of Gordon Moore is still true. The number of components on a chip can still be further increased and most likely will grow over the next years according to the International Technology Roadmap for Semiconductors [5].

When it comes to exploiting parallelism in hardware there are two paths the market is following.

3.1 *Thin Core Concept*

The concept of thin cores is mainly focused on parallelism. The basic idea is to build relatively simple cores but to build a large number of them on a single chip. Graphics processing units (GPU) heavily influenced this concept. Some solutions actually evolve originally from GPUs, which were modified in order to meet the requirements of high speed computing.

The thin core concept is based on the reasonable idea that hardware designed for high speed computing—which is typically measured in floating-point operations—should be focusing on floating-point units only. The perfect solution would be a concept in which a core is not much more than a floating-point unit.

The concept as described above is based on low clock frequencies and large numbers of cores on a chip. Very often manufacturing technology is not leading edge, as mass production for GPUs does not require high-end manufacturing technology.

Increased speed can therefore be expected for this concept from two sources mainly:

- **Higher clock frequencies:** Thin core concepts are typically based on a relatively moderate clock frequency in the range of 1 GHz or less. This allows to keep power consumption relatively low and hence squeeze more cores on a chip. Theoretically, clock frequencies still have the potential to grow by about a factor of 2–4 over the coming years. They will then have reached the level of current state-of-the-art standard processors and a further increase would lead to similar cooling and power consumption problems as with standard processors. But to increase clock frequency—even if only moderately—is an option to go for higher total performance of a chip.
- **Higher core numbers:** Using more advanced manufacturing technology, the number of components on a chip will be increased. Keeping the design of an individual core as simple as possible, the additional components can be used to increase the number of cores on a chip. This could potentially increase the number of cores on a chip by a factor of 4–8 in the coming decade.

Putting together these two trends it seems to be possible to both increase clock frequencies slightly and to further increase the number of cores on a chip. For a General Purpose GPU or for similar accelerator concepts we can expect to see a factor of 2–32 in peak performance over the coming years.

Current developments, however, seem to indicate another trend. With an increasing demand for these kinds of accelerators designers are trying to turn these cores into floating-point machines that better fit the requirements for standard simulations. Having learned that cheap GPUs can be used to speed up high-end computing systems companies increasingly see the potential in the HPC market for their product. However, the complexity of the cores has to be increased to meet the requirements of the HPC user community. Extrapolating the trend one might expect to see a stagnation in number of cores on a chip while complexity and clock frequency are increased harvesting the potential of new manufacturing technologies.

What might be expected is a moderate growth in speed only but an increase in potential for standard applications.

3.2 Fat Core Concept

The concept of fat cores could be described as the “classical” approach to high performance computing architectures. The increased speed required for HPC is delivered by increasing the complexity of processor architectures. A lot of complexity is added for example to overcome the limitations of slow memory subsystems. Additional

available components are also used to add further functional units. By doing this we gain an additional level of parallelism directly on the chip. Performance is increased by having 4 or 8 ADD-Mult per clock cycle rather than speeding up each individual ADD-Mult.

A number of processor could be described as “fat” with each of them following a different path of development.

- **X86 architectures:** The x86 processor family is the standard architecture in HPC for several years now. The number of cores is relatively low with currently 8–16 cores for a single CPU. Each of the cores itself is a highly tuned architecture with a number of sophisticated features that—if adequately programmed—turn these processors into high-end computing engines.
- **IBM Power:** The IBM Power processors is the last surviving specialized standard processor in the arena of HPC. The first standard multi-core chip for HPC was a Power processor and since then the Power processor has proven to be always of the leading edge of processor technology.
- **Vector processors:** Vector processors seemed to be extinct when NEC dropped out of what later was to become the Japanese K Computer project. However, they reappeared in 2014 when NEC introduced its SX-ACE line. The concept follows a traditional approach with introducing vector pipes as the core to achieve performance and a relatively sophisticated memory subsystem that allows pushing sustained performance to a level hardly reachable by standard processors. However, the prize for such systems is still comparably high and hence they hardly make an appearance in the TOP500 list.

3.3 Memory

At this point it makes sense to talk about memory technology. High Performance Computing hit the memory wall about 20 years ago. Increased processor speed was not matched by memory speed—neither in terms of latency nor in terms of bandwidth. Modern architectures have become increasingly imbalanced. As a result, users can expect a sustained level of performance that varies widely. The more a code is limited by memory speed the lower its sustained performance. Experts speak of about 3–5% of performance that can be achieved when working without cache aware programming.

Caches were seen to be the way to overcome memory speed limitations. Introducing small but fast caches on-chip, vendors were hoping to break the memory wall. Over time cache levels were introduced and as of today we expect to see three levels of caches in a high-end HPC processor. But as memory and cache systems get more complex users are facing two further problems.

- **Complexity:** With the growing complexity of cache hierarchies it gets increasingly difficult to optimize a code for a given hardware architecture. Once optimized for

one architecture the changes in cache hierarchy of another architecture may cause a drop in performance by as much as a factor of 10 or more. In order to fully exploit an architecture, programmers would have to be aware of the architecture—which changes rapidly and increases its complexity continuously.

- **Imbalance:** At the same time as users struggle with cache hierarchies and their complexities architect are faced with handling the memory subsystem side effects. As a result users are increasingly facing imbalances between memory and cache hierarchies on the one hand and architectural features on the other side. The most simple problem is that for some cache hierarchies to work properly a processor needs enough registers to handle all the traffic. If this is ignored it is in the end the processor architecture that kills the memory and cache subsystem.

3.4 TOP500 Trends

The HPC community criticized the TOP500 [3] list for many years. There is at least one large-scale installation that refused to participate in the list claiming that the Linpack benchmark has no whatsoever justification to be used as a yardstick for HPC systems. Although there is some truth to this claim the TOP500 list has shown to be an interesting collection of statistical data from which at least trends can be extracted [6].

Exploring recent developments in the list allows getting a better understanding of trends and markets in HPC. Over the last years, the most striking feature is that the replacement rate of systems in the list for high-end systems is slowing down. A brief analysis of the TOP10 systems over the last years shows the following:

- In November 2011 three new systems were in the TOP 10 compared to November 2010
- In November 2012 eight new systems were in the TOP 10 compared to the previous year
- In November 2013 three new systems were in the TOP 10 compared to the previous year
- In November 2014 one system was new compared to the previous year—and even this system was not a full replacement but an upgraded version of an existing system
- In November 2015, we were back to three new systems.

When we look at the five fastest systems we see no change since 2013.

What is more interesting is the trend line that can be retrieved from the last 21 years of collecting information about the fastest systems in the world. Figure 1 shows the trend lines for the performance of




Fig. 1 Trend lines of the TOP500 list. Basic data from www.top500.org with added trend lines by the author

- (a) The number 500 system (lower line),
- (b) The fastest system (middle line), and
- (c) The sum of all systems on the list

The figure was taken from the TOP500 webpage. The authors added trend lines.

The figure indicates that the number 500 system—the slowest system on the list—is unable to follow the general trend since about 2009/2010. A similar trend cannot yet be seen for the sum of performance of all systems. It looks though as if for the last four years the slope of the trend is smoother. It is too early to say that this is a general trend. However, it remains to be seen what is going to happen. The most recent version of the list—as published in June 2015—indicates that we may see a smoother slope for the total performance too.

There is an optimistic scenario for the trend which claims that especially the slower systems have not yet adopted the accelerator technology that allows faster systems to still keep pace with Moore’s law. Following this scenario, the market should catch up over the coming 2 years and the trend of the number 500 systems should go back to what it used to be. It remains to be seen though whether the owners of smaller system are able to exploit the potential of accelerators. For those among them that work in a research environment—like universities or research labs—this should not be too difficult. However, the many industrial users of low-end HPC systems may not see an incentive in investing into a technology for which there is not yet a defined

standard and for which there are not too many industrial codes that can easily run on accelerator systems. Looking back into the history of parallel computing we find that industry did catch up on parallelism but with the growing number of processors, industrial usage was increasingly decoupled from research trends.

There is also a pessimistic scenario, which claims that accelerators are a work-around for the problems of stagnating processor performance. The pessimistic scenario assumes that we will start to see a changing trend line also for the number one systems in the years ahead. The pessimistic scenario would suggest that the rule of Moore's law is over.

An interim report of the "Committee on Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science in 2017–2020" released in November 2014 states: *It is an accepted truth today that Moore's Law will end sometime in the next decade, causing significant impact to high-end systems.* And the report continues: *The transition implied by the anticipated end of Moore's Law will be even more severe—absent development of disruptive technologies; it could mean, for the first time in over three decades, the stagnation of computer performance and the end of sustained reductions in the price–performance ratio.* [7]

If we believe in the end of Moore's law we need to face the consequences and prepare for the time after.

4 What to Expect?

First, and foremost, we would have to accept that the end of Moore's law has been reached. This does not come entirely as a surprise—it was on the contrary rather surprising to see technology follow such an impressive path for more than three decades. However, it is not a reason for pessimism but rather a reason to step up the research efforts in HPC. Three main consequences follow from what we have found so far.

4.1 We Need More Investment in Better Technology

Hardware development is going to address a number of new issues beyond performance. It is reasonable to expect to see processors that are not built for floating-point performance but rather for the growing needs of data analytics. Furthermore power consumption will become a growing issue for processor architecture design. Even more than now, hardware designers will put their focus on reducing power consumption thus providing the user with lower operational costs of systems. How much this in turn will trigger a further increase in number of cores or processors remains to be seen. We may see a moderate growth into the billions of cores after a while.

We may, for example, expect to see some sort of follow-on project of the IBM BlueGene. It would be interesting to see an architecture built from billions of

low-power embedded processors. As much as this would be a challenge for programmers it could yield to interesting architectural concepts.

In any case investment will not stop at hardware. There is a growing need for better programming tools. Handling millions of cores is counterintuitive for the human being. All concepts that are able to reduce this complexity—like hybrid programming models that merge MPI and OpenMP—will be extremely useful for the user. However, such concepts are in their infancy and will require a lot of effort before they can be turned into standards and supported by all necessary tools.

4.2 Convergence and Segmentation

Thin core concepts will increase clock frequencies and will—for the sake of being useful for HPC—increase complexity of each core. Hence, they will grow fatter and hit the frequency barrier. Fat core concepts will reduce complexity in order to reduce power consumption and in order to be able to increase the number of cores. Therefore, it has to be expected that the thin core concept and the fat core concept will somehow converge. What we see already today in the market is a trend to merge accelerator technology onto standard processor parts. Sometimes these are called “fused parts.” We also see technologies like the Aurora concept of the Japanese vendor NEC [8]. The basic concept of Aurora is to turn the traditional vector processor concept (a typical fat core concept) into an accelerator that could be used like existing accelerators.

Given the wide variety of options in the design space for future processor architectures we can expect to see a market evolving that is similar to other mature markets. We can expect to see cheap solutions with a reasonable performance for a low price. We can also expect to see tailored high-end solutions for which niches will have to be carved out to survive. In any case we currently see a growing number of different solutions that all follow similar lines of architectural concepts but with the exception of the x86 architecture there is currently no solution available that could claim to be a standard solution.

4.3 Software Beats Hardware

Regardless of the directions that hardware design takes, software will become more important. As of today, we expect to see single digit sustained performance figures for large-scale systems. The recently initiated HPCG benchmark initiative [9] reveals that even for a standardized benchmark—which should by now be highly optimized—sustained performance numbers are embarrassingly low. In the future, optimization of codes is going to be a major issue. Given that hardware architecture development will slow down software developers will be given more time to get the most out of a given hardware concept. With hardware stagnation, it also makes sense to rethink

many of the old models that are used by HPC users for decades now. In the future software will make the difference between a standard system and a high performance computing system.

5 Summary

HPC is facing the end of the basis for its success story over the last three decades. With Moore's law ending, peak performance is no longer something that just so happens. This will have implications for users, vendors, and HPC centers. Centers will have to invest much more in quality of services and will have to work intensively with software developers to be able to provide high quality services. Hardware vendors will have to focus more on improved quality of hardware rather than on speed. They will have to explore and carve out niches in which HPC as a business can create a reasonable ecosystem. This will bring industrial users much more into the focus of activities than has been the case over the last decades. HPC users will have a much harder time improving their simulations. The focus will have to move from speed to quality. This is the time when software has to be improved. This is the time when models have to be improved. This is the time when algorithms have to be improved. This is going to be a great time for HPC experts—computer scientists, mathematicians, computational scientists, and engineers.

References

1. ISC 2015, Dresden, Germany. www.supercomp.de
2. Moore, G.E.: Cramming more components onto integrated circuits. *Electronics* **38**(8), 114–117 (1965)
3. www.top500.org
4. Resch, M.M.: Trends in architectures and methods for high performance computing simulation. In: Topping, B.H.V., Iványi, P. (eds.) *Parallel Distributed and Grid Computing for Engineering*, pp 37–48. Saxe-Coburg Publications, Stirlingshire, Scotland (2009)
5. <http://www.itrs.net/>
6. Resch, M., *Citius Altius Fortius—The TOP500 and Why Speed Matters*, ISC Conference 2013, Leipzig, Germany, June 17, 2013
7. Committee on Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017–2020, *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017–2020: Interim Report*. 2014. <http://www.nap.edu/catalog/18972/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-and-engineering-in-2017--2020>
8. IDC, NEC's HPC Vision: Bringing Vector Supercomputing to a Broader Data-Intensive User Base. <http://www.marketresearch.com/IDC-v2477/NEC-HPC-Vision-Vector-Supercomputing-8573697>
9. <http://hpcg-benchmark.org/>