

PARALLEL ADVANCING FRONT GRID GENERATION

Rainald Löhner and Juan R. Cebal

Institute for Computational Sciences and Informatics
MS 4C7, George Mason University, Fairfax, VA 22030, USA
Email: rlohner@science.gmu.edu

Abstract. A parallel advancing front scheme has been developed. The domain to be gridded is first subdivided spatially using a relatively coarse octree. Boxes are then identified and gridded in parallel. A scheme that resembles closely the advancing front technique on scalar machines is recovered by only considering the boxes of the active front that generate small elements. The procedure has been implemented on the SGI Origin class of machines using the shared memory paradigm. Timings for a variety of cases show speedups similar to those obtained for flow codes. The procedure has been used to generate grids in excess of a hundred million elements.

Keywords. Unstructured Grid Generation, Parallel Computing, CFD.

1. INTRODUCTION

The widespread availability of parallel machines with large memory, solvers that can harness the power of these machines, and the desire to model in ever increasing detail geometrical and physical features has led to a steady increase in the number of points used in field solvers. Grids in excess of 10^7 elements have become common for production runs in Computational Fluid Dynamics (CFD) [Bau93, Bau95, Jou98, Yos98, Mav99] and Computational Electromagnetics (CEM) [Dar97, Mor97]. The expectation is that in the near future grids in excess of $10^8 - 10^9$ elements will be required. While many solvers have been ported to parallel machines, grid generators have lagged behind. For applications where remeshing is an integral part of simulations, e.g. problems with moving bodies [Löh90, Mes93, Mes95, Bau96, Kam96, Löh98a, Has98] or changing topologies [Bau98, Bau99], the time required for mesh regeneration can easily consume more than 50% of the total time required to solve the problem. Faced with this situation, a number of efforts have been reported on parallel grid generation [Löh92, dCo94, Sho95, dCo95, Oku96, Che97, Oku97, Sai99].

The two most common ways of generating unstructured grids are the Advancing Front Technique (AFT) [Per87, Per88, Löh88a,b, Per90, Per92, Jin93, Fry94, Löh96] and the Generalized Delaunay Triangulation (GDT) [Bak89, Geo91, Wea92,

Wea94, Mar95]. The AFT introduces one element at a time, while the GDT introduces a new point at a time. Thus, both of these techniques are, in principle, scalar by nature, with a large variation in the number of operations required to introduce a new element or point. While coding and data structures may influence the scalar speed of the ‘core’ AFT or GDT, one often finds that for large-scale applications, the evaluation of the desired element size and shape in space, given by background grids, sources or other means [Löh96] consumes the largest fraction of the total grid generation time. Unstructured grid generators based on the AFT may be parallelized by invoking distance arguments, i.e. the introduction of a new element only affects (and is affected by) the immediate vicinity. This allows for the introduction of elements in parallel, provided that sufficient distance lies between them.

Several years ago, the author and his colleagues introduced a parallel AFT based on the subdivision of the background grid [Löh92, Sho95]. While used for some demonstration runs, this scheme was not general enough for a production environment. The background grid had to be adapted in order to be sufficiently fine for a balanced workload. As only background grid elements covering the domain to be gridded were allowed, complex in/out tests had to be carried out to remove refined elements lying outside the domain to be gridded. Furthermore, element size specified at CAD entities could not be ‘propagated’

edge-length of the finest octree box is of the order of 20 to 50 times the specified element size. This implies that its construction is very fast, and can be accomplished on a single processor without discernable CPU penalty.

3. WORK ESTIMATION AND BALANCE

The procedure outlined above will work optimally if each box requires approximately the same CPU time to complete its grid. This implies that a good work estimate should be provided. Given that the boxes are not body conforming, even for uniform grids the volume to be gridded can vary drastically from box to box. A balanced workload can be obtained by starting with many boxes, estimating the work to be done for each of them, and then gridding in parallel groups of boxes with similar workload.



Figure 3 Estimation of Volume to be Gridded

The volume to be gridded is estimated by the marching cubes procedure shown schematically in Figure 3. Given the dimensions of the box, and the list of active faces, the box is first subdivided into voxels (i.e. small cubes). In a first pass over the faces, the voxels cut by faces are marked as ‘inside the domain’, and an average normal is computed for each cut voxel. This normal information is used in a subsequent pass over the voxels in order to mark the neighbours of cut voxels as either inside or outside the domain to be gridded. The remaining voxels are then marked as inside or outside in several sweeps over the voxels. These sweeps are carried out until no further voxels can be marked. Finally, a work estimate is obtained by summing the expected number of elements in each of the voxels marked as inside the

domain to be gridded. This work estimation procedure is done in parallel.

Given the estimated work in each of the boxes, the load is balanced in such a way that each processor receives a similar amount of work. The assumption is made that the number of boxes is always larger than the number of processors. Should this not be the case, the boxes are subdivided further (8 new boxes for each box). If any given box has a work estimate that lies above the average work per processor, this box is also subdivided further. This ‘greedy’ work balancing algorithm may be summarized as follows:

- a) Obtain a minimum nr. of active boxes:
WHILE: The number of boxes is smaller than the number of processors:
 - Subdivide boxes (1:8);**ENDWHILE**
- b) Balance the work:
WHILE: Work unbalanced
 - Estimate, in parallel, the work for each box;
 - Obtain average work per processor;
 - **IF:** A box has a work estimate above average: Subdivide it further (1:8) and balance again;
 - **ENDIF**
 - Attempt to group boxes in such a way that the work in each group is close to average;
 - **IF:** Good work balance impossible: Subdivide boxes with highest work estimate (1:8) and balance again;
 - **ENDIF****ENDWHILE**

In many instances, boxes within a group will share a common face. In order to avoid the (unnecessary) buildup of many faces at the borders of boxes, an attempt is made to agglomerate these neighbouring boxes within groups. This is done recursively by checking, for any pair of boxes, if two of the dimensions are the same and the boxes are coincident in the remaining dimension. If so, the boxes are merged. Boxes are merged recursively, until no pair of boxes passes the test outlined above.

4. IMPLEMENTATION

The procedure described above was implemented using the shared memory (i.e. `c$doacross`) paradigm on the SGI Origin 2000. Although this is a distributed-memory machine, it can be programmed as a shared-memory machine. This choice was adopted for the following reasons:

Register for free at <https://www.scipedia.com> to download the version without the watermark

- Coding for a shared-memory environment is much simpler than for a distributed-memory environment. The operating system takes care of most of the inherent message passing, communication conflicts, etc., relieving the user from this task;
- The production codes used in conjunction with the grid generator scale well using the shared-memory paradigm [Löh98b]; Even on 32 processors, more than 50% of the theoretical speedup is achieved.
- For the last years, all of the large-scale production runs carried out by the author and his colleagues [Löh98a, Löh98c, Bau98, Bau99] were performed on these machines, using the shared-memory paradigm;
- The SGI Origin 2000 has become the dominant platform within the High-Performance Computing sites in the US; at present, there are more SGI Origin 2000 processors than those of all other vendors combined; the expectation is that this trend will not change drastically in the foreseeable future.

While some of the reasons stated above are clearly personal in nature, the basic ideas of the proposed algorithms are general, and may be coded within a distributed memory framework with explicit message passing.

5. EXAMPLES

to see how the procedure works in the 'best case scenario'. The unit cube is to be gridded with a uniform mesh of approximately 1 million tetrahedra. Although this is not a large grid, the timings shown in Figure 4 are illustrative.

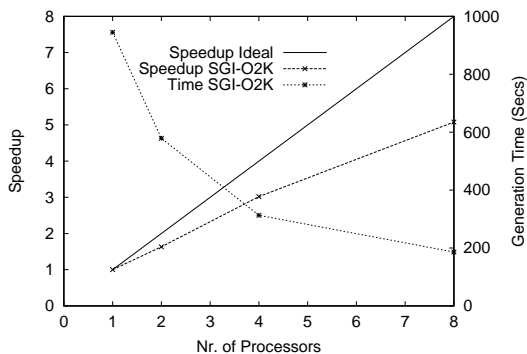


Figure 4 Cube: Speedups Obtained

5.2 Aneurism: This example is taken from a hemodynamic analysis recently conducted for this particular geometry. The outline of the domain, grid, as well as some sample results are shown in Figure 5.1.




Figure 5.1 Aneurism

The final uniform mesh had approximately 2.7 million tetrahedra. The speedup obtained is shown in Figure 5.2. As one can see, the parallel mesher was more efficient for this finer grid than for the unit cube, even though the volume to be gridded in each box can vary widely due to geometric features.

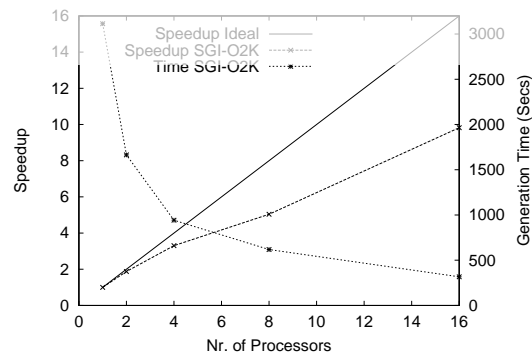


Figure 5.2 Aneurism: Speedups Obtained

5.3 Garage: This example was taken from a blast simulation recently carried out for an office complex. The outline of the domain is shown in Figure 6.1.




Figure 6.1 Garage: Wireframe

The final uniform mesh had approximately 9.2 million tetrahedra. The speedup obtained is shown in Figure 6.2. As before, the parallel mesher was more efficient for this finer grid than for the unit cube, even though the volume to be gridded in each box can vary widely due to geometric features.

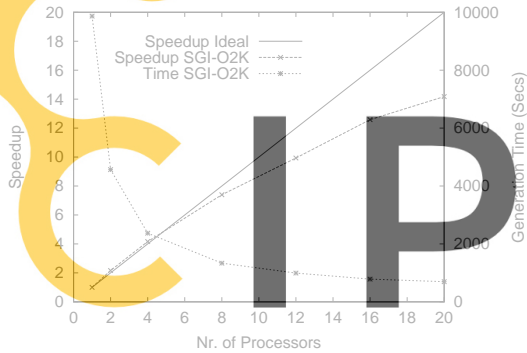


Figure 6.2 Garage: Speedups Obtained

5.3 Tysons Corner: This example was taken from a dispersion simulation recently carried out for this well-known shopping center. The outline of the domain is shown in Figures 7.1,7.2.




Figure 7.1 Tysons Corner: Wireframe (W)




Figure 7.2 Tysons Corner: Wireframe (N)

The final mesh had approximately 16 million tetrahedra. The smallest and largest specified element side lengths were 230 cm and 1400 cm respectively. The speedup obtained is shown in Figure 7.3.

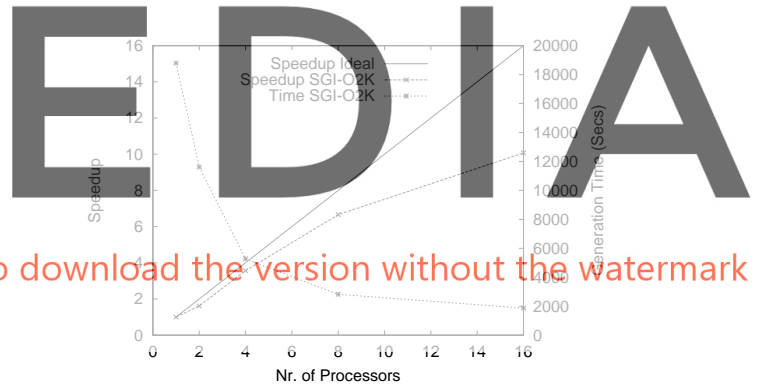


Figure 7.3 Tysons Corner: Speedups Obtained

5.4 Space Shuttle: This example is included here because it is typical of many aerodynamic data sets. The outline of the domain is shown in Figure 8.1. The surface triangulation of the final mesh, which had approximately 4 million tetrahedra, is shown in Figure 8.2. The smallest and largest specified element side lengths were 5.08 cm and 467.00 cm respectively, i.e. an edge-length ratio of approximately $1 : 10^2$ and a volume ratio of $1 : 10^6$. The spatial variation of element size was specified via approximately 200 sources [Löh96].




Figure 8.1 Space Shuttle: Outline of Domain




Figure 8.4 Space Shuttle: Surface Pressures



Figure 8.2 Space Shuttle: Surface Mesh

As could be seen from the previous examples, the grid generator certainly scales well with the number of processors. As with many other areas where parallel processing is being attempted, scalability improves with the amount of work required. The larger the grids, the better the scalability. One can also observe that for each case the ‘scalability slope’ is somewhat different, ranging from almost perfect for the garage to a 1:3 asymptote for the Shuttle. Closer inspection revealed that for the Shuttle, although the final mesh contains 4 million elements, the number of elements created in each parallel pass over increasing element sizes was rather modest, never exceeding 0.5 million elements. Thus, the inherent parallelism of this problem is rather low.

Register for free at <https://www.scipedia.com> to download the version without the watermark




Figure 8.3 Space Shuttle: Speedups Obtained

The speedup obtained is shown in Figure 8.3. Finally, the results of an Euler run for an incoming Mach-nr. of $M = 2.0$ and angle of attack of $\alpha = 1.1^\circ$ are shown in Figure 8.4.

6. CONCLUSIONS AND OUTLOOK

A parallel advancing front scheme has been developed. The domain to be gridded is first subdivided spatially using a relatively coarse octree. Boxes are then identified and gridded in parallel. A scheme that resembles closely the advancing front technique on scalar machines is recovered by only considering the boxes of the active front that generate small elements. The procedure has been implemented on the SGI Origin class of machines using the shared memory paradigm. Timings for a variety of cases show speedups similar to those obtained for flow codes. The procedure has been used to generate grids for a large variety of cases, and is nearing production maturity.

Current work is focusing on improved work prediction algorithms, as it is found that in going to a larger number of processors, any small imbalance incurs a heavy CPU penalty.

7. ACKNOWLEDGEMENTS

This research was partially supported by AFOSR, with Dr. Leonidas Sakell as the technical monitor.

8. REFERENCES

[Bak89] T.J. Baker - Developments and Trends in Three-Dimensional Mesh Generation. *Appl. Num. Math.* **5**, 275-304 (1989).

[Bau93] J.D. Baum, H. Luo and R. Löhner - Numerical Simulation of a Blast Inside a Boeing 747; *AIAA-93-3091* (1993).

[Bau95] J.D. Baum, H. Luo and R. Löhner - Numerical Simulation of Blast in the World Trade Center; *AIAA-95-0085* (1995).

[Bau96] J.D. Baum, H. Luo, R. Löhner, C. Yang, D. Pelessone and C. Charman - A Coupled Fluid/Structure Modeling of Shock Interaction with a Truck; *AIAA-96-0795* (1996).

[Bau98] J.D. Baum, H. Luo and R. Löhner - The Numerical Simulation of Strongly Unsteady Flows With Hundreds of Moving Bodies; *AIAA-98-0788* (1998).

[Bau99] J.D. Baum, H. Luo, E. Mestreau, R. Löhner, D. Pelessone and C. Charman - A Coupled CFD/CSD Methodology for Modeling Weapon Detonation and Fragmentation; *AIAA-99-0794* (1999).

[Ch97] J. P. Chou, N. Chrysoschides, F. Schur - Parallel Constrained Delaunay Meshing; *Proc. 1997 Workshop on Trends in Unstructured Mesh Generation*, June (1997)

[dCo94] H.L. de Cougny, M.S. Shephard and C. Ozturan - Parallel Three-Dimensional Mesh Generation; *Computing Systems in Engineering* **5**, 311-323 (1994).

[dCo95] H.L. de Cougny, M.S. Shephard and C. Ozturan - Parallel Three-Dimensional Mesh Generation on Distributed Memory MIMD Computers; *Tech. Rep. SCOREC Rep. # 7*, Rensselaer Polytechnic Institute (1995).

[Dar97] E. Darve and R. Löhner - Advanced Structured-Unstructured Solver for Electromagnetic Scattering from Multimaterial Objects; *AIAA-97-0863* (1997).

[Fry94] J. Frykestig - Advancing Front Mesh Generation Techniques with Application to the Finite

Element Method; *Pub. 94:10*, Chalmers University of Technology; Göteborg, Sweden (1994).

[Geo91] P.L. George, F. Hecht and E. Saltel - Automatic Mesh Generator With Specified Boundary; *Comp. Meth. Appl. Mech. Eng.* **92**, 269-288 (1991).

[Has98] O. Hassan, L.B. Bayne, K. Morgan and N. P. Weatherill - An Adaptive Unstructured Mesh Method for Transient Flows Involving Moving Boundaries; pp. 662-674 in *Computational Fluid Dynamics '98* (K.D. Papailiou, D. Tsahalis, J. Périaux and D. Knörzer eds.) Wiley (1998).

[Jin93] H. Jin and R.I. Tanner - Generation of Unstructured Tetrahedral Meshes by the Advancing Front Technique; *Int. J. Num. Meth. Eng.* **36**, 1805-1823 (1993).

[Jou98] W. Jou - Comments on the Feasibility of LES for Commercial Airplane Wings; *AIAA-98-2801* (1998).

[Kam96] A. Kamoulakos, V. Chen, E. Mestreau and R. Löhner - Finite Element Modelling of Fluid/Structure Interaction in Explosively Loaded Aircraft Fuselage Panels Using PAMSHOCK/ PAMFLOW Coupling; *Conf. on Spacecraft Structures, Materials and Mechanical Testing*, Noordwijk, The Netherlands, March (1996).

[Löh88a] R. Löhner - Some Useful Data Structures for the Generation of Unstructured Grids; *Comm. Appl. Num. Meth.* **4**, 123-135 (1988).

[Löh88b] R. Löhner, J. P. Chou, J. P. Chou - Three-Dimensional Grid Generation by the Advancing Front Method; *Int. J. Num. Meth. Fluids* **8**, 1135-1149 (1988).

[Löh90] R. Löhner - Three-Dimensional Fluid-Structure Interaction Using a Finite Element Solver and Adaptive Remeshing; *Comp. Sys. in Eng.* **1**, 2-4, 257-272 (1990).

[Löh92] R. Löhner, J. Camberos and M. Merriam - Parallel Unstructured Grid Generation; *Comp. Meth. Appl. Mech. Eng.* **95**, 343-357 (1992).

[Löh96] R. Löhner - Progress in Grid Generation via the Advancing Front Technique; *Engineering with Computers* **12**, 186-210 (1996).

[Löh98a] R. Löhner, C. Yang, J. Cebral, J.D. Baum, H. Luo, D. Pelessone and C. Charman - Fluid-Structure-Thermal Interaction Using a Loose Coupling Algorithm and Adaptive Unstructured Grids; *AIAA-98-2419* [Invited] (1998).

Register for free at <https://www.scipedia.com> to download the version without the watermark

[Löh98b] R. Löhner - Renumbering Strategies for Unstructured- Grid Solvers Operating on Shared-Memory, Cache- Based Parallel Machines; *Comp. Meth. Appl. Mech. Eng.* 163, 95-109 (1998).

[Löh98c] R. Löhner, C. Yang and E. Oñate - Viscous Free Surface Hydrodynamics Using Unstructured Grids; *Proc. 22nd Symp. Naval Hydrodynamics*, Washington, D.C., August (1998).

[Mar95] D.L. Marcum and N.P. Weatherill - Unstructured Grid Generation Using Iterative Point Insertion and Local Reconnection; *AIAA J.* 33, 9, 1619-1625 (1995).

[Mav99] D.J. Mavriplis and S. Pirzadeh - Large-Scale Parallel Unstructured Mesh Computations for 3-D High-Lift Analysis; *ICASE Rep.* 99-9 (1999).

[Mes93] E. Mestreau, R. Löhner and S. Aita - TGV Tunnel-Entry Simulations Using a Finite Element Code with Automatic Remeshing; *AIAA-93-0890* (1993).

[Mes95] E. Mestreau and R. Löhner - Airbag Simulation Using Fluid/Structure Coupling; *AIAA-95-0798* (1996).

[Mor97] K. Morgan, P.J. Brookes, O. Hassan and N.P. Weatherill - Parallel Processing for the Simulation of Problems Involving Scattering of Electromagnetic Waves; in *Proc. Symp. Advances in Computational Mechanics* (L. Demkowicz and J.N. Reddy eds.), (1997).

[Oku96] T. Okusanya and J. Peraire - Parallel Unstructured Mesh Generation; *Proc. 5th Int. Conf. Num. Grid Generation in CFD and Related Fields*, Mississippi, April (1996).

[Oku97] T. Okusanya and J. Peraire - 3-D Parallel Unstructured Mesh Generation; *Proc. Joint ASME/ASCE/SES Summer Meeting* (1997).

[Per87] J. Peraire, M. Vahdati, K. Morgan and O.C. Zienkiewicz - Adaptive Remeshing for Compressible Flow Computations; *J. Comp. Phys.* 72, 449-466 (1987).

[Per88] J. Peraire, J. Peiro, L. Formaggia K. Morgan and O.C. Zienkiewicz - Finite Element Euler Calculations in Three Dimensions; *Int. J. Num. Meth. Eng.* 26, 2135-2159 (1988).

[Per90] J. Peraire, K. Morgan and J. Peiro - Unstructured Finite Element Mesh Generation and Adaptive Procedures for CFD; *AGARD-CP-464*, 18 (1990).

[Per92] J. Peraire, K. Morgan, and J. Peiro - Adaptive Remeshing in 3-D; *J. Comp. Phys.* (1992).


[Sai99] R. Said, N.P. Weatherill, K. Morgan and N.A. Verhoeven - Distributed Parallel Delaunay Mesh Generation; to appear in *Comp. Meth. Appl. Mech. Eng.* (1999).

[Sho95] A. Shostko and R. Löhner - Three-Dimensional Parallel Unstructured Grid Generation; *Int. J. Num. Meth. Eng.* 38, 905-925 (1995).

[Wea92] N.P. Weatherill - Delaunay Triangulation in Computational Fluid Dynamics; *Comp. Math. Appl.* 24, 5/6, 129-150 (1992).

[Wea94] N.P. Weatherill and O. Hassan - Efficient Three-Dimensional Delaunay Triangulation with Automatic Point Creation and Imposed Boundary Conditions; *Int. J. Num. Meth. Eng.* 36, 2029 (1994).

[Yos98] S. Yoshimura, H. Nitta, G. Yagawa and H. Akiba - Parallel Automatic Mesh Generation Method of Ten-Million Nodes Problem Using Fuzzy Knowledge Processing and Computational Geometry; *Proc. 4th World Cong Comp. Mech.*, Buenos Aires, Argentina, July (1998).



Register for free at <https://www.scipedia.com> to download the version without the watermark