GPUs and SPH technique, growing in parallel

Alejandro Crespo*, José Domínguez, Benedict Rogers† and Moncho Gómez-Gesteira†

* EPHYSLAB, Universidade de Vigo
Campus As Lagoas s/n 32004, Ourense, Spain
e-mail: alexbexe@uvigo.es, jmdominguez@uvigo.es, mggesteira@uvigo.es
† School of Mechanical, Aerospace and Civil Engineering (MACE), University of Manchester
M13 9PL, United Kingdom
e-mail: benedict.rogers@manchester.ac.uk

ABSTRACT

Smoothed Particle Hydrodynamics (SPH) is an ideal technique to simulate free-surface flows. SPH is ideally suited to fluid and solid mechanics with highly nonlinear deformation and is opening new avenues of activity in several areas, notably fluid-structure interaction, multi-phase flows and importantly, engineering application and design. SPH describes a fluid by replacing its continuum properties with locally smoothed quantities at discrete Lagrangian locations. Thus, the domain can be multiply-connected with no special treatment of the free surface, making it ideal for examining complicated flow situations. SPH has become increasingly popular in recent years as a novel technique to model the violent hydrodynamics in wave breaking, wave-structure interaction, etc. The technique can now be used for engineering purposes in those problems involving the complex interaction between water and structures. In general, all these problems involve large domains that should be solved with fine resolution, which makes the model expensive in terms of computational requirements. This is the reason why these codes needs to be optimized and accelerated as much as possible. High Performance Computing was used to develop a SPH code capable of performing simulations of real-life applications at a reasonable time.

The open-source code DualSPHysics (http://dual.sphysics.org) has been developed (Crespo et al., 2015), which can be used both on classic CPUs (Central Processing Unit) and novel GPUs (Graphics Processing Units). DualSPHysics has been designed to be run on multi-core CPUs, which is a relatively common resource, but also on GPUs. The GPU technology has experienced a rapid development during the last few years and constitutes a fast and cheap alternative to classical computation on CPUs. Nevertheless, a single GPU is not enough to run large domains due to memory requirements and huge execution times. Thus, a multi-GPU version of the code has also been developed. In addition, pre-processing and post-processing tools have been developed to take advantage of DualSPHysics capabilities.

CUDA (Compute Unified Device Architecture) was used to exploit the huge parallel power of present-day GPUs and several optimizations are presented for the GPU implementations; maximization of occupancy to hide memory latency, reduction of global memory accesses to avoid non-coalesced memory accesses, simplification of the neighbour search, optimization of the interaction kernel and division of the domain into smaller cells to reduce code divergence. The GPU parallel computing developed here can accelerate serial SPH codes with a speedup of 148.8x using the GPU GTX Titan. The multi-GPU approach includes CUDA and MPI (Message Passing Interface) programming languages to combine the parallel performance of several GPUs in a host machine or in multiple machines connected by a network. The multi-GPU implementation has shown an efficiency close to 100% using 128 GPUs of the Barcelona Supercomputing Center, when 8 million particles per GPU have been simulated. Moreover, an application with more than 10^9 particles is presented to show the capability of the code to handle simulations that would require large CPU clusters or supercomputers otherwise.

REFERENCES

[1] A.J.C. Crespo, J.M. Domínguez, B.D. Rogers, M. Gomez-Gesteira, S.M. Longshaw, R. Canelas, R. Vacondio, A. Barreiro, O. Garcia-Feal, “DualSPHysics: Open-source parallel CFD solver based on Smoothed Particle Hydrodynamics (SPH)”, *Computer Physics Communications*, **187**, 204-216 (2015).