Performance Optimisation and Productivity

Your parallel code: better

What?

A Centre of Excellence for Performance Optimisation and Productivity.

Promoting best practices in parallel programming, funded by the EU.

FREE Services to EU industry and academia providing precise understanding of application and system behaviour.

Suggestions and support to refactor your code in the most productive way.

Why?

Why use POP?

The complexity of HPC machines and parallel codes:

- There is frequently a lack of quantified understanding of the actual behaviour of applications
- It is not always clear what the most productive direction is for code refactoring.

It is important to maximise the computational efficiency of applications as well as the code development effort.

The codes we work with

POP works with parallel codes written using all languages and paradigms: C, C++, Fortran, Python, CUDA, MPI, OpenMP …

Oil and Gas success story

improvements halve the run time

Shearwater Reveal is a seismic processing code that does time and depth analysis for land and maritime applications. Shearwater Reveal land processing tools cover all aspects of land processing from refraction statistics to final time and depth imaging.

The POP assessment made the following recommendations:

- Reduce file I/O by re-using data
- Remove I/O from deep inside loops
- Combine file I/O to improve performance
- Implement dynamic load balancing
- Identified serial code for parallelisation.

The final modified code, based on our analysis, is shown in yellow on the graph. It shows a region of interest with a performance gain of 44% when compared to the original.

In a full scale production run the computational cost was almost halved.

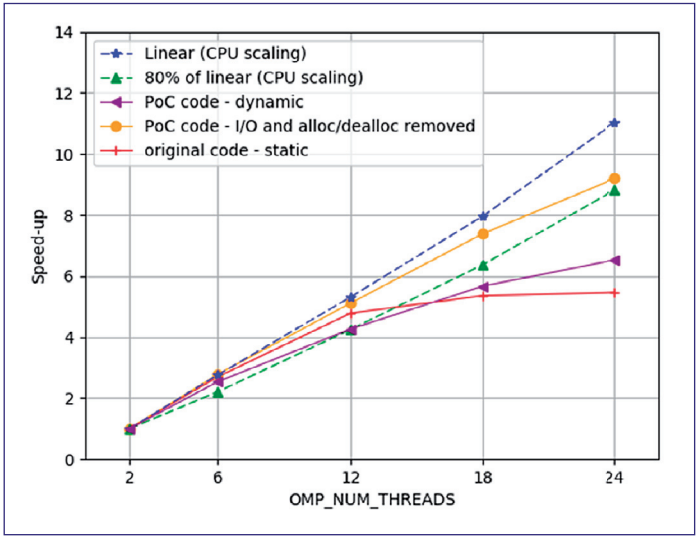

CAE success story

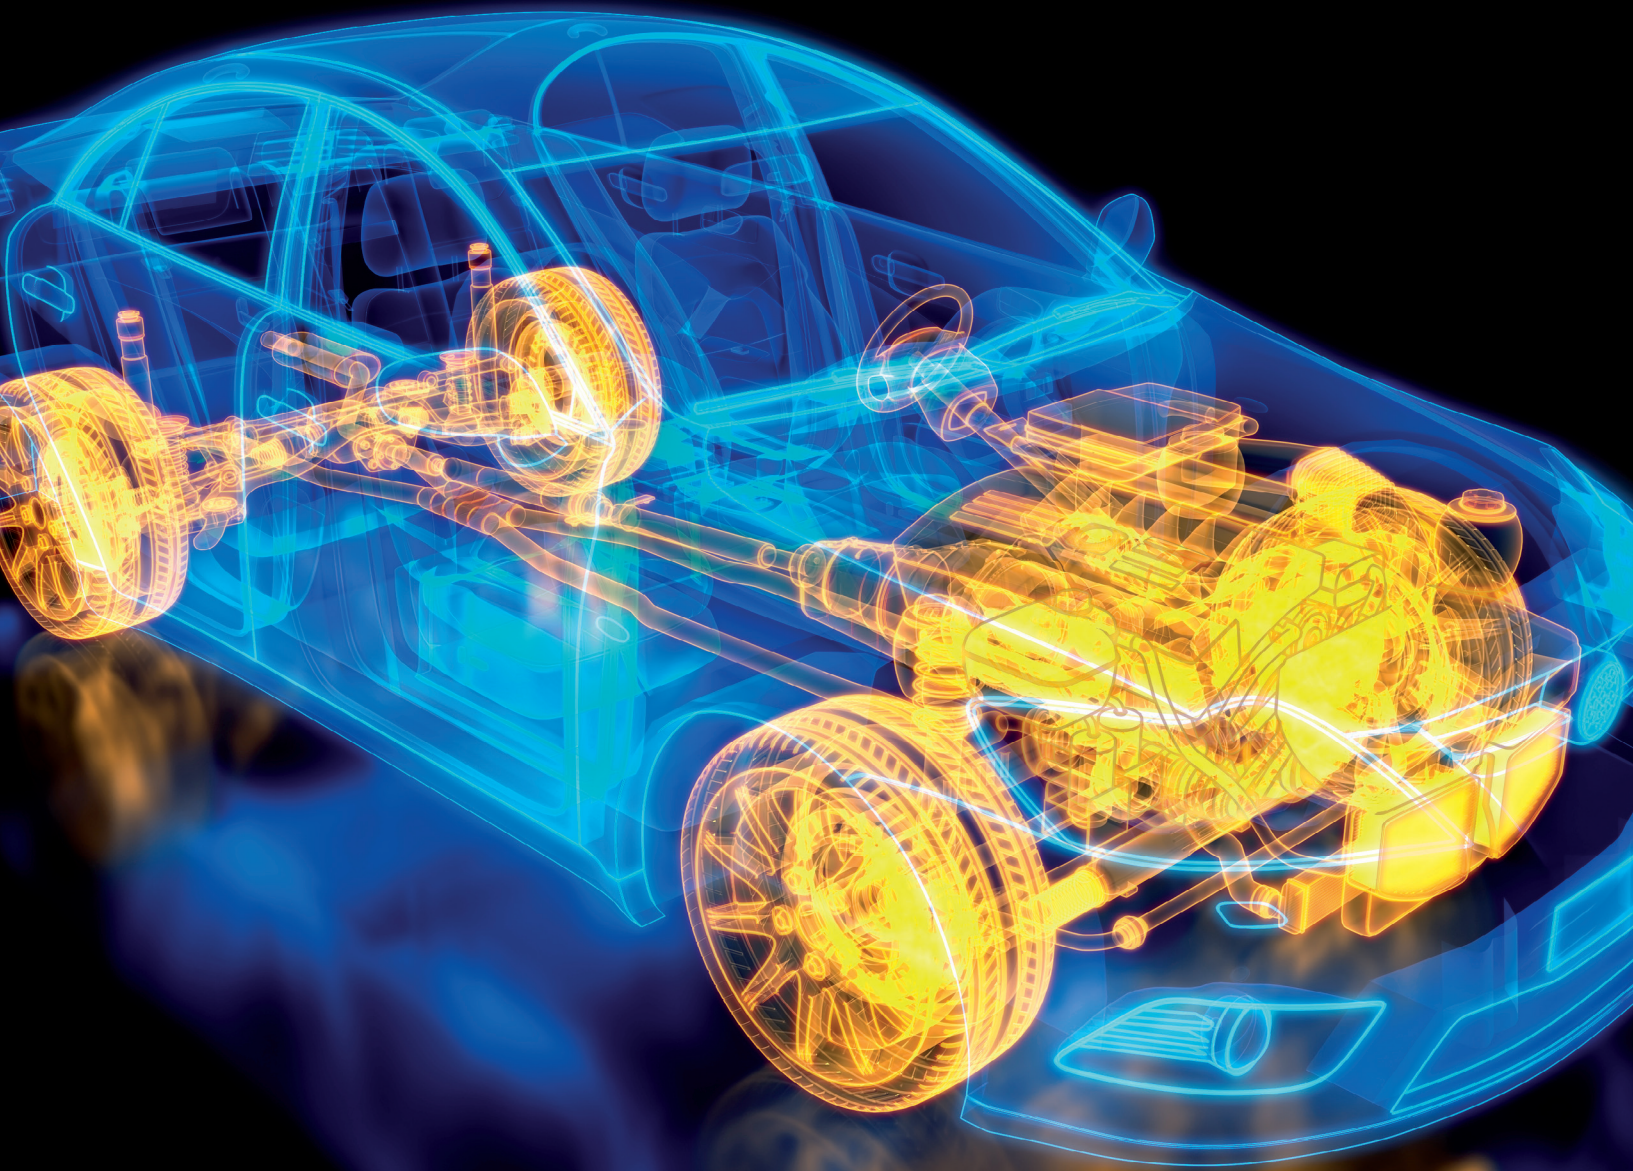

Three times speed up in Computational Fluid Dynamics

zCFD by Zenotech is a computational fluid dynamics (CFD) solver that is written in Python and C++ and parallelised with OpenMP and MPI.

As a result of a POP Proof-of-Concept study zCFD ran 3x faster on a representative input case. This was achieved by:

- Parallelising serial portions of code, specifically those which were not already running in parallel because of a compiler problem
- Improving computational load balance across OpenMP threads by avoiding very slow maths function calls
- Changing execution environment settings to boost CPU performance.

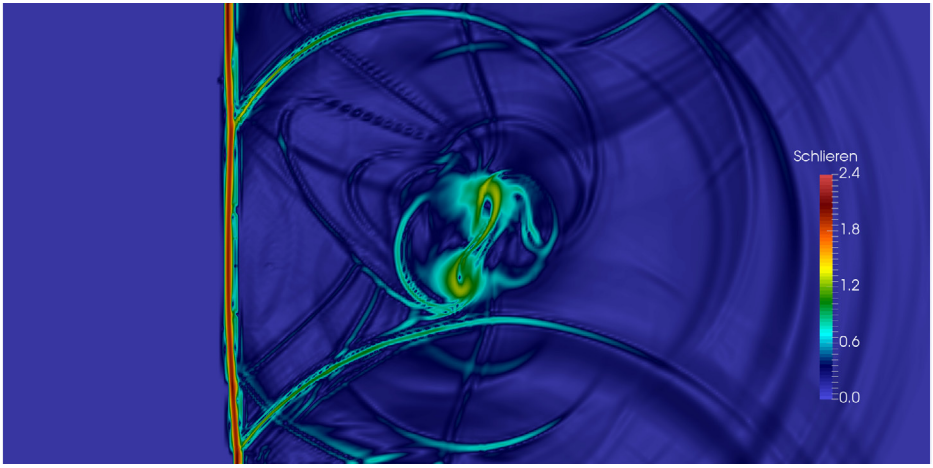

ISVs success story

ISV code computation kernel now four times faster

SCM's Amsterdam Modeling Suite (AMS) is a popular computational chemistry software package used in industrial and academic research, and includes a range of software products, many of which are based on density functional theory.

POP undertook performance assessments to investigate parallel inefficiencies in three of the AMS codes: ADF, BAND and DFTB. One study demonstrated how to improve performance for parallel multiplications of small complex matrices. We:

- Reduced the idle time by overlapping communication with computation
- Halved the time in collective communications
- Halved the time in computation.

The improvement can be seen on the trace data below, where the horizontal axis is time and vertical axis is process. These compare original and optimised computation on 288 cores where the optimised code is around four times faster.

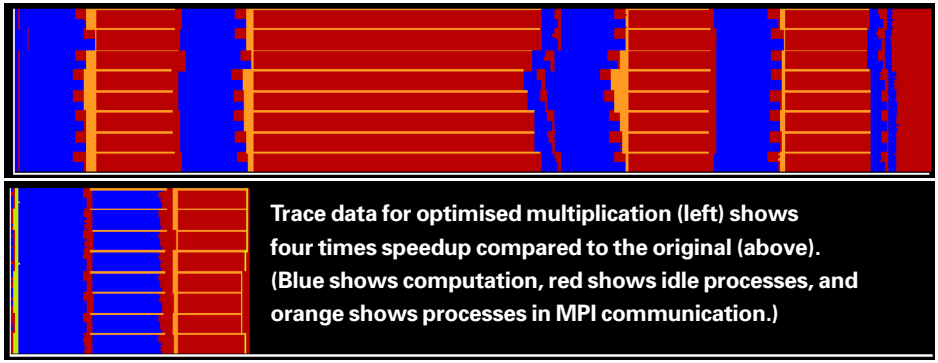

Trace data for optimised multiplication (left) shows four times speedup compared to the original (above). (Blue shows computation, red shows idle processes, and orange shows processes in MPI communication.)

Time ⟶

Sectors and Clients

POP works across a range of sectors including: Oil & Gas, Engineering, Finance, Pharma, High Performance Data Analysis and many others.

POP works with both the developers and the end-users of applications, whether in industry or in academia. We've helped Independent Software Vendors to improve their products and developers of in-house solutions to get better solutions. We've also improved the performance of a range of open-source applications for the benefit of their entire user community.

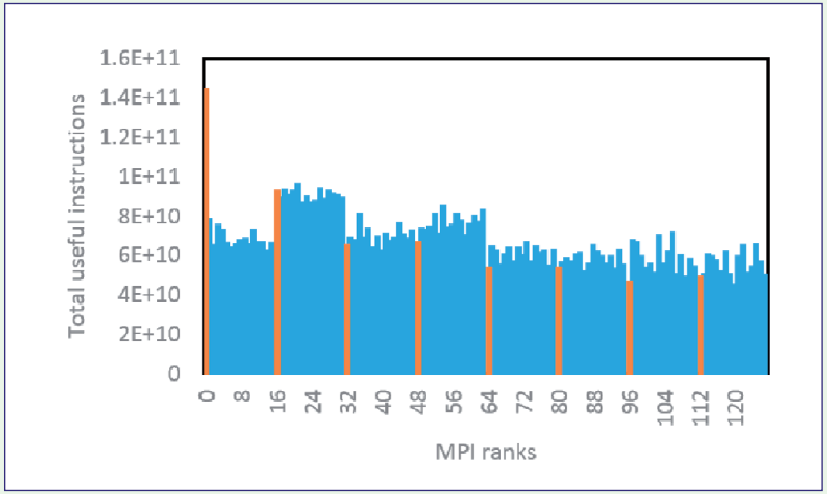
1.6E+11
1.4E+11
1.2E+11
1E+11
8E+10
6E+10
4E+10
2E+10
0
Total useful instructions
0
8
16
24
32
40
48
56
64
72
80
88
96
104
112
120
MPI ranks

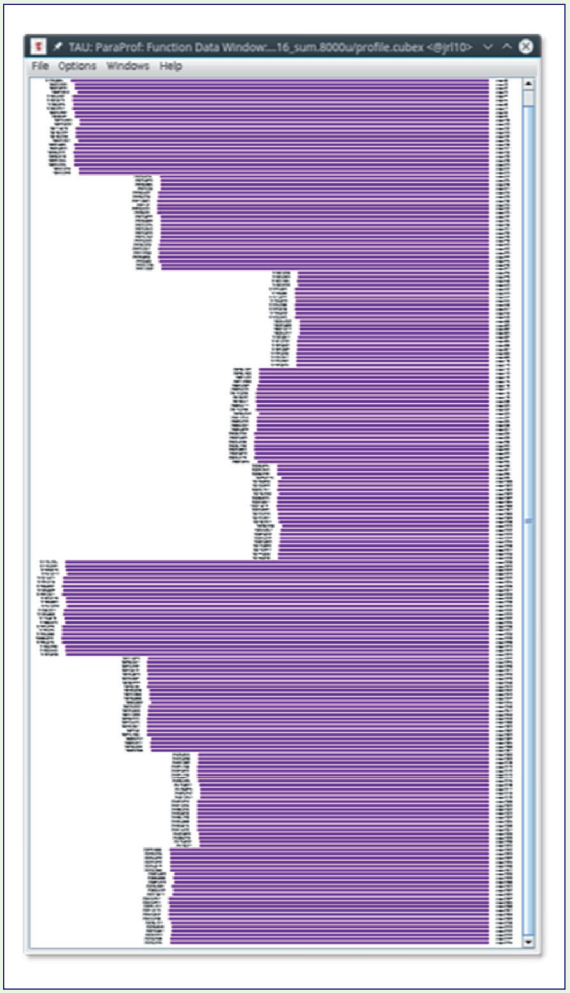
TAU: ParaProf: Function Data Window:...16_sum.8000u/profile.cubex <@jrl10>
File Options Windows Help

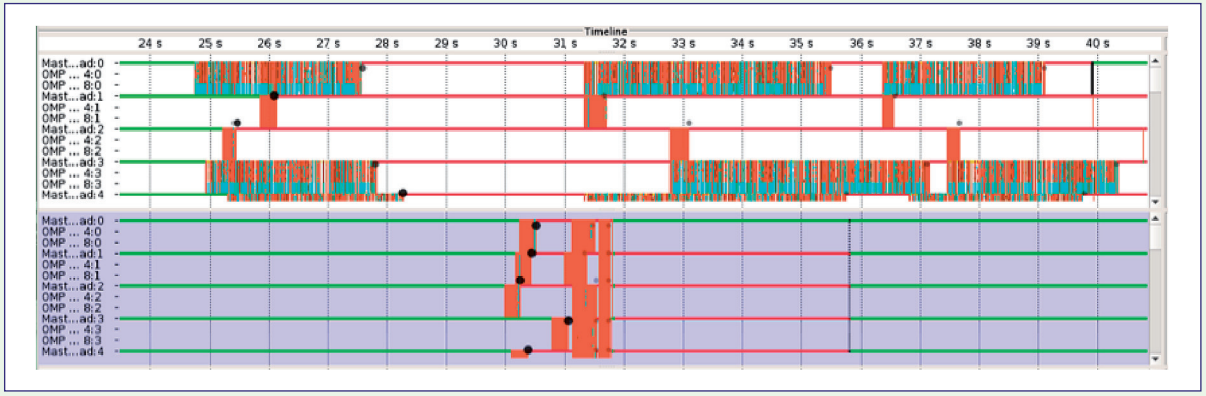
Timeline
24 s
25 s
26 s
27 s
28 s
29 s
30 s
31 s
32 s
33 s
34 s
35 s
36 s
37 s
38 s
39 s
40 s
Mast...ad:0
OMP ... 4:0
OMP ... 8:0
Mast...ad:1
OMP ... 4:1
OMP ... 8:1
Mast...ad:2
OMP ... 4:2
OMP ... 8:2
Mast...ad:3
OMP ... 4:3
OMP ... 8:3
Mast...ad:4

Process


A collaboration between YOU and POP
POP services are designed to complement your code development
Working together with POP tools we analyse your application
The POP team write a report with recommendations on the next step
If needed POP can help you improve the performance of the code
You get faster code, lower compute costs and increased productivity

Parallel Application Performance Assessment

The primary service

Performance issues, if any, are identified using a variety of POP analysis tools. Customers can install and use these themselves, if preferred, so there is no need to provide us with source code.

A report will identify the root causes of any issues found and make recommendations on how to address them.

Proof-of-Concept

A follow-up, more in-depth, service

Refactoring of selected computational kernels, I/O routines or other suboptimal code sections will be undertaken by POP staff.

Experiments will be run to show the performance improvements and indicate how these can be rolled out across the application.


mod = modifier_ob.modifiers.new("
object to mirror_ob
mod.mirror_object = mirror_ob
ion == "MIRROR_X":
mod.use_x = True
mod.use_y = False
mod.use_z = False
ation == "MIRROR_Y":
mod.use_x = False
mod.use_y = True
mod.use_z = False
ation == "MIRROR_Z":
mod.use_x = False
mod.use_y = False
mod.use_z = True
the end -add back the dese
select=1
scene.objects.activ
ted" + str(modifier_ob)) # mod
context.selected_objects[0]
objects[one.name].select = 1
("please select exactly two objects,
OPERATOR CLASSES
Operator):
mirror to the selected object"""
mirror_mirror_x"
R X"
context):
context.active_object is not None

Metrics

A standard way to characterise the performance of a parallel code so as to identify the factors which limit its scalability.

	MPI Processes 6	12	24	48
Global Efficiency	92%	95%	62%	49%
Parallel Efficiency	92%	89%	50%	40%
Load Balance	95%	94%	88%	88%
Communication Efficiency	96%	95%	57%	45%
Serialisation Efficiency	97%	96%	76%	62%
Transfer Efficiency	99%	99%	75%	73%
Computational Scalability	100%	106%	124%	123%
IPC Scalability	100%	106%	127%	130%
Instruction Scalability	100%	99%	98%	96%
Frequency Scalability	100%	101%	97%	95%

A set of metrics can be derived for any region of interest in the code using the profiling tools developed by the POP partners.

Each metric provides a measure of a different aspect of the parallel performance, helping to highlight the issues with the code and where improvements might be made.

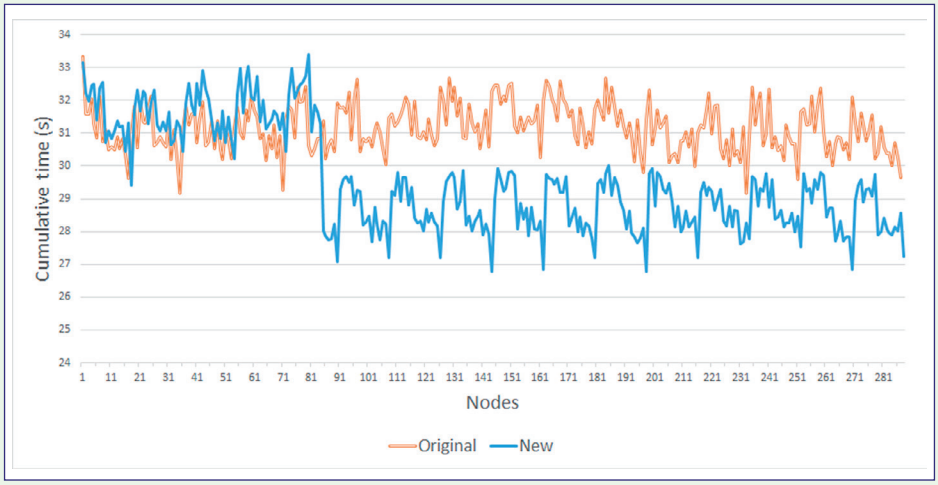
Cumulative time (s)
34
33
32
31
30
29
28
27
26
25
24
1 11 21 31 41 51 61 71 81 91 101 111 121 131 141 151 161 171 181 191 201 211 221 231 241 251 261 271 281
Nodes
Original
New

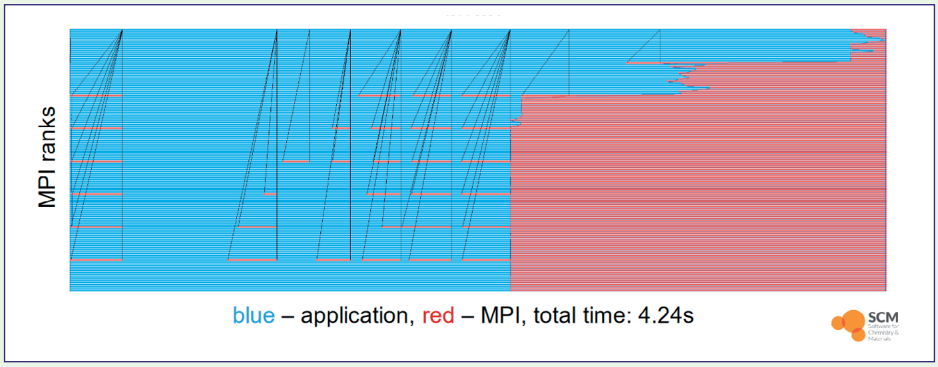
MPI ranks
blue – application, red – MPI, total time: 4.24s
SCM

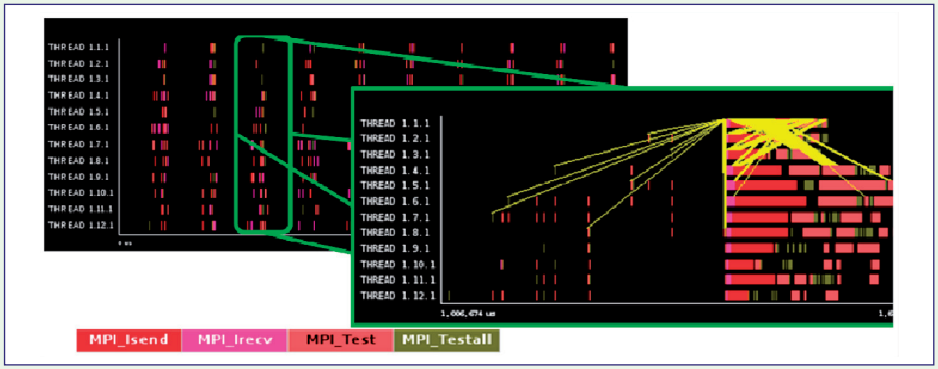
MPI_Isend
MPI_Irecv
MPI_Test
MPI_Testall


Apply Now

Apply

Getting started with a POP service is easy. Simply fill out the small questionnaire describing your application and needs:

https://pop-coe.eu/request-service-form

Got a question?

pop@bsc.es

Partners

Excellence in performance tools and tuning
Excellence in programming models and practices

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC is a globally recognised HPC Research and Service center. It operates a world-class facility in Marenostrum and has developed powerful performance analysis tools and methodologies. BSC is a key contributor to OpenMP and many other European projects, promoting best parallel programming practices towards exascale.

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

IT4Innovations (IT4I) conducts excellent research in high-performance computing and data analysis, runs a national supercomputing infrastructure and mediates its efficient utilisation to increase the competitiveness and innovativeness of Czech science and industry. IT4I is a member of PRACE, ETP4HPC, EuroHPC, EUDAT, BDVA, and I4MS.

HLRS

The High Performance Computing Center Stuttgart (HLRS) is one of three national HPC centers in Germany. The mission of the centre is to support researchers and industrial partners at National and European level in exploiting leading-edge supercomputing systems. HLRS is a member of the MPI standardisation body MPI-Forum.

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

Jülich Supercomputing Centre (JSC) is one of the leading HPC competence centres in Europe and provides HPC infrastructure both on the EU level, supporting PRACE, and the German national level, with the Gauss Centre for Supercomputing. JSC is also an internationally recognised parallel performance tool provider, including Score-P and Scalasca.

The Numerical Algorithms Group (NAG) is a global leader in HPC and Cloud consulting and services, numerical software engineering, and world-class algorithm development – solving numerical, technical and HPC problems every day, all over the world.

Teratec is a European Center of Excellence for Numerical Simulation and High-Performance Computing with over eighty technological and industrial members. Teratec offers numerous services to French and European companies, and helps them master the technology and access HPC.

RWTH Aachen University is Germany's largest technical university. The IT Center operates national HPC and Visualisation infrastructure and conducts research in these fields. RWTH is a member of the OpenMP and MPI standardisation bodies.

UVSQ is an academic leader in HPC. They have established the first master's degree entirely devoted to HPC in France. UVSQ develop performance evaluation and optimisation tools which are used by both academia and industry. Since 2009 they have collaborated with CEA and Intel to operate an Exascale Computing Research laboratory.

Partnerships

Cloud and Hardware Providers

POP can help you and your customers with:

- Assessment of performance in production conditions
- Improvements from modifying environment setup
- Training of support staff
- Benchmarking
- Customer support
- System dimensioning and design.

Professional Bodies

Do you need to promote computing innovation and best practice to your community? POP can help, please get in touch.

Testimonials

GS2 (STFC)

"POP analysis elegantly reveals in detail how our application's algorithm is running on HPC architectures. It is an extremely useful optimisation tool. Our POP contact was very knowledgeable and enthusiastic. An excellent service!"

Joseph Parker, GS2 Developer

OpenNN (Artelnics Ltd)

"High performance computing is an extremely interesting topic to our application. The POP project has helped Artelnics to speed-up Neural Designer up to five times, when compared to the serial version. And we can still improve a lot more by implementing MPI processing in computer clusters."

Dr Roberto Lopez, CEO

90% of customers **very satisfied** or **satisfied**

VAMPIRE (University of York)

"The audit of the VAMPIRE code has been extremely helpful in identifying the hot spots and specific areas to focus on performance improvements. Preliminary results suggest this may give a factor of two performance improvement on modern CPUs. I would highly recommend the service for the speed and usefulness of the audit."

Richard Evans, VAMPIRE Developer

ADF (SCM BV)

"I can honestly say your analysis gave us a new insight into performance of one of the newer features available in ADF. What is more important, it clearly showed us the limitations of the current implementation and pointed us to the ways to improve it."

Alexei Yakovlev, ADF Developer

100% considered the quality was **excellent** to **good**

Visit the POP website for learning materials, our blog, and details of how to sign-up for our newsletter and webinar programme.

https://www.pop-coe.eu

pop@bsc.es

@POP_HPC

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