

SPXXL - IBM/Lenovo User Group Meeting @ LRZ

SPXXL was founded as a worldwide user group for large-scale, High Performance Computing centers using IBM hardware. With the divestiture of IBM's x86-based server business to Lenovo in 2014, SPXXL dropped the vendor requirement from its bylaws and opened up its membership to any High Performance Computing (HPC) organization. SPXXL now provides advanced content and collaboration among multiple vendors – including IBM, Lenovo, Intel, Mellanox, and others – along with many of the world's largest scientific and technical computing centers.

Unlike other vendor-organized user groups, SPXXL is a self-organized, self-sufficient entity, registered in the United States and the State of California as a 501(c)(6) non-profit corporation. Members and affiliates participate actively in SPXXL meetings and cover their own costs for participating. At the moment SPXXL has 45 member sites. The goal of the organization is to work together with vendors to increase the capabilities and advance the technology of large-scale, parallel technical computing hardware and software and to provide guidance to vendors on essential development and support issues for HPC at scale. Some of the areas covered are: Applications, Code Development Tools, Communication, Networking, Parallel I/O, Resource Management, System Administration, Security, and Training. The group addresses topics across a wide range of issues that are important to maximizing performance

and energy efficiency of scientific/technical computing on highly scalable parallel systems.

The last SPXXL Winter Meeting took place from February 15 - 19 at the Leibniz Supercomputing Centre in Garching near Munich and was focussed around the Lenovo HPC roadmap. It included presentations from Mellanox about the latest Infiniband developments and Intel about their upcoming CPUs and OmniPath interconnect. Member sites also gave presentations on research projects, best practices, and the latest developments in system administration and user support to round out the week-long programme.

With the rotational elections at the end of the meeting Dr. Michael Stephan (JSC) was elected as new president of SPXXL.

References

[1] www.spxxl.org

contact:
Michael Stephan,
m.stephan@fz-juelich.de

• Michael Stephan

Jülich
Supercomputing
Centre (JSC),
Germany

New PATC Course: Introduction to hybrid Programming in HPC

Most HPC systems are clusters of shared memory nodes. These SMP nodes can be small multi-core CPUs up to large many-core CPUs. Parallel programming may combine the distributed memory parallelization on the node interconnect (e.g., using MPI) with the shared memory parallelization inside of each node (e.g., using OpenMP or MPI-3.0 shared memory).

As such hybrid programming techniques are getting more and more important in HPC, GCS as one of the 6 European PRACE Advanced Training Centres (PATC) has extended its curriculum by a new PATC course on hybrid programming techniques, which took place at LRZ on January 14, 2016 for the first time. The new course attracted around 40 international participants.

Similar tutorials about hybrid programming have been very successfully presented by the two lecturers Dr. habil. Georg Hager (RRZE, winner of the "Informatics Europe Curriculum Best Practices Award: Parallelism and Concurrency") and Dr. Rolf Rabenseifner (HLRS, member of the steering committee of the MPI-3 Forum) during various supercomputing conferences in the past, but have never been presented as a course with hands-on sessions in a GCS centre before.

The course analysed the strengths and weaknesses of several parallel programming models on clusters of SMP nodes. These models were compared



Figure 1: Participants of the new PATC Course: "Introduction to hybrid programming in HPC" @ LRZ, January 14, 2016

with various hybrid MPI+OpenMP approaches and pure MPI. Numerous case studies and micro-benchmarks were presented to demonstrate the performance-related aspects of hybrid programming. Tools for hybrid programming such as thread/process placement support and performance analysis were presented in a "how-to" session. Hands-on exercises gave attendees the opportunity to try the new MPI shared memory interface and explore some pitfalls of hybrid MPI+OpenMP programming.

Due to the great success of the course, it will be repeated at HLRS on June 13, 2016 and at LRZ on January 12, 2017.

References

[1] https://www.lrz.de/services/compute/courses/x_lecturenotes/hybrid_programming_hpc/

contact:
Volker Weinberg,
Volker.Weinberg@lrz.de

• Volker Weinberg

Leibniz
Supercomputing
Centre (LRZ),
Germany