Future Data Centers for Energy-Efficient Large-scale Numerical Simulations

Markus Geveler*, Martin Köhler†, Dirk Ribbrock*, Jens Saak†, Gerry Truschkewitz†, Peter Benner†, Stefan Turek*

Abstract

Besides energy harvesting based on renewables and energy nets adapting to this decentralised energy production, *energy consumers* have to adapt to the needs of the German Energiewende, which in principle means a necessary increase of *energy efficiency*. Mathematical Modeling, Simulation and Optimization (MSO) are rightfully considered offering powerful tools for conserving energy in industrial processes. Still, this automatically puts the focus on how energy efficient the corresponding computations can be performed.

For energy-efficient simulation, both hardware and software issues have to be addressed. Numerical simulation, for instance Computational Fluid Dynamics, has large computational needs: Technical flows arising in industrial processes are characterised by complex (multi-)physics, high accuracy requirements, large systems of locally and globally coupled equations as well as large amounts of data. The computational hardware consists of massively parallel distributed memory systems where the compute nodes are again multi-level parallel and heterogeneous. They normally comprise a high-end server CPU for instance based on the x86 architecture, accelerated by GPUs or other floating point arithmetic hardware. HPC sites of this type have huge energy requirements. Often a financial strain of the same order of magnitude as the initial total system cost is imposed over their lifetime. This makes them unaffordable especially for small and medium enterprises. There is a huge gap between pure performance and energy-efficiency: The Top500 list's best performing HPC system (dissipating power in the mega-watts range) is only ranked 84th on the corresponding Green500 list, whereas the most energy-efficient system in place only performs 160th in the metric of pure floating point performance [Meuer et al. 2015] [Feng et al. 2015]. This means, that energy efficiency on the hardware side is not a matter of pure floating point performance. Scientific software often fails to cover this fact and does not internalise power dissipation, energy consumption or energy efficiency and therefore often remains solely tuned for high performance. Fortunately, performance engineering is starting to cover the power and energy metrics [Hager et al. 2014] [Benner et al. 2013] [Benner et al. 2014].

The most obvious feature all Green500 top ten systems share is, that they rely on accelerators - mostly GPUs, but the top three even on an unconventional ARM- and FPGA- based microarchitecture. We have already shown, that the energy to solution in simulation can be reduced by using ARM-based processors designed for mobile computing exploiting the fact, that these originate from the domain of embedded computing and that they are historically optimised for energy efficiency as opposed to the compatibility-optimised commodity CPUs [Göddeke et al. 2013]. With the NVIDIA Tegra K1, a programmable embedded low power GPU becomes accessible alongside the ARM cores on one System-on-Chip (SoC), providing a jump in theoretical peak performance whilst preserving low power requirements. In this talk, we demonstrate a prototype of a future data center, combining high-end photovoltaic farming with a supercomputer consisting of low-energy components. We utilise NVIDIA Tegra K1 Processors in each compute node and add a specially designed low-energy data storage subsystem based on the Banana-Pi board. In addition, we discuss the performance engineering issues for simulation software on mobile SoC-based supercomputers in order to validate our approach. We compare the performance, energy-to-solution and energy efficiency of integral parts of PDE-based simulation on our system to the corresponding values on x86 and desktop GPUs. These benchmarks include BLAS, sparse matrix-vector multiply and direct as well as iterative solvers up to geometric multigrid.

Our system comprises 60 NVIDIA Tegra K1 SoC, that is 240 ARM CPU cores and 60 GPUs offering a theoretical peak performance of more than 21 TFlop/s at a total power dissipation of 2kW with no additional energy-cost due to an insular solar power supply and battery system. The storage backend provides up to 10TBytes of distributed filesystem space with a total energy consumption that is lower than that of a single standard server processor. In addition, the system is cheap: Including housing and power supply it only requires less than 100k EUR. Our design is autonomous and scalable. We show, that compared to commodity CPUs, -accelerators, and storage solutions, energy-efficiency is enhanced to a great extend.

With our approach, we show that the power of the mass market-driven mobile computing industry offers a way to change hosting of simulations. This, on the one hand, lowers the energy requirements of future HPC centers, and on the other hand, makes MSO accessible to more enterprises.

Keywords: energy efficient HPC, performance engineering, ARM, GreenIT

References

BENNER, P., EZZATTI, P., QUINTANA-ORTÍ, E., AND REMÓN, A. 2013. On the impact of optimization on the time-power-energy balance of dense linear algebra factorizations. In *Algorithms and Architectures for Parallel Processing*, R. Aversa, J. Kolodziej, J. Zhang, F. Amato, and G. Fortino, Eds., vol. 8286 of *Lecture Notes in Computer Science*. Springer International Publishing, 3–10.

BENNER, P., EZZATI, P., QUINTANA-ORTÍ, E. S., AND REMÓN, A. 2014. Trading off performance for energy in linear algebra operations with applications in control theory. *CLEI Electronic Journal 17*, 1, Paper 3.

FENG, W., CAMERON, K., SCOGLAND, T., AND SUBRAUMANIAM, B., 2015. Green500 List, July. http://www.green500.org/lists/green201506.

GÖDDEKE, D., KOMATITSCH, D., GEVELER, M., RIBBROCK, D., RAJOVIC, N., PUZOVIC, N., AND RAMIREZ, A. 2013. Energy efficiency vs. performance of the numerical solution of PDEs: an application study on a low–power ARM–based cluster. *Journal of Computational Physics*, 237 (Mar.), 132–150.

HAGER, G., TREIBIG, J., HABICH, J., AND WELLEIN, G. 2014. Exploring performance and power properties of modern multicore chips via simple machine models. *Concurrency and Computation: Practice and Experience*.

MEUER, H., STROHMEIER, E., DONGARRA, J., SIMON, H., AND MEUER, M., 2015. Top500 List, July. http://top500.org/lists/2015/06/.

*TU Dortmund,
email:{markus.geveler, dirk.ribbrock, stefan.turek}@math.tu-dortmund.de

†MPI for Dynamics of Complex Technical Systems, Magdeburg,
email:{koehlerm, saak, truschkewitz, benner}@mpi-magdeburg.mpg.de