

ITEA Gold Award for ParMA

GCS is happy to announce that the European ITEA2 research project ParMA (Parallel Programming for Multicore Architectures) received the "ITEA Achievement Award 2010 in Gold". With this award, ITEA rewards high-level technical contributions based on European collaborations providing significant results while promoting ITEA and its aims. The ITEA2 Board was impressed by the innovation and fast exploitation in ParMA which developed advanced

technologies to exploit multicore architectures for High Performance Computing. For GCS, teams from HLRS and JSC are participating in ParMA. HLRS was work package leader and acted as German coordinator for the project. An important contribution to ParMA was the UNITE development tool package designed and implemented by JSC which includes a full set of interoperable tools for advanced debugging and performance analysis including the JSC tool Scalasca.

Understanding the Formation of Wait States in parallel Programs

• Bernd Mohr

Jülich
Supercomputing
Centre

With today's supercomputers featuring tens of thousands of cores, writing efficient codes that exploit all the available parallelism becomes increasingly difficult. Load and communication imbalance, which frequently occurs during simulations of irregular and dynamic domains – typical of many engineering codes – presents a key challenge to achieving satisfactory scalability. Even delays of single processes may spread wait states across the entire machine, and their accumulated duration can constitute a substantial fraction of the overall resource consumption. In general, wait states may propagate across process boundaries along far-reaching cause-effect

chains before they materialize at a synchronization point much later in the program.

To better understand how the effects of such imbalanced behaviour slow down program execution, David Böhme, a Ph.D. candidate at the Jülich Supercomputing Centre, and his colleagues developed a scalable technique that analyzes the formation of wait states and attributes their costs in terms of resource waste to their original cause. Building on earlier work by Meira, Jr. et al. [1], the researchers take execution traces of parallel programs and replay the recorded communication. A first replay in forward direction

identifies wait states and measures their duration. A second replay, performed in backward direction, traces these wait states back to the imbalance responsible for their occurrence, letting their costs travel along the reversed cause-and-effect chain until they can eventually be mapped onto their root cause. Since the replay occurs in parallel, it was possible to demonstrate the new approach with up to 65,536 processes. An article describing the idea along with experimental results [2] won the Best Paper Award at the International Conference on Parallel Processing (ICPP) 2010 in San Diego, California. To allow more target-oriented optimizations of imbalance phenomena in the daily practice of application tuning at HPC centres, the new method is currently being integrated into Scalasca [3] (Fig. 1), a performance-analysis tool developed at the Jülich Supercomputing Centre and the German Research School for Simulation Sciences in Aachen.

David Böhme's dissertation project is funded through a scholarship from the Aachen Institute for Advanced Study in Computational Engineering Science (AICES), a graduate school at RWTH Aachen University established in November 2006 under the auspices of the Excellence Initiative of the German state and federal governments.

References

- [1] Wagner Meira, Jr., LeBlanc, T. J., Virgílio, A. F.
Almeida: Using cause-effect analysis to understand the performance of distributed programs, in: Proc. of the SIGMETRICS Symposium on Parallel and Distributed Tools (SPDT'98), Welches, OR, USA, pp. 101-111, ACM, August 1998
- [2] Böhme, D., Geimer, M., Wolf, F., Arnold, L.
Identifying the root causes of wait states in large-scale parallel applications, in: Proc. of the 39th International Conference on Parallel Processing (ICPP), San Diego, CA, USA, pp. 90-100, IEEE Computer Society, September 2010
- [3] Scalasca: www.scalasca.org

• Felix Wolf

German
Research School for
Simulation Sciences,
Aachen

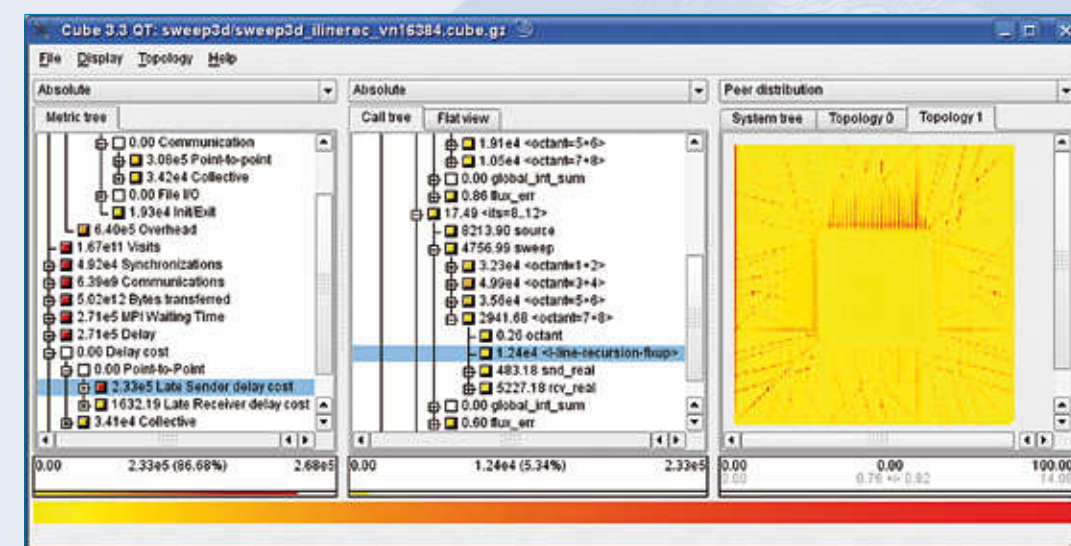


Figure 1: Using the new technique, future versions of the performance-analysis software Scalasca will allow a precise analysis of the sources and costs of wait states that occur in the wake of load imbalance.