EDITOR DEJAN MILOJICIC
Hewlett Packard Labs;
dejan.milojicic@hpe.com

More Real Than Real: The Race to Simulate Everything

Rosa M. Badia, Barcelona Supercomputing Center

Ian Foster, Argonne National Laboratory and University of Chicago

Dejan Milojicic, Hewlett Packard Labs

Extreme times require extreme measures. In this column, we discuss how high-performance computing embraces artificial intelligence and data analytics to address global challenges.

Humanity is confronted with unprecedented challenges, such as pandemics and global warming. Our future is at stake, and we require powerful tools to explore solutions to these challenges. High-performance computing (HPC) is needed more than ever to guide us through the many possible paths that are in front of us. Exascale computers are being deployed around the world, and new artificial intelligence (AI) and data analytics techniques are being adopted to complement and enhance traditional HPC techniques focused on number crunching.

Suddenly, using HPC, we are simulating and predicting many aspects of our lives: the spread of pandemics, climate change, gaming, the metaverse, digital twins, supply chains, and much more. We are almost in a race to simulate anything and everything, and that new reality is threatening to become even more real than our real life. Demands are becoming much higher; tolerance for latency to get results much smaller.

In its quest, HPC has many challenges, such as the end of Moore's law, the cost of moving data, energy limitations, and legacy programming models. Cloud computing offers some advantages but comes with its own challenges for HPC. It offers substantially higher (auto) scaling, modern development tools, and natural integration points with AI. However, today's cloud platforms cannot support large-scale, tightly coupled applications, due to jitter and unoptimized interconnects.

To further discuss opportunities and challenges of HPC in modern science, we have engaged two veterans from the HPC and parallel/distributed computing world: Rosa M. Badia, of the Barcelona Supercomputing Center

Digital Object Identifier 10.1109/MC.2022.3173359
Date of current version: 4 July 2022

(BSC), and Ian Foster, of Argonne National Laboratory and the University of Chicago. Well traveled across HPC and world locations, with broad and deep perspectives on HPC from an international point of view, Rosa and Ian will help us predict the future of HPC.

DEJAN MILOJICIC: What are the most challenging problems of HPC today?

ROSA M. BADIA: In terms of the infrastructure, in my opinion, the challenges are in coping with the new technologies for memory and in the heterogeneity of the new devices in the memory storage. While these new devices have the opportunity of solving the input–output (I/O) bandwidth challenge, they also introduce new problems. In the software area, in my opinion, the challenges come from complexity of the applications that try to leverage the heterogeneous infrastructure. New solutions to develop dynamic workflows and complex applications that involve huge simulations are necessary.

IAN FOSTER: Everything is changing all at once: hardware, software, and applications. And more things are happening at once within HPC systems, both in the sense of there being more processors (which brings its own challenges) and due to the coupling of more and different physics, new data-driven applications, AI agents, and online data analysis and reduction, among other factors. Add in growing demand from industry for precisely the skills that HPC specialists have spent so long developing, and you have a difficult situation.

MILOJICIC: How do you see regional competition in various races in HPC (petascale, exascale, and so on)? Are they helpful in advancing the technology?

BADIA: The European High Performance Computing Joint Undertaking (EuroHPC JU) aims to coordinate efforts and pool resources to make Europe a world leader in supercomputing. This will boost Europe's scientific excellence and industrial strength, and support the digital transformation of its economy while ensuring its technological sovereignty. The road map includes the deployment of two exascale systems, three pre-exascale systems (two already under construction) and five petascale systems (four already deployed and one under construction). I do not see this as much as a race but as each region having its own infrastructures and not needing to rely on the others.

FOSTER: The aggressive development of extreme-scale systems in multiple regions is great for science and engineering and great for technology. Furthermore, competition and diversity of ideas are clearly driving progress. However, I am also concerned that we are missing opportunities for collaboration. The vast scope of the problems to be solved means that every region should be looking carefully at what it must build itself versus what it can build in collaboration with, or borrow from, others. One wrinkle here, discussed by Moshe Vardi in a recent Communications of the ACM column,[2] is that HPC may be perceived as so strategic that useful cooperation is hindered. We'll see how that plays out.

MILOJICIC: Accelerators helped overcome and at least offset the slowdown if not the end of Moore's law. This is especially true for AI/machine learning (ML)/deep learning (DL). How important are they in HPC?

BADIA: If you look at the TOP500 list, most systems are equipped with GPUs. They have been very important to sustain the end of Moore's law, and have been very important for workloads that include some form of AI and, especially, DL. It is clear that accelerators are very important for HPC, but still there are challenges in the programmability of these devices and their integration in the overall system architecture.

FOSTER: When GPUs first appeared, I expected them soon to be subsumed by more usable extensions to conventional processors. But I was not accounting for DL and the resulting immense demand for linear algebra accelerators. It is ironic that those developments occurred in parallel with computational science moving beyond dense matrices. Now DL is rediscovering sparsity and adaptivity, so perhaps the next generation of accelerators will be more usable for sophisticated HPC algorithms. Or maybe we'll work out how to leverage accelerators better, for example, by rethinking algorithms and repurposing AI/ML methods?

MILOJICIC: Rumor has it that there are over 100 accelerator start-ups in the world and at least 60 in Silicon Valley. How do these 1,000 blooming flowers help or not help our community? For example, driving innovation versus increasing efforts into system integration and programming them due to nonstandard solutions.

BADIA: I think that at this moment, all these start-ups are driving innovation but also making more diverse the programming environments. At some moment, the number of alternatives will get reduced. The efforts for converging into standard solutions for programming have been there for a long time, but still there is work to do to offer high-level, simple interfaces empowered by toolchains to support performance portability.

FOSTER: I'm with Rosa that high-quality toolchains and standardized interfaces are the keys to broad impact. And while TensorFlow and PyTorch are suitable interfaces for DL, we need far more for HPC applications.

MILOJICIC: There is a lot of discussion about quantum computing and a lot of money being invested in this area. What is the take of the HPC community on quantum computing?

BADIA: BSC is also investing efforts in quantum computing. Currently, we are coordinating Quantum Spain, a project that includes the construction and commissioning of the first quantum computer in southern Europe. We are also involved in initiatives at the European level. In particular, my group is involved in the parallelization of a classical quantum computing simulator aiming to perform large-scale simulations of quantum systems, what is called *classical simulation*. Since quantum computers are neither cheap nor easy to build, classical simulation is a valuable method for efficient simulation of quantum algorithms. Classical simulation tools are a must to understand noise sources and improve the performance of quantum algorithms. These are hungry simulations both in terms of data and compute, able to fill an exascale supercomputer.

FOSTER: There is much exciting work underway in quantum computing, communications, and sensing. At Argonne and the University of Chicago, for example, there are efforts ranging from physical testbeds (for example, for quantum networking) to foundational work in algorithms and materials. However, I am not betting on quantum computing playing an important role in HPC any time soon. I suspect that view is widely shared within the community.

MILOJICIC: What is your take on the confluence of HPC with high-performance data analytics (HPDA) and AI?

BADIA: The community has identified the convergence of HPC, HPDA, and AI as critical for the new type of applications that involve the three aspects. I am personally the principal investigator of the eFlows4HPC EuroHPC JU project that aims to provide a software stack to enable the development of workflows that include HPC simulations or modeling together with AI and data analytics. In the project, we have use cases that base their workflows in simulations for manufacturing, climate prediction and modeling, and urgent computing for natural hazards (for instance, earthquakes and their subsequent tsunamis). We aim at providing tools to make the development of such workflows easier. At the core of the software stack, we find the environment developed by my team, PyCOMPSs, which provides the glue for integrating the HPC, HPDA, and AI components. The project also aims at making the deployment and execution of such workflows easier in HPC infrastructures through the new methodology of HPC workflows as a service, inspired in cloud practices.

FOSTER: It's so important and, indeed, overdue. I'll return later to the opportunities inherent in the integration of HPC and AI, but a key point is that this convergence greatly expands the number of people that can benefit from HPC technologies. And it has implications for just about every aspect of our computing infrastructure, not least in programming models and tools.

MILOJICIC: How important are novel interconnects in large-scale supercomputers? With the increasing adoption of AI and less tightly coupled code, how are interconnects emerging?

BADIA: Simulations based on the message passing interface (MPI) are still dominant in HPC. For this type of application, an efficient, low-latency, and high-bandwidth interconnect is a must. In addition, in my opinion, with larger data sets and more demanding workloads, AI will evolve to distributed computing also, where efficient interconnects will also be needed. While Infiniband is present in a large number of systems, specialized networks, such as the Tofu D interconnect in the Top1 Fugaku, have appeared.

FOSTER: What Rosa said.

MILOJICIC: Storage-class memories and high-bandwidth memories are finding their adoption in many areas, including in HPC. Where do you see additional needs in memories for HPC?

BADIA: As I answered in question 1, I think this is one of the challenges of HPC today. These new types of memories are needed due to the demands in terms of latency and bandwidth of the data-intensive applications, such as AI workloads or those that combine AI and HPC. However, at the same time, this new type of memory and the new memory hierarchy is a challenge.

FOSTER: These are really hard problems. More dynamic and data-intensive applications need places to stash increasing volumes of input, intermediate, and output data on widely varying timescales. New memory technologies certainly can help, but much work will be needed to integrate them efficiently with the full HPC stack, from applications to schedulers and beyond. Maybe we need new data abstractions, like policy-aware key value stores, to enable integration with more complex application flows. What might we do with such memories if they could also perform certain types of computation?

MILOJICIC: How is adoption of AI changing HPC? Will it enable new applications, new verticals? Broader adoption of HPC beyond science and core engineering?

FOSTER: The opportunities here are numerous and exciting. First, HPC is, of course, fundamental to modern AI. In particular, DL depends heavily on massive floating-point computation, large-scale linear algebra, high-speed communications, and parallel I/O—all things that the HPC community has been working on for years. (It's no coincidence that Jack Dongarra won the 2021 Turing Award.) So HPC has a new role. These new "HPC for AI" applications are proving to be a fascinating source of new problems that will drive many advances in algorithms, software, and hardware, for example, accelerators and reduced-precision arithmetic.

Meanwhile, "AI for HPC" is enabling a generational revolution in how HPC is employed to solve intractable problems. HPC had become almost dull as researchers worked increasingly to achieve only incremental improvements in model resolution or physics fidelity. Now we see researchers developing new AI methods that learn from computations, for example, to generate fast surrogate climate models, construct machine-learned force fields in computational chemistry, or choose the next region to explore in a molecular dynamics simulation. These new methods, when combined with extreme-scale HPC, can deliver results that are transformative rather than incremental. We're still in the early stages of exploring these opportunities. Not everything will work as expected, but there will also be unanticipated discoveries. For example, what will happen when exascale computers are used to train foundation models on data from large numbers of simulations and experiments?

BADIA: I agree with all that Ian said. In fact, in my group, we are interested in providing solutions for applications or workflows that combine, at the same time, the traditional HPC simulation and modeling with AI models. Some of the use cases of eFlows4HPC are leveraging this idea, for example, by performing AI-driven pruning of ensemble members in a large Earth System Models simulation experiment to better use the computational and storage resources and releasing computational resources accordingly or by applying reduced order modeling techniques to a large number of HPC simulation results to generate a digital twin that can be deployed in edge devices to be applied in manufacturing scenarios.

MILOJICIC: AI has increased adoption of Python and new software frameworks. Are they penetrating the HPC community?

BADIA: Yes, sure. Python is the dominating language for AI and data analytics software. In this sense, it is used in HPC systems for these types of workloads but also as programming environments for workloads that combine more traditional HPC with AI and HPDA. An example of this type of environment would be PyCOMPSs from my group, which supports the development of workflows that include MPI simulations and invocations to AI components.

FOSTER: I love the trailer for Geert Jan Bex's massive open online course[1] on Fortran for scientific computing, in which (at 1:30) the monster Fortran drags away a lifeless Python as its breakfast. But in practice, Python is being used at ever-larger scales, driven, in part, by the more dynamic and heterogeneous task-parallel applications that we have already discussed. Packages like Parsl, PyCOMPSs, Radical Pilot, Colmena, and DeepDriveMD make it easy to implement such applications, including those in which computational structures evolve over time, for example, by starting, monitoring, reconfiguring, and stopping subcomputations as a simulation proceeds. The subcomputations themselves are typically coded in low-level languages, but the flexibility of Python, if used appropriately, increases overall productivity.

MILOJICIC: Virtualization has helped in both developer productivity and provider efficiency. After virtual machines and containers, we now have serverless, also known as *function as a service (FaaS)*. Are they convenient for HPC, and what are their advantages versus opportunities for improvement?

BADIA: Containers have been largely adopted in HPC. However, the concept of serverless, or FaaS, requires some openness in the systems that are possible. For example, most HPC systems will not easily enable communication between the computing nodes and external servers. In this sense, then the FaaS is limited to the login nodes and, most of the time, under certain conditions.

FOSTER: Serverless is part of a move to a world in which tasks can flow readily to wherever is most accessible, convenient, or efficient. In such a computing continuum, HPC systems can serve as high-powered compute engines, enabling new applications, such as smart experimental facilities that engage large-scale, on-demand computing to drive decisions. We've also found serverless to be a useful abstraction within HPC systems. We've had success building both classes of applications (for example, via use of the funcX federated FaaS system), but it is certainly true, as Rosa notes, that aspects of current HPC architectures can get in the way.

MILOJICIC: Will there be a new HPC programming model introduced [for instance, in addition to MPI and partitioned global address space (PGAS)/multithreading]?

BADIA: BSC has been proposing a task-based programming model as an alternative to MPI for around 15 years now. BSC has fostered the adoption of tasks and tasks with dependencies in OpenMP as well as its use with accelerators. We have promoted portable solutions with simple, high-level interfaces. In addition to the multicore/accelerator solution that is OmpSs and its successor OmpSs-2, we have also been working on solutions for distributed computing (from the grid to the cloud and now in the continuous edge to cloud, or HPC). This distributed solution has been based on the COMPSs runtime, which recently has been reengineered from a centralized to a distributed design to cope with the needs of the computing continuum.

FOSTER: As Rosa says, there already is. Task-parallel programming, once a niche alternative to single program, multiple data (whether MPI or PGAS) suitable for a few specialized applications, is increasingly mainstream and, indeed, fundamental to important new applications. There remain barriers to the most effective use of this new model, some of which we've discussed. I expect that we'll be working for most of the next decade to make such applications truly first-class citizens on HPC platforms.

MILOJICIC: Do you expect the continued growth of commercial cloud providers ultimately to subsume conventional HPC? If so, what are the pros and cons of that happening?

BADIA: You can run some MPI workloads in clouds if the latency and number of communications is not very high. Other workloads that do not require this level of interconnection should be able to run well. I personally do not think that security is an issue here.

FOSTER: If by HPC we mean "computers with hardware and software that support large-scale, fine-grained, often data-intensive parallel computations," then the question is, "Will commercial cloud providers provide computers with such characteristics?" They certainly could, given their vast resources. Whether they will is largely a question of economics. High-end computation is a niche market that HPC centers support very well, so my expectation is that while we'll see increasing use for small-to-medium-scale science and engineering applications, clouds won't be a factor for extreme-scale computations in the immediate future.

If we take a broader view of HPC as "the computing infrastructure used to solve challenging problems in science and engineering," then there is a big, largely unexploited, opportunity, namely, to use the cloud to host "science services" that, for example, reduce data-sharing friction, enable collaborative analysis of large data sets, and train scientific foundation models. As an example, the Globus research data management service (https://globus.org), running on Amazon Web Services (AWS) for more than a decade, today supports more than 300,000 registered users at thousands of institutions. We can and should be building many more such services.

MILOJICIC: What are the key benefits of HPC's cloud deployment? Cost/affordability, DevOps (for example, toolchains), something else?

FOSTER: For HPC as access to specialized computers, I'd point to the elastic capacity as a key potential benefit of cloud deployment. So much of computational science demands rapid response, whether to align with human thinking processes or to meet external demands for decisions. And particularly for smaller computations, commercial cloud can provide more rapid response than conventional HPC. Various HPC centers are deploying on-demand Kubernetes clusters to provide similar capabilities; it remains to be seen whether such deployments can support growing on-demand workloads cost effectively.

For HPC as "scientific computing infrastructure," the incredible richness of cloud ecosystems—far exceeding anything that can reasonably be deployed and operated at an HPC center—can be a major benefit. Using Globus as an example once again, AWS support for geographical replication, scalable logging, elastic scaling, reliable storage, security controls, and many other features has made it possible for a relatively small team to construct and operate a service that could never have been built at an HPC center.

BADIA: I agree with what Ian has said, although with COMPSs, we are able to apply elasticity in Slurm-based clusters. I think HPC is also learning from these practices in clouds, like the use of containers to ease and enable automatic deployment. Adopting these methodologies will widen the use of HPC resources in new user communities.

MILOJICIC: Workflows are an increasingly important abstraction in HPC. As HPC reaches out of supercomputer centers and into the cloud, how are workflows evolving?

BADIA: Depending on the nature of the workflows, they can naturally run in clouds the same way that they can run in HPC systems. If the components of the workflow are not very demanding in terms of the interconnection network, they can run well in clouds. In case of the type of workflows I was describing earlier, some of the components of the workflow can be MPI simulation or modeling codes that require the interconnection network of an HPC system.

FOSTER: We've already talked about the growing importance of task-parallel computing, which engages some aspects of workflow, within HPC systems. Also of growing importance are workflows that link scientific instrumentation with HPC to cope with exploding data rates that far exceed human cognitive capacities. Recurring variants of this pattern include on-demand HPC for data analysis (for example, to reconstruct large data sets), HPC for training ML models that are then deployed at the instrument, and integration of experimental data with simulations to form digital twins of experimental processes.

MILOJICIC: Is there any other prediction in HPC you would like to make, Rosa and Ian?

BADIA: There is a huge potential of almost-here exascale systems that will enable us to tackle grand challenges in science and engineering. These infrastructures will be integrated in the so-called computing continuum, with instruments and sensors in the edge and Internet of Things devices. In this sense, HPC will evolve toward a more usable instrument to the general scientist through the adoption of cloud-inspired methodologies for deployment and execution.

FOSTER: I've already noted the profound transformation that I see happening in the nature and role of HPC. In the grand scheme of things, it isn't so long ago that HPC was a niche technology, important only for a few specialized applications. Now HPC technologies and applications are becoming quasi-universal. It seems likely that this trend will accelerate.

I also posit that AI and HPC will combine in more surprising ways. Specifically, I think that scientists and engineers are likely to learn, sooner rather than later, how to construct AI "foundation models" that capture the current state of understanding in specific domains (for example, the climate system, soft materials, and cellular biology). As such models emerge, we will find that simulations and experiments are increasingly driven by observations of internal inconsistencies, uncertainties, and gaps in the models' representations of the world (and associated data) rather than by leaps of human intuition. This move to model-driven simulations and experiments (and simulation- and experiment-driven models) will be immensely powerful, if disconcerting, for individual researchers and will place new demands on both HPC and experimental systems.

REFERENCES

1. G. J. Bex, *Trailer for the MOOC on Fortran for Scientific Programming*. (Jul. 7, 2021). Accessed: May 19, 2022. [Online Video]. https://www.youtube.com/watch?v=l6pEaUttWo8
2. M. Vardi, "War and tech (and ACM)," *Commun. ACM*, vol. 65, no. 5, p. 9, May 2022, doi: 10.1145/3528570.

ROSA M. BADIA is a group manager at the Barcelona Supercomputing Center, Barcelona, 08034, Spain. Contact her at rosa.m.badia@bsc.es.

IAN FOSTER is a director of the Data Science and Learning Division and distinguished fellow at Argonne National Laboratory, Lemont, 60439, Illinois, USA, and a professor of computer science at the University of Chicago, Chicago, Illinois, 60637, USA. Contact him at foster@anl.gov.

DEJAN MILOJICIC is a distinguished technologist at Hewlett Packard Labs, Palo Alto, California, 94306, USA. Contact him at dejan.milojicic@hpe.com.

Digital Object Identifier 10.1109/MC.2022.3178649