HPC Forecast: Cloudy and Uncertain

DANIEL REED, University of Utah, USA
DENNIS GANNON, Indiana University, USA
JACK DONGARRA, University of Tennessee, USA, Oak Ridge National Laboratory, USA, and University of Manchester, UK

The world of computing is in rapid transition, driven by the growth of smartphones, cloud services, and embedded devices, all while the future of semiconductors is in great flux due to the slowing of Moore's Law and increasing semiconductor foundry costs. Concomitantly, the future of advanced scientific computing (aka supercomputing or high-performance computing (HPC)), is at an important inflection point. For the last 60 years, the world's fastest computers were almost exclusively produced in the United States on behalf of scientific research in the national laboratories. Change is now in the wind. While costs now stretch the limits of U.S. government funding for advanced computing, Japan and China are now leaders in the bespoke HPC systems funded by government mandates. However, another, perhaps even deeper, fundamental change has occurred. The major cloud vendors have invested in global networks of massive scale systems that dwarf today's HPC systems. Driven by the computing demands of AI, these cloud systems are increasingly built using custom semiconductors, reducing the financial leverage of traditional computing vendors, while also reshaping how we think about the nature of scientific computation. Building the next generation of leading edge HPC systems will require rethinking many fundamentals and historical approaches by embracing end-to-end co-design; custom hardware configurations and packaging; large-scale prototyping, as was common thirty years ago; and collaborative partnerships with the dominant computing ecosystem companies. Universities, industry and governments ll need to reinvest in the basics of collaborative co-design, chiplet systems, quantum technology and new fabrication technologies. We need to reinvest in training the next generation of computer architects and collaborate on experimental prototypes.

CCS Concepts: • **Hardware**; • **Computer systems organization**; • **Social and professional topics → Computing / technology policy**;

Additional Key Words and Phrases: high performance computing, cloud computing, data centers, semiconductors

ACM Reference Format:
Daniel Reed, Dennis Gannon, and Jack Dongarra. 2022. HPC Forecast: Cloudy and Uncertain . 1, 1 (July 2022), 12 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Today, computing pervades all aspects of our society, in ways once imagined by only a few. Within science and engineering, computing has often been called the third paradigm, complementing theory and experiment, with big data and AI often called the fourth paradigm [15]. Spanning both data analysis and disciplinary and multidisciplinary modeling, scientific computing systems have,

Authors' addresses: Daniel Reed, University of Utah, Salt Lake City, Utah, USA, 84112, dan.reed@utah.edu; Dennis Gannon, Indiana University, Bloomington, Indiana, USA, 46202, dennis.gannon@outlook.com; Jack Dongarra, University of Tennessee, Knoxville, Tennessee, USA, 37996 and Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA, 37830 and University of Manchester, Manchester, UK, dongarra@icl.utk.edu.

XXXX-XXXX/2022/7-ART $15.00
https://doi.org/XXXXXXX.XXXXXXX

like their commercial counterparts, grown ever larger and more complex, and today's exascale scientific computing systems rival global scientific facilities in cost and complexity. However, not all is well, in the land of scientific computing.

In the initial decades of digital computing, government investments and the insights from designing and deploying supercomputers often shaped the next generation of mainstream and consumer computing products. Today, that economic and technological influence has increasingly shifted to smartphone and cloud service companies. Moreover, the end of Dennard scaling [3], slowdowns in Moore's Law, and the rising costs for continuing semiconductor advances, have made building ever-faster supercomputers more economically challenging and intellectually difficult.

As Figure 1 suggests, our thesis is that current approaches to designing and constructing leading edge high-performance computing (HPC) systems must change in deep and fundamental ways, embracing end-to-end co-design; custom hardware configurations and packaging; large-scale prototyping, as was common thirty years ago; and collaborative partnerships with the dominant computing ecosystem companies, smartphone and cloud computing vendors. We distinguish *leading edge HPC* – the very highest performing systems – from the broader mainstream of midrange HPC. For the later, market forces continue to shape the expansion of that market. Let's begin by examining how all of this has happened, then examining possible future directions for high-performance computing innovation and operations.

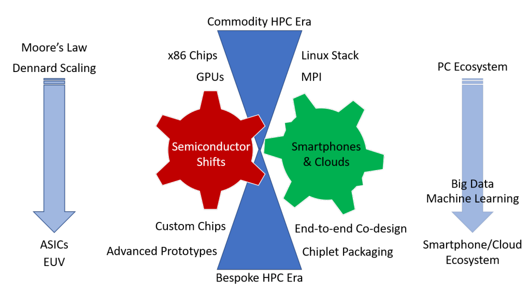

Fig. 1. Technical and Economic Forces Reshaping HPC

2 ECOSYSTEM SHIFTS

To understand the potential future of high performance computing, one must examine the fundamental shifts in computing technology. These shifts have occurred along two axes: the rise of massive scale commercial clouds and the economic and technological challenges associated with the evolution of semiconductor technology.

2.1 Cloud Innovations

Apple, Samsung, Google, Amazon, Microsoft, and the other cloud service companies are now major players the computing hardware and software ecosystem, both in scale and in technical approaches. Initially, these companies purchased standard servers and networking equipment for deployment in traditional collocation centers (colos). As scale increased, they began designing purpose-built data centers, optimized for power usage effectiveness (PUE), deployed at sites selected via multifactor optimization – inexpensive energy availability, tax incentives and political subsidies, political and geological stability, network access, and customer demand.

As cloud scale, complexity, and operational experience continued to grow, additional optimization and leverage opportunities emerged, including software defined networking, protocol offloads, and custom network architectures (greatly reducing dependence on traditional network hardware vendors) [7]; quantitative analysis of processor [16], memory [39, 47], network [8, 11] and disk failure modes [34, 38], with consequent redesign for reliability and lower cost (dictating specifications to vendors via consortia like Open Compute [32]); custom processor SKUs, custom accelerators (FPGAs and ASICs), and finally, complete processor design (e.g., Apple silicon, Google TPUs [20] and AWS Gravitons). In between, the cloud vendors deployed their own global fiber networks.

This virtuous cycle of insatiable consumer demand for rich services, business outsourcing to the cloud, expanding data center capacity, and infrastructure cost optimization has had several effects. Most importantly, it has dramatically lessened – and in many cases totally eliminated – their dependence on traditional computing vendors. One need look no further than cloud service provider and smartphone vendor market capitalizations, each near or in excess of $1T, to see the dramatic shifts in influence and scale. Put another way, the locus of innovation and influence has shifted from chip vendors and system integrators to cloud service providers.

2.2 Semiconductor Evolution

Historically, the most reliable engine of performance gains has been the steady rhythm of semiconductor advances – smaller, faster transistors and larger, higher performance chips. However, as chip feature sizes have approached 5 nanometers and Dennard scaling ended [3], the cadence of new technology generations has slowed, even as semiconductor foundry costs have continued to rise. With the shift to extreme ultraviolet (EUV) lithography [4] and gate-all-around FETs [5], the "minimax problem" of maximizing chip yields, minimizing manufacturing costs, and maximizing chip performance has grown increasingly complex for all computing domains, including HPC.

Chiplets [1, 27, 30] have emerged as a way to address these issues, while also integrating multiple functions in a single package. Rather than fabricating a monolithic system-on-a-chip (SoC), chiplet technology combines multiple chips, each representing a portion of the desired functionality, possibly fabricated using different processes by different vendors and including IP from multiple sources. Chiplet designs are part of the most recent offerings from Intel and AMD, where the latter's EPYC and Ryzen processors have delivered industry-leading performance via chiplet integration [30]. Similarly, Amazon's Graviton3 uses a chiplet design with seven different chip dies.

3 AN HPC CHECKPOINT

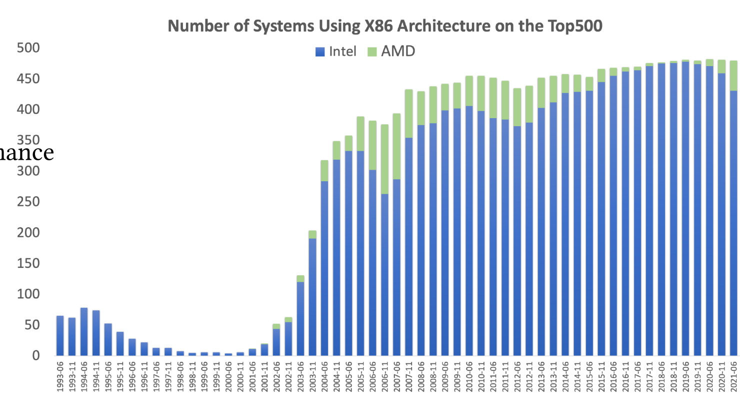

Fig. 2. Systems Using the x86-64 Architecture on the TOP500 [43]

Given the rise of cloud services and increasing constraints on commodity chip performance increases, it is useful to examine the current state of high-performance computing (HPC) and how the HPC ecosystem evolved to reach its current structure. From the 1970s to the 1990s, HPC experienced a remarkably active period of architectural creativity and exploration. In the late 1970s, the Cray series of machines [36] introduced vector processing. Companies like Denelcor and Tera then explored highly multi-threaded parallelism via custom processor design. Universities and companies were also active in exploring new shared memory designs (e.g., NYU Ultracomputer [10], Illinois Cedar [23], Stanford DASH [25], and BBN Butterfly [24]).

Finally, distributed memory, massively parallel computer designs (e.g., the Caltech Cosmic Cube [41], Intel iPSC/2 [18], and Beowulf clusters [42]) established a pattern for hyperscaled performance growth. Riding Moore's law, the ever-increasing performance of standard microprocessors, together with the cost advantage of volume production, led to the demise of most bespoke HPC systems, a shift often termed the "Attack of the Killer Micros" [31]. What followed was academic and industry standardization based on x86-64 processors (see Figure 2) and predominantly gigabit Ethernet and Infiniband networks, the Linux operating system, and message passing via the MPI standard.

By 2000, architectural innovation was limited to node accelerators (e.g., the addition of GPUs), high-bandwidth memory, and small network improvements. This monoculture of processor, operating system, and network have become standard interfaces that now define the market boundaries for innovation. At one time, dozens of high performance computing companies offered competing products. Today, only a few companies build HPC systems at the largest scales; see Figure 3.

While incremental performance improvements continue, with new x86-64 processors and GPU accelerators, basic innovation at the architectural level for supercomputers has been largely lost. However, in the last two years, sparks of architectural creativity are again re-emerging, driven by the needs to accelerate AI deep learning. Hardware startups, including Graphcore [19], Groq [2], and Cerebras [13] are exploring new architectural avenues. Concurrently, the major cloud service and smartphone providers have also developed custom processor SKUs, custom accelerators (FPGAs and ASICs), and finally, complete processor designs (e.g., Apple A15 SoCs, Google TPUs [20] and AWS Gravitons).

Against this HPC backdrop, the larger computing ecosystem itself is in flux:

- Dennard scaling [3] has ended and continued performance advances increasingly depend on functional specialization via custom ASICs and chiplet-integrated packages.
- Moore's Law is also at or near an end, and transistor costs are likely to increase as feature sizes continue to decrease.
- Advanced computing of all kinds, including high-performance computing, requires ongoing non-recurring engineering (NRE) investment (.i.e., endothermic) to develop new technologies and systems.
- The smartphone and cloud services companies are cash rich (i.e., exothermic), and they are designing, building, and deploying their own hardware and software infrastructure at unprecedented scale.
- AI is fueling a revolution in how both businesses and researchers think about problems and their computational solution.
- Talent is following the money and the intellectual opportunities, which are increasingly in a small number of very large companies or creative startups.

With this backdrop, what is the future of computing? Some of it is obvious, given the current dominance of smartphone vendors and cloud service providers. However, it seems likely that continued innovation in advanced high-performance computing will require rethinking some of our traditional approaches and assumptions, including how, where, and when academia, government laboratories, and companies spend finite resources and how we expand the global talent base.

4 LEADING EDGE HPC FUTURES

It now seems self-evident that supercomputing, at least at the highest levels, is endothermic, requiring regular infusions of non-revenue capital to fund the non-recurring engineering (NRE) costs to develop and deploy new technologies and successive generations of integrated systems. In turn, that capital can come from either other, more profitable divisions of a business or from external sources (e.g., government investment). Although most basic research is conducted in universities, several large companies (e.g., IBM, Microsoft, and Google) still conduct long-term basic research in addition to applied research and development.

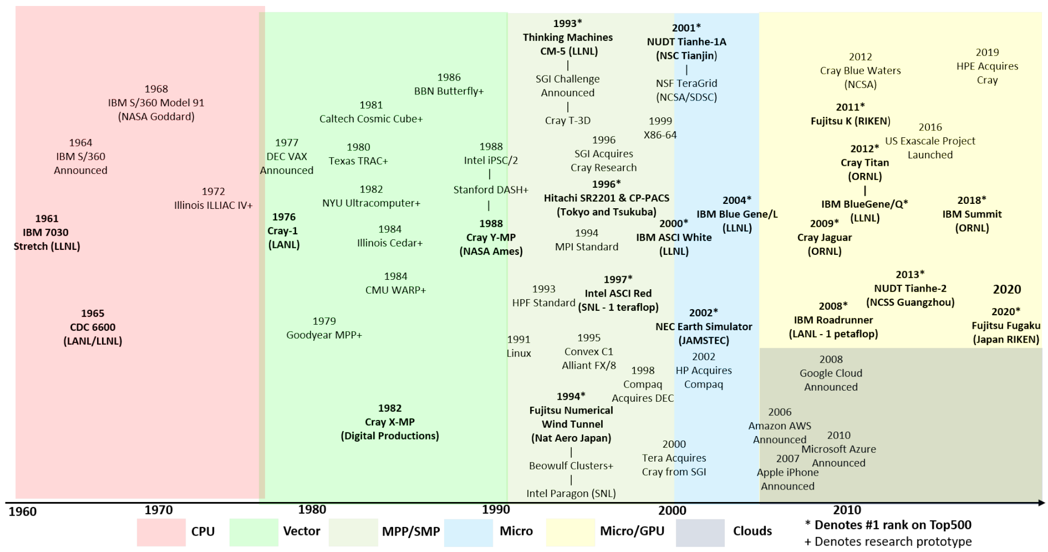

Fig. 3. Timeline of Advanced Computing

Cloud service companies now offer a variety of HPC clusters, of varying size, performance, and price. Given this, one might ask why cloud service companies are not investing even more deeply in the HPC market? Any business leader must always look at the opportunity cost (i.e., the time constant, the talent commitment, and cost of money) for any NRE investments and the expected return on investments. The core business question is always how to make the most money with the money one has, absent some other marketing or cultural reason to spend money on loss leaders, bragging rights, or political positioning. The key phrase here is“the most money;” simply being profitable is not enough, which is why HPC is rarely viewed as a primary business opportunity.

The NRE costs for leading edge supercomputing are now quite large relative to the revenues and market capitalization of those entities we call “computer companies,” and they are increasingly out of reach for most government agencies, at least under current funding envelopes. The days are long past when a few million dollars could buy a Cray-1/X-MP/Y-MP/2 or a commodity cluster and the resulting system would land the top ten of the TOP500 list. Today, hundreds of millions of dollars are needed to deploy a machine near the top of the TOP500 list, and at least similar, if not larger, investments in NRE are needed. In addition, the energy and cooling costs for operating such systems are now substantial and continuing to rise. What does this brave new world mean for leading edge HPC? We believe **five maxims** must guide future HPC government and private sector research and development strategies, for all countries.

Maxim One: *Semiconductor constraints dictate new approaches.* The “free lunch” of lower cost, higher performance transistors via Dennard scaling [3] and faster processors via Moore’s Law is at an end. Moreover, the *de facto* assumption that integrating more devices onto a single chip is always the best way to lower costs and maximize performance no longer holds. Individual transistor costs are now flat to rising as feature sizes approach one nanometer, due to the interplay of chip yields on 300nm wafers and increasing fabrication facility costs. Today, the investment needed to build state of the art facilities is denominated in billions of dollars per facility.

As recent geopolitical events have shown, there are substantial social, political, economic, and national security risks for any country or region that lacks a robust silicon fabrication ecosystem.

Fabless semiconductor firms rightly focus on design and innovation, but manufacturing those designs depends on reliable access to state of the art fabrication facilities, as the ongoing global semiconductor shortage has shown. In the U.S., the the U.S. CHIPS Act [28] and its successors, which would provide government support, are topics of intense political debate, with similar conversations underway in the European Union. Finally, Intel, TSMC, and GlobalFoundries recently announced plans to build new chip fabrication facilities in the U.S., each for different reasons.

Optimization must balance chip fabrication facility costs, now near $10B at the leading edge, chip yield per wafer, and chip performance. This optimization process has rekindled interest in packaging multiple chips, often fabricated with distinct processes and feature sizes. Such chiplets [27, 30] are more than a way to mix capabilities from multiple sources, they are an economic and engineering reaction to the interplay of chip defect rates, the cadence of feature size reductions, and semiconductor manufacturing costs. However, this approach requires academic, government, and industry collaborations to establish interoperability standards (e.g., the Open Domain-Specific Architecture (OSDA) project [45] within the Open Compute Project [32] and the Universal Chiplet Interconnect Express (UCIe) [1] standard). *Open chiplet standards can allow the best ideas from multiple sources be integrated effectively, in innovative ways, to develop next-generation HPC architectures.*

Maxim Two: *End-to-end hardware/software co-design is essential.* Leveraging the commodity semiconductor ecosystem has led to an HPC monoculture, dominated by x86-64 processors and GPU accelerators. Given current semiconductor constraints, substantially increased system performance will require more intentional end-to-end co-design [29], from device physics to applications. China and Japan are developing HPC systems outside of the conventional path, as seen by the Top500. The supercomputer Fugaku[37] (Post-K Computer), which was developed jointly by RIKEN and Fujitsu Limited, based on Arm technology with vector instructions, has taken the top spot on the Top500 List, a ranking of the world's fastest supercomputers. It also swept the other rankings of supercomputer performance (i.e., HPCG, HPL-AI, and Graph500). The supercomputer Fugaku is designed for versatile use based on a co-design approach between an application team and a system development. Similarly, the Chinese government, academic community, and domestic HPC vendors have made great efforts in the last few years to build a mature, self-designed software ecosystem and promote the possibility of running large and complex HPC applications on large, domestically produced supercomputers. It has been reported that China has two exaflops systems (OceanLight and Tianhe-3); several Gordon Bell prize submissions were run on OceanLight [26].

Similar application driven co-designs were evident in the recent batch of AI hardware startups mentioned above, as well as the cloud vendor accelerators. Such co-design means more than encouraging tweaks of existing products or product plans. Rather, it means looking holistically at the problem space, then envisioning, designing, testing, and fabricating appropriate solutions. *In addition to deep partnership with hardware vendors and cloud ecosystem operators, end-to-end co-design will require substantially expanded government investment in basic research and development, unconstrained by forced deployment timelines.* In addition to partnerships with x86-64 vendors, the ARM license model and the open source RISC-V [12] specification offer intriguing possibilities.

Maxim Three: *Prototyping at scale is required to test new ideas.* Semiconductors, chiplets, AI hardware, cloud innovations – the computing system is now in great flux, and not for the first time. As Figure 3 shows, the 1980s and 1990s were filled with innovative computing research projects and companies, many aided by government funding, that built novel hardware, new programming tools, and system software at large scale. *To escape the current HPC monoculture and build systems better suited to current and emerging scientific workloads at the leading edge, we must build real hardware and software prototypes at scale, not just incremental ones, but ones that truly test new ideas using custom silicon and associated software.* Implicitly, this means accepting the risk of failure, including at substantial scale, drawing insights from the failure, and building lessons based on those insights. A prototyping project that must succeed is not a research project; it is a product development.

Building such prototypes, whether in industry, national laboratories, or academia, depends on recruiting and sustaining integrated research teams — chip designers, packaging engineers, system software developers, programming environment developers, and application domain experts — in an integrated, end-to-end way. Such opportunities make it intellectually attractive to work on science and engineering problems, particularly given industry partnerships and opportunities to translate research ideas into practice. Implicit in such teams is coordinated funding for workforce development, basic research, and the applied R&D needed to develop and test prototype systems.

Maxim Four: *The space of leading edge HPC applications is far broader now than in the past.* Leading edge HPC originated in domains dominated by complex optimization problems and solution of time-dependent partial differential equations on complex meshes. Those domains will always matter, but other areas of advanced computing in science and engineering are of high and growing importance. As an example, the *Science 2021 Breakthrough of the Year* [40] was for AI-enabled protein structure prediction [21], with transformative implications for biology and biomedicine.

Even in traditional HPC domains, the use of AI for data set reduction and reconstruction and for PDE solver acceleration, is transforming computational modeling and simulation. The deep learning methods developed by the cloud companies are changing the course of computational science, and university collaborations are growing. The University of Washington, with help from Microsoft Azure on protein-protein interaction [17], is part of a bioscience revolution. In other areas, OpenAI is showing that deep learning can solve challenging Math Olympiad problems. In astrophysics, deep learning is being used to classify galaxies [22], generative adversarial networks (GANs) [9] have been used to understand groundwater flow in superfund sites [46], and deep neural networks have been trained to help design non-photonic structures [33]. This past year, the flagship conference of supercomputing (SC2021) had over 20 papers on neural networks in its highly selective program. *The HPC ecosystem is expanding and engaging new domains and approaches in deep learning, creating new and common ground with cloud service providers.*

Maxim Five. *Cloud economics have changed the supply chain ecosystem.* The largest HPC systems are now dwarfed by the scale of commercial cloud infrastructure and social media company deployments. A $500M supercomputer acquisition every five years provides limited financial leverage relative to the billions of dollars spent each year by cloud vendors. Driven by market economics, computing hardware and software vendors, themselves small relative to the large cloud vendors, now respond most directly to cloud vendor needs.

In turn, government investment (e.g., the U.S. Department of Energy (DOE) Exascale DesignForward, FastForward, and PathForward programs [44], and the European Union's HPC-Europa3 [6]) are small compared to the scale of commercial cloud investments and their leverage with those same vendors. For example, HPC-Europa3, funded under the EU's Eighth Framework Programme, better known as Horizon 2020, has a budget of only €9.2M [6]. Similarly, the U.S. DOE's multiyear investment of $400M via the FastForward, DesignForward, and PathForward programs as part of the Exascale Computing Project (ECP) targeted reduced power consumption, resilience, improved network and system integration. The DOE only supplied approximately $100M in NRE for each of the exascale systems under construction, while the cloud companies invested billions. Market research [35] suggests that China, Japan, the United States, and the European Union may each procure 1-2 exascale class systems per year, each estimated at approximately $400M.

The financial implications are clear. The government and academic HPC communities have limited leverage and cannot influence vendors in the same ways they did in the past. *New, collaborative models of partnership and funding are needed that recognize and embrace ecosystem changes and their implications, both in use of cloud services and collaborative development of new system* *architectures.* The cloud is evolving as a platform where specialized services such as attached quantum processors, specialized deep learning accelerators and high-performance graph database servers, can be configured and integrated into a variety of scientific workflows. However, that is not the whole HPC story. Massive scale simulations require irregular sparse data structures and the best algorithms are extremely inefficient on the current generation of supercomputers. The commercial cloud is part of the future of HPC, but it is by no means all. New architecture research and advanced prototyping are also needed.

As we have emphasized, the market capitalizations of the smartphone and cloud services vendors now dominate the computing ecosystem, and the overlap between commercial AI application hardware needs and those of scientific and engineering computing are creating new opportunities.. We realize this may be heretical to some, but there are times and places where commercial cloud services can be the best option to support scientific and engineering computing needs.

The performance gaps between cloud services and HPC gaps have lessened substantially over the past decade, as shown by a recent comparative analysis [14]. Moreover, HPC as a service is now both real and effective, both because of its performance and the rich and rapidly expanding set of hardware capabilities and software services. The latter is especially important; cloud services offer some features not readily available in the HPC software ecosystem.

Some in academia and national laboratory community will immediately say, "But, we can do it cheaper, and our systems are bigger!" Perhaps, but those may not be the appropriate perspectives. Proving such claims means being dispassionate about technological innovation, NRE investments, and opportunity costs. In turn, this requires a mix of economic and cultural realism in making build versus use decisions and taking an expansive view of the application space, unique hardware capabilities, and software tools. Opportunity costs are real, though not often quantified in academia or government. Today, capacity computing (i.e., solving an ensemble of smaller problems) can easily be satisfied with a cloud-based solution, and on-demand, episodic computing of both capacity and large-scale scientific computing can benefit from cloud scaling.

5 CONCLUSIONS

The computing ecosystem is in enormous flux, creating both opportunities and challenges for the future of advanced scientific computing. For the past twenty years, the most reliable engine of HPC performance gains has been the steady improvement in commodity CPU technology due to semiconductor advances. But with the slowing of Moore's Law and the end of Dennard scaling, improved performance of supercomputers has increasingly relied on larger scale (i.e., building systems with more computing elements) and GPU co-processing. Concurrently, the computing ecosystem has shifted, with the rise of hyperscale cloud vendors who are themselves developing new hardware and software technologies.

Looking forward, it seems increasingly unlikely that future high-end HPC systems will be procured and assembled solely by commercial integrators from only commodity components. Rather, future advances will require embracing end-to-end design, testing and evaluating advanced prototypes, and partnering strategically with both traditional chip and HPC vendors but also with the new cloud ecosystem vendors. These are likely to involve (a) collaborative partnerships among academia, government laboratories, chip vendors, and cloud providers, (b) increasingly bespoke systems, designed and built collaboratively to support key scientific and engineering workload needs, or (c) a combination of these two.

Put another way, in contrast to midrange systems, leading edge, HPC systems are increasingly similar to large-scale scientific instruments (e.g., the Vera Rubin Observatory, the LIGO gravity wave detector, or the Large Hadron Collider), with limited economic incentives for commercial development. Each contains commercially designed and constructed technology, but each also contains large numbers of custom elements for which there is no sustainable business model. Instead, we build these instruments because we want them to explore open scientific questions, and we recognize that their design and construction requires both government investment and innovative private sector partnerships.

Like many other large-scale scientific instruments, where international collaborations are an increasingly common way to share costs and facilitate research collaborations, leading edge computing would benefit from increased international partnerships, recognizing that national security and economic competitiveness issues will necessarily limit sharing of certain ”dual use” technologies. *Subject to those very real constraints, we believe greater government investment in semiconductor futures – both basic research and foundry construction – along with an integrated, long-term research and development program that funds academic, national laboratory, and private sector partnerships to design, develop, and test advanced computing prototypes will be needed if we are to build more performant leading edge high-performance computing systems. These investments must be tens, perhaps hundreds of billions of dollars, in scale.*

We have long relied on the commercial market for the building blocks of leading edge HPC systems. Although this has leveraged commodity economics, it has also resulted in systems ill-matched to the algorithmic needs of scientific and engineering applications. With the end of Moore's Law, we now have both the opportunity and the pressing need to invest in first principles design.

Investing in the future is never easy, but it is critical if we are to continue to develop and deploy new generations of high-performance computing systems, ones that leverage economic shifts, commercial practices, and emerging technologies to advance scientific discovery. Intel's Andrew Grove was right when he said "only the paranoid survive", but paranoia alone is not enough – successful competitors also need substantial financial resources and a commitment to technological opportunities and scientific innovation.

ACKNOWLEDGMENTS

Jack Dongarra was supported in part by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

The authors are grateful for thoughtful feedback on earlier drafts of this article from the anonymous reviewers, as well as Doug Burger (Microsoft), Tony Hey (UK Science and Technology Facilities Council), and John Shalf (Lawrence Berkeley National Laboratory).

REFERENCES

[1] 2022. Universal Chiplet Interconnect Express (UCIe). https://www.uciexpress.org.

[2] Dennis Abts, Jonathan Ross, Jonathan Sparling, Mark Wong-VanHaren, Max Baker, Tom Hawkins, Andrew Bell, John Thompson, Temesghen Kahsai, Garrin Kimmell, Jennifer Hwang, Rebekah Leslie-Hurd, Michael Bye, E. R. Creswick, Matthew Boyd, Mahitha Venigalla, Evan Laforge, Jon Purdy, Purushotham Kamath, Dinesh Maheshwari, Michael Beidler, Geert Rosseel, Omar Ahmad, Gleb Gagarin, Richard Czekalski, Ashay Rane, Sahil Parmar, Jeff Werner, Jim Sproch, Adrian Macias, and Brian Kurtz. 2020. Think Fast: A Tensor Streaming Processor (TSP) for Accelerating Deep Learning Workloads. In *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture* (Virtual Event) *(ISCA '20)*. IEEE Press, 145–158. https://doi.org/10.1109/ISCA45697.2020.00023

[3] Mark Bohr. 2007. A 30 Year Retrospective on Dennard's MOSFET Scaling Paper. *IEEE Solid-State Circuits Society Newsletter* 12, 1 (2007), 11–13. https://doi.org/10.1109/N-SSC.2007.4785534

[4] Yao-Wen Chang, Ru-Gun Liu, and Shao-Yun Fang. 2015. EUV and E-Beam Manufacturability: Challenges and Solutions. In *Proceedings of the 52nd Annual Design Automation Conference* (San Francisco, California) *(DAC '15)*. Association for Computing Machinery, New York, NY, USA, Article 198, 6 pages. https://doi.org/10.1145/2744769.2747925

[5] S. Dey, T. P. Dash, E. Mohapatra, J. Jena, S. Das, and C. K. Maiti. 2019. Performance and Opportunities of Gate-All-Around Vertically-Stacked Nanowire Transistors at 3nm Technology Nodes. In *2019 Devices for Integrated Circuit (DevIC)*. 94–98. https://doi.org/10.1109/DEVIC.2019.8783385

[6] European Commission. 2022. Transnational Access Programme for a Pan-European Network of HPC Research Infrastructures and Laboratories for scientific computing. https://cordis.europa.eu/project/id/730897.

[7] Daniel Firestone et al. 2018. Azure Accelerated Networking: SmartNICs in the Public Cloud. https://www.usenix.org/system/files/conference/nsdi18/nsdi18-firestone.pdf. In *15th USENIX Symposium on Networked Systems Design and Implementation*, Vol. 15. Accessed: 8-Feb-2022.

[8] Phillipa Gill, Navendu Jain, and Nachiappan Nagappan. 2011. Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications. In *Proceedings of the ACM SIGCOMM 2011 Conference* (Toronto, Ontario, Canada) *(SIGCOMM '11)*. Association for Computing Machinery, New York, NY, USA, 350–361. https://doi.org/10.1145/2018436.2018477

[9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf

[10] Allan Gottlieb, Ralph Grishman, Clyde P. Kruskal, Kevin P. McAuliffe, Larry Rudolph, and Marc Snir. 1982. The NYU Ultracomputer—Designing a MIMD, Shared-Memory Parallel Machine (Extended Abstract). In *Proceedings of the 9th Annual Symposium on Computer Architecture* (Austin, Texas, USA) *(ISCA '82)*. IEEE Computer Society Press, Washington, DC, USA, 27–42.

[11] Ramesh Govindan, Ina Minei, Mahesh Kallahalla, Bikash Koley, and Amin Vahdat. 2016. Evolve or Die: High-Availability Design Principles Drawn from Google's Network Infrastructure. In *Proceedings of the 2016 ACM SIGCOMM Conference* (Florianopolis, Brazil) *(SIGCOMM '16)*. Association for Computing Machinery, New York, NY, USA, 58–72. https://doi.org/10.1145/2934872.2934891

[12] Samuel Greengard. 2020. Will RISC-V Revolutionize Computing? *Commun. ACM* 63, 5 (apr 2020), 30–32. https://doi.org/10.1145/3386377

[13] Patrick Groeneveld. 2020. Wafer Scale Interconnect and Pathfinding for Machine Learning Hardware (Invited). In *Proceedings of the Workshop on System-Level Interconnect: Problems and Pathfinding Workshop* (San Diego, California) *(SLIP '20)*. Association for Computing Machinery, New York, NY, USA, Article 7, 1 pages. https://doi.org/10.1145/3414622.3432992

[14] Giulia Guidi, Marquita Ellis, Aydin Buluç, Katherine Yelick, and David Culler. 2021. *10 Years Later: Cloud Computing is Closing the Performance Gap*. Association for Computing Machinery, New York, NY, USA, 41–48. https://doi.org/10.1145/3447545.3451183

[15] Tony Hey, Stewart Tansley, and Kristin Tolle (Eds.). 2009. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Microsoft Research, Redmond, Washington. http://research.microsoft.com/en-us/collaboration/fourthparadigm/

[16] Peter H. Hochschild, Paul Turner, Jeffrey C. Mogul, Rama Govindaraju, Parthasarathy Ranganathan, David E. Culler, and Amin Vahdat. 2021. Cores That Don't Count. In *Proceedings of the Workshop on Hot Topics in Operating Systems* (Ann Arbor, Michigan) *(HotOS '21)*. Association for Computing Machinery, New York, NY, USA, 9–16. https://doi.org/10.1145/3458336.3465297

[17] Eric Horvitz. 2022. A Leap Forward in Bioscience. https://erichorvitz.com/Leap_forward_bioscience.htm. Accessed: 9-Feb-2022.

[18] Intel. 1988. The Intel IPSC/2 System: The Concurrent Supercomputer for Production Applications. In *Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications: Architecture, Software, Computer Systems, and General Issues - Volume 1* (Pasadena, California, USA) *(C³P)*. Association for Computing Machinery, New York, NY, USA, 843–846. https://doi.org/10.1145/62297.62412

[19] Zhe Jia, Blake Tillman, Marco Maggioni, and Daniele Paolo Scarpazza. 2019. Dissecting the Graphcore IPU Architecture via Microbenchmarking. *CoRR* abs/1912.03413 (2019). arXiv:1912.03413 http://arxiv.org/abs/1912.03413

[20] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. 2017. In-Datacenter Performance Analysis of a Tensor Processing Unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture* (Toronto, ON, Canada) *(ISCA '17)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3079856.3080246

[21] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* (Aug. 2021), 1476–4687. https://doi.org/10.1038/s41586-021-03819-2

[22] Edward J. Kim and Robert J. Brunner. 2016. Star-galaxy Classification Using Deep Convolutional Neural Networks. *Monthly Notices of the Royal Astronomical Society* 464, 4 (2016).

[23] David J. Kuck, Edward S. Davidson, Duncan H. Lawrie, and Ahmed H. Sameh. 1986. Parallel Supercomputing Today and the Cedar Approach. *Science* 231, 4741 (1986), 967–974. https://doi.org/10.1126/science.231.4741.967 arXiv:https://www.science.org/doi/pdf/10.1126/science.231.4741.967

[24] Thomas J. LeBlanc, Michael L. Scott, and Christopher M. Brown. 1988. Large-Scale Parallel Programming: Experience with BBN Butterfly Parallel Processor. *SIGPLAN Not.* 23, 9 (jan 1988), 161–172. https://doi.org/10.1145/62116.62131

[25] Daniel Lenoski, James Laudon, Kourosh Gharachorloo, Wolf-Dietrich Weber, Anoop Gupta, John Hennessy, Mark Horowitz, and Monica S. Lam. 1992. The Stanford Dash Multiprocessor. *Computer* 25, 3 (mar 1992), 63–79. https://doi.org/10.1109/2.121510

[26] Yong (Alexander) Liu, Xin (Lucy) Liu, Fang (Nancy) Li, Haohuan Fu, Yuling Yang, Jiawei Song, Pengpeng Zhao, Zhen Wang, Dajia Peng, Huarong Chen, Chu Guo, Heliang Huang, Wenzhao Wu, and Dexun Chen. 2021. Closing the "Quantum Supremacy" Gap: Achieving Real-Time Simulation of a Random Quantum Circuit Using a New Sunway Supercomputer. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (St. Louis, Missouri) *(SC '21)*. Association for Computing Machinery, New York, NY, USA, Article 3, 12 pages. https://doi.org/10.1145/3458817.3487399

[27] Gabriel H. Loh, Samuel Naffziger, and Kevin Lepak. 2021. Understanding Chiplets Today to Anticipate Future Integration Opportunities and Limits. In *2021 Design, Automation Test in Europe Conference Exhibition (DATE)*. 142–145. https://doi.org/10.23919/DATE51398.2021.9474021

[28] Rep. Michael T. McCaul. 2020. H.R.7178 – CHIPS for America Act. https://www.congress.gov/bill/116th-congress/house-bill/7178.

[29] Cherry Murray, Supratik Guha, Dan Reed, Gil Herrera, Kerstin Kleese van Dam, Sayeef Salahuddin, James Ang, Thomas Conte, Debdeep Jena, Robert Kaplar, Harry Atwater, Rick Stevens, Dushan Boroyevich, William Chappell, Tsu-Jae King Liu, Justin Rattner, Michael Witherell, Khurram Afridi, Simon Ang, Jon Bock, Srabanti Chowdhury, Suman Datta, Keith Evans, Jack Flicker, Mark Hollis, Noble Johnson, Ken Jones, Peter Kogge, Sriram Krishnamoorthy, Matthew Marinella, Todd Monson, Sreekant Narumanchi, Paul Ohodnicki, Ramamoorthy Ramesh, Michael Schuette, John Shalf, Shadi Shahedipour-Sandvik, Jerry Simmons, Valerie Taylor, Tom Theis, Eric Colby, Robinson Pino, Andy Schwartz, Katie Runkles, Joseph Harmon, Michele Nelson, and Vicki Skonicki. 2018. Basic Research Needs for Microelectronics: Report of the Office of Science Workshop on Basic Research Needs for Microelectronics, October 23 – 25, 2018. (10 2018). https://doi.org/10.2172/1616249

[30] Samuel Naffziger, Noah Beck, Thomas Burd, Kevin Lepak, Gabriel H. Loh, Mahesh Subramony, and Sean White. 2021. Pioneering Chiplet Technology and Design for the AMD EPYC™ andRyzen™ Processor Families : Industrial Product. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. 57–70. https://doi.org/10.1109/ISCA52012.2021.00014

[31] New York Times. 1991. The Attack of the Killer Micros. https://www.nytimes.com/1991/05/06/business/the-attack-of-the-killer-micros.html. Accessed = 1991-05-06.

[32] Open Compute Project. 2022. Open Possibilities. https://www.opencompute.org/. Accessed: 8-Feb-2022.

[33] Peurifoy et al. 2022. Nanophotonic particle simulation and inverse design using artificial neural networks. *Science Advances* 8, 5 (2022).

[34] Eduardo Pinheiro, Wolf-Dietrich Weber, and Luiz André Barroso. 2007. Failure Trends in a Large Disk Drive Population. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies* (San Jose, CA) *(FAST '07)*. USENIX Association, USA, 2.

[35] Hyperion Research. 2021. HPC Market Update Briefing During SC21. https://hyperionresearch.com/hpc-market-update-briefing-during-sc21/. Accessed = 2022-02-28.

[36] Richard M. Russell. 1978. The CRAY-1 Computer System. *Commun. ACM* 21, 1 (jan 1978), 63–72. https://doi.org/10.1145/359327.359336

[37] Mitsuhisa Sato, Yutaka Ishikawa, Hirofumi Tomita, Yuetsu Kodama, Tetsuya Odajima, Miwako Tsuji, Hisashi Yashiro, Masaki Aoki, Naoyuki Shida, Ikuo Miyoshi, Kouichi Hirai, Atsushi Furuya, Akira Asato, Kuniki Morita, and Toshiyuki Shimizu. 2020. Co-Design for A64FX Manycore Processor and "Fugaku". In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Atlanta, Georgia) *(SC '20)*. IEEE Press, Article 47,

15 pages.

[38] Bianca Schroeder and Garth A. Gibson. 2007. Understanding Disk Failure Rates: What Does an MTTF of 1,000,000 Hours Mean to You? *ACM Trans. Storage* 3, 3 (oct 2007), 8–es. https://doi.org/10.1145/1288783.1288785

[39] Bianca Schroeder, Eduardo Pinheiro, and Wolf-Dietrich Weber. 2011. DRAM Errors in the Wild: A Large-Scale Field Study. *Commun. ACM* 54, 2 (feb 2011), 100–107. https://doi.org/10.1145/1897816.1897844

[40] Science. 2021. Science's 2021 Breakthrough: AI-powered Protein Prediction. https://www.aaas.org/news/sciences-2021-breakthrough-ai-powered-protein-prediction.

[41] Charles L. Seitz. 1985. The Cosmic Cube. *Commun. ACM* 28, 1 (jan 1985), 22–33. https://doi.org/10.1145/2465.2467

[42] T. Sterling, D. Savarese, D. J. Becker, J. E. Dorband, U. A. Ranawake, and C. V. Packer. 1995. BEOWULF: A Parallel Workstation for Scientific Computation. In *Proceedings of the 24th International Conference on Parallel Processing*, Vol. I, Architecture. CRC Press, Boca Raton, FL, I:11–14.

[43] Erich Strohmaier, Hans W. Meuer, Jack Dongarra, and Horst D. Simon. 2015. The TOP500 List and Progress in High-Performance Computing. *Computer* 48, 11 (2015), 42–49. https://doi.org/10.1109/MC.2015.338

[44] U.S. Department of Energy. [n.d.]. Pathforward. https://www.exascaleproject.org/research-group/pathforward/. Accessed 2022-02-10.

[45] Bapiraju Vinnakota. 2021. The Open Domain-Specific Architecture: Next Steps to Production. In *Proceedings of the Eight Annual ACM International Conference on Nanoscale Computing and Communication* (Virtual Event, Italy) *(NANOCOM '21)*. Association for Computing Machinery, New York, NY, USA, Article 19, 5 pages. https://doi.org/10.1145/3477206.3477462

[46] Liu Yang, Sean Treichler, Thorsten Kurth, Keno Fischer, David Barajas-Solano, Josh Romero, Valentin Churavy, Alexandre Tartakovsky, Michael Houston, Prabhat, and George Karniadakis. 2021. Highly-scalable, physics-informed GANs for learning solutions of stochastic PDEs. *ArXiv* abs/1910.13444v1 (2021).

[47] Darko Zivanovic, Pouya Esmaili Dokht, Sergi Moré, Javier Bartolome, Paul M. Carpenter, Petar Radojković, and Eduard Ayguadé. 2019. DRAM Errors in the Field: A Statistical Approach. In *Proceedings of the International Symposium on Memory Systems* (Washington, District of Columbia, USA) *(MEMSYS '19)*. Association for Computing Machinery, New York, NY, USA, 69–84. https://doi.org/10.1145/3357526.3357558