11 partners, providing comprehensive expertise in social, technical and business domains for both, HPC and GSS.

- High Performance Computing Center Stuttgart, Universität Stuttgart (USTUTT), Germany
- Universität Potsdam (UP), Germany
- Global Climate Forum EV (GCF), Germany
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk (PSNC), Poland
- Fondazione Istituto Per L'Interscambio Scientifico (ISI), Italy
- Scuola IMT Alti Studi Di Lucca (IMT), Italy
- Consorzio TOP-IX (TOP-IX), Italy
- Chalmers Tekniska Hoegskola AB (CHALMERS), Sweden
- ATOS Spain (ATOS), Spain
- The COSMO Company SAS (COSMO), France
- Dialogik Gemeinnuetzige Gesellschaft für Kommunikations- und Kooperationsforschung (DIA), Germany

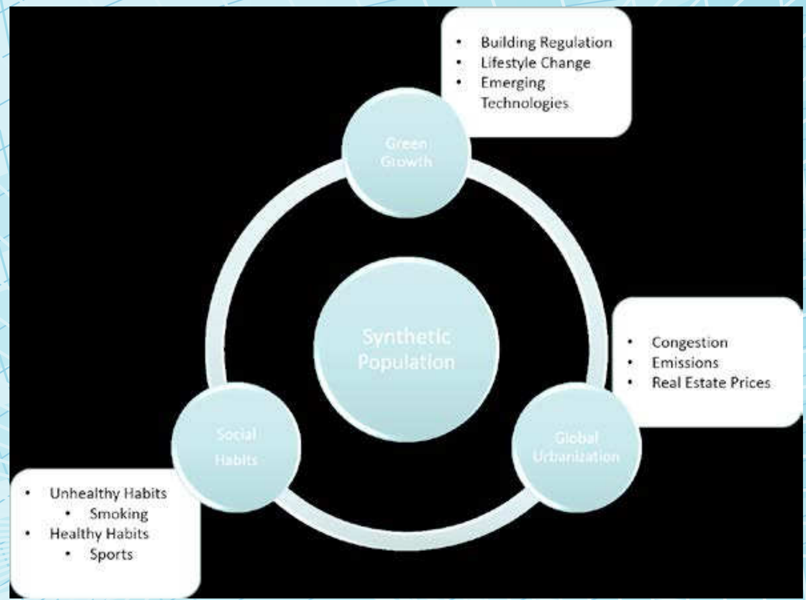


Technical realizations related to High Performance Computing are managed by the partners USTUTT, PSNC and ATOS, social research including risk assessment is performed by DIA, IMT and GCF. Furthermore, Global Systems Sciences related topics are conducted by GCF, UP, CHALMERS, ISI, IMT and COSMO and a strong business context for exploiting and creating significant services and results is provided by ATOS, TOP-IX and GCF. Especially this in general complementary structure, which still enables cooperation on research and business topics amongst the consortium partners, guarantees a sustainable operation and uptake of HPC and GSS communities.

contact:
Michael Gienger,
gienger@hlrs.de,
Bastian Koller,
koller@hlrs.de

• Michael Gienger

University of Stuttgart (HLRS), Germany





# POP: Performance Optimization and Productivity

## A Centre of Excellence in Computing Applications

Inaugurated October 1, 2015, the new EU H2020 Center of Excellence (CoE) for Performance Optimisation and Productivity (POP) provides performance optimisation and productivity services for academic and industrial codes. European's leading experts from the High Performance Computing field will help application developers getting a precise understanding of application and system behaviour.

Established codes, but especially codes never undergone any analysis or performance tuning, may profit from the expertise of the POP services which use latest state-of-the-art tools to detect and locate bottlenecks in applications, suggest possible code improvements, and may even help by Proof-of-Concept experiments and mock-up test for customer codes on their own platforms.

Today's complexity of high performance computer systems and codes makes it increasingly harder to get applications running fast and efficient on the latest hardware. Often expert knowledge and a good amount of experience is needed to figure out the most productive direction of code refactoring. Domain experts in many research areas and industry use computer simulations but lack this knowledge. Thus, their codes are often far away from using the hardware in an efficient way, using much more compute time than needed. By this, they either waste energy, require superfluously oversized and expensive hardware, or just miss research potential as their codes can only handle smaller or less complex problems in the available compute time.

The dominant practise in analysing and presenting the performance of applications and how they scale with increased core counts is to report speed-ups derived from execution timing or domain-specific performance metrics (e.g. simulated time/days). These are global observations but often do not give sufficient insight into the actual inefficiencies inside the simulation programs. In many cases the reference execution is a parallel run with a relatively small core count but can already this experience inefficiencies (e.g., load imbalance), thus masking the effect of such behaviour in the scaling study.

Beyond the global speed-up study, profiles may be collected that point to actual routines that dominate the execution, but often do not give real insight into the fundamental behaviour of the application. As computing applications become more and more complex, their performance is below the optimal levels for a number of reasons:

- Load imbalance that causes waiting on slower processes - This lack of balance can be caused by a different amount of work per process (computational imbalance) or different performance (hardware-related imbalance) or a combination of both.

- Serialisation – dependencies between code regions that cause chains of delay propagated along different processes.
- Data transfer – time cost of non-overlapped data transfer of data between processes.
- Instructions Per Cycle (IPC) – actual performance of the sequential computation. This can be significantly below the core peak performance due to issues with the memory hierarchy, instruction mix, non-pipelined instructions, or dependencies.
- Amount of instructions and type of instruction – is the algorithm optimal for the given problem in terms of computational and instruction complexity and code balance?
- I/O, storage – frequency of or the time spent on I/O operations and how it affects computational efficiency.

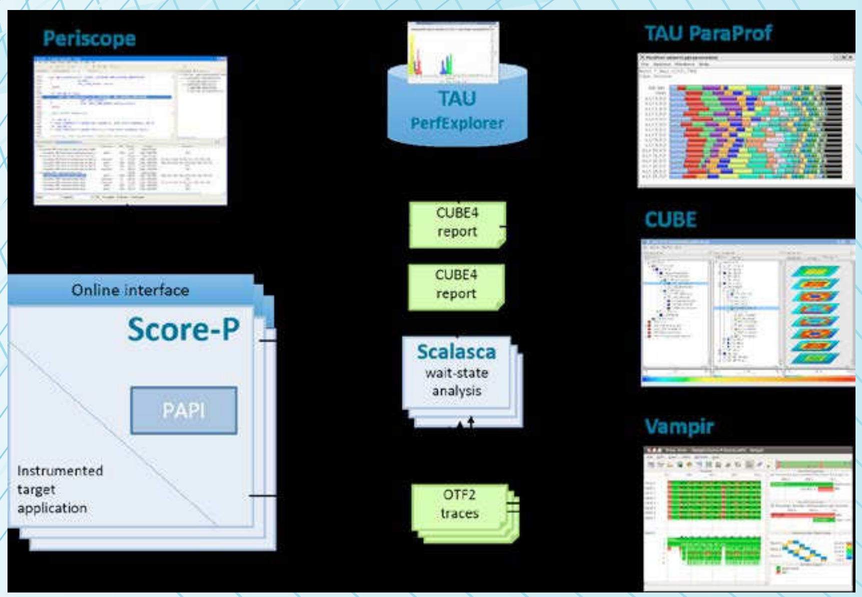


Figure 1: Performance Analysis Tools of the Score-P Ecosystem.

To overcome this situation, the POP CoE brings two tightly coupled disciplines to the user as a service, which are crucial for the efficient use of parallel computers in the future: First, powerful performance analysis tools, methodologies, and expertise needed to precisely understand and gain real insight into the actual application and system behaviour; Second, deep understanding of programming models and best practice guidance needed to express algorithms in the most flexible, maintainable and portable way, while still being able to maximise the performance achieved.

## POP-Technologies

The German project partners will provide their powerful and well-established performance analysis tools centered around the community-developed instrumentation and measurement infrastructure Score-P. Core developers are (among others) the Gauss Alliance members JSC, RWTH Aachen University, Technical University Darmstadt, Technical University Dresden, and Technical University Munich. Score-P can instrument and measure the performance of typical HPC applications written in Fortran or C/C++ and based on the MPI, SHMEM, OpenMP, OmpSs, CUDA, OpenCL and Pthread programming paradigms. It can collect flat and callpath profiles or detailed execution traces in the open CUBE4 or OTF2 formats. Score-P analysis tools include the event trace analyser Scalasca (developed by JSC) which provides very scalable wait-state, delay or root-cause analysis, the event trace analyser and visualisation tool Vampir (developed by TU Dresden), and the online performance analysis tool Periscope developed by GCS partner TU Munich. CUBE4 profiles can be analysed with the help of the Cube browser from JSC or TAU ParaProf (developed by the Univeristy of Oregon).



Projects

In addition to the Score-P tool universe, project partner BSC will provide their instrumentation and measurement package Extrae and their event visualiser Paraver. Finally, the POP experts will of course also use project-external tools like vendor products installed on the POP customer target platforms.

## POP-Services

The POP CoE team consists out of six partners which are experts in High Performance Computing with long-standing experience in performance tools and tuning as well as researchers in the field of programming models and programming practices. All partners come with a research and development background and proven commitment in application of their knowhow to real academic and industrial use cases. The POP CoE will provide three kind of service levels to its customers – depending on their background, knowledge, and demands:

### ? Application Performance Audit

This is the primary service of the POP CoE and the starting point for any further work. Applications undergoing this service will be analyzed by the POP experts after an initial discussion with respect to their best practices and provide a first impression of the code status. Within the Performance Audit, performance issues of customer code will be identified at the customer site. It will serve as a starting point for further analysis or initial code refactoring. The duration for a Performance Audit is expected to be around one month and a successful Performance Audit may be seen as a code quality certificate in HPC.

### ! Application Performance Plan

The Performance Plan Service follows the Performance Audit if the customer needs more detailed knowledge where and how to address specific issues in the code. POP experts together with the customer will develop a plan how and with which tools to analyse the issues under investigation. The POP experts will then analyse the code in detail and give quantified hints to overcome the problems so that they can be fixed by the customer. The duration for a Performance Plan is very problem-specific, but will in general take between one and three months including a closer look into the source code.

### √ Proof-of-Concept

If requested, Proof-of-Concept studies will be performed. This includes experiments and mock-up tests for customer codes. The details of the proof-of-concept study will be decided in very close collaboration with the customer and may include kernel extraction from the application, parallelisation or mini-apps experiments to show effects of the proposed optimisations of the POP experts. As this very complex task goes into deep detail, Proof-of-Concept work should be expected to require about six months.

Besides the above three key services, the POP CoE will also provide a variety of training activities in the field of performance analysis and optimisation based on the user's needs to improve their basic high performance programming knowledge and increase the awareness of performance issues and potentials in general.

## POP-Impact

The POP CoE with its service and training activities will have a wide impact within all areas of research and industry, making it a real transversal activity:

- It provides access to computing application expertise that enables researchers and industry to be more productive, leading to scientific and industrial excellence.

Autumn 2015 • Vol. 13 No.2 • inSiDE

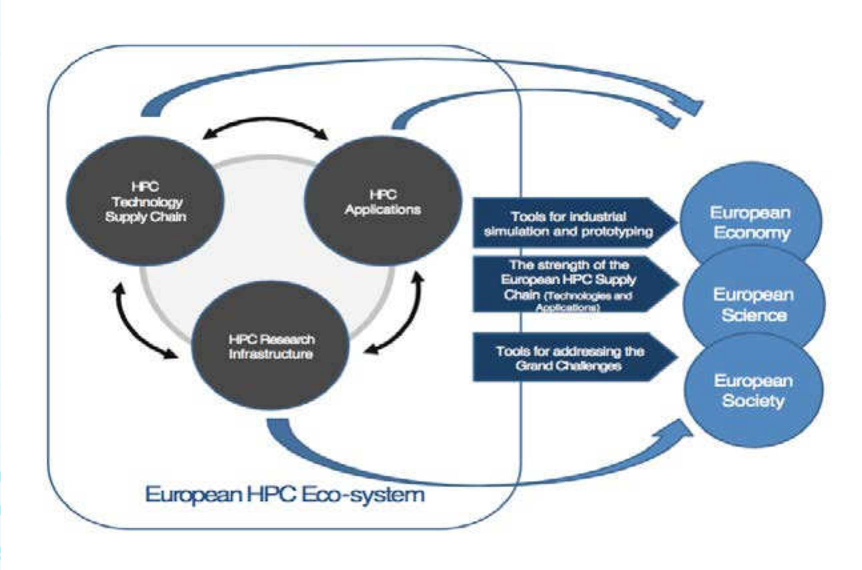


Figure 2: The European HPC Eco-system.

• Christoph Niethammer[1]
• José Gracia[1]
• Bernd Mohr[2]
• Brian J. N. Wylie[2]

[1] University of Stuttgart (HLRS), Germany

[2] Jülich Supercomputing Centre (JSC), Germany

- It improves competitiveness for the Centre's customers by generating a tangible Return on Investment (ROI) in terms of savings, elimination of waste, errors, and delays by making their applications leaner and issue-free.
- As the Centre represents the European world-class expertise in this area, its deployment will strengthen Europe's leading position in the development and use of applications that address societal challenges or are important for industrial applications through better code performance and better code maintenance and availability; POP will drive the cultural shift towards focusing on the health of applications.
- The Centre's services will include training on the use of computational methods and optimisation of applications.
- The Centre will build a repository of user cases (i.e., the computing application issues resolved) which will serve as a basis for further research in the field.

According to an IDC Report, the global HPC applications market will grow by 8% between 2013 and 2018. According to the same report, HPC is a proven accelerator of economic competitiveness. With high-end supercomputers now costing $200-500 million, their ROI can be a scientific advance or corporate profit, revenues, new jobs, or retaining jobs. Also, ROI arguments will become increasingly important for funding systems, which makes any measures improving the system's cost/performance ratio increasingly appealing. The POP CoE directly addresses this issue.

## Partners

Barcelona Supercomputing Centre (BSC), High Performance Computing Center Stuttgart of the University of Stuttgart (HLRS), Jülich Supercomputing Centre (JSC), Numerical Algorithm Group (NAG), Rheinisch-Westfälische Technische Hochschule Aachen (RWTH), TERATEC (TERATEC).

## Timeframe

October 2015 – March 2018

## POP Coordination

Prof. Jesus Labarta, Judit Gimenez
Barcelona Supercomputing Center (BSC)
Email: pop@bsc.es

http://www.pop-coe.eu

## Acknowledgements

This project is supported by the European Commission under H2020 Grant Agreement No. 676553

contact:
Christoph Niethammer,
niethammer@hlrs.de





# ESSEX – Equipping Sparse Solvers for Exascale

The Priority Programme 1648 "Software for Exascale Computing" (SPPEXA) of the German Research Foundation is approaching the end of its third of six years. 13 projects started in January 2013 to address various challenges of exascale computing. In this issue, we present project ESSEX.

The major objective of the ESSEX project is to develop an Exascale Sparse Solver Repository (ESSR) for large sparse eigenvalue problems motivated by quantum physics research and apply it, exemplarily, to graphene-based structures and topological insulators. To this end various rather general aspects of sparse eigenvalue problems have to be addressed: Computation of (i) the minimal and the maximal eigenvalue, (ii) a block of eigenpairs at the lower end or in the middle of the spectrum, and (iii) high quality approximations to the complete eigenvalue spectrum. Classic and novel numerical schemes are implemented and optimized for efficient use of heterogeneous supercomputers. Complementing related sparse solver library developments (e.g. Trilinos [1] or sparse MAGMA [2]), the ESSEX project pursues a coherent co-design of all software layers where a holistic performance engineering (PE) process guides code development across the classic boundaries of application, numerical method and basic kernel library (see Fig. 1). The ESSR will finally cover a wide range of quantum physics problems and provide blueprints of sparse solvers adapted and optimized for the exascale challenges of heterogeneity, extreme parallelism, energy/code efficiency and fault tolerance (FT).

Partners from all three software layers are actively involved in ESSEX: The application layer is represented by the group of PI Fehske (Physics, University of Greifswald), expertise on sparse eigensolvers is contributed by PI Basermann (Simulation&Software, DLR) and PI Lang (Applied Mathematics, University of Wuppertal), and the basic building block development (including the project-wide PE and FT activities) is pursued by PI Hager (Erlangen Regional Computing Center) and PI Wellein

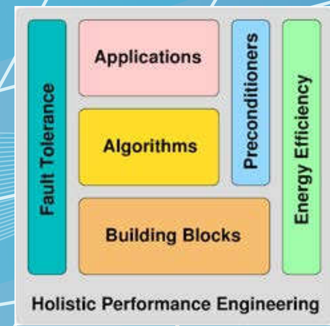


Figure 1: Overall ESSEX project structure.

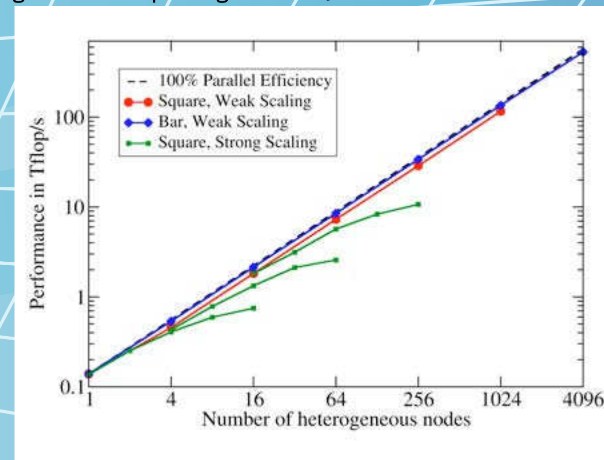


Figure 2: Weak and strong scaling performance of the full KPM solver for a topological insulator application with different geometries ("Square" vs. "Bar") on CRAY XC30. Each node comprises one Intel Xeon E5-2670 and one NVIDIA K20X card. The weak scaling runs achieve approx. 11% of LINPACK performance.

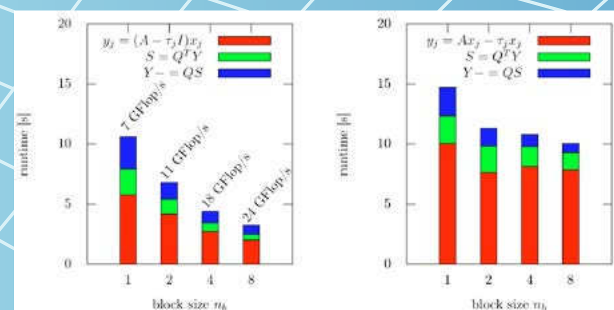


Figure 3: Runtime of the basic JADA operations for a representative ESSEX matrix with different vector block sizes on one Intel Xeon E5-2670 processor. The ESSR implementation (left panel) outperforms the Trilinos basic building block library (right panel) by up to 3x. See [8] for details.