



Extreme-Scale Parallel Mesh Generation: Telescopic Approach

Nikos Chrisochoides, Andrey Chernikov, Daming Feng and Christos Tsolakis

Center for Real-Time Computing
Computer Science Department
Old Dominion University
Norfolk VA



INTRODUCTION

Finite Element Mesh Generation is a critical component for many (bio-) engineering and science applications. This project will deliver a novel framework for highly scalable and energy efficient guaranteed quality mesh generation for the Finite Element (FE) analysis in three and four dimensions. Parallel unsteady blood flow simulations will be used as a test-bed end-to-end application. This project will combine domain-and application-specific knowledge with run-time system support to improve energy efficiency and scalability of parallel FE mesh generation codes. Traditionally, parallel FE mesh generation methods and software are developed without considering the architectural features of the supercomputer platforms on which they are eventually used for production.

The proposed approach is to abstract and expose parallel mesh generation run-time information to the underlying run-time system which can guide the execution towards the most efficient utilization of resources on the given supercomputer. The issues of performance and energy efficiency are closely related, and we will study them in tandem.

The project will focus on the following three objectives:

- O1:** Integration of multiple parallel Delaunay mesh generation methods into a coherent hierarchical frame work.
- O2:** Development of application-specific models that describe the inherent concurrency and data access patterns of this framework.
- O3:** Development of domain-specific energy-efficient race-to-halt, concurrency throttling, and component-level (core and memory) power scaling and test them on parallel mesh generation using O2.

In this poster we focus and present our preliminary results pertinent to the first objective. The goal of this project is to study the telescopic approach (see Fig 1) and to develop Delaunay-based methods to explore concurrency at all hardware layers using abstractions at

(a) medium-grain level for many cores within a single chip and (b) coarse-grain level, i.e., sub-domain level using proper error metric- and application-specific continuous decomposition method .

Parallel Optimistic Delaunay Mesh generator (PODM) In [2,8] is presented a Parallel Delaunay Refinement algorithm for polyhedral domains with good performance but also with high communication cost.

Parallel Delaunay Refinement (PDR) In [3,4] using existing meshing code we increased the scalability of the meshing procedure. In PDR the initial input is being decomposed using a quadtree and buffer zones are being created for each quadtree leaf. This enhancement enables concurrent meshing of quadtree leaves that are at a safe distance.

Parallel Constrained Delaunay Mesh generator (PCDM) [5] uses only asynchronous communication and takes advantage of the domain decomposition. The initial input is decomposed into regions which are meshed in parallel. In order to achieve consistency across the mesh, split messages are being sent between neighboring regions when a point is inserted on their common separator.

Parallel Domain Decoupling Delaunay method (PD³) [6,7] using an approximation of the medial axis, the domain is divided into regions using pre-refined separators in such a way that it is guaranteed that no other refinement on the boundaries of the regions will be performed and thus each region may be refined independently.

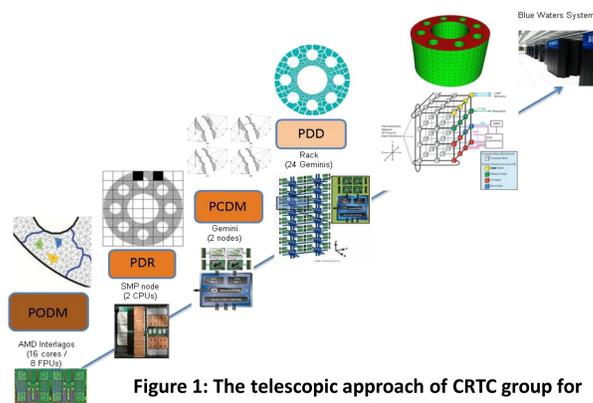


Figure 1: The telescopic approach of CRTG group for Parallel Mesh Generation.

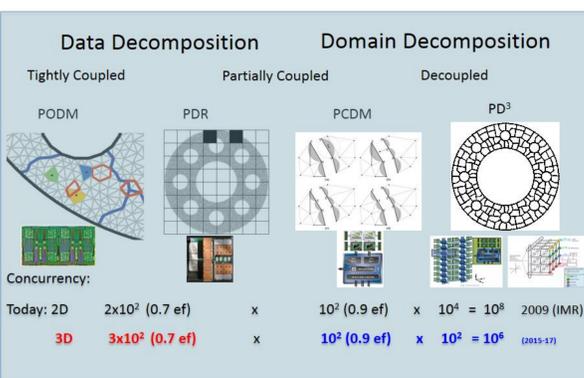


Figure 2: Today we can reach 3×10^2 concurrency, the expectation is to reach 10^6 by the end of 2017.

METHODS and RESULTS

Our recent research is mainly focused on achieving similar scalability results but this time with three-dimensional input. Our expectation is to have at least 10^6 of concurrency.

Parallel Optimistic Delaunay Mesh generator (PODM)

In [8] we developed a tightly coupled Delaunay refinement algorithm for meshing 3D medical images. PODM offers provable quality guarantees and exhibits good performance for up to 144 cores creating elements with a peak rate of 14.3 million elements per second.

In [9] improving the data locality we increased the limit to roughly 200 cores achieving almost 70% efficiency, with the Locality-Aware Parallel Delaunay Image-to-Mesh conversion algorithm (LAPD).

#Cores	1	32	64	128	144	160	176
#Tets	10.7M	3.49M	7.44M	132M	151M	167M	185M
Speedup	1.00	33.71	63.33	119.56	123.67	94.10	86.36
efficiency	1.00	1.05	0.99	0.93	0.86	0.59	0.49

Figure 3: Performance of PODM [8].

PDR.PODM

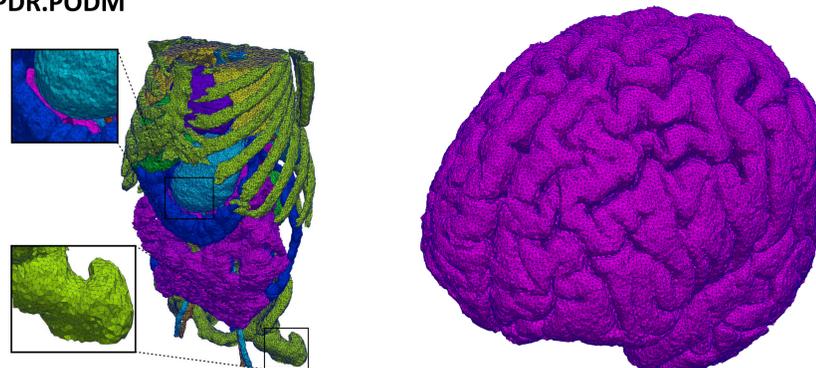


Figure 4: PDR.PODM can handle arbitrary complex geometries.

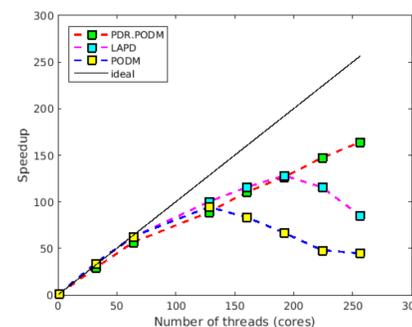


Figure 5: Weak Speedup comparison of PODM optimizations for meshes that vary from 3M tets to 745M tets on 256 cores.

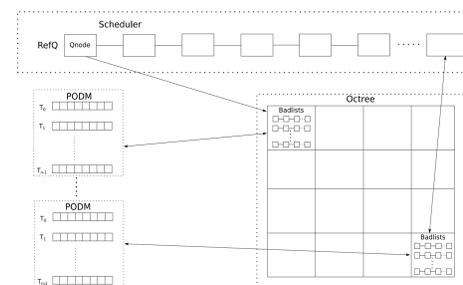


Figure 6: The high-level structure of PDR.PODM.

The data management that PDR offers, has been adapted to PODM. The implementation of PDR.PODM relies on an octree for data decomposition [10]. More precisely, the initial image is being divided into subregions using an octree. In each octree leaf a group of PODM threads is assigned. Each group has its own block of memory so the cost of communication within a block is low. Like in the original version of PDR every time a region is refined a buffer of neighboring regions is locked in order to prevent other threads from accessing elements that belong to other group of PODM threads.

FUTURE WORK

The next step of our telescopic approach for 3D input images is to create software similar to PCDM and PD³ in three dimensions and integrate our existing code. Already, we have some results for both methods using simple input images.

In addition we are planning for Extreme-scale Parallel Anisotropic Meshing to address geometry, mesh generation and high-performance computing requirements for the NASA's CFD 2030 Vision: "a study to address the long range, strategic planning required by NASA's Revolutionary Computational Aerosciences (RCA) program in the area of computational fluid dynamics (CFD), including future software and hardware requirements for High Performance Computing (HPC)."

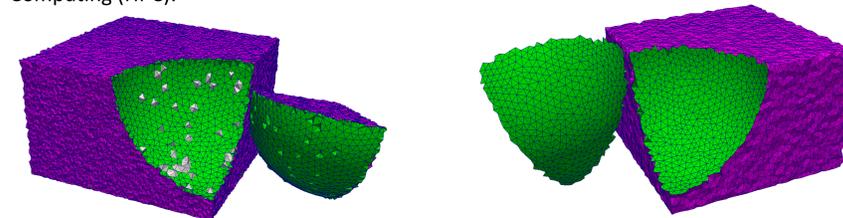


Figure 7: Left, The curved surface is treated as a constrain. Like PCDM for each encroaching element (white tetrahedron) a split message has to be sent. Right, Decoupled method. The curved surface here has been pre-refined and the two parts can be refined with no communication at all.

Finally, we are targeting brain aneurysms which are abnormal ballooning of intracranial arteries which when left untreated may result in fatal outcome for the patients. Neurovascular stents have been used extensively in stent-assisted coiling treatment of aneurysms. Treatment with flow diverters (low porosity, fine-mesh stents) is used for complicated aneurysm geometries not amenable to coiling. Angiographic analysis with mathematical modeling has been applied for comparison of different flow diversion designs as well as towards prediction of flow diversion efficacy. High resolution computational fluid dynamics (CFD) simulations would provide the opportunity to study localized alterations in hemodynamics due to different devices.

ACKNOWLEDGEMENTS

This work in part is funded by NSF grants , CCF-1439079,CCF-1139864 and CCF-1346401 with special thanks to Dr. Ramnath Rajiv and Dr. Chitkelkanova Almadena at NSF for making it possible to cover the IMR travel expenses for Christos Tsolakis. In addition, it utilized resources from the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575. We thank the systems group in the Pittsburgh Supercomputing Center for their great help and support. Finally, it is partially supported by the Richard T. Cheng Endowment. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NSF.

REFERENCES

- 1.N. Chrisochoides, A. Chernikov, A. Fedorov, A. Kot, L. Linardakis and P. Foteinos "Towards Exascale Parallel Delaunay Mesh Generation." In *Proceedings of the 18th International Meshing Roundtable*, edited by Brett W. Clark, 319–36. Springer Berlin Heidelberg, 2009. doi:http://dx.doi.org/10.1007/978-3-642-04319-2_19.
- 2.D. Nave,N. Chrisochoides ,P.Chew "Guaranteed: Quality Parallel Delaunay Refinement for Restricted Polyhedral Domains." In *Proceedings of the Eighteenth Annual Symposium on Computational Geometry*, 135–44. SCG '02. New York, NY, USA: ACM, 2002. doi:10.1145/513400.513418
- 3.A. Chernikov and N. Chrisochoides."Practical and Efficient Point Insertion Scheduling Method for Parallel Guaranteed Quality Delaunay Refinement." In *Proceedings of the 18th Annual International Conference on Supercomputing*, 48–57. ICS '04. New York, NY, USA: ACM, 2004. doi:10.1145/1006209.1006217.
- 4.A. Chernikov and N. Chrisochoides. "Parallel Guaranteed Quality Delaunay Uniform Mesh Refinement." *SIAM Journal on Scientific Computing* 28, no. 5 (January , 2006): 1907–26. doi:10.1137/050625886.
- 5.A. Chernikov and N. Chrisochoides. "Algorithm 872: Parallel 2D Constrained Delaunay Mesh Generation." *ACM Trans. Math. Softw.* 34, no. 1 (January 2008): 6:1–6:20. doi:10.1145/1322436.1322442.
- 6.L. Linardakis and N. Chrisochoides. "Delaunay Decoupling Method for Parallel Guaranteed Quality Planar Mesh Refinement." *SIAM Journal on Scientific Computing* 27, no. 4 (January , 2006): 1394–1423. doi:10.1137/030602812.
- 7.L. Linardakis and N. Chrisochoides. "Algorithm 870: A Static Geometric Medial Axis Domain Decomposition in 2D Euclidean Space." *ACM Trans. Math. Softw.* 34, no. 1 (January 2008): 4:1–4:28. doi:10.1145/1322436.1322440.
- 8.P. Foteinos and N. Chrisochoides. "High Quality Real-Time Image-to-Mesh Conversion for Finite Element Simulations." *Journal of Parallel and Distributed Computing* 74, no. 2 (February 2014): 2123–40. doi:10.1016/j.jpdc.2013.11.002.
- 9.D. Feng, A. Chernikov and N. Chrisochoides. "Two-level Locality-Aware Parallel Delaunay Image-to-Mesh Conversion." *Journal of Parallel and Distributed Computing* (submitted)
- 10.A. Chernikov and N. Chrisochoides. "Three-dimensional Delaunay refinement for multi-core processors". 22nd ACM International Conference on Supercomputing, pages 214–224. Island of Kos, Greece, June 2008.