Geophysical Research Abstracts
Vol. 21, EGU2019-5204, 2019
EGU General Assembly 2019

Chasing the limits of computing efficiency in terrestrial modeling: memory bandwidth

Michal Belda (1) and Luděk Kučera (2)

(1) Charles University, Faculty of Mathematics and Physics, Department of Atmospheric Physics, Praha, Czech Republic (michal.belda@mff.cuni.cz), (2) Charles University, Faculty of Mathematics and Physics, Department of Applied Mathematics, Praha, Czech Republic (ludek@kam.mff.cuni.cz)

The presentation shows that modern computers aren't well suited for simulations in meteorology and climatology, being able to use just few percents of their computing power. The community should clearly pronounce the need of accelerated research to increase unsufficient processor memory bandwidth, the main reason of the inefficiency.

Modern supercomputers execute tens or hundreds of PFlop/s (10^{15} floating point operations per second). However, there is a growing feeling that they use their immense computing power in an inefficient way when running grid-based simulations of physical processes. This is why J. Dongarra et al. prepared the HPCG (High Performance Conjugate Gradients) benchmark as a method to check the ability of a computer to perform sparse grid based simulations.

The HPCG results are very disappointing (www.hpcg-benchmark.org). With few exceptions, existing supercomputers use only about 1.5-2% of their peak computing power when running the HPCG(!!!).

The HPCG solves elliptic differential equations of the heat diffusion that are discretized as systems of linear equations with very sparse matrices. The main reason of the computing inefficiency is that the extremely high memory bandwidth of modern processors is still too small compared to their computing power.

One of the main parts of the HPCG is SpMV (Sparse Matrix - Vector) product, where only two arithmetic operations are performed per each 8 byte (double precision) matrix element, loaded from the external memory (multiplying by a vector element and adding to an accumulator) – 0.25 operations per 1 byte from the memory.

The most powerful accelerator of the date, Nvidia Volta-100, has the memory bandwidth 900 GB/s, enough to perform not more than 225 SpMV Gflop/s. However, the peak DP computing power of V-100 is 7800 GFlop/s, and hence the computing efficiency is less than 225/7800 $\approx$ 2.88%. The HPCG preconditioning is computationally similar.

Processes of the atmospheric physics are much more complex than the heat diffusion: the air is a complex mixture of gases including water in different phases, the state of a node is described by a wide vector involving temperature, pressure, humidity, etc. Thus, it has been expected that the amount of compution for one grid node would be considerably higher than in the case of the HPCG, and hence the computing efficiency would be substantially higher as well.

Our research, unfortunately, kills this hope. We investigated three important parts of the WRF (Weather Research and Forecast) program, namely the microphysics driver, the long wave radiation, and the advect, and we have found that (in the less memory-bandwidth demanding single precision (SP) configuration) the flop/byte ratio of the three WRF components is below 1.00 for realistic cache sizes.

Thus, considering Volta-100 again, the memory bandwidth is sufficient for not more than 900 WRF SP GFlop/s, while the SP peak computing power is 15,600 Gflops: the theoretical computing efficiency below 5.8%.

Since WRF is not known to be less performant than other NWP programs, it can be conjectured that similar figures apply to other programs of the atmospheric physics as well.