

ANTAREX – AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems

Cristina Silvano, Giovanni Agosta, Andrea Bartolini, Andrea Beccari, Luca Benini, João M. P. Cardoso, Carlo Cavazzoni, Jan Martinovič, Gianluca Palermo, Martin Palkovič, et al.

▶ **To cite this version:**

Cristina Silvano, Giovanni Agosta, Andrea Bartolini, Andrea Beccari, Luca Benini, et al.. ANTAREX – AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems. 18th IEEE International Conference on Computational Science and Engineering, Oct 2015, Porto, Portugal. hal-01235713

HAL Id: hal-01235713
https://inria.hal.science/hal-01235713

Submitted on 3 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ANTAREX – AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems*

Cristina Silvano*, Giovanni Agosta*, Andrea Bartolini†, Andrea Beccari‡, Luca Benini†, João M. P. Cardoso§, Carlo Cavazzoni¶
Jan Martinovič‖, Gianluca Palermo*, Martin Palkovič‖, Erven Rohou**, Nico Sanna¶, Katerina Slaninova‖
*DEIB – Politecnico di Milano, †IIS – Eidgenössische Technische Hochschule Zürich, ‡Dompé S.p.A., §FEUP – Universidade do Porto,
¶CINECA, ‖IT4Innovation National Supercomputing Center, **INRIA Rennes
Email: *name.surname@polimi.it †{barandre,lbenini}@iis.ee.ethz.ch, ‡andrea.beccari@dompe.it, §jmpc@acm.org,
¶n.surname@cineca.it, ‖name.surname@vsb.cz, **erven.rohou@inria.fr

***Abstract*—The main goal of the ANTAREX project is to express by a Domain Specific Language (DSL) the application self-adaptivity and to runtime manage and autotune applications for green and heterogeneous High Performance Computing (HPC) systems up to the Exascale level. Key innovations of the project include the introduction of a separation of concerns between self-adaptivity strategies and application functionalities. The DSL approach will allow the definition of energy-efficiency, performance, and adaptivity strategies as well as their enforcement at runtime through application autotuning and resource and power management.**

***Keywords*—High Performance Computing, Autotuning, Adaptivity**

I. Introduction

High Performance Computing (HPC) has been traditionally the domain of grand scientific challenges and a few industrial domains such as oil & gas or finance, where investments are large enough to support massive computing infrastructures. Nowadays HPC is recognized as a powerful technology to increase the competitiveness of nations and their industrial sector, including small scale but high-tech businesses – *to compete, you must compute* has become an ubiquitous slogan [1]. The current road-maps for HPC systems aim at reaching exascale levels (10^{18} FLOPS) within 2020 – a $\times 1000$ improvement over petascale, which was reached in 2009, and a $\times 100$ improvement over current systems. Reaching exascale poses the additional challenge of significantly limiting the energy envelope while providing massive increases in computational capabilities – the target power envelope for future exascale system ranges between 20 and 30 megawatts. Thus, "Green" HPC systems, are being designed aiming at maximizing a FLOPS/Watt metric, rather than just FLOPS, and employing increasingly heterogeneous architectures with GPGPU or MIC accelerators. On average, the efficiency of heterogeneous systems is more than double that of homogeneous systems (i.e., $1,938$ MFLOPS/Watt vs 743 MFLOPS/Watt[1]). This level of efficiency is still two orders of magnitude lower than that needed for supporting exascale systems at the target power envelope of 20 megawatts. To this end, European efforts have been focused towards building supercomputers out of the less power-hungry ARM cores and GPGPUs [2]. On the other hand, the wide margin provided by modern chip manufacturing techniques, combined with the inability to exploit this space to produce faster, more complex cores due to the breakdown of Dennard scaling [3], has given rise to a pervasive diffusion of a number of parallel computing architectures, up to the point where embedded systems are also characterized by multicore processors. The large design effort has led to a variety of approaches in terms of core interconnection and data management. Thus, the ability to port applications designed for current platforms, based on GPGPUs like the NVIDIA Kepler or Tesla families, to heterogeneous systems such as those currently designed for embedded systems is critical to provide software support for future HPC.

Designing and implementing HPC applications is a difficult and complex task, which requires mastering many specialized languages and tools for performance tuning. This is incompatible with the current drive to opening HPC infrastructures to a much wider range of users. The current model of having the HPC center staff directly supporting the development of the application will become unsustainable in the long term. Thus, the availability of effective standard programming languages and APIs is critical to provide migration paths towards novel heterogeneous HPC platforms as well as to guarantee the ability of developers to work effectively on these platforms. To fulfil the 20MWatt target, energy-efficient heterogeneous supercomputers need to be coupled with a radically new software stack capable of exploiting the benefits offered by heterogeneity at all the different levels (supercomputer, job, node).

ANTAREX will address these challenging problems through a holistic approach spanning all the decision layers composing the supercomputer software stack and exploiting effectively the full system capabilities (including heterogeneity and energy management). The main goal of the ANTAREX project is to express by a DSL the application self-adaptivity and to runtime manage and autotune applications for green heterogeneous HPC systems up to Exascale.

One key innovation of the proposed approach consists of introducing a separation of concerns (where self-adaptivity and energy efficient strategies are specified aside to application functionalities) promoted by the definition of a DSL inspired by aspect-oriented programming concepts for heterogeneous systems. The new DSL will be introduced for expressing at compile time the adaptivity/energy/performance strategies and to enforce at runtime application autotuning and resource and power management. The goal is to support the parallelism, scalability and adaptivity of a dynamic workload by exploiting the full system capabilities (including energy management) for emerging large-scale and extreme-scale systems, while reducing the Total Cost of Ownership (TCO) for companies and public organizations. ANTAREX approach will be based on: (1) introducing a new DSL for expressing adaptivity and autotuning strategies; (2) enabling the performance/energy control capabilities by introducing software knobs (including application parameters, code transformations and code variants); (3) designing scalable and hierarchical optimal control-loops capable of dynamically leveraging them together with performance/energy control knobs at different time scale (compile-, deploy- and run-time) to always operate the supercomputer and each application at the maximum energy-efficient and thermally-safe point. This can be done by monitoring the evolution of the supercomputer as well as the application status and requirements and bringing this information to the ANTAREX

*ANTAREX is supported by the EU H2020 FET-HPC program
[1]www.green500.org, June 2014

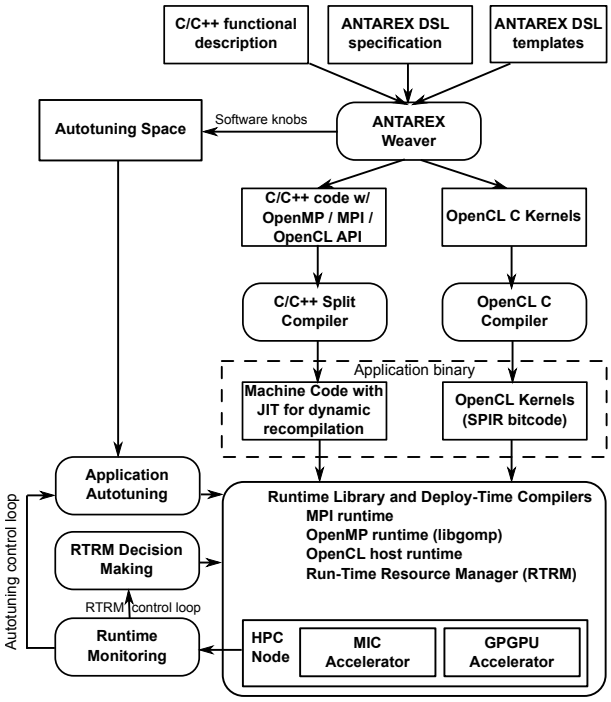


Fig. 1. The ANTAREX Tool Flow

energy/performance-aware software stack.

The ANTAREX project is driven by two use cases from highly relevant HPC application scenarios: (1) a biopharmaceutical HPC application for drug discovery deployed on the 1.2 PetaFlops heterogeneous NeXtScale Tier-1 Intel-based IBM system based at CINECA and (2) a self-adaptive navigation system for smart cities deployed on the server-side on the 1.457 PetaFlops heterogeneous Intel® Xeon Phi™ based system provided by IT4Innovations National Supercomputing Center. All the key ANTAREX software innovations will be designed and engineered since the beginning to be scaled-up to the Exascale level. Performance metrics extracted from the two use cases will be modelled to extrapolate these results towards Exascale systems expected by the end of 2020.

The ANTAREX Consortium comprises a wealth of expertise in all pertinent domains. Four top-ranked academic and research partners (Politecnico di Milano, ETHZ Zurich, University of Porto and INRIA) with extensive experience in the salient research topics to be explored will be complemented by the Italian Tier-0 Supercomputing Center (CINECA), the the Tier-1 Czech National Supercomputing Center (IT4Innovations) and two industrial application providers, one of the leading biopharmaceutical companies in Europe (Dompé) and the top European navigation software company (Sygic).

II. The ANTAREX Approach

The ANTAREX approach and related tool flow, shown in Figure 1, operates both at design-time and runtime. The application functionality is expressed through C/C++ code (possibly including legacy code), whereas the non-functional aspects of the application, including parallelisation, mapping, and adaptivity strategies are expressed through the DSL developed in the project. One of the benefits consists of facilitating the reuse of legacy code. In the definition of these strategies, the application developer or system integrator can leverage DSL templates that encapsulate specific mechanisms, including how to generate code for OpenCL or OpenMP parallelisation, and how to interact with the runtime resource manager. The DSL weaver and refactoring tool will then enhance the C/C++ functional specification with the desired adaptivity strategies, generating a version of the code that includes the necessary libraries as well as the partitioning between the code for the GPPs and the code for the accelerators (such as GPGPUs and MIC accelerators [4]). A mix of off-the-shelf and custom compilers will be used to generated code, balancing between development effort and optimization level.

Thus, the ANTAREX compilation flow leverages a runtime phase with compilation steps, through the use of split-compilation techniques. The application autotuning is delayed to the runtime phase, where the software knobs (application parameters, code transformations and code variants) are configured according to the runtime information coming from the execution environment. Finally the runtime resource and power manager are used to control the resource usage for the underlying computing infrastructure given the changing conditions. At runtime, the application control code, thanks to the design-time phase, now contains also runtime monitoring and adaptivity strategies code derived from the DSL extra-functional specification. Thus, the application is continuously monitored to guarantee the required Service Level Agreement (SLA), while communication with the runtime resource-manager takes place to control the amount of processing resources needed by the application. The application monitoring and auto-tuning will be supported by a runtime layer implementing an application level collect-analyse-decide-act loop.

III. Technology Challenges

In this section, we introduce the main technological challenges addressed by ANTAREX.

A. Programming Languages & Compiler Technology

1) A Domain Specific Language for Adaptivity and Specialization: HPC applications may profit by adapting to operational and situational conditions, such as changes in contextual information (e.g., workloads), in requirements (e.g., deadlines), and in resources availability (e.g., energy, connectivity, number of processor nodes available). A simplistic approach to both adaptation specification and implementation (see, e.g., [5]) employs hard coding of, e.g., conditional expressions and parameterizations. In our approach the specification of runtime adaptability strategies will rely on an extended version of the LARA DSL [6], [7]. Our approach will consider runtime strategies to adapt applications to the resources in the target computing system, e.g., controlling resource allocation and task mapping, subject to specific requirements [8], and dynamic compilation (see, e.g., ADAPT [9]). The separation of concerns to be followed will allow developers to retain the original application source-code while exploiting the benefits of an automatic approach for applying strategies. At compile time, the DSL will be used to specify strategies to guide and control code transformations, and to specify specific code translation schemes. At runtime, DSL strategies will be executed to provide runtime adaptivity, including specialization, tuning of parameters, and task migration and mapping decisions keeping into account SLA, energy budget, and system resources.

2) Parallel Programming Model and Language Runtime: Currently, directive-driven programming models (e.g., MPI and OpenMP) are used to map computations to multi-cores and multiprocessors. These models are also limited in scope and do not allow to express complex, crosscutting adaptation strategies. The ANTAREX DSL approach aims at reaching a higher abstraction level, to separate and express data communication and computation parallelism, and to augment the capabilities of existing programming models by passing hints and metadata to the compilers. Some DSL higher-abstractions are then translated to lower level directive-driven programming

models for the target HPC system. Moreover, heterogeneity is a key factor for reaching Exascale level. OpenCL is the industry standard for programming heterogeneous many-core architectures combining GPUs, CPUs and DSPs. It achieves efficiency by fully exposing the memory hierarchy and device heterogeneity to the programmer, who is burdened with fine-tuning for performance, which is very sensitive to even minimal variation in the architectural parameters [10], [11]. Through the ANTAREX DSL, we aim at improving performance portability while integrating the management of heterogeneity (e.g., through specialized code versions) into the same mechanisms used to express and manage other non-functional constraints.

3) Split Compilation: Iterative compilation [12] techniques are attractive, since they allow to identify the best compiler optimizations for a given program/code fragment considering possible trade-offs. Given the diversity of heterogeneous multiprocessors and the potential for optimizations provided by runtime information, runtime optimization is also desirable. To combine the two approach, a *split compilation* will be used. The key idea is to split the compilation process in two steps - offline, and online - and to offload as much of the complexity as possible to the offline step, conveying the result to runtime optimizers [13]. We will encode code generation strategies to drive a dynamic code generator in response to particular hardware features as well as dynamic information. This combination of iterative- and split-compilation will have a significant impact on the performance of applications, but also on the productivity of programmes as it will relieve programmers from the burden of repeatedly optimizing, tuning, compiling and testing.

B. Self-Adaptivity & Autotuning

The management of system adaptivity and autotuning is a key issue in HPC systems, where the computing infrastructure can easily evolve and the system needs to react promptly to changing workloads and events, without impacting too much its extra-functional characteristics, such as energy and thermal features. The motivation behind this trend can be easily explained by the requirement to meet the maximum performance/power ratio across all the possible deployments of the applications. This is especially important when considering the rapid grow of computing infrastructures that continue to evolve on one hand by increasing computing nodes, while on other hand by increasing the performance exploiting heterogeneity in terms of accelerators/co-processors. Thus, there is a requirement on applications to become adaptive with respect to the computing resources. In this direction, another interesting effect is that there is a growing need of guaranteeing Service Level Agreement (SLA) both at the server- and at the application-side. This need is related to the performance of the application, but also to the maximum power budget that can be allocated to a specific computation. In this context, our efforts are mainly focused on two main paths: i) the development of an autotuning framework capable to configure and adapt application-level parameters and ii) the possibility to export the concept of precision autotuning to HPC applications.

Application Autotuning. Two types of framework have been investigated so far to support application autotuning depending on the knowledge about the target domain: white-boxes and black-boxes. White-box techniques are those approaches based on autotuning libraries (thus not general) that deeply use the domain specific knowledge to fast surf the parameter space. On the other side, black-box techniques are always-applicable frameworks since they do not require any knowledge on the underlying application, however suffering of long convergence time and less custom possibilities. The proposed framework falls in the area of grey-box approaches since, starting from the idea of non-domain knowledge, it can be fed code annotations to shrink the search space by focusing the auto-tuner on a certain subspace. Moreover, the framework includes a monitoring loop to be used to monitor the application and to trigger the application adaptation. The monitoring, together with application properties/features, represents the main support to the decision-making during the application autotuning phase since it is used to perform statistical analysis related to system performance and other SLA aspects. Continuous on-line learning techniques is adopted to update the knowledge from the data collected by the monitors, giving the possibility to auto-tune the system always according to the more recent operating conditions. Machine learning techniques are also adopted in the decision-making engine to support the autotuning by predicting the most promising set of parameter settings.

Precision Autotuning. In the recent years, customized precision has emerged as a promising approach to achieve power/performance trade-offs. It derives from the fact that many applications can tolerate some loss of quality during computation, as in the case of media processing (audio, video and image) and data mining. These applications are increasingly common in the emerging field of real-time HPC. In ANTAREX, the benefits of customized precision HPC applications will be investigated in tight collaboration with the domain experts of the two use cases. Additionally, we also plan to apply fully automatic dynamic optimizations, based on profiling information, and data acquired at runtime, e.g. dynamic range of function parameters.

In both cases the usage of the ANTAREX DSL will be a key-point to decouple the functional specification of the application from the definition of software knobs (such as code variants or application parameters) and from the precision tuning phase.

C. Runtime Resource & Power Management

ANTAREX follows a holistic approach toward next-generation energy-efficient exascale supercomputers. While traditional design of green supercomputer relies the integration of best-in-class energy-efficient components [14], recent works [15], [16] show that as effect of this design practice supercomputers are nowadays heterogeneous system. Indeed supercomputers are not only composed by heterogeneous computing architectures (GPGPUs and CPUs), but different instances of the same nominal component execute the same application with 15% of variation in the energy-consumption. Different applications on the same resources show different performance-energy trade-offs, for which an optimal selection of operating point can save in between 18% and 50% of node energy w.r.t. the default frequency selection of the Linux OS power governor. Moreover it has been recently shown that environmental conditions, such as ambient temperature, can significantly change the overall cooling efficiency of a supercomputer, causing more than 10% *Power usage effectiveness* (PUE) loss when transitioning from winter to summer [17]. These sources of heterogeneity, coupled with current worst-case design practices, lead to significant less in energy-efficiency, and to missed opportunities for run-time resource and power management (RTRM, RTPM). ANTAREX leverages RTRM and RTPM by combining: (1) novel introspection points, application progress and dynamic requirements; (2) autotuning capabilities enabled by the DSL in the applications; and (3) information coming from the processing elements of the Exascale machine and IT infrastructure and the respective performance knobs (resource allocation, DVFS, cooling effort, room temperature). These information flows converge in a scalable multilayer resource management infrastructure. The information will be used to allocate to each application the set of resources and their operating point which maximize the entire supercomputer energy-efficiency, while respecting SLA and safe working conditions. The latter will be ensured by the resource management solution by

optimal selection of the cooling effort and by a distributed optimal thermal management controller. The ANTAREX power management and resource allocation approach is based on: (1) Expanding the energy/performance control capabilities by introducing novel software control knobs (i.e. software reconfigurability and adaptability); (2) Designing scalable and hierarchical optimal control-loops capable of dynamically leveraging the control knobs together with classical performance/energy control knobs (job dispatching, resource management and Dynamic Voltage and Frequency Scaling) at different time scale (compile time, deployment time and run-time); (3) Monitoring the HW supercomputing evolution, the application status and requirements, bringing this information to the energy/performance-aware software stack. This approach will allow to always operate the supercomputer and each application at the most energy-efficient and thermally-safe point.

IV. Application Scenarios & Target Platforms

The ANTAREX project is driven by two industrial HPC applications chosen to address the self-adaptivity and scalability characteristics of two highly relevant scenarios towards the Exascale era.

a) Use Case 1: Computer Accelerated Drug Discovery: Computational discovery of new drugs is a compute-intensive activity that is critical to explore the huge space of chemicals with potential applicability as drugs. Typical problems include the prediction of properties of protein-ligand complexes (such as docking and affinity), and the verification of synthetic feasibility. These problems are massively parallel, but demonstrate unpredictable imbalances in the computational time, since the verification of each point in the solution space requires a widely varying time. Moreover, different tasks might be more efficient on different type of processors, especially in a heterogeneous system. Dynamic load balancing and task placement are critical for the efficient solution of such problems [18], [19].

b) Use Case 2: Self-Adaptive Navigation System: To solve the growing automotive traffic load, it is necessary to find the best utilization of an existing road network, under a variable workload. The basic idea is to provide contextual information from server side to traditional mobile navigation users and vice versa. The approach will help to overcome the major shortcomings of the currently available navigation systems exploiting synergies between server-side and client-side computation capabilities. The efficient operation of such a system depends strongly on balancing data collection, big data analysis and extreme computational power [20], [21].

Prototypes of these two use cases will be developed, integrated and validated in relevant and realistic environments to practically assess the benefits of the self-adaptive holistic approach, as well as the scalability of the proposed approach towards Exascale systems.

The target platforms are the CINECA's Tier-1 IBM NeXtScale hybrid Linux cluster, based on Intel TrueScale interconnect as well as Xeon Haswell processors and MIC accelerators [4], and IT4Innovations Salomon supercomputer, which is a PetaFlop class system consisting 1008 computational nodes. Each node is equipped with 24 cores (two twelve-core Intel Haswell processors). These computing nodes are interconnected by InfiniBand FDR and Ethernet networks. Salomon also includes of 432 compute nodes with MIC accelerators (two Intel® Xeon Phi™ 7120P per node).

V. Conclusions

Exascale HPC systems will need the definition of new software stacks to fully exploit heterogeneity while meeting power efficiency requirements. The goal of ANTAREX is to provide a holistic system-wide adaptive approach for next generation HPC systems. Our long-term vision is to explore an innovative application programming paradigm and description methodology to decouple functional and extra-functional aspects of the application. The impact and benefits of such technology are far reaching, beyond traditional HPC domains.

References

[1] J. Curley, "HPC and Big Data," *Innovation*, vol. 12, no. 3, Jul. 2014.

[2] N. Rajovic, P. M. Carpenter, I. Gelado, N. Puzovic, A. Ramirez, and M. Valero, "Supercomputing with Commodity CPUs: Are Mobile SoCs Ready for HPC?" in *Proc. Int'l Conf. on High Performance Computing, Networking, Storage and Analysis*. ACM, 2013, pp. 40:1–40:12.

[3] M. Bohr, "A 30 Year Retrospective on Dennard's MOSFET Scaling Paper," *IEEE SSCS Newsletter*, vol. 12, no. 1, pp. 11–13, Winter 2007.

[4] G. Chrysos, "Intel® Xeon Phi™ Coprocessor-the Architecture," Intel Whitepaper, 2014.

[5] J. Floch, S. Hallsteinsen, E. Stav, F. Eliassen, K. Lund, and E. Gjorven, "Using architecture models for runtime adaptability," *IEEE Softw.*, vol. 23, no. 2, pp. 62–70, Mar. 2006.

[6] J. M. Cardoso, T. Carvalho, J. G. Coutinho, W. Luk, R. Nobre, P. Diniz, and Z. Petrov, "Lara: An aspect-oriented programming language for embedded systems," in *Proc. 11th Annual Int'l Conf. on Aspect-oriented Software Development*. ACM, 2012, pp. 179–190.

[7] J. M. Cardoso, P. C. Diniz, J. G. Coutinho, and Z. Petrov, *Compilation and Synthesis for Embedded Reconfigurable Systems: An Aspect-Oriented Approach*. Springer, 2013.

[8] G. Mariani, G. Palermo, C. Silvano, and V. Zaccaria, "Arte: An application-specific run-time management framework for multi-core systems," in *Proc. 2011 IEEE 9th Symp. on Application Specific Processors*. IEEE Comp. Soc., 2011, pp. 86–93.

[9] M. J. Voss and R. Eigemann, "High-level adaptive program optimization with adapt," in *Proc. Eighth ACM SIGPLAN Symp. on Principles and Practices of Parallel Programming*. ACM, 2001, pp. 93–102.

[10] G. Agosta, A. Barenghi, G. Pelosi, and M. Scandale, "Towards Transparently Tackling Functionality and Performance Issues across Different OpenCL Platforms," in *2nd Int'l Symp. on Computing and Networking (CANDAR)*, Dec 2014, pp. 130–136.

[11] G. Agosta, A. Barenghi, A. Di Federico, and G. Pelosi, "OpenCL Performance Portability for General-purpose Computation on Graphics Processor Units: an Exploration on Cryptographic Primitives," *Concurrency and Computation: Practice and Experience*, 2014.

[12] F. Bodin, T. Kisuki, P. Knijnenburg, M. O'Boyle, and E. Rohou, "Iterative compilation in a non-linear optimisation space," 1998.

[13] A. Cohen and E. Rohou, "Processor virtualization and split compilation for heterogeneous multicore embedded systems," in *Proc. 47th Design Automation Conference*. ACM, 2010, pp. 102–107.

[14] B. Subramaniam, W. Saunders, T. Scogland, and W.-c. Feng, "Trends in Energy-Efficient Computing: A Perspective from the Green500," in *4th Int'l Green Computing Conference*, June 2013.

[15] F. Fraternali, A. Bartolini, C. Cavazzoni, G. Tecchiolli, and L. Benini, "Quantifying the Impact of Variability on the Energy Efficiency for a Next-generation Ultra-green Supercomputer," in *Proc. 2014 Int'l Symp. on Low Power Electronics and Design*. ACM, 2014, pp. 295–298.

[16] A. Auweter, A. Bode, M. Brehm, L. Brochard, N. Hammer, H. Huber, R. Panda, F. Thomas, and T. Wilde, "A Case Study of Energy Aware Scheduling on SuperMUC," in *Supercomputing*. Springer, 2014, vol. 8488, pp. 394–409.

[17] A. Borghesi, C. Conficoni, M. Lombardi, and A. Bartolini, "MS3: a Mediterranean-Stile Job Scheduler for Supercomputers - do less when it's too hot!" in *2015 Int'l Conf. on High Perf. Comp. & Simulation*. IEEE, 2015.

[18] C. Beato, A. R. Beccari, C. Cavazzoni, S. Lorenzi, and G. Costantino, "Use of experimental design to optimize docking performance: The case of ligendock, the docking module of ligen, a new de novo design program," *J. Chem. Inf. Model.*, vol. 53, no. 6, pp. 1503–1517, 2013.

[19] A. R. Beccari, C. Cavazzoni, C. Beato, and G. Costantino, "LiGen: A High Performance Workflow for Chemistry Driven de Novo Design," *J. Chem. Inf. Model.*, vol. 53, no. 6, pp. 1518–1527, 2013.

[20] R. Tomis, J. Martinovič, K. Slaninová, L. Rapant, and I. Vondrák, "Time-dependent route planning for the highways in the czech republic," in *Proc. 14th Int'l Conf. on Computer Information Systems and Industrial Management, CISIM 2015*, 2015.

[21] D. Fedorčák, T. Kocyan, M. Hájek, D. Szturcová, and J. Martinovič, "viaRODOS: Monitoring and Visualisation of Current Traffic Situation on Highways," in *Computer Information Systems and Industrial Management*. Springer, 2014, vol. 8838, pp. 290–300.