ParCFD'2020
32^{nd} International Conference on Parallel Computational Fluid Dynamics
May-11-13 2020, Nice FRANCE

MASSIVELY PARALLEL ANISOTROPIC MESH ADAPTATION, PERFORMANCES AT SUPERCOMPUTER SCALE AND APPLICATION TO EARTH LAND ELEVATION

H. DIGONNET* AND L. SILVA†

* École Centrale de Nantes
Institut de Calcul Intensif
1 rue de la Noë, 44300 Nantes, France
e-mail: hugues.digonnet@ec-nantes.fr, web page: https://ici.ec-nantes.fr/

†École Centrale de Nantes
Institut de Calcul Intensif
1 rue de la Noë, 44300 Nantes, France
e-mail: luisa.rocha-da-silva@ec-nantes.fr, web page: https://ici.ec-nantes.fr/

Key words: Massively parallel computation, Mesh adaptation, Massive input data, Application to Earth

Abstract. In this paper, we present work and results obtained during a "Grand Challenges" session on the French supercomputer Irène Joliot-Curie. The goal of such a project was to increase the scalability of our mesher [1] at full supercomputer scale (several dozens of thousands of cores) and so validate the good scalability of the approach. A second objective was to evaluate the possibility of graphically exploit such huge results and compare several IO strategies. An application to Earth land elevation is shown, where we use a modified level-set method to represent the elevation by introducing "real data" (the huge SRTM data [2]) as the input for our anisotropic mesh adaptation code.

1 INTRODUCTION

During the "Grand Challenges" session on the new French supercomputer Irène Joliot-Curie, we have mainly two partitions: one with 79 488 Intel Xeon SkyLake Cores for 320 TB of memory; the other with 56 304 Intel Xeon Phi KNL cores for 80 TB of memory. A third one is under construction, with 293 376 AMD Epyc cores for 573 TB of memory and will be tested as soon as possible, providing a step forward to have a next generation Exascale supercomputer. Thanks to GENCI, that provided us access to such huge supercomputers, we were able to evaluate the scalability of our mesher [1] at a large scale. Running over several dozens of thousand of cores validated the good scalability of the application, but also allowed looking deep into some more practical aspects, such as the IO strategies and the possibility of graphically exploit such huge results.

2 PARALLEL PERFORMANCE AT A SUPERCOMPUTER SCALE

On new hardwar,e we start by evaluating the scalability of our mesh adaptation application using an isotropic size benchmark. Taking into account the hardware of the Skylake partition of Irène Joliot-Curie, the goal was to be able to build a 1000 billion-element mesh. Building such a mesh requires the use of at least 65 536 cores and up to 200 TB of memory (almost the full supercomputer). During this bench, two output strategies have been tested: the standard one consists in writing one file per core; the distributed MPI-IO one consists in writing one large MPI-IO file per group of 256 cores. If the first one is still the most efficient, with a 130 GB/s, it uses a large amount of Inodes (here, 65 536). The second one is a little less efficient, but reduces the number of Inodes used (256) and produces a "well sized" file of 120GB.

Figure 1 shows the weak scalability test done in 2d and up to 65 536 cores. Since the workload per core remains the same regardless the number of cores used, we may have, in the ideal case, a constant computational time during the mesh adaptation process. In practice, we notice three regions in the plot:

- the first is from the beginning and until the almost steady state at 128 cores, where the time spent increases slowly with the number of cores used. This may be due to two factors: the increase in the number of communications per core; the underpopulated use of the CPU for a small number of cores.

- then, time spent to refine the mesh is almost constant between 128 cores and 6 384 cores. this is when we we fully use the CPU and reach an almost steady number of communications per core.

- finally, there is a degradation of the parallel performance at the full system scale, mainly due to the poor performance of IntelMPI setup and IO and this scale. MPI_Init takes more than 20 minutes and IO time becomes not negligible, with around 5 minutes at 130 GB/s.

When using almost the full supercomputer (65 536 cores), we get a 2d mesh containing 509 billion nodes and 1018 billion elements, whick took around 40 TB of disk storage. It's a well adapted mesh, with a scaled quality (from 0, flat element, to 1, equilateral triangle with perfect mesh size) of 0.41 for the worst element, 1 for the best and 0.93 in average. The partition of the mesh is also good enough, with the best subdomain containing only 2 neighbours, the worst 22 neighbours (on one unconnected domain), and an average of 6.3 neighbours per subdomain.

3 APPLICATION TO EARTH LAND ELEVATION

An anisotropic application has also been done by considering the capture of the Earth's surface elevation at full scale. We want here to use huge "real data" as the input of our anisotropic mesh adaptation procedure. For that, we chose to use huge SRTM data [2] that represent the Earth land elevation with a latitude/longitude delta space of 8.33e-4,

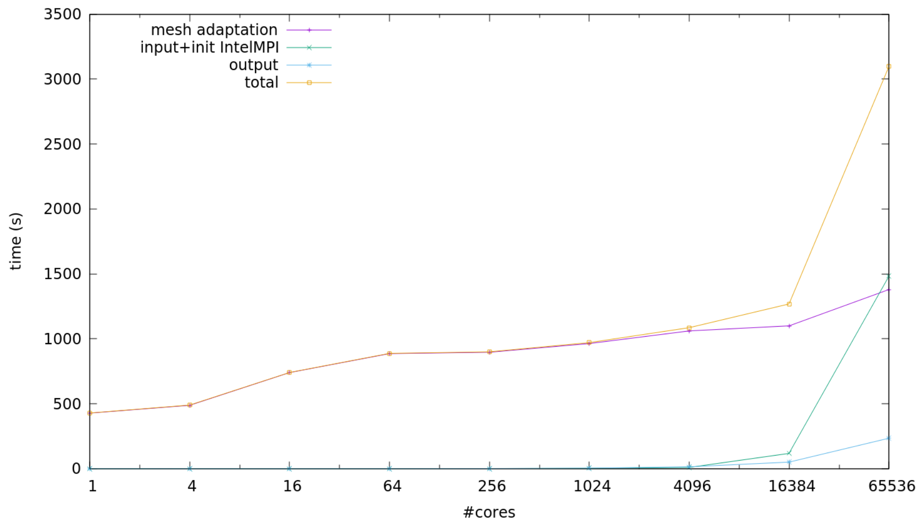

Figure 1: Weak scalability benchmark done from 1 to 65 536 cores on Irène Joliot-Curie: time spent to refine a 2d mesh by a factor two. The workload per core is constant and we end with a mesh containing around 15 million of elements per core. The largest mesh generated contains 1018 billions of elements.

which represents a 12.7 GB of compressed images. The global image contains 26 billions of points, torepresent the -60 to +60 on latitude region of the Earth globe. This high resolution land elevation is completed by a coarser one at the poles.

The first step is to translate these 2d images into a large set of 3d faces that will represent the land elevation. Then, this high resolution surface is decimated using the above mesh adaptation strategy to keep high resolution where needed (in the mountains) and have quite coarse facets elsewhere (in the sea parts). All this processing is done in parallel on each 5x5 latitude-longitude image, to end with a final set of 800-million adapted faces representing the Earth surface. All these faces are reordered in boxes to only load the faces needed to compute the level-set that will represent the land elevation surface.

Once this work accomplished, we may start the anisotropic mesh adaptation procedure on a background volume mesh, including computation of the signed distance to the set of faces representing the Earth elevation, an error estimator function, the metric construction and the topological mesh adaptation. This set of instructions iterates until a good quality mesh has been generated.

For taht, computations have been run using 1 384 cores on our Tier2 supercomputer Liger (6 384 cores) and have generated an anisotropic adapted mesh containing 2.5 billions of elements, keeping details of less than 200 meters, with an element size that varies between 1 m and 50 km, inside a global 4 096 billions of km^3 cube. Figures 2 and 3 show pictures of the land elevation field on that mesh, at the full scale and on some particular details.

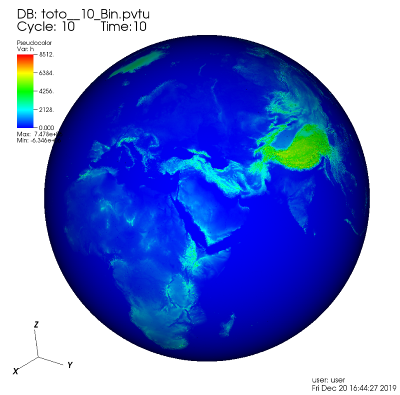

Figure 2: The zero isovolume of the LevelSet function representing land elevation on Earth at its full scale. The color represents the land elevation (in relative to the sea level.

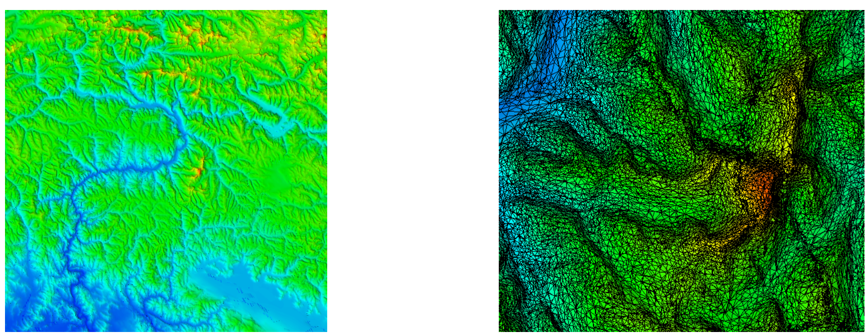

Figure 3: Two details concerning the Himalaya mountain chain region. On the left, we show the elevation and, on the right, we show elevation and the adapted mesh close to the Nanga Parbat mountain.

4 CONCLUSIONS

Our mesher parallelization strategy has shown a good parallel performance at the full national supercomputer size, up to 65 536 cores, 200 TB of RAM including parallel IO and visualization. At that scale, we have been able to generate a 1 000 billion of elements mesh. We have also shown an anisotropic application to well capture the land elevation on Earth by using huge real data taken from STRM **??** measurements. We used only 1 384 cores to generate an anisotropic adapted mesh of 2.5 billion elements, but we are able to that keep details of less than 200 m inside a global 4 096 billion km^3 cube, with a directional mesh size that varies between 1 m and 50 km.

REFERENCES

[1] Digonnet, H. and al. *Massively parallel anisotropic mesh adaptation.* Int. J. High Perf. Comp. App., Vol. **33**(1) pp. 3-24, (2019).

[2] Jarvis A., H.I. Reuter, A. Nelson, E. Guevara, *Hole-filled seamless SRTM data V4*, International Centre for Tropical Agriculture (CIAT), available from http://srtm.csi.cgiar.org. (2008)