

HPC Application Optimisation in SODALITE

Daniel Vladusic, Dragan Radolovic, Joao Pita Cost
XLAB
daniel.vladusic,dragan.radolovic,joao.pitacosta,@xlab.si

Kalman Meth
IBM Research - Haifa
meth@il.ibm.com

Elisabetta Di Nitto
Politecnico di Milano
elisabetta.dinitto@polimi.it

Alfio Lazzaro, Nina Mujkanovic
HPE HPC/AI EMEA Research Lab
alfio.lazzaro,nina.mujkanovic@hpe.com

Maria Carbonell
ATOS
Maria.carbonell@atos.net

Published in [SYSTOR'21] ACM International Systems and Storage Conference. June 2021, Haifa, Israel.

ACM ISBN 978-1-4503-8398-1/21/06

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825480.

HPC Application Optimisation in SODALITE

Daniel Vladusic, Joao Pita Costa,
XLAB
{daniel.vladusic,joao.pitacosta}@xlab.si

Dragan Radolovic, XLAB
dragan.radolovic@xlab.si

Kalman Meth, IBM Research -
Haifa
meth@il.ibm.com

Elisabetta Di Nitto, Politecnico di
Milano
elisabetta.dinitto@polimi.it

Alfio Lazzaro, Nina Mujkanovic,
HPE HPC/AI EMEA Research Lab
{alfio.lazzaro,nina.mujkanovic}@hpe.com

Maria Carbonell, ATOS
Maria.carbonell@atos.net

ABSTRACT

We propose to tackle the complexity of deploying and operating modern applications onto heterogeneous HPC and cloud-based systems by providing application developers and infrastructure operators with tools to abstract their application and infrastructure requirements.

CCS CONCEPTS

• **Software and its engineering** → *Model-driven software engineering; Development frameworks and environments*; • **Computer systems organization** → *Heterogeneous (hybrid) systems*; Cloud computing.

KEYWORDS

HPC, Cloud computing, application optimisation

1 PROBLEM STATEMENT

Software application developers and users are now targeting a wide range of diverse computing platforms: clusters and supercomputers with homogeneous or heterogeneous node architectures for heavy batch computations, including resources available on the Cloud. HPC jobs with requirements for specialized execution environments, such as compute accelerators or specialized hardware for data movement, are becoming more and more popular. However, the support of specific software technology to manage the heterogeneous

infrastructures in compute centers requires standards that are not yet properly defined.

2 PROPOSED SOLUTION

SODALITE [1], an Horizon 2020 project, proposes to tackle the complexity of deploying and operating modern applications onto heterogeneous HPC and cloud-based systems. It does so by providing application developers and infrastructure operators with tools to abstract their application and infrastructure requirements. Models are built to abstract performance and other characteristics of HPC and cloud-based environments. Models may also be built to capture the behaviour of a distributed application. This enables simpler and faster development, deployment, operation, and execution of heterogeneous applications reflecting diverse circumstances over heterogeneous high-performance and cloud infrastructures. SODALITE attempts to bring the vast knowledge of performance optimization accrued by the HPC industry over decades into the cloud computing area. Using application and infrastructure Performance Models, the SODALITE deployment framework enables automated performance optimization before deployment (static) and at runtime (dynamic). For static optimizations, three broad application types are supported: AI training/Inference, Big Data Analytics, and Traditional HPC. Using the SODALITE Integrated Development Environment (IDE), application experts can select optimizations for these applications, which resolve to native instantiations of the application tuned for execution on the given hardware. Additionally, application runtime parameters can also be autotuned, further improving performance and the efficiency of the system.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825480.

REFERENCES

[1] Sodalite. 2019. *Software Defined AppLIcations Infrastructures management and Engineering*. <https://www.sodalite.eu/>

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 the author(s) 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.
SYSTOR '21, June 14–16, 2021, Haifa, Israel
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8398-1/21/06...\$15.00

<https://doi.org/10.1145/3456727.3463829>