MAXI – Multi-System Application Extreme-Scaling Imperative

Dirk BRÖMMEL [a], Wolfgang FRINGS [a], Brian J. N. WYLIE [a]
[a] *Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, Germany*

Future supercomputer systems are expected to employ very large numbers of processors/cores and possibly to become more heterogeneous. Application scalability has therefore become a critical requirement, and with significant differences between supercomputer architectures experience with multiple computer systems is imperative.

In preparation, research is done on how to make the coming resources more accessible to users, e.g. new programming models, performance and debugging tools, and algorithmic improvements are developed. At the same time, numerous application code teams have already demonstrated the ability to exploit massive process/thread parallelism or adapt to limited compute-node memory.

While workshop and conference contributions from the aforementioned more theoretical research side are abundant, the view of application developers is scarce in comparison. The goal of this mini-symposium was to alleviate this situation and give application code teams a venue to report their outstanding achievements in both strong and weak scalability. In particular, we encouraged the participants to also report on the tools they used to get insight into performance, the issues encountered which needed to be resolved, and open issues that still need to be addressed. Special focus on application performance on multiple state-of-the-art systems and a comparison between those systems was desired.

Ten presentations were accepted for this mini-symposium — seven from code developers on experience scaling their applications from the fields of Earth modelling, computational fluid dynamics, engineering and neuroscience, complemented by three from supercomputer centres surveying highly-scaling applications on their systems — from which eight contributed papers for the proceedings.

In his presentation "Earthquake simulations at Petascale – Optimizing high-order ADER-BG in `SeisSol`," Michael Bader from the Technical University of Munich gave a comprehensive overview of the end-to-end optimisation of the `SeisSol` earthquake simulation code, which performs multiphysics simulations of the dynamic rupture process couples to seismic wave propagation, reporting on experience and scalability tests on *Tianhe-2*, *Stampede* and *SuperMUC*. `SeisSol` was a finalist for the 2014 Gordon Bell Prize reaching 8.6 PFLOPS on 8192 compute nodes of *Tianhe-2* based on Intel Xeon Phi processors.

Computational fluid dynamics on IBM Blue Gene/Q systems was addressed by two codes. Panagiotis Hadjidoukas of ETH Zürich presented "High throughput simulations of two-phase flows on Blue Gene/Q" with the `CUBISM-MPCF` code which won the 2013 ACM Gordon Bell Prize for peak performance, used to simulate cavitation collapse dynamics with 13 trillion finite-volume cells on 1.6M cores of *Sequoia* at 14.4 PFLOPS.

Jens Henrik Göbbert from the Jülich Aachen Research Alliance presented "Direct numerical simulation of fluid turbulence at extreme scale with psOpen" featuring scalability results on the 458,752 cores of *JUQUEEN* obtained from the 2015 JUQUEEN Extreme Scaling Workshop.

Two other workshop participants also presented their successful scaling to 458,752 cores of *JUQUEEN*. Aleksandr Ovcharenko of the EPFL Blue Brain Project presented "Simulating morphologically detailed neuronal networks at extreme scale" with up to 155 million neurons using CoreNeuron software isolated from the NEURON simulator and optimized with reduced memory footprint and three-level node, thread and vector parallelisation. Martin Lanser from the University of Cologne presented "FE2TI: Computational scale bridging for dual-phase steels," including recent scaling results with up to 786,432 cores on the *Mira* BG/Q at ANL of simulations of two- and three-dimensional nonlinear and micro-heterogeneous hyperelasticity problems.

The performance of two codes with different simulation capabilities using Lattice Boltzmann methods were compared on the three German tier-0 computer systems *JUQUEEN*, *Hermit/Hornet* and *SuperMUC* phases 1 & 2 at the largest scales. "waLBerla, an ultra-scalable multi-physics simulation framework for piecewise regular grids" was presented by Christian Godenschwager from the University of Erlangen-Nuremberg. "Performance of a Lattice Boltzman solver Musubi on various HPC systems" presented by Jiaxing Qi from the University of Siegen additionally included comparison with the *Kabuki* NEC SX-ACE multi-core vector supercomputer at HLRS which achieved the best single node performance.

The results of the seven international application code teams that participated in the 2015 *JUQUEEN* Extreme Scaling Workshop, and which were all successful within 24 hours in scaling to its full 458,752 cores, were reviewed by Dirk Brömmel of JSC in his presentation "Extreme-scaling applications 24/7 on *JUQUEEN* Blue Gene/Q." The associated High-Q Club now has a wide variety of over 24 codes (five in the mini-symposium) demonstrating extreme scalability.

Nicolay Hammer of Leibniz-Rechenzentrum (LRZ) presented "Extreme scale-out of *SuperMUC* phase 2, lessons learned," showing that of the 14 code teams given a month of early access to the phase 2 of *SuperMUC* at LRZ, most were able to improve scalability achieved in prior years' scaling workshops to successfully run on its 86,016 Intel Haswell cores, including SeisSol and three High-Q Club codes. Comparison of this form of scaling activity with that of short dedicated extreme scaling workshops was discussed, identifying its benefits when a new system is installed with a significant existing community of scalable applications.

Finally, in the presentation "K-scale applications on the *K* computer and co-design effort for the design and development of post-K" by Miwako Tsuji of RIKEN AICS, initial access to the *K* computer was similarly dedicated to selected code teams with scalable applications, two of which became ACM Gordon Bell Prize winners and were examined in detail. Experience with applications from nine priority research areas resulted in the selection of nine application codes which are now being used to co-design the successor to the *K* computer.

The organisers want to thank all who contributed a presentation and paper as well as everyone taking part in the discussions during the mini-symposium.