



Performance Optimisation and Productivity

A Centre of Excellence in HPC



POP Newsletter 9 – Issue December 2018

Welcome to the ninth newsletter from the EU [POP](#) Centre of Excellence. This is the first newsletter for the second phase of POP as funding for this phase started on 1st December 2018. For new requests, please see section Free Code Optimisation Help at the bottom of this newsletter.

This issue includes:

- POP Project Restarted 1st December 2018;
- POP Partner Profiles: The Performance Tools team at UVSQ and IT4Innovations at VSB-TUO;
- POP Performance Analysis Tool: MAQAO;
- Apply for free help with code optimisation;
- The POP Helpdesk;

For information on our services and past editions of the newsletter see the [POP website](#).

POP Project Restarted 1st December 2018

After a very successful first phase of the POP project from October 2015 to March 2018, where we performed over 160 performance audit, performance plan, and proof-of-concept services for our customers, the project secured funding for a second 3-year phase starting 1st December 2018. For our past and potential future customers nothing major changed: We still will provide free performance optimisation and productivity services for academic and industrial codes in all domains! The services are still free of charge to organisations / SMEs / ISVs / companies in the EU! We will also continue our successful training and tuning workshops programme including our performance analysis webinars.

To improve the service, the second phase of the project includes a few minor changes and additions which are listed below.

New Service Structure

In the first phase of the project we offered three types of services: Performance Audit, Performance Plan and Proof-of-Concept (PoC). The performance audits and plans represented different depths in the analysis of observed behaviour and prediction of the impacts that would result from proposed refactoring or architectural retargeting. Initially, Performance Audits should include simpler and more generic analyses and take less effort from the POP analyst than Performance Plans which should be very deep and targeted to analysing specific issues and potential improvements. Proof-of-Concept services aim at demonstrating to the customer on a subset of their code, mini-apps or kernels extracted from them, how our proposed techniques should be applied and measure the gain obtained. Experience in the first phase has shown the blurred boundary between Performance Audits and Plans. Based on this experience and also from our users' feedback, we decided to integrate Performance Audits and Plans into one new service called **Performance Assessment**.

Internally we will consider the audit as the identification of the region to focus the analysis and the measurement of the efficiencies. These results will be reported to the user as a first initial feedback unless the results indicate there is no performance issue or the user declines to continue the study. Otherwise, the next step will be a deeper analysis of the components that degrade the performance providing concrete recommendations for each of them.

New Project Partners

For the second phase, teams from the Performance Tools team at UVSQ in France and IT4Innovations at VSB-TUO in the Czech Republic will join our group of performance experts. The other members are Barcelona Supercomputing Center, High Performance Computing Center Stuttgart, Jülich Supercomputing Centre, Numerical Algorithms Group, Rheinisch-Westfälische Technische Hochschule Aachen, and Teratec. See below for more details on our new project partners.

New Co-design Data Repository

We will create a co-design data repository which includes statistics about common performance issues of HPC applications as well as micro-kernels extracted from real applications each characterising fundamental performance behaviour. Hardware architects or system software designers from other EU projects will be able to get quantitative information of how to estimate the potential impact of an architectural or system software approach they may be developing. The micro-kernels will also be useful within POP (but also outside) as training material examples, and for demonstrating benefits of the programming model features and practices that POP promotes in dissemination activities.

POP Partner Profiles

The POP CoE consists of [8 partners](#) and in this newsletter, we take a closer look at the two new partners who joined the project for its second phase (December 2018 to November 2021) : The performance tools (PT) team at UVSQ ([LI-PaRAD Laboratory](#)) and [IT4Innovations](#) at VSB-TUO.

The Performance Tools team at UVSQ

The performance tools (PT) team at [University Versailles / Saint-Quentin-en-Yvelines](#), France ([LI-PaRAD Laboratory](#)) has been working on performance optimization tools and methodologies since 2004 starting with Itanium based clusters. UVSQ started working on real applications within a joint Laboratory (LRC-ITACA) with the French Department of Energy (CEA DAM). Today UVSQ is a partner of the [Excascale Computing Research Laboratory](#) together with Intel and CEA since 2009.



UVSQ pursues the following goals with respect to HPC:

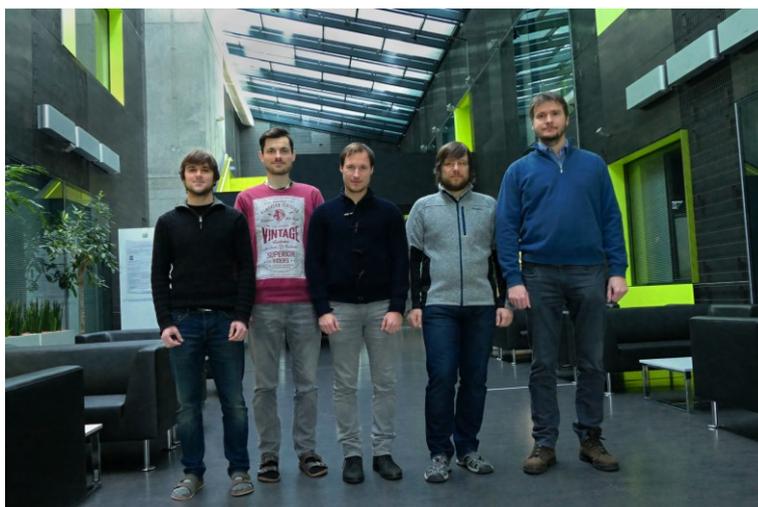
- Develop performance evaluation tools (MAQAO) targeting both core/node level and parallel issues
- Develop program behaviour analysis tool (performance, energy, numerical accuracy) allowing to capture interaction between hardware and HPC applications
- Develop a methodology and a User experience (UX) to guide non-expert users to restructure their code to improve performance
- Develop partnerships with industrial and academic partners to test and improve our MAQAO Toolset
- Educate tomorrow's HPC experts through a specialized HPC Master programme
- Take part in French and International HPC projects
- Training: Help users analyzing and optimizing their codes in Workshops / Events

POP services are provided by the Performance Tools team based on a strong experience acquired through success stories involving leading Industrial companies.

IT4Innovations at VSB-TUO

[IT4Innovations National Supercomputing Center](#), a university institute at VSB – Technical University of Ostrava, assures the operability of the most powerful supercomputing technology in the Czech Republic. With respect to its uniqueness and significance, it ranks highest among large national research infrastructures of the Czech Republic. It is currently operating two powerful supercomputers – Anselm and Salomon, the latter of which still ranks among the top European supercomputers. At the beginning of the year 2019, a new system with an approximate theoretical peak performance of 800 Tflop/s is planned to support our current supercomputers. This new system is to be equipped with the latest available technology including processors with the AVX-512 instruction set, Nvidia Tesla V100 accelerators, a fat compute node with a memory of up to 6 TB, 200 Gb/s interconnect, NVMe memory-based storage, and BurstBuffer technology for accelerating data access.

IT4Innovations is also a research and development centre with strong international links and as such, it is currently involved in good number of international projects funded primarily by the Horizon 2020 programme. Since its foundation in 2011, IT4Innovations has been a member of the prestigious pan-European PRACE (Partnership for Advanced Computing in Europe) research infrastructure, where it represents the Czech Republic. Since 2016, it has also been involved in the European Technology Platform for High-Performance Computing (ETP4HPC), which focuses on defining technology and research priorities in the field of high-performance computing in Europe. Since the beginning of the year 2018, IT4Innovations has also been participating in the preparation of the EuroHPC joint undertaking.



Its activities are also focused on supporting the deployment of computationally intensive numerical simulations and advanced data analysis primarily in small and medium-sized enterprises. IT4Innovations is registered by the European Commission as a Digital Innovation Hub in the fields of HPC, artificial intelligence, and advanced data analysis, see <http://s3platform.jrc.ec.europa.eu>. Moreover, we have become an important centre for education in HPC. We are proud to be one of the PRACE Training Centres, offering a comprehensive training programme. We have also established a brand new Computational Sciences PhD study programme.

Within the POP project the IT4I team is mainly involved in Performance Assessment and Proof of Concept services where it brings the experience in development, performance analysis, and optimization of the HPC applications achieved from EXA2CT (H2020: Exascale Algorithms and Advanced Computational Techniques) and IPCC (Intel Parallel Computing Center) projects. For the new Co-design activity, IT4I will participate in kernel extraction and definition. IT4I will also carry out promotion activities to find out potential users for the services especially in the central and eastern part of Europe in both academic and industrial environments. Finally, IT4I will participate in POP Dissemination and Training activities using its experience in organizing various training events and workshops.

POP's Performance Analysis Tools

The POP CoE uses various [performance tools](#) developed by project partners for the performance optimisation services offered by the project. Here we introduce the MAQAO tool set from our new project partner USQV.

UVSQ Performance Tools

The performance tools (PT) team at UVSQ ([LI-PaRAD Laboratory](#)) has been working on performance optimization tools and methodologies since 2004 starting with Itanium based clusters.

[MAQAO](#) (Modular Assembly Quality Analyzer and Optimizer) is a performance analysis and optimisation toolsuite operating at binary level. It can run single and multi-node parallel applications. When compared to other tools It provides an advanced focus on core/node performance.

MAQAO
Global
Application
Functions
Loops
Topology

Help is available by moving the cursor above any ? symbol or by checking [MAQAO website](#).

Experiment Summary		Configuration Summary	
Application	/home/mtribalat/YALES2/3d_cylinder/3D_cylinder	Dataset	/home/mtribalat/YALES2/3d_cylinder/dataset
Timestamp	2018-12-11 14:04:26	Run Command	<binary>
Experiment Type	MPI	MPI Command	mpirun -n <number_processes>
Machine	localhost.localdomain	Number Processes	4
Architecture	x86_64	Filter	(type = number ; value = 10 ;)
Micro Architecture	HASWELL		
Model Name	Intel(R) Core(TM) i7-4770 CPU @ 3.40GHz		
Cache Size	8192 KB		
Number of Cores	4		
Compilation Options	binary: Intel 15.0.1.133 -I/home/eosere/toto1/hdIS-1.8.8/hdl5_intel/include -I/opt/intel/15.0.1/mkl/include/ftw -I. -I/opt/apps/YALES2/yaes2/lib_optim -I/opt/intel/mpi5.0.1.035/intel64/include/gfortran -I/opt/intel/mpi5.0.1.035/intel64/include -fpp -DFFTW3 -DHDF5 -DLAPACK -DCVODE -DSLATEC -module /opt/apps/YALES2/yaes2/lib_optim -g -O2 -xCORE-AVX2 -DFORTRAN2008 -c -o /opt/apps/YALES2/yaes2/lib_optim/linear_solver_mat_op_m.o		
Number of processes observed	4		
Number of threads observed	4		

Global Metrics		CQA Potential Speedups Summary																																													
Total Time (s)	41	<table border="1" style="font-size: 0.7em; margin-top: 5px;"> <caption>CQA Potential Speedups Summary Data</caption> <thead> <tr> <th>Process Count</th> <th>Clean</th> <th>FP Vectorized</th> <th>Fully Vectorized</th> </tr> </thead> <tbody> <tr><td>1</td><td>1.00</td><td>1.00</td><td>1.00</td></tr> <tr><td>2</td><td>1.00</td><td>1.03</td><td>1.08</td></tr> <tr><td>3</td><td>1.00</td><td>1.04</td><td>1.11</td></tr> <tr><td>4</td><td>1.00</td><td>1.05</td><td>1.13</td></tr> <tr><td>5</td><td>1.00</td><td>1.05</td><td>1.15</td></tr> <tr><td>6</td><td>1.00</td><td>1.06</td><td>1.17</td></tr> <tr><td>7</td><td>1.00</td><td>1.06</td><td>1.18</td></tr> <tr><td>8</td><td>1.00</td><td>1.06</td><td>1.19</td></tr> <tr><td>9</td><td>1.00</td><td>1.06</td><td>1.20</td></tr> <tr><td>10</td><td>1.00</td><td>1.06</td><td>1.20</td></tr> </tbody> </table>		Process Count	Clean	FP Vectorized	Fully Vectorized	1	1.00	1.00	1.00	2	1.00	1.03	1.08	3	1.00	1.04	1.11	4	1.00	1.05	1.13	5	1.00	1.05	1.15	6	1.00	1.06	1.17	7	1.00	1.06	1.18	8	1.00	1.06	1.19	9	1.00	1.06	1.20	10	1.00	1.06	1.20
Process Count	Clean			FP Vectorized	Fully Vectorized																																										
1	1.00			1.00	1.00																																										
2	1.00			1.03	1.08																																										
3	1.00			1.04	1.11																																										
4	1.00			1.05	1.13																																										
5	1.00			1.05	1.15																																										
6	1.00			1.06	1.17																																										
7	1.00			1.06	1.18																																										
8	1.00			1.06	1.19																																										
9	1.00	1.06	1.20																																												
10	1.00	1.06	1.20																																												
Compilation Options	OK																																														
Flow Complexity	1.00																																														
Array Access Efficiency (%)	26.06																																														
Clean	Potential Speedup: 1.01																																														
	Nb Loops to get 80%: 4																																														
FP Vectorized	Potential Speedup: 1.06																																														
	Nb Loops to get 80%: 5																																														
Fully Vectorized	Potential Speedup: 1.20																																														
	Nb Loops to get 80%: 7																																														
Data in L1 Cache	Potential Speedup: 1.11																																														
	Nb Loops to get 80%: 5																																														

Its main goal is to guide application developers along the optimization process through synthetic reports and hints. The tool mixes both dynamic and static analyses based on its ability to reconstruct high level structures such as functions and loops from an application binary. Since MAQAO operates at binary level, it is agnostic regarding the language used in the source code and does not require recompiling the application to perform analyses. Another key feature of MAQAO is its extensibility. Users can easily write their own modules using the Lua scripting language, allowing fast prototyping of new MAQAO modules. MAQAO has also been designed to concurrently support multiple architectures. At the moment, the Intel64, Xeon Phi and ARM architectures are implemented.

The main modules of MAQAO are:

- ONE View, a supervising module responsible for invoking other modules and aggregating their results, which are then presented as synthetic reports in HTML or XLS format.
- LProf (Lightweight Profiler), a sampling-based lightweight profiler offering results at both function and loop levels and capable of categorizing its results depending on their source. LProf is agnostic with regard to the MPI or OpenMP runtime used by the application.
- CQA (Code Quality Analyzer), a static analyser assessing the quality of the code generated by the compiler and producing a set of reports describing potential issues, estimations of the gain if fixed, and hints on how to achieve this through modifications to the source code or the compilation chain. Since it relies on static analysis, CQA can also provide projections of code performance on different architectures (cross-analysis).
- VProf is a value profiler relying on instrumentation through binary rewriting. For example, it can be used to retrieve the number of iterations and execution time of loops and grouping the results by loop instances with a similar behaviour. It can also be used to gather typical parameter values for subroutine arguments.
- DECAN (DECremental ANalyzer), is an advanced MAQAO module using differential analysis on innermost loops to locate performance issues. To achieve this, DECAN generates several variants of an assembly loop through binary rewriting, by removing or modifying groups of instructions and adding probes to measure time or read hardware counters for the modified loop. The comparison of the results obtained after executing the different variants generated then allows to infer impact of the modified instructions on the overall loop performance. This also allows DECAN to predict the application behaviour if certain conditions are met (e.g. all data in L1 cache) or some transformations have been performed.
- ASSIST is a prototype code restructuring tool implementing advanced Profile Guided Optimization techniques: it first uses more refined performance analysis provided by MAQAO modules and second it performs complex code restructuring such as specialization.

Apply for free help with code optimisation

We offer a range of [free services](#) designed to help EU organisations improve the performance of parallel software. If you're not getting the performance you need from parallel software, please apply for help via the short [Service Request Form](#), or [email us](#) to discuss further.

These services are funded by the European Union Horizon 2020 research and innovation programme - there's no direct cost to our users! If you're interested in our services, please contact us soon to express an interest.

The POP Helpdesk

Past and present POP users are eligible to use our [email helpdesk \(pop-helpdesk@bsc.es\)](mailto:pop-helpdesk@bsc.es). Please contact our team of experts for help analysing code changes, to discuss your next steps, and to ask questions about your parallel performance optimisation.



<https://pop-coe.eu>



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 676553 and 824080.

