Exascale-Era Mesh Generation For Medical Images*

Nikos Chrisochoides, Andrey Chernikov and Christos Tsolakis

Center for Real-Time Computing, Computer Science Department

Old Dominion University, Norfolk, VA

There have been numerous advances in the computer hardware over the last 15 years that drove the current and emerging architectures for exascale-era computing. Networking techniques have been integrated into the memory hierarchy resulting in multi-layered (w.r.t. memory and network) supercomputing architectures. In the past 15 years we developed the theory for parallel mesh generation capable of taking advantage of multi-layered architectures [1, 2]. By combining our experience in designing algorithms and software, for each of the layers, we have developed a novel **telescopic approach which explores concurrency at different levels of granularity**: (1) medium-grain level for multiple cores within a single chip, and (2) coarse-grain level for multiple chips within a node or a cluster, both using proper error metrics and application-specific discrete data and continuous domain decomposition methods. By leveraging careful mapping of work units of each parallel meshing method to each level of the memory hierarchy, the telescopic approach is capable of creating guaranteed quality aerospace simulations. For example, in order to accurately model the entirety of the neuron network in the human brain model, one will need more than 100 trillion elements, this requires a massive amount of memory. Likewise, in aerospace engineering, NASA's CFD vision for 2030 addresses the need for a highly efficient and scalable parallel mesh generation method, which is essential for simulating the physics in pertinent complex systems . Such data sets can be handled only by taking advantage of the full potential of the upcoming HPC systems. Our preliminary data are encouraging, to the best of our knowledge, we were the first to generate 18M elms/sec on 256 cores [3] while mathematically guarantee quality and fidelity of the meshes.

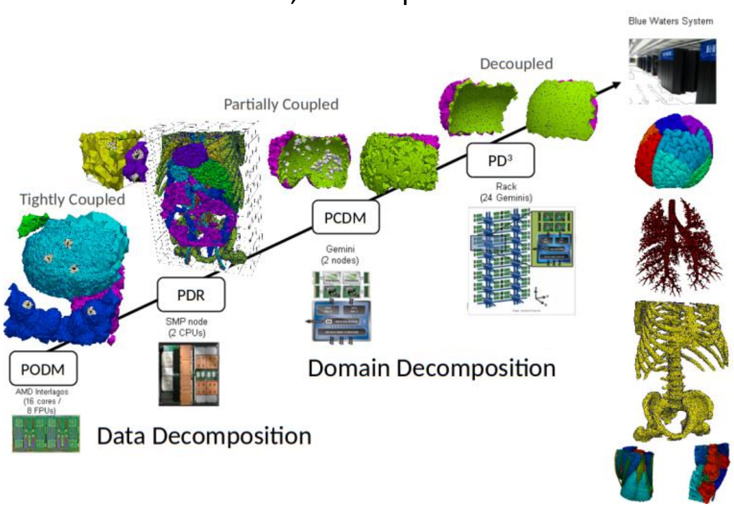

References

[1] Parallel Mesh Generation. N. Chrisochoides. In Numerical Solution of Partial Differential Equations on Parallel Computers, Springer-Verlag, (eds. A. M. Bruaset, P. Bjorstad, A. Tveito), Vol. 51, p 237-259, 2005.

[2] Towards Exascale Parallel Delaunay Mesh Generation. N. Chrisochoides, A. Chernikov, A. Fedorov, A. Kot, L. Linardakis and P. Foteinos. In 18th Intern. Meshing Roundtable, p 319 - 336, Utah, Oct., 2009.

[3] Scalable 3D Hybrid Parallel Delaunay Image-to-Mesh Conversion Algorithm for Distributed Shared Memory Architectures. D. Feng, C. Tsolakis, A. Chernikov and N. Chrisochoides. In 24th International Meshing Roundtable, Austin, Texas, Oct., 2015.

(*)This work is funded in part by NSF grant no. CCF-1439079, NASA grant no. NNX15AU39A, DoD's PETTT Special Project PP-CFD-KY07-007, John Simon Guggenheim Foundation and Richard T. Cheng Endowment. Dr. L. Leonidas, Dr. P. Foteinos and Daming Feng, contributed in this effort with their PhD work at CRTC.