

Optimization Approaches and Performance Characteristics of Lattice Boltzmann Kernels

Gerhard Wellein^{1,2}, Thomas Zeiser¹, Georg Hager¹, Stefan Donath¹,
Frank Deserno³, Peter Lammers⁴, and Monika Wierse⁵

Within the past decade, the lattice Boltzmann method (LBM) [1] has evolved into a promising alternative for the numerical simulation of (time-dependent) incompressible flows. Owing to the high scientific potential of LBM for large scale applications, it is the aim of this paper to demonstrate architecture-dependent optimization strategies of a Lattice Boltzmann kernel to achieve high performance on commodity “off-the-shelf” (COTS) technology, e.g. Intel or AMD processors, as well as on vector processors.

Delivering high sustained performance for scientific, memory-intensive applications is a core task in high performance computing (HPC). The main objective in the design of vector computers is to meet this challenge and bridge the continuously increasing gap between memory and processor speed [2]. Commodity “off-the-shelf” architectures do not mainly focus on HPC requirements, but dominate the HPC market due to their (often) moderate price-performance ratio.

Using an LBM benchmark kernel we point out different architecture-dependent optimization strategies and discuss single processor performance numbers for a wide range of different processor architectures ranging from IA32 compatible (Intel Xeon, AMD Opteron), superscalar RISC (IBM Power4) and IA64 (Intel Itanium 2) to classical vector (NEC SX6) and novel vector (Cray X1) architectures.

Our results demonstrate that vector systems can outperform COTS architectures by more than one order of magnitude. The NEC SX6 and Cray X1 achieve comparable performance levels on large problem sizes. Comparing different programming models of the LBM kernel shows that the Cray X1 can deliver good performance even on the standard implementation of this kernel.

References

- [1] D. A. Wolf-Gladrow, *Lattice-Gas Cellular Automata and Lattice Boltzmann Models*, Vol. 1725 of *Lecture Notes in Mathematics*, Springer, Berlin, 2000.
- [2] L. Oliker, J. C. A. Canning, J. Shalf, D. Skinner, S. Ethier, R. Biswas, J. Djomehri, R. V. d. Wijngaart, *Evaluation of cache-based superscalar and cacheless vector architectures for scientific computations*, in: *Proceedings of SC2003*, CD-ROM, 2003.

¹Regionales Rechenzentrum Erlangen (RRZE), Erlangen, Germany

²gerhard.wellein@rrze.uni-erlangen.de

³Lehrstuhl für Systemsimulation, Erlangen, Germany

⁴Höchstleistungsrechenzentrum Stuttgart, Stuttgart, Germany

⁵Cray Computer GmbH, Germany