Score-E – Scalable Tools for Energy Analysis and Tuning in HPC

For some time already, computing centers feel the severe financial impact of energy consumption of modern computing systems, especially in the area of High-Performance Computing (HPC). Today, the share of energy already accounts for a third of the total cost of ownership, see Fig. 1, and is continuously growing.

The main objective of the Score-E project, funded under the third "HPC software for scalable parallel computers" call of the Federal Ministry of Education and Research (BMBF), is to provide user-friendly analysis and optimization tools for the energy consumption of HPC applications. These tools (Scalasca [2], Vampir [3], Periscope [4], and TAU [5]) will enable software developers to investigate the energy consumption of their parallel programs in detail and to identify program parts with excessive energy demands, together with suggestions on how to make improvements and to evaluate them quantitatively. In addition, the project will develop models to describe not directly measurable energy-related aspects and a powerful visualization of the measurement results.

Measurement of energy consumption will take advantage of hardware counters like Intel's "Running Average Power Limit" (RAPL) counters introduced with the Sandy Bridge architecture, IBM's

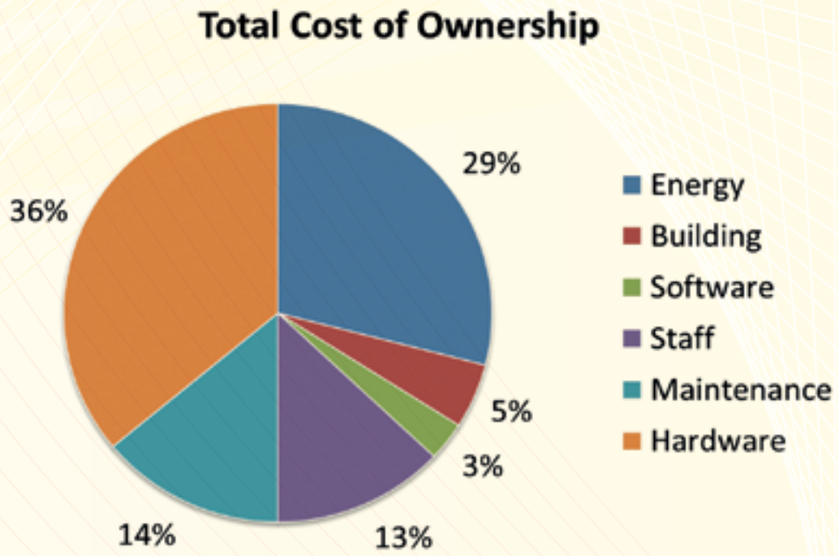

Figure 1: Total Cost of Ownership of a typical supercomputer according to a study by RWTH [1].

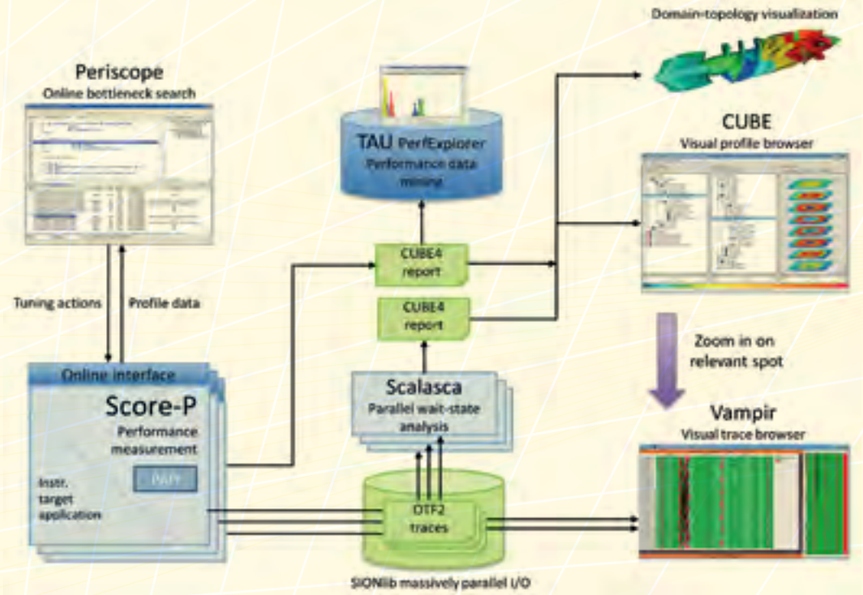

Figure 2: Interaction of the tools Persicope, Scalasca, Vampir, TAU and Scalasca's profile browser CUBE via the data exchange formats OTF2 for traces and CUBE4 for profiles.

Blue Gene/Q application programming interfaces to query power consumption on a node-board level, as well as system specific power measurement infrastructure on the node and rack level e.g., on SuperMUC. The new energy-related metrics can not only be visualized by the performance tools mentioned above, but also be used to trigger specific tuning actions during runtime by utilizing Score-P's online access interface.

With regards to visualization the Score-E project will investigate domain-topology visualizations as well as linked, multiple-view data presentation to gain deeper insight of the performance and energy behavior of complex applications.

At the same time, the project will further develop and maintain the community instrumentation and measurement system Score-P, see Fig. 2 [6], which forms the common base of all four tools mentioned above. To serve a broad user base, both, within the Gauss Alliance and beyond, most of the performance tools are released free of charge to the community under an open-source license. Only Vampir, due to its sophisticated user interface, is distributed commercially.

One particular task in the development will be the extension of the "Scalable I/O library for parallel access to task-local files" (SIONlib [7]) to support not only MPI plus basic OpenMP but arbitrary programming models and heterogeneity by providing a generic callback-based interface as well as a key-value scheme to handle a variable number of tasks per process. This extension widens the applicability of SIONlib to all

programming models currently in use in HPC.

The software products are accompanied by training and support offerings through the Virtual Institute–High Productivity Supercomputing (VI-HPS [8]), and will be maintained and adapted to emerging HPC architectures and programming paradigms beyond the lifetime of the Score-E project itself.

Besides the evident economic and environmental benefits in terms of energy, Score-E will also empower the optimized programs to unlock new scientific and commercial potentials.

The academic project partners in LMAC are the Jülich Supercomputing Centre, the German Research School for Simulation Sciences, RWTH Aachen University, TU Dresden and TU Munich. The industrial partner GNS mbH, a private company that specializes in services related to metal forming simulations, such as mesh generation for complex structures and finite element analyses, coordinates the project.

In addition, the University of Oregon, an associated partner, complements the Score-E objectives with corresponding extensions to the performance tool TAU. Further associated partners are Engys UG, who specializes in the application, support and development of Open Source Computational Fluid Dynamics (CFD) software and Munters Euroform which its expertise in engineering droplet separation systems for various industrial purposes.

For more information see:
http://www.vi-hps.org/projects/score-e
and http://www.score-p.org

- Christian Rössel [1]
- Bernd Mohr [1]
- Michael Gerndt [2]

[1] Jülich Supercomputing Centre (JSC)

[2] Technische Universität München

Acknowledgements

The Score-E project is funded by the German Federal Ministry of Research and Education (BMBF) under Grant No. 01IH13001.

References

[1] Bischof, C., an Mey, D., Iwainsky, C. Brainware for green HPC Computer Science–Research and Development, 27(4):227–233, 2012.

[2] http://scalasca.org/

[3] http://www.vampir.eu/

[4] http://www.lrr.in.tum.de/~periscop/

[5] https://www.cs.uoregon.edu/research/tau/home.php

[6] http://www.score-p.org

[7] http://www.fz-juelich.de/ias/jsc/EN/Expertise/Support/Software/SIONlib/_node.html

[8] http://www.vi-hps.org

Christian Rössel,
c.roessel@fz-juelich.de

UNICORE 7 Released

The UNICORE middleware suite is well-established as one of the major solutions for building federations and e-infrastructures, with a history going back to 1996 [1]. It is in worldwide use in HPC oriented infrastructures, for example PRACE, the US-project XSEDE [2] and in national grid initiatives such as PL-Grid. This spring, UNICORE 7 was released, which is the first major release since UNICORE 6.0 in August 2007.
It is no paradigm change as was the change from UNICORE 5 to UNICORE 6. Instead, UNICORE 7 is built on the same ideas and principles as UNICORE 6, and the two versions are compatible. We decided to make this a major release due to a number of improvements that serve to put the software on an updated technological basis.

As the most prominent change, we updated the internal web services stack to use the Apache CXF framework [3], which is the most advanced and mature Java services stack available today. This allows to build both WS/SOAP services as are currently used in UNICORE and RESTful services that will become more and more important in the future.

As the major new feature, we have added the possiblitity to deploy and run UNICORE in a way that end-users do not need certificates, using the Unity group management and federated identity solution [4]. Instead of using X.509 certificates to identify themselves, UNICORE clients request a signed Security Assertion Markup Language (SAML) [5] document from the Unity service, which is validated by the UNICORE services to assert the user's identity. Nevertheless, the strong client authentication based on X.509 certificates is still available and will continue to be supported in future releases.

Apart from the enhanced web services container and improved security stack, there are a number of other new features. For example, a new data-oriented processing feature allows to define data processing via user-defined rules. Jobs can be restarted easily, and data staging now supports wildcards.

Several changes have been made to improve the performance of UNICORE 7. For example, security sessions have been introduced to reduce the amount of XML data transferred between client and server, also reducing the CPU time required to process the XML messages. Several new batch operations have been added, for example allowing to delete multiple files or to check the status of many jobs using a single request/reply web service call. In data staging, the transfer of directories or multiple files has been optimized. Now, multiple files can be transferred in a single session, greatly using the overhead. This works especially well in conjunction with the UFTP high-performance data transfer protocol.

Together with the UNICORE 7.0 release, a first version of the new UNICORE Portal component was made available. This serves the increasing demand of users and infrastructure operators for a web-based access to UNICORE