Applied Analysis and Computational Mathematics

DOE/ER/25053--3

DE91 002696

Review of Previous Work

Our adaptive methodology has made significant advances in the past few years. We have developed an efficient method of adaptive mesh refinement, henceforth AMR, which concentrates the computational effort where it is most needed. The method is designed for problems with multiple length scales that need high resolution over small portions of the computational domain. Our method has been combined with several higher-order versions of Godunov's method [Berger and Colella]. The increased resolution of this combined approach has enabled the study of fluid flow phenomena not previously possible. For example, AMR was used in a study of mach reflection of a blast wave off an oblique wedge. This has been used as a test problem for many years, in particular, see the survey paper by Woodward and Colella [1984]. However, previous calculations were apparently not fully converged. With the extra resolution provided by AMR, a Kelvin-Helmholtz-type instability was computed along the slip line culminating in the jet at the bottom boundary. This was not observed in the previous calculations, although experiments confirm its existence. The AMR code is now being used in several laboratories across the country. The two dimensional AMR has been extended to incorporate moving quadrilateral meshes, and work is in progress to simulate incompressible fluid flow. AMR has been combined with other techniques, such as the front-tracking algorithm of [Chern and Colella] and volume of fluid (SLIC) techniques. Using AMR, my collaborators have been able to obtain close agreement with shock tube experiments, verify or disprove theoretical conjectures, and discover new phenomena in regimes not easily accessible to laboratory experiment [Colella and Henderson].

AMR is based on the use of uniform, local grid refinements superimposed on an underlying coarse grid. These embedded grid refinements can be recursively nested to achieve a fixed level of accuracy in the calculation. Unlike other embedded grid refinement methods [e.g. Dannenhoffer], in our method the embedded fine grid cells are grouped together into grid patches containing many refined grid cells, but also some coarse cells that may not have needed to be refined. This is done so that all grids are uniform and logically rectangular, which allows vectorization without using gather/scatter operations. It also allows for a simple user interface; a finite difference scheme can be written for a uniform rectangular grid without concern for the connectivity of each cell. The use of fine grid patches also reduces the storage overhead, which is on a per grid basis for our method, rather than the overhead per grid point found in unstructured mesh calculations

DISCLAIMER

[Mavripilis and Jameson; Lohner]. The fine grids use a smaller time step as well, further concentrating the work on the fine grids. Other adaptive methods for transient flows use the same time step for the whole mesh [Lohner; Brackbill and Saltzman]. This is less efficient, since the Courant number on the adapted coarse cells can be tiny. The additional complications introduced by this method occur only at the interfaces of the fine and coarse grids.

AMR uses an automatic error estimation procedure, based on Richardson extrapolation, to determine the regions in the domain where the resolution in the solution is insufficient. These coarse grid cells are "flagged" as needing refinement. This past year, we have developed a new automatic grid generation algorithm that groups these flagged cells into rectangular grid patches [Berger and Rigoutsos]. Our new algorithm uses ideas from pattern recognition and computer vision to find near optimal groupings of flagged points into rectangles. Especially for three dimensional calculations, it is important to minimize the area that is unnecessarily refined, while keeping vector lengths as long as possible.

Great attention was paid to the efficient vectorization of AMR. For example, our initial implementation of the boundary interpolation procedures for the fine grids took 40% of the CPU time. A careful restructuring that minimized subroutine calls and performed block memory transfers reduced that to approximately 10% of the run time. The interior integration of a grid, which uses approximately 75% of the run time, has been clocked at 75 megaflops on the XMP. AMR has developed from an experimental algorithm to a highly efficient code that has been used for parameter studies and for production runs.

We have previously studied the parallel processing of AMR, so far only on shared memory multiprocessors. The interesting question here concerned load balancing. The computational load changes dynamically in time, as well as being a function of space. Since AMR already uses a domain decomposition approach to arrive at an accurate solution, it was natural to try to use a domain decomposition approach to multiprocessing AMR. The data structures that were developed for doing grid management and the inter-grid communication were already in place, and were easily adapted to the new task of inter-processor communication. Thus, I used a high level of parallelism based on grids, rather than the DO loop level of parallelism that is frequently employed. In general, a coarser granularity of parallelism incurs less synchronization and overhead costs than do fine grained approached.

The binary decomposition algorithm [Berger and Bokhari] was used to load balance the computation. It was implemented within AMR by partitioning the grids themselves so that each processor would entirely "own" the grid in its region of space. If a grid extended across a partition line, it was divided into two grids. Processors communicated with each

other through the boundary of their respective grids.

By parallelizing the code at such a high level, only five subroutines, accounting for 98% of the CPU time, needed to be modified. Using four processors, these five subroutines were sped up by a factor of 3.6, giving an overall speedup of 3.4, or 87%. This 10% imbalance in the computational load is due partly to the neglect of boundary work in the work estimates (3%), and partly to the inherent imbalance of partitioning fine grids only at coarse grid lines (7%), a structural requirement of AMR. Despite this, and including the serial degradation caused by creating more grids than are needed in the serial case, the speedup is still a respectable 84.5%.

Proposed Research

Our goal is to extend and combine these advanced numerical methods to tackle the complex physics and geometry found in realistic three dimensional problems.

Cartesian mesh methods

Our principal project will be the development of a Cartesian mesh algorithm with adaptive refinement to compute flows around arbitrary geometries. We are interested in both time-accurate and steady-state calculations in two and three space dimensions. In this approach, we keep a uniform rectangular (or Cartesian) grid and allow the boundary to intersect the grid cells in an essentially arbitrary way (see Figure 1).

This treatment of boundaries can be combined naturally with our AMR adaptive strategy of using locally uniform meshes. We retain the advantages (efficiency and accuracy) of uniform grids and are able to resolve fine scale flow features induced by complex geometries. Previous calculations using Cartesian meshes have either been underresolved, at the leading edge of an airfoil, for example, or inefficient, using tensor product meshes to refine the leading edge, but leaving the grid lines extending wastefully out to the far field [Choi and Grossman]. We are using our previously developed adaptive mesh refinement algorithm to achieve accuracy comparable to the body-fitted meshes, where grid points can be bunched in an a priori manner to improve the accuracy of the solution.

Figure 1 shows some preliminary examples of the method applied to the Euler equations to compute the interaction of a shock impinging on two cylinders. The only grid cells drawn in the figures are the irregular cells formed where the grid intersects the body. As the flow field becomes more complex, refined grids are used to increase the resolution. The location of these grids is indicated along with the density contours. The contour plots themselves show the information taken from the finest grid in each region. Note that the contour lines go smoothly across the grid interfaces between finer and coarse grids. The time steps used correspond to a CFL number of roughly 0.5 on the regular grid cells, though up to a thousand times larger on the smallest irregular cells.

END

DATE FILMED

12/04/90