JSC hosted the 4th JLESC Workshop

• Robert Speck

Jülich Supercomputing Centre (JSC), Germany

From December 2 to 4, the 4th JLESC Workshop took place at the Gustav-Stresemann-Institut in Bonn-Bad Godesberg, organized this time by JSC. This event was the second in 2015 of the biannual meetings of the Joint Laboratory on Extreme Scale Computing (JLESC) and the first one hosted by JSC. The Joint Laboratory brings together researchers from the Institut National de Recherche en Informatique et en Automatique (Inria, France), the National Center for Supercomputing Applications (NCSA, USA), Argonne National Laboratory (ANL, USA), Barcelona Supercomputing Center (BSC, Spain), RIKEN AICS (Japan) and JSC.

The key objective of JLESC is to foster international collaborations on state-of-the-art research related to computational and data focused simulation and analytics at extreme scales. Within JLESC, scientists from many different disciplines as well as from industry address the most critical issues in advancing from petascale to extreme scale computing. The collaborative work is organized in projects between two or more partners. This includes mutual research visits, joint publications and software releases. Every six months, all JLESC partners meet during a workshop to discuss the most recent results and to exchange ideas for further collaborations.

With more than 100 scientists and students from the six JLESC partners, the meeting in Bad Godesberg covered a broad range of topics crucial for today's and tomorrow's supercomputing. Together with the other participants, scientists and PhD students from JSC and German partner universities could catch up on cutting-edge research from the fields of resilience, I/O and programming models as well as numerical methods, applications, big data and performance tools. Besides their talks, the participants had time for fruitful discussions about their on-going and future research during project meetings, open microphone sessions and a social event. In addition, PhD students and postdocs could attend the first JLESC "young scientist dinner" to exchange ideas in a less formal (and very well-received) setting.

Organized by Inria, the next JLESC workshop will take place at ENS Lyon from June 27 to 29, 2016 continuing this successful series of internationally recognized and valued meetings. This workshop is followed by a two-day summer school on resilience. For more information on JLESC, the workshops and JSC's participation, visit http://www.fz-juelich.de/ias/jsc/jlesc

contact:
Robert Speck,
r.speck@fz-juelich.de

7th Blue Gene Extreme Scaling Workshop

Feedback from last year's very successful workshop motivated the organization of a three-day extreme scaling workshop February 1-3, 2016 at Jülich Supercomputing Centre (JSC), continuing the series started in 2006. The entire 28-rack JUQUEEN Blue Gene/Q with 458,752 cores was reserved for over 50 hours to allow eight selected code teams to investigate and improve their application scalability, assisted by staff of JSC Simulation Laboratories and Cross-Sectional Teams. Scalasca/Score-P and Darshan were employed to profile application execution performance, focussing on MPI/OpenMP communication and synchronisation as well as file I/O.

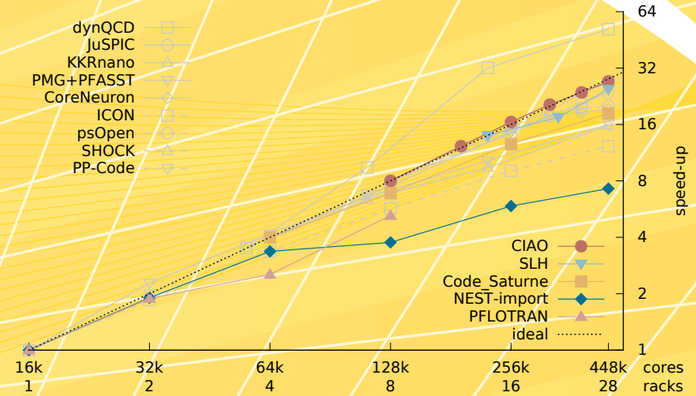

Figure 1: Strong scaling of workshop application codes compared to existing High-Q Club members

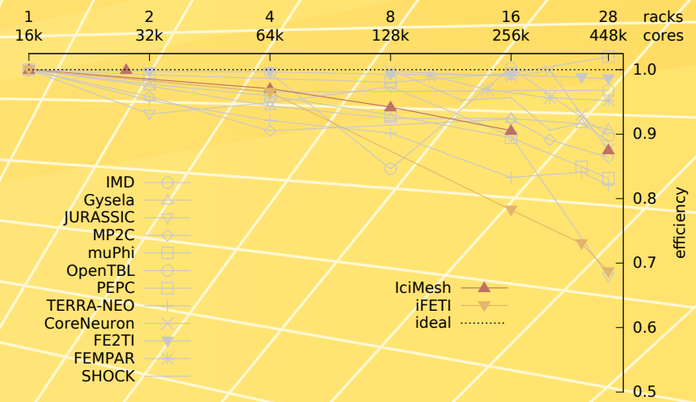

Figure 2: Weak scaling of workshop application codes compared to existing High-Q Club members

Code_Saturne from STFC (Daresbury) and **Seven-League Hydro** from HITS (Heidelberg) were both able to display good strong scalability and thereby become candidates for High-Q Club membership. Both used 4 OpenMP threads per MPI process, over 1.8 million threads in total. Existing members, **CIAO** from RWTH-ITV and **iFETI** from University of Cologne and TU Freiberg, were able to show that they had additional solvers which also scaled acceptably using purely MPI, and in-situ visualization was demonstrated with a CIAO/JUSITU/VisIt simulation using 458,752 MPI processes running on 28 racks.

Two adaptive mesh refinement libraries, **p4est** from University of Bonn and **IciMesh** from Ecole Centrale de Nantes, showed that they could respectively run with 917,504 and 458,752 MPI ranks, but both encountered problems loading large meshes. Parallel file I/O limitations also prevented large-scale executions of the FZJ IEK-6/Amphos21 **PFLOTRAN** subsurface flow and reactive transport code, however, a **NEST-import** HDF5 module developed by the EPFL Blue Brain Project could be optimized to use collective MPI file reading calls to load and connect 1.9 TB of neuron and synapse data and enable large-scale

data-driven neuronal network simulations with 458,752 threads.

Detailed workshop reports provided by each code-team, and additional comparative analysis to the 25 High-Q Club member codes, are available in a technical report. Figures 1 and 2 show the strong and/or weak scalability of this year's workshop codes. Despite more mixed results than the previous workshop, we learnt more about application file I/O limitations and inefficiencies which continue to be the primary inhibitor to large-scale simulations, and all of the participants found the workshop to have been very valuable.

- Dirk Brömmel
- Wolfgang Frings
- Brian J.N. Wylie

Jülich Supercomputing Centre (JSC), Germany

References

[1] The High-Q Club, http://www.fz-juelich.de/ias/jsc/high-q-club

[2] Brömmel, D., Frings, W., Wylie, B.J.N. (eds.), JUQUEEN Extreme Scaling Workshop 2015, Tech. Report FZJ-JSC-IB-2015-01, https://juser.fz-juelich.de/record/188191

[3] Brömmel, D., Frings, W., Wylie, B.J.N. (eds.), JUQUEEN Extreme Scaling Workshop 2016, Tech. Report FZJ-JSC-IB-2016-01, https://juser.fz-juelich.de/record/283461

contact:
Brian Wylie,
b.wylie@fz-juelich.de

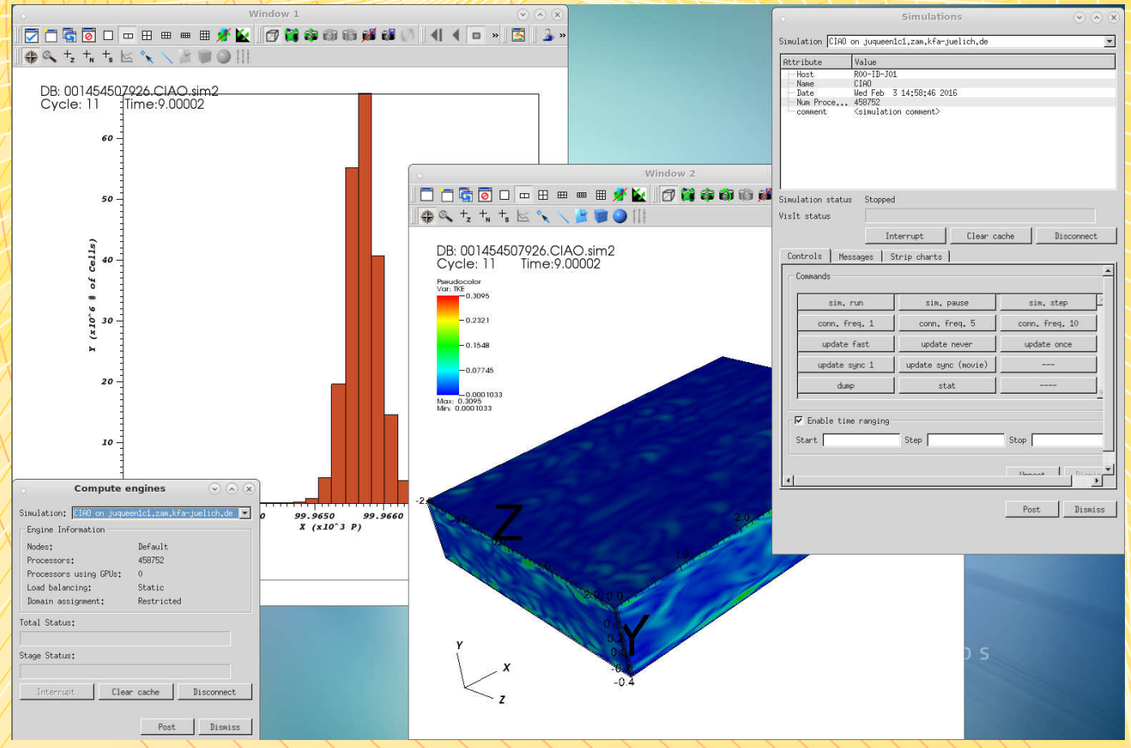

Figure 3: Visualization client showing CIAO/JUSITU/VisIt with 458,752 MPI processes on JUQUEEN

JURECA Porting and Tuning Workshop

In the past couple of years, Jülich Supercomputing Centre (JSC) has been running Porting and Tuning Workshops on its highly scalable BlueGene/Q system JUQUEEN. These workshops attracted up to 47 participants from our wide, European user base and focussed on the specialized hardware and software of BlueGene/Q.

The expected trend for the near future is a continued increase in complexity for the supercomputers to come. A precursor of this trend is the latest system installed in Jülich, the general purpose cluster JURECA, the "Jülich Research on Exascale Cluster Architectures" system. While it is comprised of standard Intel processors and thus appears to be a off-the-shelf computer, with a high-speed Mellanox EDR InfiniBand network, nodes with varying amounts of memory and additional 75 GPU nodes with 2 NVIDIA K80 accelerators each it is not. The complexity starts with the Intel Xeon chips supporting simultaneous multi-threading with up to 48 threads per node, each of them being able to use FMA instructions and wide SIMD vectors. The latter being essential to reach peak performance. Ideally, this can be combined with the available GPUs, adding another way to parallelize codes and increasing the complexity when it comes to coordinating distributed and shared memory parallelization with executing kernels on the accelerators and transferring the necessary data.

The workshop will take place June 6-8., its goal will be making the users aware of the increasing effort to reach peak performance and suggesting possible routes for this task. The topics will include best practices for JURECA, possibilities for visualization, and scientific big data analytics. We will also cover efficient I/O and ways to achieve multi-threading and vectorization. To facilitate GPU programming, we will compare OpenCL, OpenACC and CUDA. Special focus will be put on node-level performance and performance analysis.

Since this is a very broad range of topics, we will only have time to introduce the ideas and highlight their applicability. At the heart of the workshop will be extensive hands-on sessions with the participants' codes, aimed at helping with porting applications to JURECA and understanding performance bottlenecks. This will be supervised by members of staff from JSC's Simulation Laboratories and cross-sectional teams Application Optimization, Performance Analysis, and Mathematical Methods and Algorithms. At the end of the workshop the participants should have their codes running on JURECA and have a clear picture on how to improve the performance.

- Dirk Brömmel

Jülich Supercomputing Centre (JSC), Germany

contact:
Dirk Brömmel,
d.broemmel@fz-juelich.de