



# Performance Optimisation and Productivity

A Centre of Excellence in Computing Applications



## POP Newsletter 3 – Issue December 2016

Welcome to the third newsletter from the [POP \(Performance Optimisation and Productivity\) CoE \(Centre of Excellence\)](#). We're an EU funded project and offer free of charge support to EU organisations to help improve performance of parallel software.

This issue includes:

- Details of a new training workshop in Southampton, UK.
- Profiles of two POP partners
  - The High Performance Computing Center Stuttgart (HLRS)
  - TERATEC
- Recent analysis highlights
- A set of standard metrics for parallel performance analysis

Past editions of our newsletter are available via the [POP website](#) which also includes a description of all our services plus further information about the project.

---

## Apply for free help with code optimisation

We offer a range of [free services](#) designed to help EU organisations improve 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!

---

## 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.

# @POP\_HPC Twitter Account

[Follow us on Twitter](#) for news, info and interesting content on performance optimisation and parallel computing. Share your experiences with us as well, we're always interested to learn and discuss different approaches and tools.

---

## A set of standard metrics for parallel performance analysis

Attempting to optimise performance of a parallel code can be a daunting task, and often it is difficult to know where to start. For example, we might ask if the way computational work is divided is a problem? Or perhaps the chosen communication scheme is inefficient? Or does something else impact performance? To help address this issue, POP has defined a methodology for analysis of parallel codes, to provide a quantitative way of measuring the relative impact of the different factors inherent in parallelisation.

A feature of the methodology is that it uses a hierarchy of metrics, each metric reflecting a common cause of inefficiency in parallel programs. These metrics then allow comparison of parallel performance (e.g. over a range of thread/process counts, across different machines, or at different stages of optimisation and tuning) to identify which characteristics of the code contribute to inefficiency.

<ADD LINK HERE> [Click here to read](#) <ADD LINK HERE> our article that introduces the metrics, explains their meaning, and provides insight into the thinking behind them.


---

## Recent analysis highlights

### Communication inefficiency in Hybrid MPI + OpenMP

A popular Fortran CFD code using Hybrid MPI + OpenMP was investigated in a POP Performance Audit for varying numbers of MPI processes. The performance analysis using Extrae & Paraver identified specific regions of execution with high parallel inefficiency, and also identified that data transfer was a significant performance bottleneck.

A communications analysis (see adjacent matrix) identified imbalances in the total number of MPI point-to-point calls and in the amount of data transferred, and also found potential endpoint contention issues. Various recommendations for code refactoring were made, these issues and potential solutions were suggested for further investigation as a POP Performance Plan or Proof of Concept study.



### Load imbalance due to idle OpenMP threads

A Linear-Scaling Density Functional Theory code, using Fortran MPI + OpenMP, was investigated to find potential improvements using Score-P and Scalasca. These tools

allowed easy identification of the specific routines that exhibit poor parallel performance. Our analysis identified a large load imbalance; although individual OpenMP regions were well balanced there were many of them, with large regions between where only the master thread was active. This resulted in large CPU idle times and large OpenMP overheads. It was also identified that time in useful computation reduced from 45% to 15% for the largest core count studied.

Our recommendations included combining small parallel regions to reduce overheads and make more efficient use of cores, plus further in-depth investigation of the routines identified as problematic.

### Communication bottleneck

A C + MPI code was investigated in a POP Performance Audit which identified poor parallel scaling. Our analysis showed the cause was primarily communication issues; in particular, poor load balance (caused by load imbalance in total useful instructions per process) and low values for transfer efficiency (i.e. a large percentage of time transferring data).

Interestingly the audit also identified that although computational scalability was good (i.e. total useful computation is largely independent of process count) the number of instructions does increase significantly. However, this instructions increase is masked by a corresponding increase in IPC (instructions per cycle). As the number of processes increases further this increase in IPC will reach its upper limit, at this point it is likely computation scalability will reduce and significantly impact performance.

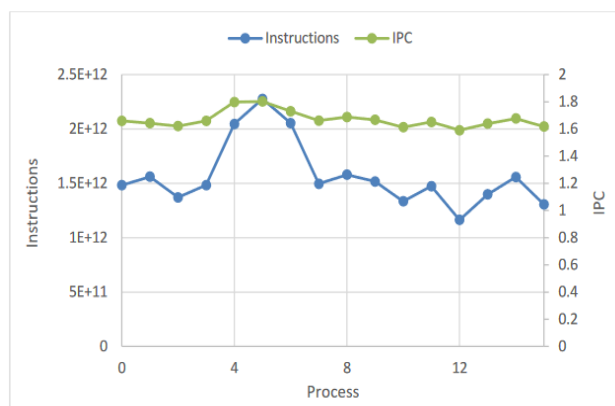
### Computational inefficiencies

A commercial Fortran + MPI code was investigated using Extrae & Paraver, to identify causes of a known parallel inefficiency.

Significant computational imbalance was identified, the underlying cause being imbalance in the number of instructions per process (see figure) which results in process idle times.

Increased time in computation for increasing process count was also identified, caused by decreasing IPC. This is being investigated in a POP Performance Plan.

This work nicely demonstrates the power of these tools in helping identify exact causes of parallel inefficiencies.



---

## POP Training & Learning

We offer our Performance Analysis Workshop to developers of parallel applications; our next training session is 8-10<sup>th</sup> February 2017 in Southampton, UK. This covers concepts and principles of profiling parallel applications and gives a practical introduction to the Scalasca and Paraver open source performance analysis tools. There will be opportunities for attendees to use these tools on their own codes.

Click [here](#) for details of registration.

Future training will be announced on [our website](#). For further information on POP training please [email us](#).

In addition, a range of learning material is available on the [POP website](#).

---

## Meet POP at upcoming events

### POP @ PRACE Days 2017

[PRACE Days 2017](#) are in Barcelona between 16-18<sup>th</sup> May. This forum is for domain experts from industry and research to discuss how they use HPC to further their research.

### POP @ ISC 2017

The [International Supercomputing Conference \(ISC\)](#) is 18-22<sup>nd</sup> June in Frankfurt. ISC is a major HPC conference held annually in Germany, and is attended by users and vendors who showcase the latest technologies and developments in HPC. The conference also hosts many presentations on novel uses of HPC to advance scientific and industrial productivity.

### POP @ Teratec Forum 2017

The [Teratec Forum](#) is a two-day event dedicated to HPC on the 27-28<sup>th</sup> June in Paris. It includes keynotes, technical sessions and an exhibition, with a strong emphasis on industrial usage of HPC.

---

## POP Partner Profiles

The POP project comprises [6 partners](#); in past newsletters we introduced The Numerical Algorithms Group, the HPC Group at RWTH Aachen, the Barcelona Supercomputer Center, and the Jülich Supercomputing Centre.

We now introduce the final two partners: the High Performance Computing Center Stuttgart (HLRS) and TERATEC.


### The SPMT Group of the High Performance Computing Center Stuttgart (HLRS)

The origins of the [High Performance Computing Center \(HLRS\)](#) go back to the computing centre at the University of Stuttgart, which by 1986 already had an operating Cray -2 system. HLRS as known today was officially established in 1996, it was the first national German High Performance Computing (HPC) centre, and in 2007 HLRS was one of the founding members of the [German Gauss Center for Supercomputing](#).

Today HLRS offers HPC resources and services via a wide range of partnerships with national and European research and industry. For some time, the supercomputing systems operated by HLRS have been listed amongst the top systems in the world, ranked within the top 20 in the TOP 500 list.

Alongside its strong background in the field of engineering, the nearly 100 experts working at HLRS, together with their contacts, cover every aspect of supercomputing and simulation. HLRS is also building new connections to important topics like sustainability, the impact of simulation on science, technological development, and the social handling of simulation results.

POP services are provided by the Scalable Programming Models and Tools (SPMT) group at HLRS. The SPMT group undertakes research into parallel programming models and tools to assist development for parallel programming languages. Additionally, the group maintains part of the software stack on the various HLRS platforms as a service to HLRS users.



The group has an excellent history as a partner within various national and EU funded projects. These demonstrate its European and international HPC expertise in exascale application development and hardware (CRESTA, MONT-BLANC) as well as in new programming models and tools (TEXT, DASH). Furthermore, the SPMT team develops the [Temanejo](#) debugger (for all task based programming models), maintains test suites for the MPI and OpenMP standards, and organises (with ZIH Dresden) the annual International Parallel Tools Workshop.

This experience and expertise means SPMT is perfectly placed to help customers identify performance issues in their applications and to identify solutions to overcome them.

## **TERATEC - a European Industrial Initiative for HPC, Big Data, and Simulation**

[TERATEC](#) was founded in 2005 as the result of a joint initiative of the French CEA (The Alternative Energies and Atomic Energy Commission) along with several industrial partners to form a European competence centre for numerical simulation and big data.



Its mission is to federate all industrial and academic HPC players (whether providers or users) and to offer access to the most powerful systems, to promote and increase attractiveness, and to foster economic development in this field.

Thereby TERATEC gathers leading [HPC, Simulation and Big Data players](#) from industry and research with the ambition to contribute to the development and use of technologies in these areas, to accelerate the design and implementation of the most powerful systems, and to



foster the emergence of new technologies and associated tools. Thus TERATEC creates expertise and highly qualified jobs in these areas.

**The TERATEC mission includes:**

**Mastering Technology** - TERATEC participates actively in initiatives to improve industrial mastery in the numerical simulation and HPC sector, which is crucial to keeping companies competitive and innovative. One such initiative is [ETP4HPC](#), a European platform to accelerate the development of European technologies at every step of the HPC chain, these technologies now being essential to almost all industry sectors.

**Industrial Research** - TERATEC also helps set up and promote [French and European research projects](#) involving industrial companies, technology suppliers, and research centers via R&D programs in France (Competitiveness Clusters, French National Research Agency) and in Europe (H2020, ITEA). Thus the participation of TERATEC in the POP project is well aligned with its mission and is reinforced by two Third Parties: the [INRIA](#) and [CNRS](#) laboratories that share in the technical program.

**Dissemination across Industries and Services** - TERATEC helps companies, especially small and medium-sized businesses, gain access to the high-performance computing technologies they need to develop new products and services. TERATEC and [GENCI](#) jointly run the French national program [SiMSEO](#) for the dissemination of numerical simulation among small and medium-sized businesses throughout the country, offering awareness sessions, training activities, sector-specific services and local support. In addition, the [TERATEC Forum](#) is the annual leading event for HPC, Big Data and Simulation in Europe, a gathering of international experts who meet to discuss technological and economic issues.

**Support for Technological SMEs** - TERATEC supports technological SMEs and start-ups in their actions to optimize their expertise, facilitates their access to industry leaders and helps them set up and finance their R&D projects.

**Teaching and Training** - TERATEC has joined forces with universities and major engineering schools to design initial and continuing education programs that cover the entire spectrum of high performance simulation and modeling. These initiatives will be expanded and reinforced in Europe.

One of the TERATEC successes is realising the first European Technopole dedicated to high-performance computing on the [TERATEC Campus](#). This is the home to an incubator, business centre, technology companies & industrial laboratories, and provides service platforms providing industrial stakeholders with the computing resources, software and technical expertise needed to carry out HPC projects. For more information about TERATEC see their [website](#).



<https://pop-coe.eu>



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

