LA-UR-97-2577

Approved for public release;
distribution is unlimited

Title: Adaptive Mesh Refinement Algorithm
Development and Dissemination

RECEIVED

AUG 27 1997

OSTI

Author(s): J. S. Saltzman, CIC-19
D. L. Brown, CIC-19
K. D. Brislawn, CIC-12
G. S. Chesshire, CIC-19
D. J. Quinlan, CIC-19
M. Berger, Depts. App. Math. and Comp. Sci., NYU

Submitted to: DOE Office of Scientific and Technical Information (OSTI)

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Los Alamos
NATIONAL LABORATORY

29
MASTER

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

Adaptive-Mesh-Refinement Algorithm Development and Dissemination

Jeffrey S. Saltzman,* David L. Brown, Kristi D. Brislawn,
Geoffrey S. Chesshire, and Daniel J. Quinlan
Computing, Information, and Communications Division
Los Alamos National Laboratory

Marsha Berger
Departments of Applied Mathematics and Computer Science
New York University

Abstract

This is the final report of a three-year, Laboratory Directed Research and Development (LDRD) project at the Los Alamos National Laboratory (LANL). The project objective was to develop and disseminate adaptive mesh refinement (AMR) algorithms for structured and unstructured meshes. Development of AMR algorithms will continue along several directions. These directions include algorithms for parallel architectures, techniques for the solution of partial differential equations on adaptive meshes, mesh generation, and algorithms for nontraditional or generic applications of AMR. Dissemination of AMR algorithms is also a goal of the project. AMR algorithms are perceived as difficult to meld to current algorithms. We are developing tools that diminish this perception and allow more computational scientists to use AMR within their own work.

Background and Research Objectives

Our objective was to develop and disseminate adaptive mesh refinement (AMR) algorithms for structured meshes. Development of AMR algorithms continued along several directions. These directions included algorithms for parallel architectures, techniques for the solution of partial differential equations on adaptive meshes, mesh generation, and algorithms for nontraditional or generic applications of AMR. Dissemination of AMR algorithms was also a goal of our work. AMR algorithms are perceived as difficult to meld to current algorithms. We are developing tools that diminish this perception and allow more computational scientists to use AMR within their own work.

Prior to the work stemming from this project, we have worked on development of adaptive techniques for concentrating computational effort where it is most needed as a function of space, time and data described in Refs. 1 and 2. There are two settings in

* Principal Investigator, E-mail: jss@lanl.gov

which we have developed such algorithms. In local adaptive mesh refinement (AMR), subgrids in a finite-difference calculation are dynamically inserted and removed to maintain a uniform level of accuracy and resolve important time-dependent features in the solution to a problem. The second type of technique addresses resolving the complex geometries within which solutions of partial differential equations are being approximated. A structured overlapping composite (SCOM) grid describing the region around a complex object can be created, for example, by making a grid for each component of the object and then combining these grids as described in Ref. 3. This method allows practical application of algorithms to real world problems where geometry becomes an important issue. All of these adaptive methods have been used very successfully for time dependent problems in two and three dimensions.

A SCOM grid is a collection of structured curvilinear meshes that physically overlap and communicate through interpolation. The loose requirement of simple overlap allows complicated geometries to be gridded without the usual mesh generation issues such as smooth line continuity across mesh components. SCOM meshes directly address the issue of complicated geometries that often arise in industrial applications. AMR is a technique where a mesh is locally refined by embedding a finer mesh within a given mesh to enhance the local resolution of the calculation. The finer meshes may have even finer meshes embedded within them and so on. AMR has been shown to enhance resolution efficiency of unstructured and structured algorithms by factors anywhere from ten to one hundred.

Our preliminary work on parallel AMR is described in Ref. 4. We have simplified the AMR algorithm in two-dimensional (2-D) geometry by making all refinement patches a fixed size. This leads to a simple distribution scheme and minimal communication above and beyond what is called for in a uniform grid computation. In addition to completion of the 2-D work, we extended this work to 3-D. The results are described below.

Importance to LANL's Science and Technology Base and National R&D Needs

The research relates well to Laboratory core competencies and capabilities through the high-performance computing strategies and related tactics of the Laboratory. The primary mission of the Laboratory, reducing the global nuclear danger, depends on science-based stockpile stewardship (SBSS), the purpose of which is to assess the readiness and reliability of the weapons designed at LANL and previously tested at the test site in Nevada. Because of the current ban on testing of nuclear weapons, the Laboratory must rely on a combination of non-nuclear testing and simulation. The challenge LANL

faces today is that code development must move towards first-principle computations. In order to do first-principle computations, resolution of unprecedented detail must be obtained. The development of AMR technologies is one of the keystones in reaching this goal of finely detailed resolution. Already it is seen that work is proceeding on adaptive mesh schemes using both structured and unstructured meshes in a variety of production codes at LANL.

In the spirit of dual use technologies, AMR and SCOMs are also being used to solve problems of mutual interest to industry and the Laboratory. In particular, moving adaptive SCOMs are being applied to combustion-fluid-dynamics problems found in diesel engine design. Here the complex geometry of the diesel engine cylinder, with its opening and closing valves and moving piston, poses a particular computational challenge. Combining AMR with SCOMs within a parallel computing environment poses an extreme programming challenge but much progress has been made.

Scientific Approach and Accomplishments

During the first year of our project we supported work found in Ref. 4 and extended the work to three dimensions. The 2-D and 3-D codes solved a set of hyperbolic conservation equations, called the Euler equations, using a patch-based AMR algorithm. The extension of the actual code to 3-D was relatively straightforward. This was primarily due to the internal AMR algorithms being tensor products of 1-D techniques. However, several implementation difficulties were encountered during the development cycle. The most trivial of the difficulties was visualization of the data. There are few visualization packages that can handle the combination of patched-based structured data and the enormity of the datasets generated by the code. Figure 1 shows an example of the type of test computations that we ran and visualized. The computational domain is a unit size box with another smaller box attached at the bottom. The smaller box has heated fluid that is allowed to expand into the colder fluid in the larger box.

Because of the special data layout scheme and the large patch size ($32 \times 32 \times 32$ cells), the code ran efficiently (approximately 75%) compared to the uniform grid case. Memory saving and the resulting CPU savings were on the order of a factor of 100. Coupled with the efficiency of 75%, the overall saving was a factor of 75 over the uniform grid code. In order to achieve this type of efficiency, we also had to do machine specific optimizations. Communication between processors within the AMR algorithm only takes place when there is inter-refinement level communication such as interpolation, refluxing

(making sure mass is conserved between levels). There are several ways within the CM-5 programming language CMF (CM-FORTRAN) to achieve these results. We found that using unstructured communication was actually *faster* than using structured communication. Initially we found this not very intuitive until we discovered that reducing the number of communication steps in a computation was far more important than using a lower bandwidth style of communication.

The other half of the project focused on applying AMR methods to codes that were developed to find solutions of partial differential equations whose domains are complex. The structured composite overlapping meshes (SCOMs), described above, were targeted for the project. Again, because the component grids of the SCOM were structured, we found it straightforward to apply AMR techniques. This was done by applying AMR to each component and then examining what needed to be done at component grid boundaries. Work that needed to be done at grid boundaries included ensuring that error estimation algorithms that flagged points near a boundary for refinement also flagged points on physically adjacent grids. Interpolation from one component to another turned out not to be difficult in that only additional interpolation points were needed with no modification of the interpolation algorithms themselves.

Development of adaptive SCOM algorithms was performed in both 2- and 3-dimensions. Figure 2 shows a visualization of a shock impinging on a spherical obstacle. There are three component grids. The entire computational region is covered by a cubic mesh called the background grid. The sphere in the middle of the cubical domain is covered by two hemispherical mesh shells that overlap at an equator. The hemispherical meshes and physical interior boundary cut away at the background cubical mesh. There are no geometric singularities in the spherical meshing so large time steps can be used to run the computation.

Adaptive mesh refinement is used in addition to the composite grids. From the shading of the image in Figure 2, it is apparent that a shock is moving away from the observer and from left to right. The adaptive meshing highlights the shock as well as resolves it. In this computation only a factor of 2 refinement was used in conjunction with the constraint of a single level of refinement. However, larger computations were run using multiple levels of refinement and refinement factors of 4.

In developing the aforementioned algorithms and codes much experience was gained in how to make these techniques easier to use. Two library development efforts that are a result of this experience are currently being carried on by members of the Scientific Computing Group at LANL. The Overture Library⁵ is a collection of C++ objects that allow the code developer to construct moving composite grid codes through reuse and

selective modification of library code. The AMR++ library, currently under development, will be used in conjunction with Overture to add AMR capabilities. The use of Overture and AMR++ will facilitate development of new algorithms on complex domains with adaptive meshes so that the developer can focus on the central numerical algorithm instead of the infrastructure.

The final area of development that is still under way is finding conservative interpolation techniques for overlapping boundaries. Using Lagrangian interpolation with overlapping meshes in either an adaptive or nonadaptive way results in a loss of material between grids. By using Cartesian or cutout methods, where the overlapping grids cut away at underlying grids, single grid algorithms with global conservation but irregular zoning can be used. Work has begun, as indicated in Ref. 6, in addressing these irregular zones at overlap boundaries by treating them implicitly. By treating the zones implicitly the time steps for a simulation can still be constrained by the regular zoning away from boundaries. However, challenges remain in characterizing the irregular zoning at overlap boundaries when there is dynamic regridding or three-dimensional geometry or both. Finally, parallel algorithms for this kind of treatment are still not known.

Publications

1. Brislawn, K. D., Brown, D. L., Chesshire, G. S., and Saltzman, J. S., "Adaptively-Refined Overlapping Grids for the Numerical Solution of Systems of Hyperbolic Conservation Laws," Los Alamos National Laboratory document LA-UR 95-257, Proceedings of the ICASE/Langley Conference on Adaptive Meshes.
2. Berger, M. J., and Saltzman, J. S., "AMR on the CM-2," Applied Numerical Mathematics, 14, pp. 238-253, 1994.
3. Pao, K. I., and Saltzman, J. S., "An Implicit Method for Hyperbolic Conservation Laws on Meshes with Small Cells in One Dimension," Los Alamos National Laboratory document LA-UR-96-2932.

References

1. Berger, M. J., and Colella, P., "Local Adaptive Mesh Refinement for Shock Hydrodynamics," *Journal of Computational Physics*, **82**, pp. 64-84, 1989.
2. Bell, J. B., Berger, M. J., Saltzman, J. S., and Welcome, M. L., "Three Dimensional Adaptive Mesh Refinement for Hyperbolic Conservation Laws," *SIAM Journal on Scientific Computation*, **15**, pp. 127-138, 1994.
3. Chesshire, G. S., and Henshaw, W. D., "Composite Overlapping Meshes for the Solution of Partial Differential Equations," *Journal of Computational Physics*, **90**, p. 1, (1990).
4. Berger, M. J., and Saltzman, J. S., "AMR on the CM-2," *Applied Numerical Mathematics*, **14**, pp. 238-253, 1994.
5. Henshaw, W. D., and Brown, D. L., "Overture: An Advanced Object-Oriented Software System for Moving Overlapping Grid Computation," *Proceedings of the Computational Aerosciences Workshop at NASA Ames Research Center, Los Alamos National Laboratory document LA-UR-96-2931*.
6. Pao, K. I., and Saltzman, J. S., "An Implicit Method for Hyperbolic Conservation Laws on Meshes with Small Cells in One Dimension," *Los Alamos National Laboratory document LA-UR-96-2932*.

Figure 1. A CM-5 computation of a very hot fluid blasting out of a cubical hole into a cold region at a late time showing slices of the 3-D region adaptively resolved with an intermediate isosurface of density added. The isosurface shows rollup taking place due to the vorticity generation of the flow.

Figure 2. A shock computation that has proceeded around a sphere. The mesh is an adaptive composite overlapping grid with three components. There are two hemispherical components around the sphere and a single Cartesian background grid. The hemispherical grids avoid the classical polar singularity. Each grid surface has shading denoting density. Lighter shades are higher in density.