The 8th European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS Congress 2022
5-9 June 2022, Oslo, Norway

A novel approach for anisotropic mesh adaptation on massively parallel distributed-memory systems

Bastien Andrieu*, **Bruno Maugars**, **Eric Quémerais**

ONERA
29 avenue de la Division Leclerc, Châtillon, France
bastien.andrieu@onera.fr

Keywords: *Mesh adaptation, Parallel meshing algorithms, High Performance Computing.*

Anisotropic mesh adaptation has become a powerful tool for making Computational Fluid Dynamics (CFD) simulations more reliable and efficient. In order to enable simulations on very large scales while optimizing the workflow, mesh adaptation must be performed with the same computational means as the CFD solver. Most parallel mesh adaptation algorithms proposed in the literature[1] follow the same approach as the CFD codes, namely a coarse-grained parallelism based on domain decomposition. It is worth noting that some proposed algorithms also make use of multithreading to parallelize remeshing operators at the subdomain level. The general approach thus consists in adapting the interior of each subdomain while constraining mesh elements at partition boundaries. The domain decomposition is then modified so that some of these constrained elements can now be adapted. This procedure is carried on until all elements are adapted. Work prediction heuristics are implemented in order to yield well load-balanced partitions, which is critical to achieve good scalability. However, the repartitioning step still represents a significant part of the total running time, and the shape (connectedness, . . .) of the partitions can also have a strong impact on the performance of these algorithms.

Our approach does not rely on domain decomposition. Instead, it adopts a fine-grained parallelism, based on decomposition into elementary tasks distributed over all MPI processes. All elementary operations (point insertion/deletion, edge/face swap, . . .) are performed by a unique cavity operator[2]. Candidate cavities are first constructed according to the operation to perform. Next, a parallel graph coloring technique is used to identify a large set of mutually disjoint cavities. The selected cavities are then evenly redistributed on all processes and remeshed in a massively parallel fashion. Rejected cavities are enlarged in subsequent passes, thus increasing their chances of success.

All mesh entities are evenly distributed on all processes without building local mesh partitions. This distribution is dynamically updated as entities are deleted/inserted during the adaptation procedure. This way, an optimal load and memory balance is acheived at each step of the algorithm, thus compensating for the communication overhead. By design, the adapted mesh produced by the proposed algorithm is independent of the number of cores.

Preliminary tests on two-dimensional triangular meshes have been performed with isotropic and anisotropic metric fields, showing promising results with good strong scalability up to a thousand cores. Insertion rates of about $8 \cdot 10^5$ points per second are achieved, producing adapted meshes of up to one billion points.

REFERENCES

[1] C. Tsolakis, N. Chrisochoides, M. Park, A. Loseille, and T. Michal, *Parallel Anisotropic Unstructured Grid Adaptation. AIAA Journal*, Vol. **59**, pp. 4764–4776, 2021.

[2] A. Loseille, and V. Menier. *Serial and parallel mesh modification through a unique cavity-based primitive*, *Proceedings of the 22nd International Meshing Roundtable*, Springer, pp. 541–558, 2014.