



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/164225/>

Version: Accepted Version

Proceedings Paper:

Djemame, K and Carr, H (2020) Exascale Computing Deployment Challenges. In: Djemame, K, Altmann, J, Bañares, JÁ, Agmon Ben-Yehuda, O, Stankovski, V and Tuffin, B, (eds.) Lecture Notes in Computer Science. GECON2020: 17th International Conference on the Economics of Grids, Clouds, Systems, and Services, 15-17 Sep 2020, Izola, Slovenia. Springer, Cham, Switzerland, pp. 211-216. ISBN: 9783030630577. ISSN: 0302-9743.

https://doi.org/10.1007/978-3-030-63058-4_19

© Springer Nature Switzerland AG 2020. This is an author produced version of a conference paper published in Lecture Notes in Computer Science. Uploaded in accordance with the publisher's self-archiving policy. This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use

(<https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: https://doi.org/10.1007/978-3-030-63058-4_19 .

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Exascale Computing Deployment Challenges

Karim Djemame¹[0000-0001-5811-5263] and Hamish Carr¹[0000-0001-6739-0283]

School of Computing, University of Leeds, UK K.Djemame, H.Carr@leeds.ac.uk
<https://eps.leeds.ac.uk/computing/staff/187/professor-karim-djemame>
<https://eps.leeds.ac.uk/computing/staff/499/dr-hamish-carr>

Abstract. As Exascale computing proliferates, we see an accelerating shift towards clusters with thousands of nodes and thousands of cores per node, often on the back of commodity graphics processing units. This paper argues that this drives a once in a generation shift of computation, and that fundamentals of computer science therefore need to be re-examined. Exploiting the full power of Exascale computation will require attention to the fundamentals of programme design and specification, programming language design, systems and software engineering, analytic, performance and cost models, fundamental algorithmic design, and to increasing replacement of human bandwidth by computational analysis. As part of this, we will argue that Exascale computing will require a significant degree of co-design and close attention to the economics underlying the challenges ahead.

Keywords: Exascale computing · High Performance Computing · Holistic Approach · Economics.

1 Introduction

The Exascale project will accelerate exciting advances and scientific discovery in many diverse fields such as genomics, market economics, astrophysics as well as contribute to the economic competitiveness [6]. Modern High-Performance Computing (HPC) is already at the petascale, using distributed techniques pioneered in the 1990's in response to shortcomings of the previous massive vector machine models. Thirty years later, the pattern is now repeating, as the limitations of distributed computing drive the development of hybrid models where distributed and vector techniques are layered on top of each other. This generational transition in underlying hardware models means that effective transition to Exascale computing and later to yottascale, requires not just development of immediately applicable techniques, but fundamental attention to the underlying models of not computational science, but computer science. Consequently, a separation of concerns is necessary if Exascale is to become widely available in a sustainable manner. This paper argues that algorithmic development to harness the power of Exascale cannot be achieved without considering the fundamental computer science issues that Exascale computing raises. Its primary contribution is the formal identification of the major issues that must be overcome, the hardware

and software environments inside which Exascale development will happen, the software abstractions necessary to cope efficiently with ever-larger computations, the programming methodologies based on the abstractions that will enable more efficient application development, the interpretation methods necessary at Exascale, and the underlying economics issues to be addressed to ensure effective adoption.

The paper is structured as follows: Section 2 gives an overview of the landscape in the shadows of the petascale cluster and presents the fundamental computer science challenges that Exascale computing raises. The associated economics issues to productively deploy Exascale are identified in Section 3. Section 4 summarises and concludes the paper.

2 Grand Challenges and Research Agenda

The development of the extreme scale supercomputers arena such as Argonne National Laboratory’s *Aurora*, Oak Ridge National Laboratory’s *Frontier* and Lawrence Livermore National Laboratory *El Capitan* is under way [1]. These three CRAY supercomputers being fielded in the 2021-2023 timeframe will all employ the Cray Shasta architecture, its Slingshot interconnect and a new software platform. Vector and distributed parallelism often demand different approaches with unique optimisations, workflow and shortcuts. Hybrid parallelism, such as will be needed for the next generation, demands the best of both worlds, and poses a challenge that must be overcome in order to realise the potential for not just Exascale computation, but for the yottascale computation that will come after it. Substantial research efforts have been put into Exascale computing over the last few years. The Exascale Computing Project (ECP) in the USA is bringing together research, development, and deployment projects as part of a capable Exascale computing ecosystem to ensure an enduring Exascale computing capability [7]. The EU plans to develop and reinforce the European high-performance computing and data processing capabilities to achieve Exascale capabilities by 2023, and has funded a number of technology projects [2]. What becomes urgent is to draw a complete picture of the state of the art and the challenges ahead to exploit the Exascale. Instead of tackling research into the areas of *input, programming, systems* and *output/analysis* separately, the proposed way forward is to take a holistic approach that considers the entire Exascale software stack to: 1) identify the missing functionalities, bottlenecks, unsolved problems, best practice and trends to support Exascale computing across the stack, and 2) to define and integrate new requirements into the design and development process for software. We acknowledge that as Exascale hardware evolves, software and applications need to adapt, and as Exascale application requirements evolve, hardware and software design need to adapt as well. This mutual adaptation process is a problem that needs addressing.

Challenge 1 – Exploiting Exascale Systems A one million processors supercomputer can run unprecedentedly large scale simulations and offer orders of magnitude greater performance than conventional processors. Moreover, the

future direction of Exascale systems is extreme heterogeneity, often of potentially reconfigurable systems. Nodes will be a collection of heterogeneous processors: some general purpose (CPUs), some specialised (such as SIMD units, e.g. GPUs) and some configurable cores (FPGA fabrics) [9]. Moreover, the shift to hybrid symmetric multiprocessing and distributed parallelism represents one of the biggest challenges that software developers will face. Research is needed to investigate how to automate efficient mapping across the complex and heterogeneous Exascale landscape, and how to adapt workflows to meet infrastructure-level objectives needs addressing, e.g. through machine learning approaches, as well as the identification of the building blocks for run-time reconfigurable Exascale systems.

Challenge 2 – Application design and development Research is needed to identify the best practice to increase programmer efficiency considering the emerging requirements for Exascale algorithms. Techniques such as constraint analysis and program synthesis to discover code patterns that match heterogeneous hardware needs investigation. Research will identify the required features to improve the performance and interoperability between the current programming models, and how to provide the tool(s) below the Programming Model and interfaces that exploit the best practices in terms of scalability and performance. Domain Specific Languages (DSLs) as higher level abstractions can be used to increase programmer efficiency, hide the hardware complexity and allow hardware independence [8], but must be designed first. Generic computing patterns that allow the implementation of customised algorithms to optimise load balancing, data handling, and support fault tolerance can also ease application development [5]. The Programming Model itself should support ease of programming and migration for legacy code.

Challenge 3 – Middleware, Resource Management and Performance One of the most critical aspects in the evolution of High Performance Exascale systems is a dedicated middleware to manage the enormous complexity of such systems where deep heterogeneity is needed to handle the wide variety of applications. Intelligent resource management thanks to automated methods using machine learning / data science is the way forward for handling the complexity of Exascale systems and optimise their performance. With current Petascale storage systems support data processing in, or close to, the storage location to improve performance, the identification of the technologies supporting data access for High Performance Exascale Data Analytics becomes key.

Challenge 4 - Data analysis and interpretation The human visual system allows humans to grapple with data in the megabyte to gigabyte range, and recent innovations allow artificial intelligence to operate with similar amounts of data. Beyond this, interpretation involves some form of abstraction, reduction or summarisation of the data [10]. For a terabyte of data, this has meant tools that allow humans to look at 0.1% of the data through slicing, statistical summarisation, custom domain-dependent tests, or visualisation, i.e. a three orders of magnitude reduction in the data presented to the human. At the Exascale, computational analysis of all forms will become increasingly important, but the

analytic tools in each field tend to be isolated. Research is needed to map out the tools currently used for data interpretation and analysis, and to define a framework that can be carried forward for tools with wide application at the Exascale.

3 The Economics

The diffusion of Exascale computing will follow the pattern of (old) High Performance Computing but will create visionary ideas and a process of generalised adoption. This gives a very broad set of options that will emerge due to economic constraints, critical mass issues, industrialisation aspects, legacy and usability problems.

Investment. The deployment of Exascale is driven by economic and societal needs and takes into account the changes expected in the technologies and architectures of the expanding underlying hardware and software infrastructure. Data acquisition will continue to fuel demand for more processing as well as enabling workloads that combine HPC, advanced analytics and IoT at scale. However, the investments needed to productively deploy Exascale are substantial. Future business models will be categorised based on criteria that take into account their value propositions, their technological and economic incentives and emerging trends in the market of Exascale.

Software (Re)Engineering. In the discipline of software (re)engineering, making applications Exascale-ready will incur costs, but the resulting software itself will have economic attributes as well. The inherent re-engineering cost factors will include the quality of the software to be re-engineered, the availability of the expertise and tools support for re-engineering as well as the extent of the (possible) data conversion which is required.

New business models. Industry and SMEs are currently increasingly relying on the power of supercomputers to work on innovative solutions, reduce costs and decrease time to market for products and services. An example of HPC provision is Cray on Azure to answer the demand for a large-scale, dedicated compute instance in the cloud [4]. They will continue to do so with Exascale, and may provide previously unimaginable scenarios in upstream processing that leverage the Exascale.

Application deployment. The rise of new disruptive applications such as Cyber Physical Systems, which are evolving into Cyber Physical Systems of Systems and other major technological advances in ICT are profoundly transforming their development and management processes. These are difficult to analyse because they may be deployed across a multitude of systems which includes Exascale facilities, cloud data centres, edge/fog components and networks. Their deployment should take place in an efficient manner, not only technically but economically as well.

Exascale as a Service. Cloud Service Providers such as Microsoft Azure and Amazon Web Services (AWS) are providing elastic and scalable cloud infrastructure to run applications on HPC systems [3]. With Exascale they will

face more economic challenges with the provision of costly mixed CPU, GPU and FPGA types and high-performance low-latency interconnects. For these investments to pay off, finding new accounting and pricing models as well as maximising the use of the high-cost infrastructure is a key requirement as the classic cloud service provision model will not be adequate. In an IoT context, Exascale as a service will become a part of the overall application workflow across a number of domains: IoT, edge/fog and cloud, with a complete data flow all the way from the IoT devices up to the Exascale/cloud data centres.

4 Conclusion

The ongoing deployment of Exascale computing is bringing a number of challenges. Research on both Exascale applications and related technologies will expand from the traditional fields which deploy HPC solutions to new ones, in order to not only address the requirements of (often disruptive) applications such as Artificial Intelligence and Data Analytics but the generational transition in underlying hardware models. The paper has argued that a separation of concerns is necessary if Exascale is to become widely available in a sustainable manner. This will require a truly interdisciplinary effort to not only develop applicable techniques, but address the underlying models of both computational science and computer science. Moreover, it has identified key economics issues following such paradigm shift.

References

1. Cray Exascale (2019), <https://www.cray.com/company/news-and-media/doe-el-capitan-press-release>
2. European High Performance Computing Handbook (2019), https://www.etp4hpc.eu/pujades/files/ETP4HPC_Handbook_2019_web.pdf
3. Amazon High Performance Computing (2020), <https://aws.amazon.com/hpc/>
4. Cray in Azure (2020), <https://www.cray.com/solutions/supercomputing-as-a-service/cray-in-azure>
5. Alowayyed, S., Piontek, T., Suter, J., Hoenen, O., Groen, D., Luk, O., Bosak, B., Kopta, P., Kurowski, K., Perks, O., Brabazon, K., Jancauskas, V., Coster, D., Coveney, P., Hoekstra, A.: Patterns for High Performance Multiscale Computing. *Future Generation Computer Systems* **91**, 335 – 346 (2019)
6. Gagliardi, F., Moreto, M., et al, M.O.: The International Race Towards Exascale in Europe. *CCF Transactions on HPC* **1**, 3–13 (2019)
7. Messina, P.: The Exascale Computing Project. *Computing in Science Engineering* **19**(3), 63–67 (2017)
8. Silvano, C., Agosta, G., Bartolini, A.: ANTAREX: A DSL-Based Approach to Adaptively Optimizing and Enforcing Extra-Functional Properties in High Performance Computing. In: *Proceedings of the 2018 21st Euromicro Conference on Digital System Design (DSD)*. pp. 600–607 (2018)
9. Terzo, O., Djemame, K., Scionti, A., Pezuela, C.: *Heterogeneous Computing Architectures: Challenges and Vision*. CRC Press, 1st edn. (2019)
10. Tierny, J., Carr, H.: Jacobi Fiber Surfaces for Bivariate Reeb Space Computation. *IEEE Transactions on Visualization and Computer Graphics* **23**(1), 960–969 (2017)