Modern many-core CPU architecture: a promise for future exascale super-computers.

Mathieu Lobet (1), Matthieu Haefele (1)

1) Maison de la Simulation, CEA, CNRS, Université Paris-Sud, UVSQ, Universite Paris-Saclay, F-91191 Gif-sur-Yvette, France

Exascale supercomputers will achieve more than 10^{18} floating point operations per second, 10 times the capacity of Sunway TaihuLight the current world's most powerful super-computer in China. USA and China has announced their plan to build an exascale super-computer near 2020. Achieving exascale computing facilities in the next decade will be a great challenge particularly in terms of energy consumption. They should consume no more than 20 MW which is almost the value already reached by Sunway TaihuLight. At the CPU level, many-integrated core (MIC) architectures such as Intel Xeon Phi products are now considered as a viable technology toward this goal.

Many-core processors differ from previous and current multi-core architecture in that they contain a large number of "simpler", less-powerful computing units (cores) with cache coherency. To compensate for the reduction of computing power due to clock speed, MIC processor uses wider data registers that can process multiple data in a single clock cycle (Single Instruction Multiple Data or SIMD). Recent Intel product includes a large integrated on-package shared memory called MCDRAM with a significantly higher memory bandwidth than DRAM. MIC processors have been designed for a high degree of symmetric multiprocessing (SMP) parallelism, enabling higher peak performance at lower energy consumption on a single chip than previous architecture. The Knights Landing (KNL) is the last born of the family with 72 cores at 1.5 GHz, 4 threads per core, 2 VPUs of 512-bit SIMD vector registers and 16 Gb of MCDRAM. It achieves a peak power of 3,5 Tflops in double precision for an energy consumption of 250 W. Many-core processors now equip 4 of the top 10 most-powerful supercomputers including the two world most powerful ones in China. The next most-powerful supercomputer in USA (> 100 PFlops) called Aurora will be built with the next generation of Xeon Phi of code name Knights Hill (KNH). Marconi (Italy) is currently the most powerful European KNL supercomputer ranked 14^{th} in the top500.

In order to obtain the best performances on Xeon Phi architectures, algorithms have to be designed to achieve good memory locality (at the cache and the node level) and efficient vectorization. A well-constructed hybrid multi-level parallelism is essential to deal with the high number of threads and the non-uniform memory accesses (NUMA) with the MCDRAM memory. The KNL design makes already-existing codes more sensitive to optimization issues/laziness and thus sometime inefficient on these platform. This presentation will cover the basics of the Intel Xeon Phi conception and usage focusing on KNL. We will cover the most sensitive points to address to get good performances. Optimizing code on KNL sometime requires painful changes including rethinking domain decomposition, cache usage, data-structures and communication patterns. The positive outcome is that the resulting codes will also perform better on current multi-core architectures.