



Performance Optimisation and Productivity

A Centre of Excellence in HPC



POP Newsletter 21 – Issue December 2021

Welcome to the 21st newsletter from the EU [POP](#) Centre of Excellence.

In this edition, we start with a reminder about the extension to the POP service and reveal the landmarks featured in the animation premiered in our previous issue. We invite you to attend our upcoming webinar on “Resources for Co-Design” and to catch up on the recent SME-focused one. We have a bumper selection of technical blogs this time, with lots of impressive success stories and useful technical tips. Finally, find out about the recent events we’ve contributed to and where you can learn more in our forthcoming performance tuning workshops.

If you would like to contribute technical content for this newsletter on the topic of parallel performance profiling, please contact us at pop@bsc.es.

This issue includes:

- **Reminder: POP Service Extension**
 - The POP animation landmarks revealed
- **POP Webinars**
 - Resources for Co-Design
 - POP: The SME Perspective
- **Technical Blogs**
 - 30x Speedup of the Matrix Factorization when Applying a Maths Library
 - Using the Python Extrae API to Profile Python OpenMP Codes
 - A POP proof-of-concept achieves 2x speedup of an EXCELLERAT use case
 - Using PAPI Counters to compute IPC for an MPI Fortran application
 - POP Transformed our Understanding of the Behaviour of immerFLOW
 - Near to 5x speedup of the RoSSBi astrophysics code
 - Proof-of-concept Leads to almost 2x Speedup of Atmospheric Physics Code
 - POP Performance Analysis of TensorFlow: High-Performance Deep Learning
- **Recent POP Events**
 - The NAFEMS World Congress 2021
 - POP Experts Contributed to SC 2021 Tutorial Program
- **Performance Tuning Workshops**
 - Performance Analysis Methodology Workshop
 - 41st VI-HPS Tuning Workshop
- **The POP Helpdesk**

For past editions of the newsletter, see the [POP newsletter web page](#).

Reminder: POP Service Extension

As announced in the last newsletter, the POP service has been extended by six months to run until the end of May 2022, giving you even longer to take advantage of our services. However, we suggest that you don't leave it to the last minute!

If you need a reminder of what we offer and have a spare 30 seconds, then [take a look at the animation](#) we premiered in our last newsletter, in which we set the challenge of spotting the five famous European landmarks that our POP superheroes fly over. The answers are, of course,

1. Old Trafford, Manchester, UK
2. Cologne Cathedral, Germany
3. Charles Bridge, Prague, Czech Republic
4. Sagrada Família, Barcelona, Spain
5. Palace of Versailles, France

Well done to everyone who got them all right and remember, you can apply for a performance assessment of your parallel code using our [service request form](#).

POP Webinars

Resources for Co-Design

10 December 2021, 15:00 CET / 14:00 GMT | online

[Resources for co-design](#) is a section within the POP website which gathers together a set of typical behavioural patterns seen in HPC codes, potentially resulting in some kind of performance degradation, that POP has identified in our analyses of user applications. For each of these patterns, the site links to the corresponding best-practice(s) that address their performance issues and, in many cases, also provides downloadable benchmarks to allow interested parties to compare the behaviour before and after applying a given best-practice.

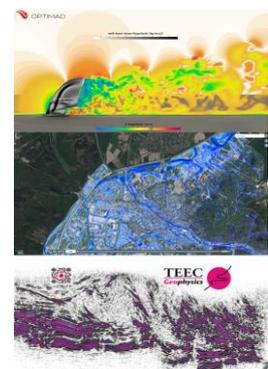
In this 30-minute webinar, Xavier Teruel of BSC will present the main motivation behind the development of the site and then explain how to navigate through the different resources for co-design we have so far created.

For further information and booking, please click [here](#).

POP: The SME Perspective

In September, we welcomed three speakers from SMEs, who have worked, or are currently working, with POP: Benedikt Rothe of [Hydrotec](#), Gerald Eisenberg-Klein of [TEEC](#) and Marco Cisternino of [Optimad](#). They each described the business they are in, the importance of HPC to it and what they have gained from their relationship with the POP service. The discussion was then opened up to cover various aspects of HPC support for SMEs. The recording, slides and handouts can all be found [here](#).

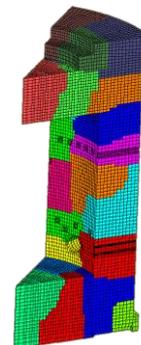
Browse the full list and catch up on all our previous webinars [here](#).



Technical Blogs

30x Speedup of the Matrix Factorization when Applying a Maths Library

[SIFEL \(Simple Finite ELeMents\)](#) is an open-source computer finite element (FE) code that has been developed since 2001 at the Czech Technical University in Prague. It is a C/C++ code parallelised with MPI. Developers requested a POP performance assessment of the PARGEF module containing the Schur complement method. It performs a matrix factorization, followed by an iterative solver. We came to the conclusion that the matrix factorization implementation is not efficient. Since the matrix factorization and the Schur complement computation is a well-studied mathematical and algorithmic problem that was already implemented in many math libraries, we decided to replace the original implementation by a call to the Pardiso library, part of the Math Kernel Library (MKL).



Find out what happened by reading [the full story](#).

Using the Python Extrae API to Profile Python OpenMP Codes

Extrae supports the collecting of trace data for Python codes. This is commonly done for codes using MPI or the multiprocessing package but, as well as this, the Intel distributions of the NumPy and SciPy Python packages both support Intel MKL which means both packages can take advantage of OpenMP parallelism.

Find out how to profile such codes in this tool time [blog post](#).

A POP proof-of-concept achieves 2x speedup of an EXCELLERAT use case

[Alya](#), one of the reference codes of the [EXCELLERAT CoE](#), is a high-performance computational mechanics code used to solve complex coupled multi-physics, multi-scale, multi-domain problems. A performance assessment was performed using an EXCELLERAT use case that simulates a Bunsen flame. The POP assessment revealed that the main factor limiting scalability for this use case is load imbalance. This load imbalance issue is inherent to the problem that is being solved, because the chemical reactions only happen, and are thus computed, in a small section of the domain. Moreover, where the chemical reactions happen cannot be predicted so the domain partition cannot be adjusted to minimize the imbalance.

Read about how POP tackled this problem in this [blog post](#).

Using PAPI Counters to compute IPC for an MPI Fortran application

[PAPI](#) (Performance Application Programming Interface) is an open-source API developed by the University of Tennessee to access hardware performance counters on most modern microprocessors. It is written in C and used by many profiling tools, including [TAU](#), [Score-P](#) and [EXTRAЕ](#). In some cases, however, it is more convenient to use PAPI directly from the application, to reduce the overhead and obtain only the desired quantities.

In [this article](#), we describe how to call this API from an MPI Fortran application.

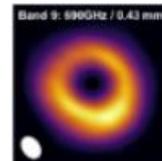
POP Transformed our Understanding of the Behaviour of immerFLOW

[Here](#) we have a guest blog post, written by Marco Cisternino of [Optimad](#) about the work carried out by POP on their immerFLOW code and how that led to a much greater understanding of the code's behaviour. Marco was also one of the participants in the webinar "POP: The SME Perspective", mentioned earlier.



Near to 5x speedup of the RoSSBi astrophysics code

RoSSBi-3D (Rotating Systems Simulation for Bi-fluids) is the 3D version of a code specifically developed to study the evolution of protoplanetary disks. During the standard POP analysis workflow, we identified several performance issues that limited the scalability and speed of computation. Read the [blog post](#) for the full story that led to a near 5x performance improvement.



Proof-of-concept Leads to almost 2x Speedup of Atmospheric Physics Code

Polar Scat is an atmospheric physics code that computes scattering cross-sections for particles in the atmosphere. It is written in Fortran and parallelized with OpenMP. A POP performance assessment was performed, showing that the parallel performance of the code was excellent, with a global efficiency of over 80% on 48 cores. However, vectorization was poor, with only 29% of vector capacity usage on 48 threads.

Find out how we tackled the vectorisation problem to give an almost 2x speedup [here](#).

POP Performance Analysis of TensorFlow: High-Performance Deep Learning

When using expensive hardware for deep learning, it is important to understand whether good use is being made of the GPU resources and to question whether the GPU Parallel Efficiency can be improved. As for any parallel computation, inefficiencies can arise due to imbalanced computation, excessive time in data transfer and idle time due to synchronisation. A POP Performance Assessment can help answer these questions about efficiency.

Find out more in our [blog post](#).

Recent POP Events

The NAFEMS World Congress 2021

POP recently participated in the [NAFEMS World Congress 2021](#), held virtually from the 25th to the 29th October. NAFEMS is the “International Association for the Engineering Modelling, Analysis and Simulation Community” and their biennial world congress is an important landmark in the analysis and simulation calendar. Ten keynote speakers, 1,000 attendees, and over 600 delegates from all over the world presented state-of-the-art results in various fields, including automotive, physics, aerospace and chemistry, as well as many more.

Further details can be found [here](#).

POP Experts Contributed to SC 2021 Tutorial Program

Experts of the POP CoE contributed to the tutorial program of the [SC 2021 conference](#), which was held as a hybrid (in-person and virtual) event. Full details of the topics covered are available [here](#).

Performance Tuning Workshops

Performance Analysis Methodology Workshop

15 December, 2021, 09:00-12:00 GMT, online

This online workshop is organized in collaboration with the [Computer Science Department](#) and [Advanced Research Computing](#), Durham University, [DiRAC](#) and [N8 CIR](#). The workshop aims to give a solid understanding of the basics of performance analysis for research and simulation codes, using high- and low-level metrics. The agenda and registration details can be found [here](#).

41st VI-HPS Tuning Workshop

7-11 February, 2022, online

POP experts will give an overview of the VI-HPS programming tools suite, explain the functionality of individual tools, and how to use them effectively, and offer hands-on experience and expert assistance using the tools. The event is organised by JSC and RWTH Aachen, Germany and the agenda and registration details can be found [here](#).

Apply For Free Help with Code Optimisation

We offer a range of [free services](#) designed to help EU/UK organisations improve the performance of parallel software. If you are not getting the performance you need from parallel software or would like to review the performance of a parallel code, please apply for help via the short [Service Request Form](#), or [email us](#) to discuss the service further and how it can be beneficial.

These services are funded by the European Union Horizon 2020 research and innovation programme so there is no direct cost to our users.

The POP Helpdesk

Past and present POP users are eligible to use our [email helpdesk](#). 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.

