



Performance Optimisation and Productivity

A Centre of Excellence in Computing
Applications



POP Newsletter 5 – Issue June 2017

Welcome to the fifth newsletter from the EU [POP](#) Centre of Excellence. Our free of charge services help EU organisations improve performance of parallel software.

This issue includes:

- **Live webinar** - How to Improve the Performance of Parallel Codes
- New **Success Stories**
- A new blog post
- More **Analysis Highlights**
- Information on where to meet POP at upcoming events

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

POP live webinar - How to Improve the Performance of Parallel Codes

Thursday 15 June 2017 - 14:00 BST | 15:00 CEST

This [30-minute live webinar](#) will discuss how to improve the performance of parallel codes. We will present a systematic approach to optimising codes while pointing out the various factors that should be considered. We will illustrate the talk with practical examples involving real codes.

Attendees will get:

- An expert practitioner's view
- A demonstration of a code profiling tool.
- Demonstrated results on real applications.
- Your questions answered, either during the session or post event.

[Register here](#)

About the presenter: Jon Gibson is a Technical Consultant at the Numerical Algorithms Group (NAG). He has 25 years' experience in scientific programming and has worked for two national HPC services in the UK.

What could you achieve if your parallel software ran faster?

A new blog post: The importance of using the right methodology for parallel code optimisation.

Improving performance of your code could mean running at the resolution you need, coupling that extra package to make your simulations more realistic, or significantly reducing your time to solution to get what you want now rather than later.

But this is not an easy task. Many connecting and interweaving factors affect the performance of your application, and unpicking them takes a serious investment of time and effort that often is not available, leading to performance optimisation with a narrow focus. But naïvely deal with only the obvious issues, and sooner or later the next problem will pop up, just like playing whack-a-mole. If you don't step back to look at the issues holistically, you may achieve little more than running a little bit faster or on a few more cores.

For the full article [click here](#).

Success story – Proof of Concept for BPFM leads to 40% runtime reduction

The public available [BPFM](#) (Bayesian Probabilistic Matrix Factorization) code is an efficient method to solve complex modelling problems. This code was improved as part of a POP Proof of Concept study which implemented:

- further optimization of the linear algebra computations
- improvement of the selection of the optimized algorithms inside BPFM
- solving load balance issues in the OpenMP parallelization

For the full story see the [POP blog](#). POP Proof of Concept studies are a free service to eligible EU organisations.

Success story – 10-fold improvement for EPW

[EPW](#) (Electron-Phonon using Wannier interpolation) is a materials science DFT code distributed within the Quantum ESPRESSO suite, and written in Fortran with MPI parallelism. A POP Performance Plan resulted in a version of the code that was **60% faster**, and a POP proof-of-concept investigation reduced file writing from over seven hours to **under one minute!**

See the [POP blog](#) for more information.

Success story – 2x speed up for k-wave

[k-Wave](#) is an open-source toolbox for time domain acoustic and ultrasound simulations in complex and tissue-realistic media, parallelised with MPI+OpenMP. The insight on code inefficiencies obtained during the POP performance audit allowed the developers to obtain a **2x speedup** whilst the audit was still live.


More information on the [POP blog](#).

Some recent analysis highlights

A POP Performance Plan for GS2

A new algorithm for the GS2 code (written using Fortran and MPI) was found to reduce communication costs at the expense of introducing additional load imbalance in a particular section of the application.

The level of the load imbalance introduced is revealed as the step change in the blue line in the profiling plot.



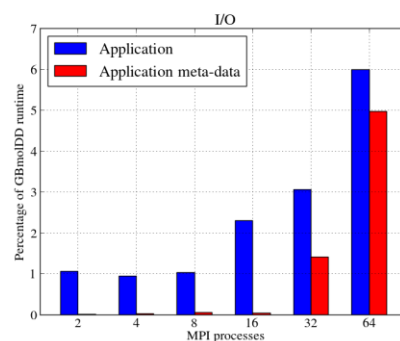
This detailed level of profiling information is invaluable for steering the developers' efforts to further optimise this application.

Analysis of MPI molecular dynamics software

The GBmolDD molecular dynamics software replicates computation over MPI processes to reduce expensive communications, but still shows poor parallel scaling. The POP audit identified that while good communications efficiency had been achieved, the cost in terms of low computational efficiency was significant (see table). We recommended the balance between replication and communication should be reconsidered, and to reduce communications by implementing hybrid MPI+OpenMP parallelism.

MPI processes	8	16	32	64
Computational efficiency	73.5	59.7	51.7	54.1
Communication efficiency	97.7	95.3	91.3	87.7

The audit also identified a significant increase in file I/O. Our analysis showed this was because at the end of the computation each MPI process wrote data to file using POSIX I/O. This results in many files for large MPI runs, with the underlying file system spending a large amount of time in the I/O meta-data phase (see figure). This could be improved using a parallel file format e.g. MPI-I/O, parallel NetCDF or parallel HDF5.



Meet POP at upcoming events

POP BoF @ ISC 2017

- Look out for a BoF titled [POP Improves HPC Applications](#) at ISC and participate!
- Talk to POP experts at the exhibition booths of BSC, JSC/HLRS and NAG.
- See the POP poster in the Project Poster session - booth L-212 in the Exhibition Hall, June 19th (pm) to June 21st.

The [International Supercomputing Conference \(ISC\)](#) (18-22nd June in Frankfurt) 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 27-28th June in Paris. It includes keynotes, technical sessions and an exhibition, with a strong emphasis on industrial usage of HPC.

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 (until end of March 2018) 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.



<https://pop-coe.eu>



HLRS



nag®



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

