H2020-FETHPC-1-2014 ANTAREX-671623

AutoTuning and Adaptivity approach for Energy efficient eXascale HPC systems

Deliverable D7.3: Lessons Learnt from ANTAREX Project

http://www.antarex-project.eu/

Horizon 2020
European Union funding
for Research & Innovation

Deliverable Title:	Lessons Learnt from ANTAREX Projects (Journal Submission)		
Lead beneficiary:	POLIMI (Italy)		
Keywords:	XX		
Author(s):	Cristina Silvano (POLIMI), Giovanni Agosta (POLIMI), Andrea Bartolini (ETHZ), Andrea R. Beccari (DOMPE), Luca Benini (ETHZ), Loïc Besnard (CNRS), João Bispo (UPORTO), Radim Cmar (SYGIC), João M. P. Cardoso (UPORTO), Carlo Cavazzoni (CINECA), Daniele Cesarini (ETHZ), Stefano Cherubin (POLIMI), Federico Ficarelli (CINECA), Davide Gadioli (POLIMI), Martin Golasowski (IT4I), Antonio Libri (ETHZ), Jan Martinovic (IT4I), Gianluca Palermo (POLIMI), Pedro Pinto (UPORTO) , Katerina Slaninová (IT4I), Emanuele Vitali (POLIMI)		
Reviewer(s):	Cristina Silvano (POLIMI);		
WP:	WP7	Task:	
Nature:	Report	Dissemination level:	Public
Identifier:	D7.3	Version:	V0.1
Delivery due date:	November 30, 2018	Actual submission date:	January 21st, 2019

Executive Summary:	This public report summarizes some lessons learnt from ANTAREX project under the leadership of POLIMI.

Approved and issued by the Project Coordinator: C Silvano	**Date:** January 21st, 2019

Project Coordinator: Prof. Dr. Cristina SILVANO – Politecnico di Milano
e-mail: silvano@elet.polimi.it - **Phone:** +39-02-2399-3692- **Fax:** +39-02-2399-3411

1 Lessons Learnt from ANTAREX project

The Horizon 2020 ANTAREX project concluded on 30 November 2018: It is time to summarize the main challenges, achievements, exploitation paths and lessons learnt.

Energy efficiency is of paramount importance for any class of systems, from embedded right the way up to high-performance computing (HPC) systems. Research on supercomputing is pushing towards the target of a 20MW Exascale supercomputer – that is, one capable of a billion of billions calculations per second – by 2021. Currently, the IBM Summit, equipped with IBM Power9 processors and NVIDIA VOLTA GV100 graphics processing units (GPUs), is the fastest computer in the world according to the **TOP500** list. This computer has reached 143.5 petaFLOPS (million billion floating operations per second) and is ranked third on the Green500 list, with 14.7 gigaFLOPS/W.

To reach a 20MW Exascale supercomputer would entail more than tripling the energy efficiency of today's supercomputers to around 50 gigaFLOPS/W. We believe that the next generation of supercomputers need a radically new software stack capable of exploiting the benefits offered by heterogeneity to meet the programmability, scalability and energy efficiency required by the Exascale.

What are the main innovations of ANTAREX?

In a nutshell, ANTAREX aims to provide a **breakthrough approach** for application self-adaptivity and to runtime manage, monitor and autotune applications for energy-efficient HPC systems up to the Exascale level.

One key innovation of our approach is the concept of *'separation of concerns'*. This enables the end-user application programmer to express adaptivity strategies and non-functional requirements, such as throughput and power constraints, in addition to application functionality. This was promoted by the design of the **LARA domain-specific language (DSL)**, inspired by aspect-oriented programming concepts for heterogeneous systems **[1].**

Leveraging the LARA DSL, ANTAREX introduced a compile-time framework composed of the **Clava** source-to-source compiler, the **LibVC** dynamic compiler **[2]**, the memoization, the custom precision and the loop splitting. The key feature of the ANTAREX DSL is the ability to separate the concerns of the end-user of the HPC infrastructure, who is generally an application domain expert and does not need to know about the extra-functional characteristics and their expression through the DSL, and the HPC expert, who on the other hand manipulates the application through the DSL itself adding its expertise to optimise the extra-functional properties. The ANTAREX DSL-based approach significantly save programming and performance tuning time and contribute to TTS (time-to-solution) reductions. Optimization steps are automatically supported by reusable strategies applied to the application and thus with a high potential to save programmers' and tuning experts' efforts. In the paper **[3],** we present the **lessons learnt** on how the key outcome of the project, the ANTAREX DSL and its associated toolbox support extra-functional requirements, as well as runtime adaptivity strategies and approximation techniques. To this end, we provide a number of examples of mechanisms and strategies encoded in the DSL, including some that demonstrate how the DSL is applied in the context of the project use cases.

In the context of ANTAREX, we also proposed the application of machine learning techniques to addressed the problem of compiler autotuning **[4]** and mitigating the compiler phase-ordering problem **[5]**. The **lesson learnt** was how a systematic approach based on machine learning techniques is needed to speedup the compile time autotuning and to achieve peak performance in HPC applications.

Another innovative technology is the ANTAREX runtime framework, based on the **mARGOt** application autotuner, to enable runtime tunable approximations at the application-level by trading off accuracy versus throughput. The **lesson learnt** was how to make HPC applications self-aware of their runtime behavior **[6].** Then, we enhanced the **mARGOt** dynamic autotuning framework with a module for the online learning of the application knowledge to drive the adaptation strategies during the production phase **[7].**

The ANTAREX runtime framework also includes the **Examon [8]** scalable, fine-grained energy-monitoring system and the **Countdown [9]** power manager, which were recently integrated into the CINECA HPC production environment. With a production application in a 3.5 K cores run, **Countdown** saves the 22.36% of energy consumption with a performance penalty of 2.88%. The **lesson learnt** was how to monitor and runtime manage the scalability and adaptability of a dynamic workload by exploiting the full system capabilities, including energy management, for emerging large-scale and extreme-scale systems.

What sort of challenges could we solve by exploiting the ANTAREX technologies?

ANTAREX addresses the autotuning and scalability characteristics of two important application scenarios in HPC: a biopharmaceutical application for accelerating drug discovery and an intelligent navigation system for traffic congestion mitigation in smart cities. These two use cases were chosen for their direct economic exploitability and their major social impact.

Scaling up and accelerating molecular docking

For the biopharmaceutical use case, we worked with the Dompé bio-pharmaceutical company, part of the ANTAREX consortium. The focus was on molecular docking, an increasingly important application for HPC-accelerated drug discovery. We started by analyzing the most computationally intensive kernels of the LiGen molecular docking application used by Dompé. We developed a runtime tunable version of the molecular docking application for use in virtual screening experiments in homogeneous **[10]** and heterogeneous nodes **[13[.** This was deployed and scaled out to the full size of the 10.4 petaFLOPS Tier-0 Marconi[1] supercomputer at CINECA to screen a database containing more than 1.2 billion ligands – biochemical substances which bind to biological molecules –with the aim of targeting unresolved infective diseases. This represents **the largest virtual screening experiment ever launched** in terms of computational threads (up to one million) and size of the compound database (one billion ligands).

Despite of system reliability issues, the **lesson learnt [D4.4]** was to demonstrate the feasibility of a molecular docking run on more than 1.2 billion of ligands by using a full-scale HPC system installed at CINECA and including 3600 nodes for a peak performance of 10 PFlop/s. Looking at future Exascale systems, besides the **lessons learnt** on scalability and performance tuning, we

[1] Marconi Intel Xeon Phi is currently ranked No. 19 in the last Top500 list, dated November 2018.

learned how fault tolerant mechanisms could play even a more relevant role, to exploit the full capability of the system. This is because Exascale systems will have much more nodes than the Marconi - KNL partition (1% of 1EFlop/s). We also believe this is a more general trend not depending on our specific application domain.

Participating in ANTAREX allowed Dompé to take advantage of HPC-accelerated and tunable solutions, thus envisioning new development paths not viable using conventional computing. Exploiting ANTAREX's HPC technologies supporting autotuning, scalability and energy efficiency, Dompé is now able to optimize molecular docking to reduce the virtual screening process for the identification of new active compounds by two orders of magnitude.

The **lesson learnt [D4.5]** is that with the autotuned application (GEODOCK) developed in ANTAREX and integrated in the LiGen code, we made it possible to explore the huge chemical space of almost all known molecules available to build drugs within one day of today supercomputer, a time that can even become less than one hour on future Exascale systems. From the point of view of the drug industry (but for the scientific research as well) this perspective open up the possibility to shorten the path from the discovery of a health threat, like the case of a sudden pandemic virus, to the prompt availability of candidate drugs. From the point of view of the European HPC, the possibility to deliver a result on huge chemical spaces in less than one hour, makes the future European Exascale systems of primary interest for the pharma industry.

From the results of experiments carried out on the Marconi system at CINECA, we are now able to estimate the amount of computational resources required for the virtual screening of the full set of ligands (in the order of one billion) over the set of protein pockets of a real pandemic virus.

Dompé virtual screening platform represents a unique asset and a competitive advantage versus leading pharmaceutical companies

The ultra-high performance enhancement of the Ligen thought the GEODOCK molecular placing tool, thanks to the fruitful cooperation among DOMPE, CINECA and POLIMI allows DOMPE having the world fastest structure based virtual screening tool. Comparing the Novartis virtual screening platform with the DOMPE platform, we reduced the time from **9 hours** (Novartis) to **10 seconds** (DOMPE) with the same number of cores. This result dramatically reduces the time and cost of the virtual screening process making feasible the screening of **billions** rather than **millions of compounds.** This makes virtual screening computationally and economically viable for all current drug discovery projects of DOMPE, our academic partners and our customers. Another important achievement obtained in ANTAREX is the capability to tune at runtime approximations and parameters. This concept enables the optimization of software parameters for each single molecule evaluated, that represents a huge optimization of time, performance and energy consumption. This approach represents a major improvement compared to the state-of-the-art where the optimization of the simulation parameters is based on a small set of molecules used in the validation process.

Moreover, reducing the conformational space sampling for low-complexity molecules enables the optimization of the calculations without affecting the quality of the results. Remarkably, this approach is unique in our platform compared to other commercial structure-based virtual screening tools, and it is one of the most relevant outcomes of the ANTAREX project for DOMPE.

Mapping traffic congestion in smart cities

Slovak navigation application company Sygic and Czech supercomputing center IT4Innovations worked on a smart navigation system **[D5.4]** to help mitigate traffic congestion in smart cities, using an innovative algorithm of road-balanced routing. The **lesson learnt** is how exploiting supercomputing power and ANTAREX code optimization technology, we reached the point of being able to calculate routes for tens of thousands of drivers simultaneously and perpetually towards a global optimum, i.e. providing traffic-flow-optimized navigation to reduce total travel time **[11] [D5.5].**

The project helped Sygic create solid foundations for a revolutionary product for municipalities, which is ready to be deployed in a pilot scheme immediately. Using the ANTAREX technologies supporting autotuning and scalability, Sygic is now able to adapt the product for a wide range of municipalities with their specific cost/performance requirements.

Sygic envisions pilot deployments in forward-thinking cities and to this end has already established contact with the municipalities of Milano, Ostrava, and Vienna. Thanks to reduced travel times and emissions, this system improves citizens' quality of life, hence we believe it will also be attractive to other cities. As the business assumption is that municipalities will bear the cost of the solution, providing the services to citizens for free, further improvements to the product will aim to improve the cost efficiency of the computation, exploiting ANTAREX optimization technologies to the full.

Sygic's Collaborative Routing Algorithm Has Become One of the Top European Innovations

The collaborative routing algorithm developed in ANTAREX thanks to the fruitful cooperation among SYGIC, IT4I and POLIMI is the key for the nomination of SYGIC among the 50 finalists for the European Commission's **"Innovation Radar Prize 2018"** in the category "Best early stage innovation". The nomination was given to the Adaptive Navigation System Software (https://www.innoradar.eu/innovation/26196).

The Innovation Radar is a European Commission initiative to identify high potential innovations in research funded by the European Union. The public is invited to vote for their favorite innovators to select the finalists from the six categories. The Collaborative Routing algorithm was nominated in the Best Early Stage Innovation, finished second in its category after public voting and qualified for the finals which took place at the ICT 2018 conference in Vienna. Each of the finalists had 3 minutes to present their innovation to the expert jury of investors and experienced entrepreneurs, which decided on the winner.

Sygic, the trusted navigation solutions architect for global businesses has become one of twenty of Europe's best innovators in Innovation Radar Prize. The expert jury at the ICT 2018 conference appreciated the solution's potential to reduce the time spent in the traffic by more than 20 percent, significantly lower costs and contribution to CO_2 reduction in cities. The expert jury has also appreciated the positive impact on the quality of life in cities.

The intelligent solution is leveraging the power of IT4Innovations supercomputer to calculate optimal route distribution for drivers in the city and save more than 20 percent of daily travel time in traffic.

„By orchestrating 500 thousand vehicles with our algorithm, we can save cities and drivers up to 400,000 euro a day," said project lead Radim Cmar from Sygic. *„Which significantly exceeds the monthly fee for using the supercomputer,"* he added.

Using ANTAREX technologies, Sygic was able to optimize critical calculations and achieved better energy efficiency with the same performance, which reflected into computation costs reduced by 25 percent.

„We are currently in talks with four cities – in Europe and in the US to start the pilot project," disclosed Cmar. The Collaborative routing algorithm is expecting to have a great influence on the quality of life in cities. The reduced number of cars stuck in traffic has the potential to cut down travel and maintenance costs of cars and public transport vehicles, lower air pollution and significantly improve commuting time.

"The fact that our project has received one of the highest counts of votes from the public, confirms how important the question of traffic for all of us in Europe is. We have all the know-how and experience to deploy the solution in any city in the World, which would show interest," said Sygic CEO Martin Strigac.

Sygic is a devoted supporter of Europe's plan to build 300 smart cities by the end of 2019. With its unique routing algorithms, developed in coordination with IT4Innovations, it can address the current biggest challenges of European cities related to air pollution and traffic. In Slovakia, the company co-founded I Want My City Smart initiative, which helps municipalities to integrate smart solutions. See also **[12].**

References

1. Pedro Pinto, Tiago Carvalho, João Bispo, Miguel António Ramalho, João M. P. Cardoso, “Aspect composition for multiple target languages using LARA”, Computer Languages Systems & Structures, Elsevier, 2018, Green Open Access: https://repositorio-aberto.up.pt/handle/10216/117169

2. Stefano Cherubin, Giovanni Agosta, “libVersioningCompiler: An easy-to-use library for dynamic generation and invocation of multiple code versions”, SofwareX, Volume 7, January-June 2018, pp. 95-100, ELSEVIER, Open Access: https://doi.org/10.1016/j.softx.2018.03.006

3. C. Silvano, G. Agosta, A. Bartolini, A. Beccari, L. Benini, L. Besnard, J. Bispo , R. Cmar, J. Cardoso, C. Cavazzoni, D. Cesarini, S. Cherubin, F. Ficarelli, D. Gadioli , M. Golasowski , A. Libri, J. Martinovic, G. Palermo, P. Pinto, K. Slaninova, E. Vitali “The ANTAREX Domain Specific Language for High Performance Computing” submitted in January 2019 to Microprocessors and Microsystems: Embedded Hardware Design (MICPRO), Elsevier: Open Access: http://arxiv.org/abs/1901.06175

4. Amir Hossein Ashouri, Giovanni Mariani, Gianluca Palermo, Eunjung Park, John Cavazos, Cristina Silvano: “COBAYN: Compiler Autotuning Framework Using Bayesian Networks”. ACM Transactions on Architecture and Code Optimization, Volume 13(2): 21:1-21:25 (2016), Green Open Access: https://re.public.polimi.it/retrieve/handle/11311/1026239/331071/COBAYN__green_open_access.pdf

5. Amir Hossein Ashouri, Andrea Bignoli, Gianluca Palermo, Cristina Silvano, Sameer Kulkarni, John Cavazos: “MiCOMP: Mitigating the Compiler Phase-Ordering Problem Using Optimization Sub-Sequences and Machine Learning”. ACM Transactions on Architecture and Code Optimization, Volume 14(3): 29:1-29:28 (2017), Green Open Access: https://re.public.polimi.it/retrieve/handle/11311/1026239/331071/COBAYN__green_open_access.pdf

6. Davide Gadioli, Emanuele Vitali, Gianluca Palermo, Cristina Silvano, "mARGOt: a Dynamic Autotuning Framework for Self-aware Approximate Computing", **IEEE Transactions on Computers,** Accepted for Publications, November 2018. Green Open Access: http://home.deib.polimi.it/silvano/Paperi_IEEE/mARGOt_GreenOpenAccess.pdf

7. Tomas Martinovic, Davide Gadioli, Gianluca Palermo, Cristina Silvano, “On-line Application Autotuning Exploiting Ensemble Models”, submitted in Dec. 2018 to IEEE Transactions on Computers. Open Access: http://arxiv.org/abs/1901.06228

8. Francesco Beneventi, Andrea Bartolini, Carlo Cavazzoni, Luca Benini, “Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools”, Design, Automation Test in Europe Conference Exhibition (DATE), 2017, Green Open Access: https://www.research-collection.ethz.ch/handle/20.500.11850/192078

9. Daniele Cesarini, Andrea Bartolini, Pietro Bonfà, Carlo Cavazzoni, Luca Benini, “COUNTDOWN - three, two, one, low power! A Run-time Library for Energy Saving in MPI Communication Primitives”, submitted in July 2018 to IEEE Transactions on Parallel and Distributed Systems, Open Access: https://arxiv.org/abs/1806.07258

10. "Tunable Approximations to Control Time-to-Solution in an HPC Molecular Docking Mini-App", submitted in December 2018 to IEEE Transactions on Computers; Open Access: http://arxiv.org/abs/1901.06363

11. Emanuele Vitali, Davide Gadioli, Gianluca Palermo, Martin Golasowski, Joao Bispo, Pedro Pinto, Jan Martinovic, Katerina Slaninova, Joao M. P. Cardoso, Cristina Silvano, "An Efficient Monte Carlo-based Probabilistic Time-Dependent Routing Calculation Targeting a Server-Side Car Navigation System", submitted in July 2018 to IEEE Transactions on Emerging Topics in Computing; Open Access: http://arxiv.org/abs/1901.06210

12. https://www.sygic.com/press/sygics-collaborative-routing-algorithm-has-become-one-of-the-top-european-innovations

13. Emanuele Vitali, Davide Gadioli, Gianluca Palermo, Andrea Beccari, Carlo Cavazzoni, Cristina Silvano, "Exploiting OpenMP & OpenACC to Accelerate a Molecular Docking Mini-App", Submitted to Journal of Supercomputing, Springer. Open Access: http://arxiv.org/abs/1901.06229