

Tsunami-HySEA: A Model to Simulate Tsunamis Generated by Earthquakes on Multiple GPUs

Marc de la Asunción, Manuel J. Castro, José M. González, Jorge Macías
University of Málaga, Spain

Key words: Tsunami simulation, Nested meshes, Earthquake, Multi-GPU

Introduction

HySEA (Hyperbolic Systems and Efficient Algorithms) is a high-performance software package developed by the EDANYA group at the University of Málaga, Spain, for the simulation of geophysical flows, such as tsunamis generated by earthquakes or landslides, river flows, sediment transport, turbidity currents, etc. Specifically, Tsunami-HySEA is the numerical model of the HySEA family to simulate tsunamis generated by earthquakes. The computations are performed using a single GPU or a multi-GPU cluster in order to take advantage of the massively parallel architecture of the modern graphics cards.

In Tsunami-HySEA, the three phases of an earthquake generated tsunami (deformation of the seafloor, propagation of the tsunami and inundation of the coastal areas) are performed in a single code. Time series and nested meshes with different spatial resolutions are also supported.

Tsunami-HySEA is currently being used by the Italian Istituto Nazionale di Geofisica e Vulcanologia (INGV) for its National Tsunami Early Warning System (TEWS), by the Joint Research Centre (JRC) of the European Commission in its TEWS, by the Spanish Instituto Geográfico Nacional (IGN) in its TEWS, and integrated in TRIDEC Cloud at the GFZ German Research Centre for Geosciences. Tsunami-HySEA has been extensively tested during several years, and validated and verified with the benchmark tests compiled in [6].

Numerical Scheme

Tsunami-HySEA implements the non-linear shallow water equations in spherical coordinates. For the propagation step, it implements a fast, two-step scheme similar to leap-frog written under a finite volume framework. In the inundation step, a second order TVD-WAF flux-limiter scheme, described in [1], is used. The combination of both schemes guarantees the mass conservation and prevents the generation of spurious high frequency oscillations near discontinuities generated by leap-frog type schemes.

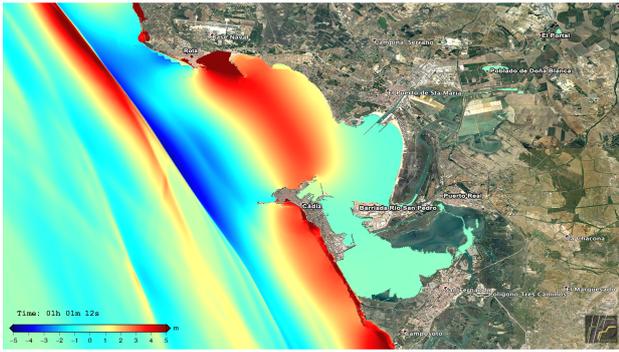
When using nested meshes, an approach based on the fluctuations of the state values (see [2] for more details on this strategy) is used to interpolate the fine ghost cells and to project the values of the fine ghost cells to the next coarser level. A flux correction step at the boundaries of the submeshes is also needed to ensure mass conservation [2].

The Okada model [5] is used to predict the initial bottom deformation caused by the earthquake, which is transmitted instantaneously to the sea surface generating the tsunami wave. Triangular faults [4] are also supported to compute the initial bottom deformation. Optionally, the Kajiura filter [3] can also be applied. It is possible to consider several fault segments and to apply the deformation on them at different times.

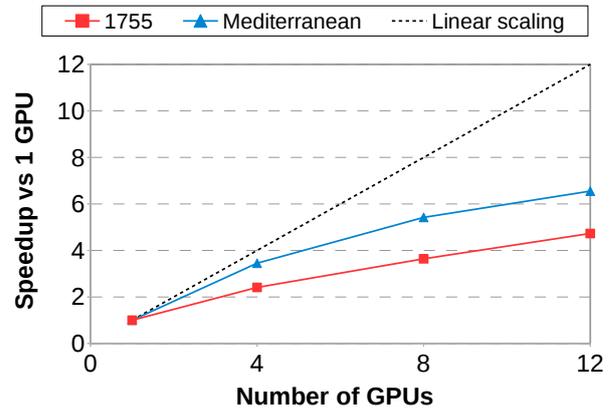
Implementation

The Tsunami-HySEA numerical scheme, the nested meshes processing, the Okada model and the Kajiura filter are implemented using CUDA and MPI. Double numerical precision is used. The initial bottom deformation can be applied to the whole spatial domain or locally in a subdomain. A 2D domain decomposition is applied by using a load balancing algorithm that takes into account the wet and dry zones and the nested submeshes, so that all the processes have a similar computational load. MPI communications and CUDA computations are overlapped to increase efficiency. CUDA streams are used to compute in parallel different submeshes in a same level of the grid hierarchy. The output is written into NetCDF files, and several variables can be output, such as time series, wave height, velocities, maximum wave height, momentum flux, arrival times of the tsunami, etc. Furthermore, the NetCDF files can be written asynchronously, so that the computations can continue while the data are being written to the disk.

It is also possible to resume a stored simulation. New grids and new points for the time series can be added when resuming a simulation. For example, this can be useful to first simulate a tsunami using a coarse mesh, and then compute the inundation in coastal areas by adding refined grids without having to start the simulation again.



(a)



(b)

Figure 1: Tsunami-HySEA results: (a) Arriving of the tsunami to Cádiz in the 1755 Lisbon earthquake problem, (b) Strong scaling obtained with both problems described in the main text using up to 12 GPUs.

Numerical Results

In this section we consider two problems using real bathymetries. All the simulations are executed in a cluster formed by 6 nodes with Intel Xeon ES-2620 processors and 2 GeForce GTX Titan Black in each node.

The first problem simulates the tsunami caused by the Lisbon earthquake occurred on November 1, 1755. This tsunami, with its epicenter in the Atlantic Ocean, caused great damages and many deaths in the south of Portugal and Spain. We simulate this tsunami using three levels with a total of 38.4 millions of volumes, where the finest level has two submeshes covering the Spanish cities of Cádiz and Huelva. Simulation time is 1.5 hours. Figure 1a shows the arriving of the tsunami at the city of Cádiz. In Figure 1b we can see that, using 12 GPUs, a speedup of 4.7 is obtained with respect to one single GPU. The runtime has been 1 hour and 23 minutes using 12 GTX Titan Black.

The second problem represents a tsunami in the Mediterranean Sea using one single level with a mesh of 10.0 millions of volumes and a resolution of 30 arc-sec. Simulation time is 4 hours. Figure 1b shows the strong scaling obtained, where it can be seen that a speedup of 6.6 has been reached with respect to one GPU when using 12 GPUs. In this problem, up to 1246 millions of volumes per second have been processed when using 12 GTX Titan Black, with a runtime of 2 minutes and 14 seconds.

References

- [1] M. de la Asunción, M. J. Castro, E. D. Fernández-Nieto, J. M. Mantas, S. Ortega, and J. M. González. Efficient GPU implementation of a two waves TVD-WAF method for the two-dimensional one layer shallow water system on structured meshes. *Computers & Fluids*, **80**, 441–452, 2013.
- [2] M. de la Asunción, and M. J. Castro. Simulation of tsunamis generated by landslides using adaptive mesh refinement on GPU. *Journal of Computational Physics*, **345**, 91–110, 2016.
- [3] K. Kajiura. The leading wave of a tsunami. *Bulletin of the Earthquake Research Institute*, **41**, 535–571, 1963.
- [4] B. J. Meade. Algorithms for the calculation of exact displacements, strains, and stresses for triangular dislocation elements in a uniform elastic half space. *Computers & Geosciences*, **33(8)**, 1064–1075, 2007.
- [5] Y. Okada. Internal deformation due to shear and tensile faults in a half-space. *Bulletin of the Seismological Society of America*, **82(2)**, 1018–1040, 1992.
- [6] C. E. Synolakis, E. N. Bernard, V. V. Titov, U. Kânoğlu, and F. I. González. Validation and verification of tsunami numerical models. *Pure and Applied Geophysics*, **165(11)**, 2197–2228, 2008.