

High Performance Computing and Industry 4.0: Experiences from the DISRUPT Project

Rizos Sakellariou
School of Computer Science
University of Manchester, UK
rizos@manchester.ac.uk

Jorge Buenabad-Chávez
School of Computer Science
University of Manchester, UK
jorge@manchester.ac.uk

Evangelia Kavakli*
School of Computer Science
University of Manchester, UK
evangelia.kavakli@manchester.ac.uk

Ilias Spais[†]
School of Computer Science
University of Manchester, UK
ilias.spais@manchester.ac.uk

Vasilios Tountopoulos
Athens Technology Center S.A.
Greece
v.tountopoulos@atc.gr

ABSTRACT

Significant effort is currently being invested on enhancing digitization in all aspects of industrial processes and manufacturing on the quest for the next transformational change towards sustainable manufacturing and the factories of the future. This vision, commonly termed as Industry 4.0, is at the core of a number of funded projects in Europe. This short paper briefly reports experiences from the EU-funded project DISRUPT and focuses on the identification of some key challenges in the interplay between Industry 4.0 and High Performance Computing.

CCS CONCEPTS

• **Computer systems organization** → *Embedded and cyber-physical systems*; • **Software and its engineering** → *Software performance*;

KEYWORDS

Industry 4.0, High Performance Computing, Industrial Internet of Things

ACM Reference Format:

Rizos Sakellariou, Jorge Buenabad-Chávez, Evangelia Kavakli, Ilias Spais, and Vasilios Tountopoulos. 2018. High Performance Computing and Industry 4.0: Experiences from the DISRUPT Project. In *2018 International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XVIII)*, July 15–19, 2018, Pythagorion, Samos Island, Greece. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3229631.3264660>

*Evangelia Kavakli is also affiliated with the Department of Cultural Technology and Communication, University of the Aegean, Greece.

[†]Ilias Spais is also affiliated with Aegis IT Research, UK.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAMOS XVIII, July 15–19, 2018, Pythagorion, Samos Island, Greece

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6494-2/18/07...\$15.00

<https://doi.org/10.1145/3229631.3264660>

1 INTRODUCTION

The term ‘Industry 4.0’ has been commonly used recently to describe the set of transformational changes in manufacturing and industrial processes that underpin the next industrial revolution and the transition towards smart manufacturing [6]. Taking advantage of enabling concepts and technologies, such as Cyber-Physical Systems (CPS) [9] and the Internet of Things [14], the overall objective is to handle efficiently and effectively the wealth of data that can be generated by modern manufacturing ecosystems through the seamless integration of low-level devices and infrastructures with high-level services and processes.

Such a seamless integration creates a ‘system of systems’ engineering challenge where different heterogeneous and independently operating stand-alone systems have to inter-operate as part of an integrated digital manufacturing ecosystem of the factory of the future [10]. To address this challenge and facilitate the transition towards Industry 4.0 and smart manufacturing, the European Union through Horizon 2020, the EU Framework Programme for Research and Innovation for 2014–2020, has already invested in a number of over 100 research projects. Among these projects, DISRUPT¹ (‘Decentralised architectures for optimised operations via virtualised processes and manufacturing ecosystem collaboration’) is a 3-year project funded through the H2020-FOF-2016 call that started in September 2016. The project consortium consists of nine partners, coordinated by Centro Ricerche Fiat, with a total budget of about 3.5 million euros. The main objective of the project is to disrupt the traditional hierarchical approach for building industrial automation software systems through decentralized self-adjustment of components, services and processes. The idea is based upon in-depth data collection via CPS and supported by complex event processing and advanced simulation, modelling and optimisation techniques that enable sophisticated decision making. The first phase of project work has focused on identifying requirements through end-user involvement and specifying a software architecture to materialize the project’s vision [3, 7, 13].

DISRUPT is designed as a cloud-based platform to accommodate the anticipated high data volume and computational needs. Following the early project experiences, this short note elaborates further on these needs on the road to help shape future agendas in the interplay between High Performance Computing and Industry 4.0.

¹<http://www.disrupt-project.eu/>

2 HIGH PERFORMANCE COMPUTING AND INDUSTRY 4.0

High Performance Computing (HPC) has long been used by the manufacturing industry to support product design, optimization or testing. The automotive industry has perhaps a dominant position with a number of applications but other industries have seen some HPC use in their production processes [8, 11]. In the USA, the current ‘High Performance Computing for Manufacturing’ (HPC4Mfg) programme aims to grant access to HPC facilities for selected industry partners to address key challenges in manufacturing [5]. A recurring characteristic of this HPC use appears to be the execution of stand-alone application codes of interest. Following the work in the DISRUPT project thus far, three major HPC-related issues have been identified that can be associated with specific challenges on the road to the factory of the future. These are listed below.

HPC on demand: The concept of ‘HPC on demand’ has long been associated with dynamic adjustment of resources for compute and visualization workloads (often through cloud infrastructures) or, additionally, with big data analytics in the context of Eulerian-Lagrangian simulations more recently [2]. An integrated smart manufacturing environment may need to run various compute- or data-intensive codes at short notice to deal with events in production that may cause some disruption and may require some rescheduling, extra (simulation) checks for product quality assurance due to lack of originally planned parts, and so on. HPC demands may vary in each case and may depend on the state of the environment.

Economical data management: The potential of the Industrial Internet of Things for generating data is practically unlimited. A recent discussion paper, identifying major challenges for data-oriented Factories of the Future, has noted that for processing the exponentially increased volumes of data HPC can be employed [1]. However, there is a risk that this approach may lead to a significant consumption of HPC resources whereas an in-depth analysis of data needs may match specific requirements and data properties with appropriate data production rates thereby reducing unnecessarily large volumes of data. The question to be addressed is how often different types of data will need to be generated and collected and how queries related to these data will need to be resolved. To give an example, it may not be necessary to read a specific plant floor sensor every second if its readings change much less often and value changes are not critical to read straightaway. Quality metrics from earlier work on Grid computing [4] may have to be enhanced to control the data deluge, yet provide meaningful information.

Objective-based system configuration: The current trend in HPC is towards heterogeneous platforms and mixed programming models. To the extent that HPC systems become tightly interconnected with smart manufacturing systems they will have to react according to overall system performance requirements. These requirements may not have as a primary or sole objective response time but may prioritize the need to reduce energy consumption, data transfer, network traffic, respect some power cap or find a trade-off between some combinations of these. Conversely, any such requirements may have to be used to adjust or reconfigure devices or other components that provide data required by HPC systems. All in

all, HPC systems will have to cope with different performance objectives and what constitutes ‘best’ performance may vary every time depending on overall system requirements and trade-offs between different objectives. The need for ‘algorithms that communicate as little as possible’, highlighted in [12] for future HPC systems, may also be particularly relevant.

3 CONCLUSION

This short paper highlighted some HPC challenges that may need to be addressed on the pathway to actualize the Industry 4.0 vision. Preliminary experiences are drawn from work in the EU-funded DISRUPT project and an appreciation of current trends in HPC and the Industrial Internet of Things and may be used to drive research priorities in future years.

ACKNOWLEDGEMENT

This work has been supported by the EU H2020-FOF-11-2016 project DISRUPT (grant no 723541).

REFERENCES

- [1] Big Data Value Association. 2018. Big Data Challenges in Smart Manufacturing. (2018). Discussion paper available from <http://www.bdva.eu>.
- [2] Rafael Ferreira da Silva, Rosa Filgueira, Ilia Pietri, Ming Jiang, Rizos Sakellariou, and Ewa Deelman. 2017. A characterization of workflow management systems for extreme-scale applications. *Future Generation Computer Systems* 75 (2017), 228 – 238. DOI: <http://dx.doi.org/https://doi.org/10.1016/j.future.2017.02.026>
- [3] Pavlos Eirinakis, Jorge Buenabad-Chavez, Rosanna Fornasiero, Haluk Gokmen, Julien-Etienne Mascolo, Ioannis Mourtos, Sven Spieckermann, Vasilios Tountopoulos, Frank Werner, and Robert Woitsch. 2017. A Proposal of Decentralised Architecture for Optimised Operations in Manufacturing Ecosystem Collaboration. In *Collaboration in a Data-Rich World*, Luis M. Camarinha-Matos, Hamideh Afsarmanesh, and Rosanna Fornasiero (Eds.). Springer International Publishing, 128–137.
- [4] Laurence Field and Rizos Sakellariou. 2017. An Evaluation of Information Consistency in Grid Information Systems. *Journal of Grid Computing* 15, 1 (01 Mar 2017), 127–137. DOI: <http://dx.doi.org/10.1007/s10723-016-9373-z>
- [5] HPC4Manufacturing. 2018. High Performance Computing for Manufacturing (HPC4Mfg). (2018). <https://hpc4mfg.llnl.gov/>
- [6] Sachin S. Kamble, Angappa Gunasekaran, and Shradha A. Gawankar. 2018. Sustainable Industry 4.0 framework: A systematic literature review identifying the current trends and future perspectives. *Process Safety and Environmental Protection* 117 (2018), 408 – 425. DOI: <http://dx.doi.org/https://doi.org/10.1016/j.psep.2018.05.009>
- [7] Evangelia Kavakli, Jorge Buenabad-Chávez, Vasilios Tountopoulos, Pericles Loucopoulos, and Rizos Sakellariou. 2018. WiP: An Architecture for Disruption Management in Smart Manufacturing. In *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*. 279–281. DOI: <http://dx.doi.org/10.1109/SMARTCOMP.2018.00057>
- [8] S. Kodiyalam, R.J. Yang, L. Gu, and C.-H. Tho. 2004. Multidisciplinary design optimization of a vehicle system in a scalable, high performance computing environment. *Structural and Multidisciplinary Optimization* 26, 3 (01 Feb 2004), 256–263. DOI: <http://dx.doi.org/10.1007/s00158-003-0343-2>
- [9] László Monostori. 2014. Cyber-physical Production Systems: Roots, Expectations and R&D Challenges. *Procedia CIRP* 17 (2014), 9 – 13. DOI: <http://dx.doi.org/https://doi.org/10.1016/j.procir.2014.03.115>
- [10] Saeid Nahavandi, D. Creighton, V. T. Le, M. Johnstone, and J. Zhang. 2015. *Future Integrated Factories: A System of Systems Engineering Perspective*. Springer International Publishing, Cham, 147–161. DOI: http://dx.doi.org/10.1007/978-3-319-15898-3_9
- [11] Anwar Osseyran and Merle Giles. 2015. *Industrial applications of high-performance computing: best global practices*. Vol. 25. CRC Press.
- [12] Daniel A. Reed and Jack Dongarra. 2015. Exascale Computing and Big Data. *Commun. ACM* 58, 7 (June 2015), 56–68. DOI: <http://dx.doi.org/10.1145/2699414>
- [13] Frank Werner and Robert Woitsch. 2018. Data Processing in Industrie 4.0. *Datenbank-Spektrum* 18, 1 (01 Mar 2018), 15–25. DOI: <http://dx.doi.org/10.1007/s13222-018-0277-x>
- [14] Li Da Xu, Wu He, and Shancang Li. 2014. Internet of Things in Industries: A Survey. *IEEE Transactions on Industrial Informatics* 10, 4 (Nov 2014), 2233–2243. DOI: <http://dx.doi.org/10.1109/TII.2014.2300753>