A new parallelized 3D SPH model: resolution of water entry problems and scalability study

Guillaume Oger [1], David Le Touzé [2], Bertrand Alessandrini [2] and Pierre Maruzewski [3]

[1] HydrOcean, 1 rue de la Noë, 44300 Nantes, France
[2] Laboratoire de Mécanique des Fluides (CNRS UMR6598), Ecole Centrale Nantes (ECN), 1 rue de la Noë, 44300 Nantes, France
[3] Laboratoire de Machines Hydrauliques, Ecole Polytechnique Fédérale de Lausanne (EPFL), Avenue de Cour 33bis, 1007 Lausanne, Switzerland

[1] guillaume.oger@hydrocean.fr; [2] david.letouze@ec-nantes.fr; [3] pierre.maruzewski@epfl.ch

Abstract

By its ability to treat efficiently problems involving more and more complex mechanics, the SPH method (Smoothed Particle Hydrodynamics) represents a very interesting alternative method to classical mesh-based methods such as Finite-Volume, Finite-Difference or Finite-Element Methods. Presently, most of the models present in the free-surface SPH related literature are two-dimensional, and thus do not really suffer from high computational cost difficulties. But actual engineering applications are indeed three-dimensional, which dramatically increases computational costs, and finally limits the use of the SPH scheme. A parallelization technique integrating some very specific optimization tools is proposed here. The implemented model SoPHy-N has been tested in terms of acceleration and efficiency of the calculation with regards to the number of processes used, with various total particle numbers. Very encouraging results have been found on the ECN Cray XD1 cluster by using up to 32 processors. In particular, a mean efficiency of about 90% has been obtained. However, the efficiency of the parallelized model had to be validated at larger scale. This scalability study is realized on a water entry problem on the 8092 processor Blue Gene cluster of EPFL, namely on a test case involving a sphere impacting the free surface at high velocity. The model capabilities are further illustrated on case of slamming of a 3D hull.

1 Introduction

SPH is a Lagrangian meshless method able to simulate complex free-surface flows of fast dynamics, including for instance jets, sprays and reconnections of the free surface. It is therefore well suited to model complex problems such as the water entry of 3D objects of complex geometry. In this method, no mesh has to be generated nor adapted to follow the fragmentations and reconnections of the free surface which remains always precisely described by the particles in their Lagrangian motion. An other advantage is that the modelling of the air is not required to perform the simulation (though it can be also be included with no difficulty), with respect to classical techniques involving Level-Set or Volume Of Fluid interface tracking.

In this context, complex 3D problems can be simulated,

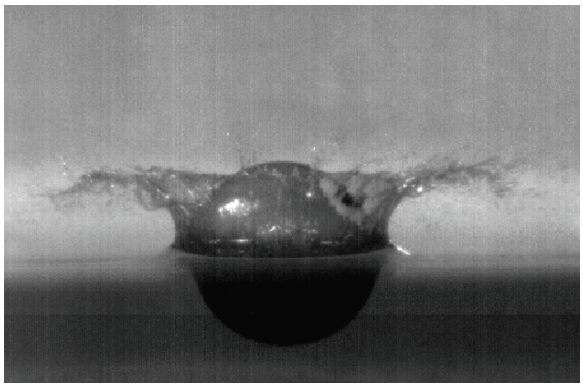

Figure 1: Billard ball impacting the free surface at 4,8 m/s (free fall), from [3].

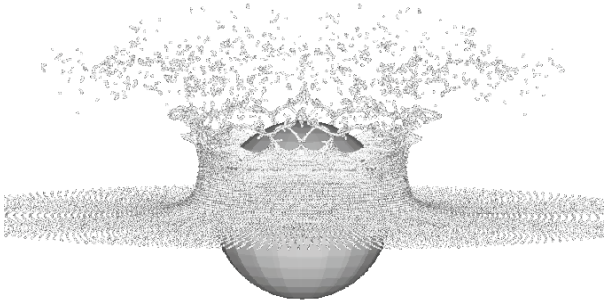

Figure 2: Illustration of a near-equivalent SPH model simulation (4.8 m/s imposed velocity).

such as the slamming of a hull shown hereafter. However, describing accurately the free surface complex evolution in such 3D problems requires the use of a large number of particles (typically millions). There is thus a need for efficiently parallelizing the models. Again, another asset of SPH makes it easier with respect to classical CFD methods: the explicit nature of the method, leading to a rather straightforward parallelization by domain decomposition.

In our model SoPHy-N developed in Nantes, significant efforts have been made in two directions (see e.g. [1]). Firstly, the accuracy of the model has been the objective of a number of developments: a Riemann solver is implemented in the model (following [2]), as well as kernel renormalization. Secondly, numerous developments have been dedicated to widen the applicability of SoPHy-N. Arbitrary complex body geometries can now be included, in free or forced motion, open problems can be treated, and a specific algorithm has been developed to be able to use particles of variable size in order to concentrate the particles in the zones of interest of the flow.

Another key of the application of the model to complex 3D problems is an efficient parallelization. In the present paper we give an outline of the developments made to reach such an objective. The main principles adopted for communications between processors are presented, followed by an evaluation of the efficiency of the proposed model, on the test case of a solid sphere impacting the free surface. This study has been realized in collaboration with EPFL, on their Blue Gene cluster. The capabilities of the model are also illustrated on the simulation of slamming of a hull, achieved on the XD1 cluster of ECN.

2 Core Implementation

The core of the SPH scheme used in the model relies on the use of a renormalized kernel and an exact Riemann solver (Godunov type scheme), extracted from the works of Vila et al (see e.g. [2]). The latter presents various advantages such as avoiding the use of the artificial viscosity required in standard SPH, decreasing numerical dissipations and increasing stability, in a way inspired from compressible finite-difference and finite-volume schemes developed from the sixties until now. Boundary conditions are imposed using ghost particles, which is the most accurate frontier treatment in our experience. Specific developments have been realized to extend this technique from a flat boundary, originally, to an arbitrary 3D shape. Further, the scheme implemented in SoPHy-N also integrates a variable smoothing length scheme, necessary to apply this model to various fields of application. Actually, this technique is particularly adapted to the test case of water entry, providing accuracy in the impact area (where particles of small size are used). The spatial distribution of particles is then slowly relaxed from the limit of this zone up to the tank border.

3 Parallelization of the SPH solver

The extension from a two-dimensional to a three-dimensional SPH model is pretty immediate. Nevertheless, for computational cost saving and memory requirement reasons, the practical implementation of such a three dimensional code implies its parallelization, in order to make possible the calculation on a cluster. This parallelization is achieved here using the standard MPI (Message Passing Interface) for inter-process communications (used conjointly with FORTRAN 90). We chose the domain decomposition as parallelization strategy, this method consisting in splitting the whole fluid domain into sub-domains, each sub-domain corresponding

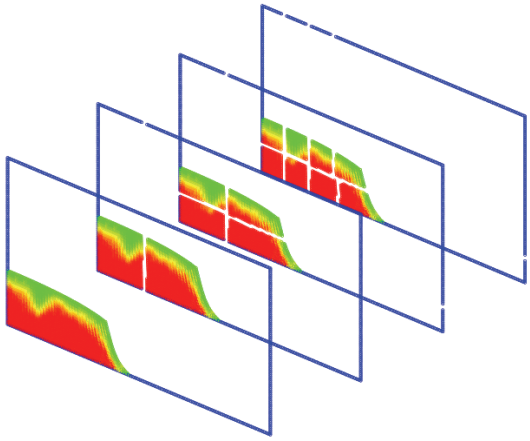

Figure 3: Successive domain splitting for a 2D calculation involving eight processors.

to one dedicated processor. In this process, a dedicated algorithm ensures to get balanced loads afterwards. An example of such a procedure is illustrated in figure 3 for the case of a two-dimensional dam breaking, for which the calculation involves eight processors. In order to make this code as simple as possible, the creation of rectangular sub-domains of optimized shape is adopted. Interactions between these various sub-domains are then achieved using MPI, by systematic particle data communications.

In this paper, for the sake of simplicity and clarity, this parallelization is discussed from a two-dimensional point of view, but its extension to the three-dimensional case is immediate and is implemented in our 3D model.

3. 1 *Parallelization description*

First, note that each domain is determined at the beginning of the calculation, its sizes being adapted during the simulation so that each process always has approximately the same number of particles to treat, in order to make the global calculation as efficient as possible. Let us consider a given process of interest. This process owns a large number of particles for which it is able to compute the interactions (discrete convolutions) everywhere except near its limits. Indeed, its neighbor processes contain the missing particles, that we call now “foreign particles”. Its neighbor processes must therefore communicate to it the foreign particle data (position, velocity, pressure...) in order to allow the process of interest to complete its calculation, namely to account for all of the neighbors of its own particles. Thus, the process of

interest has first to communicate to its neighbor processes the limits of the areas (dotted zones in figure 4) containing the foreign particles that interact with its own particles.

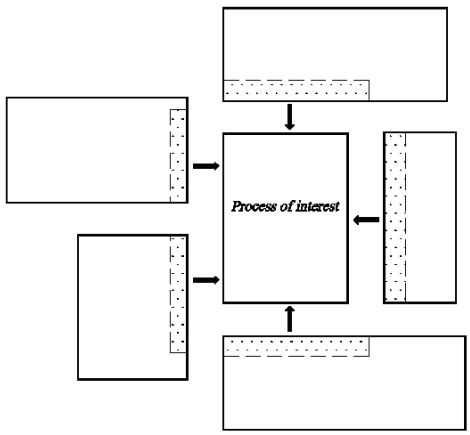

Figure 4: Process of interest with its neighbor processes.

These limits (dashed lines) are determined from the particles being the closest to the process of interest limits (see figure 5), and using the interaction radius R=2h.

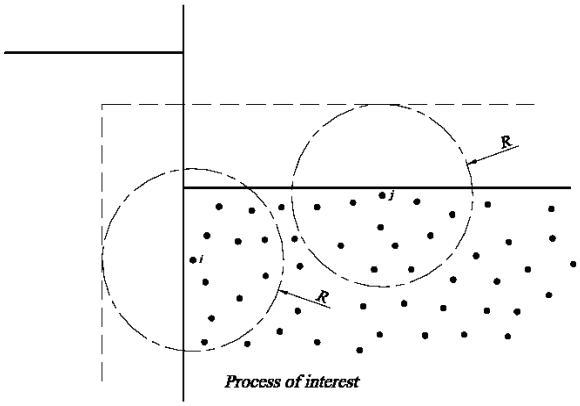

Figure 5: Limits of the foreign particle areas.

The MPI standard allows non-blocking communications. This particularity is of interest since it is possible to achieve simultaneously any operation (dependent or not on the parallelization itself) and the communications. The present scheme widely uses this advantage, which provides large optimization opportunities.

3. 2 *Scalability results*

This parallelized model gave successful results in terms of efficiency for a limited number of processes (about 90% on 32 processes) on the Cray XD1 cluster of ECN. The next step then consisted in applying this model on larger process numbers (up to 1024 processes). Such large process number tests are quite challenging, because of difficulties due to sub-domain size heterogeneities, larger foreign particle numbers (responsible for increased communication delays and memory requirements), and subtleties appearing from the variable-H scheme. These new scalability tests have been carried out on the 8092-processor Blue Gene computer of the EPFL for various particle and process numbers; acceleration and efficiency results are summarized in figures 6 and 7 respectively.

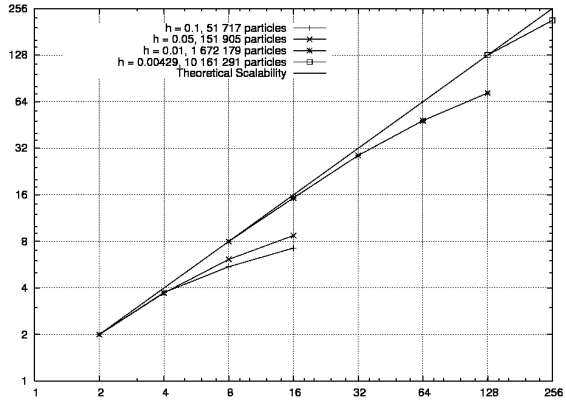

Figure 6: Acceleration of the calculation.

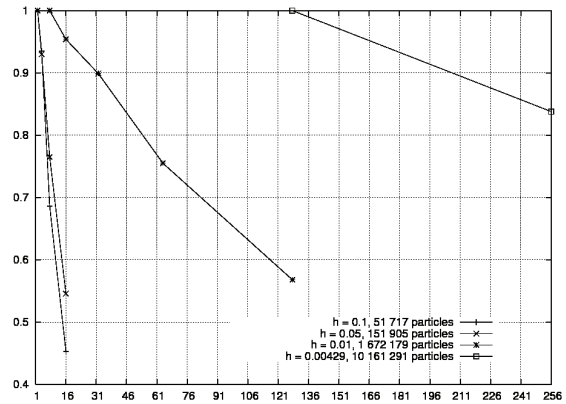

Figure 7: Efficiency of the calculation.

In these tests, the time reference value is based on 2 processes for 51,717 and 151,905 particles, on 8 processes for 1,672,179 particles, and 128 processes for 10,161,291 particles. Reference was 256 processors than, even though simulations with up to 55 million particles on 1024 processors have been performed since. As expected, for a given particle number, the efficiency decreases as the number of processes increases, due to the importance of communication delays compared to the effective calculation time, so that the use of a number of processes bigger than 16 is useless for 150,000 particles or less. For about 1,600,000, a correct efficiency is preserved with a maximum of 64 processes (about 75%). These results also reveal an efficiency higher than 80% for 256 processes with 10,000,000 particles (this efficiency should be taken into account carefully because the time reference value chosen is for 128 processes).

4 Impact of a complex 3D geometry

Our SPH model SoPHy-N is applied here to the simulation of a complex hull impacting the free surface, in the case of a forced motion. The experiments concerning this test case were performed in the Large Wave Basin (50 m x 30 m x 5 m) of the Fluid Mechanics Laboratory of ECN. 2,400,000 particles were involved in

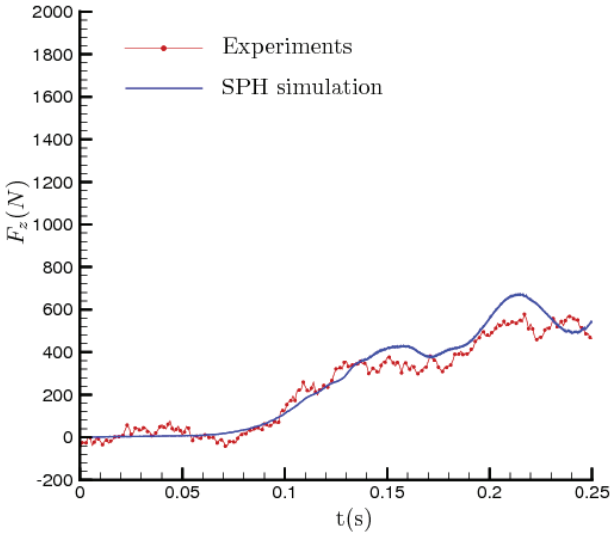

Figure 8: Vertical force time history.

the SPH simulation, achieved on 16 processes during 24 hours on the Cray XD1 cluster of ECN. Figure 9 illustrates the calculated free surface deformation evolution during the impact.

This test case represents a good opportunity to test our model on a real engineering application, while most of SPH calculations discussed in the literature remains prototypal and are usually applied on simple bi-dimensional test cases. The data of interest to be compared with the experiments concerns the vertical force time history on the ship F_z during the impact, as described in figure 9, showing a fair agreement.

5 Conclusion

In this paper, a study based on a three-dimensional parallelized SPH model applied to free surface impact simulations has been presented. A special attention has been paid to the description of the parallelization itself, before a brief analysis of the first scalability results obtained on the 8092-processor Blue Gene computer of the EPFL through the use of a maximum of 1024 processors up to now with 50,000 to 55,000,000 particles. These results are very encouraging, presenting interesting efficiency values. This work also reveals the need for improvements concerning namely the memory consumption, which remains quite subtle because of the variable-H scheme constraints. Further development is still needed, before to pursue these scalability tests, to finally reach possible calculations on 8092 processors. The resulting model should be a powerful tool based on SPH. The applicability of the model on realistic engineering problems is also illustrated on a case of slamming of a 3D hull.

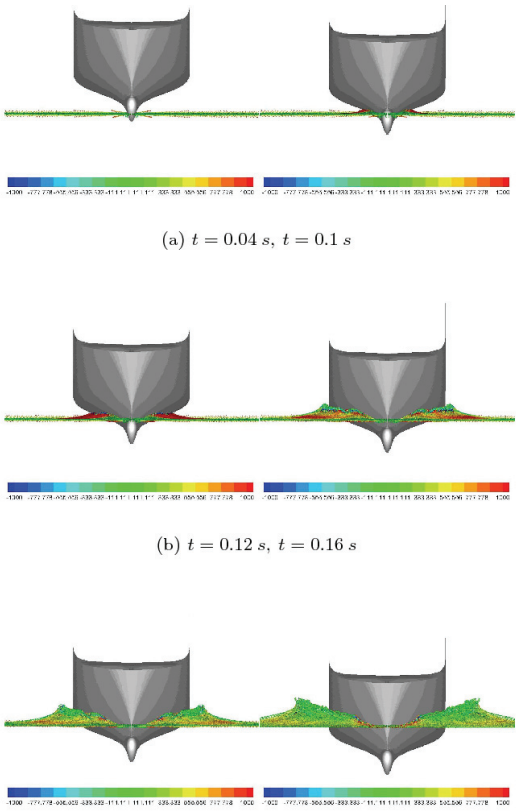

(a) $t = 0.04\ s,\ t = 0.1\ s$

(b) $t = 0.12\ s,\ t = 0.16\ s$

(c) $t = 0.18\ s,\ t = 0.24\ s$

Figure 9: Front view of the numerical free surface elevation evolution during the slamming impact.

References

[1] Oger G., Doring M., Alessandrini B., Ferrant P. An improved SPH method: towards higher-order convergence, Journal of Computational Physics, 225: 1472-1492, 2007.

[2] Vila J.P., On particle weighted methods and Smooth Particles Hydrodynamics, Mathematical Models and Methods in Applied Sciences, 9: 161-210, 1999.

[3] Laverty S.M., Experimental hydrodynamics of spherical projectiles impacting on a free surface using high speed imaging techniques, Master in science in ocean engineering at the MIT.