



Performance Optimization and Productivity

EU H2020 Center of Excellence (CoE)



1 October 2015 – 31 March 2018 (30 months)



- **A Center of Excellence**
 - On **Performance Optimization and Productivity**
 - Promoting **best practices in performance analysis and parallel programming**
- **Providing Services**
 - Precise understanding of application and system behavior
 - Suggestion/support on how to refactor code in the most productive way
- **Horizontal**
 - Transversal across application areas, platforms, scales
- **For academic AND industrial codes and users !**



Partners



• Who?

- BSC (coordinator), ES
- HLRS, DE
- JSC, DE
- NAG, UK
- RWTH Aachen, IT Center, DE
- TERATEC, FR



A team with

- Excellence in performance tools and tuning
- Excellence in programming models and practices
- Research and development background AND proven commitment in application to real academic and industrial use cases



Motivation



Why?

- Complexity of machines and codes
 - Frequent lack of quantified understanding of actual behavior
 - Not clear most productive direction of code refactoring
- Important to maximize efficiency (performance, power) of compute intensive applications and the productivity of the development efforts

Target

- Parallel programs , mainly MPI /OpenMP ... although can also look at CUDA, OpenCL, Python, ...



3 levels of services



? Application Performance Audit

- Primary service
- Identify performance issues of customer code (at customer site)
- Small Effort (< 1 month)

! Application Performance Plan

- Follow-up on the service
- Identifies the root causes of the issues found and qualifies and quantifies approaches to address the issues
- Longer effort (1-3 months)

✓ Proof-of-Concept

- Experiments and mock-up tests for customer codes
- Kernel extraction, parallelization, mini-apps experiments to show effect of proposed optimizations
- 6 months effort

Reports

Software
demonstrator

Apply @
<http://www.pop-coe.eu>

The screenshot shows a web browser displaying the 'Request Service Form' page of the Performance Optimisation and Productivity (POP) Centre of Excellence. The page features a navigation menu on the left with links to Blog, News, Partners, Services, Request Service Form (highlighted), Target Customers, Further Information, and Contact. The main content area is titled 'Request Service Form' and contains a 'Contact Details' section with fields for 'Applicant's Name', 'Institution', and 'e-mail'. Below this is a 'Code' section with a 'Name of the code' field and a dropdown menu for 'Scientific/technical area and class of problems it solves'. At the bottom, there is a 'Contribution' section with radio buttons for 'Core developer', 'Module developer', and 'User'.



Target customers



- **Code developers**

- Assessment of detailed actual behavior
- Suggestion of more productive directions to refactor code

- **Users**

- Assessment of achieved performance on specific production conditions
- Possible improvements modifying environment setup
- Evidences to interact with code provider

- **Infrastructure operators**

- Assessment of achieved performance in production conditions
- Possible improvements modifying environment setup
- Information for allocation processes
- Training of support staff

- **Vendors**

- Benchmarking
- Customer support
- System dimensioning/design



Best practices in Performance analysis



- **Powerful tools ...**

- Extrae + Paraver
- Score-P + Scalasca/TAU/Vampir + Cube
- Dimemas, Extra-P
- Other commercial tools

- **... and techniques**

- Clustering, modeling, projection, extrapolation, memory access patterns, ...
- ... with extreme detail ...
- ... and up to extreme scale

- **Unify methodologies**

- Structure
 - Spatio temporal / syntactic
- Metrics
 - Parallel fundamental factors: Efficiency, Load balance, Serialization
 - Programming model related metrics
 - User level code sequential performance
- Hierarchical search
 - From high level fundamental behavior to its causes

- **To deliver insight**

- **To estimate potentials**



Best practices in parallel programming



- **MPI and OpenMP**

- Active members of OpenMP consortium (RWTH, BSC)
- Active members of MPI Forum (JSC, RWTH)

- **Pushing application as early adopters and co-design drivers**

- **Promoting new features ...**

- **... gathering feedback**

- **Promoting a throughput oriented methodology**

- Task based programming
- Asynchrony, overlap
- Locality
- Malleability, Dynamic Load Balancing
- Nesting, recursion



Activities



• External access

- WEB (www.pop-coe.eu)
 - Request form
 - Feedback questionnaires
 - News and blog

• Internal organization

- CRM
- TRAC ticketing system
- Wiki

The collage consists of four overlapping screenshots:

- Top Left:** The POP website homepage. It features the logo and the text "Performance Optimisation and Productivity A Centre of Excellence in Computing Applications". A sidebar on the left lists navigation options like Blog, News, Partners, Services, Request Service Form, Target Customers, Further information, and Contact. The main content area has a "Mission" section and a "News" section.
- Top Right:** A screenshot of the TRAC ticketing system. It shows a search results page for "(14) Work in progress (13 matches)". A table lists tickets with columns for Ticket ID, Summary, Component, Owner, Status, Start date, POP State, and % Complete.
- Bottom Left:** A screenshot of the SuteCRM interface. It shows a search for POPUsers with a list of results including names, organizations, and primary contact information.
- Bottom Right:** A screenshot of the POP Wiki Home page. It includes a welcome message, a "Project News" section with a list of recent updates, and a "WORK PACKAGES" section.



Activities



• Services

- Completed/reporting: 15
- Codes being analyzed: 6
- Waiting user input: 13
- Cancelled: 1

• By type

- Audits: 31
- Plan: 2
- Proof of concept: 2

• Reports

- 5 -15 pages

OpenNN performance assessment report

Document Information

Reference Number	POP_AM_13
Author	Taru Oksanen (BSC)
Coauthor(s)	Jari Lehto (BSC)
Date	March 29 th , 2014

Notes:

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No "101015710".

© 2015 POP Consortium Partners. All rights reserved.

Figure 1: Run on 41 CYBOLA... (text partially obscured)

Figure 2: Run on 22 JURECA... (text partially obscured)

Figure 2: Speedup achieved with respect to threads (left) and cores (right)

The analysis of the traces reports the efficiency numbers in Table 2. The table indicates that the parallel efficiency is very good for the different exercises with a small variability strongly related with the global load balance. The analysis identifies that the main reason that it is limiting the speedup is the efficiency achieved executing the computation (computation efficiency).

	2	4	8
Parallel Efficiency	0.9887	0.9338	0.9113
Load Balance	0.9951	0.9340	0.9393
Synchronization Efficiency	0.9998	0.9998	0.9914
Computation Efficiency	1.0	0.917	0.833
Global Efficiency	0.9947	0.7423	0.5395

Table 2: Time efficiencies observed in the parallel region

The computation efficiency is determined by the number of instructions and the instructions per cycle (IPC) values efficiencies are detailed in Table 3.

	2	4	8
IPC Scaling Efficiency	1.000	0.9461	0.794
Instruction Scaling Efficiency	1.000	0.873	1.126

Table 3: Other efficiencies

4. Scalability

Figure 3 highlights the scalability of the code on 256 processors on the left and their 4 processors on the right. As a perfectly linear scaling example, processes doubled, the total execution time is halved (see the top chart at the right side of Fig. 3). In the same figure is lower a further improvement in global speedup improvement is not necessary lower than the performance issue initially reported by

Figure 3: Scalability of PCA. Timeline of computational regions and speedup chart.

5. Efficiency

Table 1 and Table 2 show metrics for fundamental factors and efficiencies from the PCA of the executions using 16 to 256 MPI processes. Values are in percentage with higher values being better.

The observed global efficiency of the application decreases steadily from 92.19% at 16 processes to 70.43% at 256 processes, with an additional drop from 128 to 256 processes. The decreasing global efficiency is mainly caused by a decreasing load balance and decreasing computation efficiency, i.e. an increasing amount of time (accounted over all processes) is spent in computation for higher process counts.

The communication efficiency, however, is fairly constant and overall in a high range. Load balance is discussed in more detail in Section 6. The decreasing computation efficiency is also influenced by a decreasing number of instructions performed per cycle (IPC), which declines to 64.4% at 256 processes.

Serial Performance

- Evolution of IPC when scaling from 16 to 256 cores
- Tending to lower IPC for higher scales
- In addition, higher dispersion

Application Structure and Focus of Analysis

- Initial Audit: Parallel efficiency drops for more than 200 cores
- Analysis for 16 to 256 cores
- Truncated to the first 50 iterations, i.e. 2.55s out of 20,000s



Other activities



- **Promotion and dissemination**

- Market and community development
- Dissemination material and events

- **Customer advocacy**

- Gather customers feedback, ensure satisfaction, steer activities

- **Sustainability**

- Explore business models

- **Training**

- Best practices on the use of the tools and programming models (MPI + OpenMP)
 - Lot of interest ... customers want to learn how to do it themselves



Answer to Questions



- **Presented what we offer, what we are doing**
- **Requests by EC**
 - Women participation: BSC: 4/6, HLRS: 1/3, RWTH: 1/3, NAG: 1/3, JSC: 0/3
 - Interaction between CoEs: Training EoCoE, Events (EsiWACE, ...). Assessments to other CoEs
- **International cooperation**
 - we do have many activities (JLESC, VI-HPS,...) as individual partners.
 - Not at project level
- **PRACE scientific case, SRA, other FETHPC projects**
 - Involved on SRA at individual partners level
 - Have customers from other CoE and FETHPC projects. Our tools technologies are also used and partially developed in them



Conclusion



- We have established our internal operation infrastructure and procedures
- We have already performed 15 assessments and 19 are in the pipeline
- We believe the effort to unify our methodologies can become the core of best practices in performance analysis and programming practices that expand at international level
- We consider the POP CoE is progressing at a fairly good pace and results are already showing up.





www.pop-coe.eu

THANKS

