ID de Contribution: **4** Type: **Non spécifié**

Matrix free conjugate gradient with Maxeler Data Flow Engine technology

lundi 4 juillet 2022 17:00 (30 minutes)

In this presentation, the implementation of a miniapp extracted from a production code in material science (Metalwalls) using Maxeler technology will be explained, after which a chip to chip comparison between a CPU, a GPU and an FPGA, as well as a scalability study on multiple FPGAs will be presented.
The core algorithm is a matrix free conjugate gradient that computes the total electrostatic energy thanks to an Ewald summation at each iteration.
The FPGA implementation using 40 bits floating point number representation outperforms the CPU implementation both in terms of computing power and energy usage resulting in an energy efficiency more than 14 times better. Compared to the GPU of the same generation, the FPGA reaches 60\% of the GPU performance while the ratio of the performance per watt is still better by a factor of 3.
Thanks to its low average power usage, the FPGA bests both fully loaded CPU and GPU in terms of number of conjugate gradient iterations per second and per watt.

Orateur: PROUVEUR, Charles (CEA)