CRAY CHANNELS

SPRING 1991 · A CRAY RESEARCH, INC., PUBLICATION

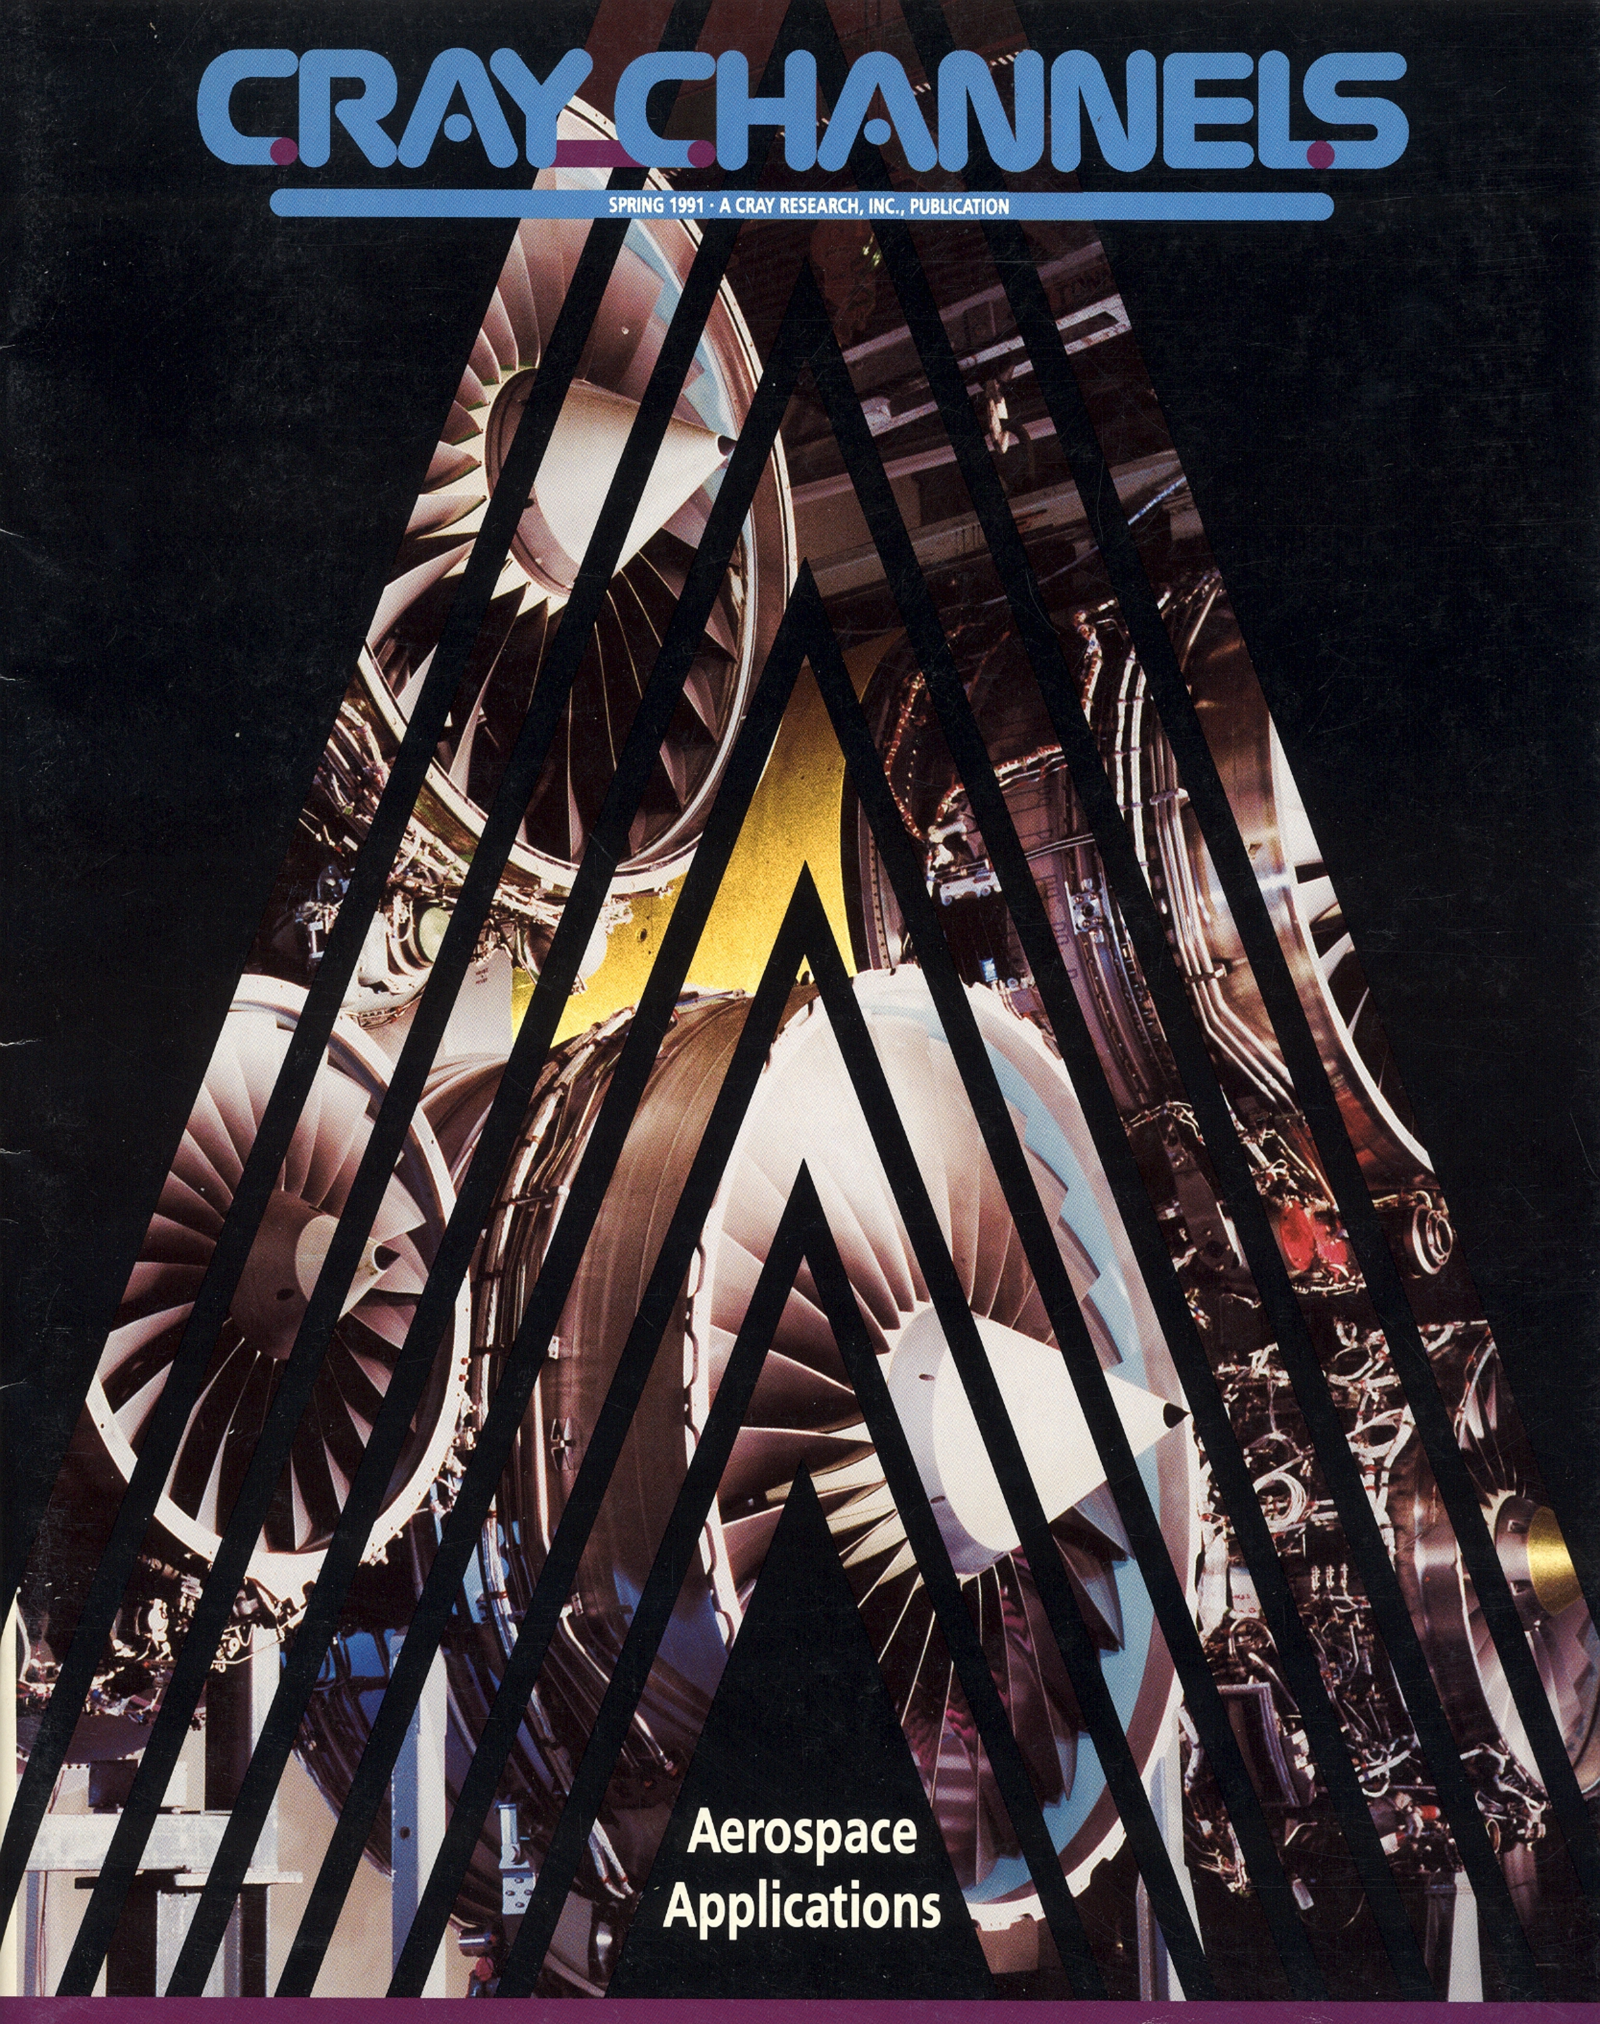

Aerospace Applications

Announcing new hardware and software products

Large-scale methods in computational electromagnetics

Daniel S. Katz and Allen Taflove, Northwestern University, Evanston, Illinois
Jeffrey P. Brooks and Evans Harrigan, Cray Research, Inc.

The numerical modeling of electromagnetic wave phenomena can be a computationally intensive task. To date, the design and engineering of aerospace vehicles has been the primary application driving the development of large-scale methods in computational electromagnetics (CEM). Efforts in this area have been aimed primarily at minimizing the radar cross section (RCS) of aerospace vehicles. RCS minimization enhances the survivability of vehicles that are subjected to precision-targeted ordnance. The physics of RCS is determined by Maxwell's equations and the constitutive properties of a vehicle's materials. As a result, the interesting situation arises in which the effectiveness and cost of state-of-the-art aerospace systems in part depends on the ability to develop an efficient engineering understanding of 120-year-old equations that describe the propagation and scattering of electromagnetic waves.

Two algorithms are of primary interest in this field: the robust, traditional, full-matrix, frequency-domain integral equation method of moments (MoM); and emerging time-domain, grid-based direct solutions of Maxwell's curl equations. Both types of algorithm make efficient use of Cray Research hardware and software capabilities.

Full-matrix MoM field computations at 2 GFLOPS

In the MoM area, one group of important codes originated with the Rao-Wilton-Glisson triangular surface patch technique for RCS analysis of arbitrarily shaped three-dimensional conducting structures.[1] Cray Research analysts determined that the primary task here involves the solution of very large, dense, complex-valued matrices (10K by 10K and larger) that exceed the available central memory. A strategy evolved to develop a complex-valued lower-upper matrix decomposition program that utilizes an efficient out-of-memory scheme and is adaptable to multiple CPU usage. The result was CLUD — Complex Lower-Upper Decomposition, with versions developed for the CRAY-2, CRAY X-MP, and CRAY Y-MP computer systems. This work rapidly gained popularity among MoM users, and Cray Research scientists have provided assistance to members of the CEM community in adapting these matrix solvers to many MoM codes.

A second group of important MoM codes originated with the Newman ESP-3 rectangular surface patch technique for RCS analysis of arbitrarily shaped three-dimensional conducting structures.[2] Cray Research analysts adapted their out-of-memory scheme for this code and subsequently developed an out-of-memory solver suitable for simultaneous solution of the monostatic RCS at a large number of illumination angles (right-hand sides). In fact, the number of right-hand sides could be in the thousands, approximating the order, *N*, of the MoM matrix. Subsequently, a parallel-processing version of the "*N* right-hand-sides" code was developed.

Although CLUD works well, two drawbacks had to be addressed for very large problems:

- ☐ The input/output (I/O) for CLUD is either synchronous to disks or synchronously staged from disk to Cray Research's SSD solid-state storage device. If the matrix is scaled to fit entirely in an SSD, this is not troublesome, and near-peak performance is achieved on the CRAY X-MP and CRAY Y-MP systems. However, a 20K-by-20K complex-valued MoM matrix requires an 800 Mword SSD, which is not currently available. In CLUD, very large problems of this size require synchronous I/O between disks and SSD, which reduces overall performance.
- ☐ The CLUD algorithm is based on a SAXPY type kernel that works on individual columns. This kernel runs at peak performance on the CRAY X-MP and CRAY Y-MP systems, but not on the CRAY-2 system because of a high ratio of memory operations to computation.

Because the Cray Research mathematical software group had optimized the BLAS-3 (Basic Linear Algebra Subroutines) to run at near-peak performance on all Cray Research computer systems, an improved algorithm was developed that was based on these kernels. A block-oriented method was adapted from LAPACK to run out-of-memory by Jeffrey Brooks of Cray Research's benchmarking department. The routine, CGETRF, made use of two BLAS-3 kernels, CGEMM (complex matrix multiply) and CTRSM (complex triangular backsolve).

To adapt CGETRF to run out-of-memory, the matrix is divided into slabs. (A slab is a matrix block consisting of a large number of adjacent columns of the matrix.) The matrix is decomposed from left to right, one slab at a time. Computation works on pairs of slabs. To compute a new leading slab, all preceding slabs need to be brought into memory, one at a time, for computation. This is an I/O pattern similar to that used in the existing CLUD code. However, three slab-sized memory buffers are used in the new code to allow for asynchronous I/O. The partial-pivoting scheme used in CGETRF is preserved in the new out-of-memory version.

A routine called CMXMA was written to take advantage of Golub's identity, which reduces the multiplication count for complex-number products. CMXMA converts complex matrix products to three real matrix multiplies and several matrix additions. (This routine is available as CGEMMS in SCILIB 6.0.) The Strassen's real matrix multiply (SGEMMS) was used to save further on operations. SGEMMS is a Strassen's algorithm extension to the standard BLAS-3 matrix multiply routine, SGEMM. SGEMMS was written by Cray Research's mathematical software group and is included in version 6.0 of Cray Research's UNICOS operating system.

When automatically multitasked to run on all eight processors on a CRAY Y-MP system, the new out-of-memory code ran at average computation rates exceeding 2.1 GFLOPS. Only 1.99 hours were required to process a 20K-by-20K matrix. During this run, 138 Gbytes of I/O were discharged to and from seven DD-40 disk drives. Yet, only 228 seconds (3.8 minutes) represented I/O wait time. In fact, 90 percent of the actual I/O operations were performed concurrently with the floating-point arithmetic by virtue of the asynchronous I/O scheme and therefore did not contribute to the observed wait time. As matrix size increased, the relative efficiency of the asynchronous scheme improved, with the I/O concurrency factor rising to 95 percent for a 40K matrix. Thus, the massive I/O associated with solving huge, dense, complex-valued MoM matrices could be buried almost completely.

When automatically multitasked to run on all eight processors on a CRAY Y-MP system, the new out-of-memory code ran at average computation rates exceeding 2.1 GFLOPS.

Multiprocessing space-grid time-domain codes

Although the LU decomposition strategy described here is highly efficient, the fundamental [order (N^3)] computational burden of LU decomposition remains dimensionally large. In fact, it is so large that there is virtually no prospect for using the traditional, full-matrix MoM to computationally model entire aerospace structures, such as fighter planes, at radar frequencies much above 150 MHz. Yet, radar frequencies of interest can greatly exceed 150 MHz, climbing to 10 GHz and higher. Much research effort, therefore, has been invested in the development of alternative iterative frequency-domain approaches, including conjugate gradient and spectral methods, that preserve the rigorous boundary-integral formulation of MoM while realizing dimensionally reduced [order (N^2) or less] computational burdens. Such methods would permit, in principle, the modeling of entire aircraft at radar frequencies above 1 GHz. However, these alternatives may not be as robust as the full-matrix MoM, insofar as they may not provide results of engineering value for a wide class of structures without the user having to wonder if the iterative algorithm has converged.

Problems involved in applying frequency-domain, full-matrix MoM technology to large scale RCS modeling have prompted much new interest in an alternative class of non-matrix approaches: direct space-grid, time-domain solvers for Maxwell's time-dependent curl equations. These approaches appear to be as robust and accurate as MoM, but have dimensionally-reduced computational burdens [approaching order (N)] such that whole-aircraft modeling for RCS can be considered in the near future. Currently, the primary approaches in this class are the finite-difference time-domain (FD-TD) and finite-volume time-domain (FV-TD) techniques.[3,4] These are analogous to existing mesh-based solutions of fluid-flow problems in that the numerical model is based upon a direct, time-domain solution of the governing partial differential equation. Yet, FD-TD and FV-TD are very nontraditional approaches to CEM for detailed engineering applications, where frequency-domain methods (primarily full-matrix MoM) have dominated.

FD-TD and FV-TD methods for Maxwell's equations are based on volumetric sampling of the unknown near-field distribution within and surrounding the structure of interest. The sampling is at sub-wavelength (λ_0) resolution to avoid aliasing of the field magnitude and phase information. Overall, the goal is to provide a self-consistent model of the mutual coupling of all the electrically-small volume cells that comprise the structure and its near field, even if the structure spans tens of λ_0 in three dimensions and there are tens of millions of space cells.

The primary FD-TD and FV-TD algorithms used today are fully explicit, second-order-accurate grid-based solvers that use highly vectorizable schemes for time-marching the six vector components of the electromagnetic near field at each of the volume cells. The explicit nature of the solvers is maintained either by leapfrog or predictor-corrector time-integration schemes. Present methods differ primarily in the set up of the space grid (almost-completely structured for FD-TD, body-fitted or unstructured for FV-TD) and the enforcement of EM field continuity at the interfaces of adjacent cells. As a result, the number of floating point operations needed to update a field vector component over one time step can vary by about 20 to 1 from one algorithm to another.

However, the choice of algorithm is not straightforward, despite this wide range of computational burdens. There is an important tradeoff decision to be made. Namely, a faster, simpler solver such as FD-TD uses meshes that may not be compatible with those used in other aerospace engineering studies, computational fluid dynamics (CFD) studies in particular. As a result, there is much "homework" to be done as researchers learn to generate a new class of three-dimensional meshes specific to Maxwell's curl equations. On the other hand, the more complex FV-TD solvers can utilize existing CFD mesh generators, but require substantially more algorithmic computer arithmetic and storage. Both FD-TD and FV-TD algorithms are highly vectorized, having been benchmarked at over 200 MFLOPS on one processor of a CRAY Y-MP system for real models. However, the attainment of even higher MFLOPS rates may be hampered by the fact that the the space grids have an unavoidable number of nonstandard cells that require either scalar or odd-lot vector operations. These nonstandard cells result from the need to program a near-field radiation condition at the outermost grid boundary (simulating the grid continuing to infinity), and the need to stitch together varying types of meshes to accommodate complex structure shapes. Despite this, it has been found possible

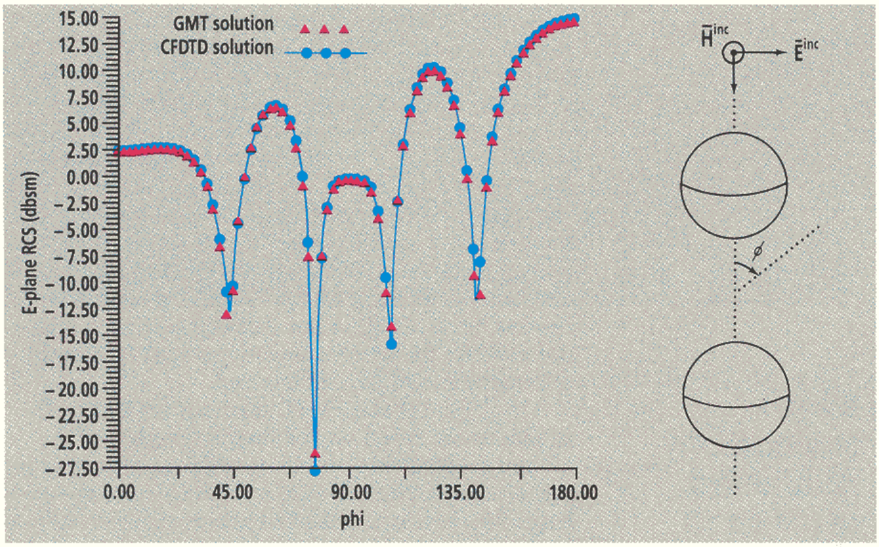

Figure 1. Agreement of FD-TD and generalized multipole technique (GMT)[6] bistatic RCS within 1 dB over a 42-dB range for a pair of $1\text{-}\lambda_0$ diameter conducting spheres separated by a $1\text{-}\lambda_0$ air gap.

to achieve nearly 100 percent concurrent utilization of all eight processors on a CRAY Y-MP system using Cray Research's Autotasking automatic multitasking software feature for three-dimensional FD-TD and FV-TD codes. Only relatively minor modifications were required to the original single-processor Fortran code.

Three-dimensional FD-TD validation example

Excellent validations of FD-TD have been obtained for three-dimensional problems that involve some of the key electromagnetic wave physics involved in RCS phenomena: near fields, monostatic RCS pattern, and bistatic RCS pattern. Here, we detail the results of a canonical, but difficult, bistatic RCS pattern validation.

Figure 1 shows the bistatic (side-scatter) RCS of a pair of $1\text{-}\lambda_0$ diameter conducting spheres separated by a $1\text{-}\lambda_0$ air gap.[5] The spheres are illuminated by a plane wave that propagates along a line connecting the centers of the spheres, and the bistatic pattern is observed in the plane of the incident electric field. (Note: when $\phi = 0^0$, the response is in the backscatter direction; that is, it is the monostatic RCS.) Here, the comparison is between FD-TD (using a mostly Cartesian, partially unstructured mesh to model conformally the spheres' surface curvatures) and an analytical approach well-suited for this problem, the generalized multipole technique (GMT).[6] Agreement between the two methods is excellent: within ±1 dB (approximately ±25 percent) over a wide 42-dB (16,000 to 1) range of RCS. This modeling accuracy occurs despite some tough electromagnetic field physics: the spheres interchange energy across the air gap in a tightly-coupled manner. For this problem, alternative FV-TD approaches using body-fitted meshes may introduce artifacts due to refraction and reflection of numerical waves propagating across global mesh distortions in the air-gap region. In fact, the two-sphere problem is a canonical example of difficult three-dimensional structures having substantial EM coupling between disjoint regions.

Electrically-large FD-TD application: jet engine inlet

The multiprocessing in-memory FD-TD code was used to model the RCS properties of an electrically large three-dimensional structure of engineering significance: a serpentine jet engine inlet (Figure 2). The overall system design problem involved sizing and shaping the engine inlet to meet specifications for both aerodynamics (thrust) and monostatic RCS at 10 GHz. The inlet was assumed embedded within a simple rectangular metal box coated with commercially available radar-absorbing material that provides approximately 30 dB (1000 to 1) suppression of electromagnetic wave reflections at 10 GHz. Thus, the FD-TD computed near-field and far-field electromagnetic response was primarily a function of the inside wall shaping of the inlet and not any exterior embedding.

As shown in Figure 2, the incident wave was assumed to propagate from right to left and be polarized with its electric field pointing across the narrow gap dimension (y direction) of the inlet. In this figure, the aperture of the inlet is located at the right, and the inlet is shorted by a conducting wall that represents the turbine assembly at the far left. With the box dimensions set at 30" x 10.5" x 10", the overall inlet and box target configuration spanned $25.4\lambda_0$ x $8.89\lambda_0$ x $8.47\lambda_0$ at 10 GHz. For this target, the FD-TD space cell size was 1/8" $(\lambda_0/9.43)$; and the overall lattice had 270 x 122 x 118 cells that spanned $4608\lambda_0^3$ and contained 23,321,520 unknown vector field components. Starting with zero-field initial conditions, 1800 time steps were used (95.25 cycles of the incident wave) to march the field components to the sinusoidal steady state. The computer running time was only 3 minutes and 40 seconds per monostatic RCS calculation on the CRAY Y-MP system using automatic multiprocessing across eight processors (7.97/8 processor concurrency), yielding an average computation rate of 1.6 GFLOPS.

In addition to simple data for the RCS pattern, the FD-TD modeling provided details of the complex near field. Figure 3 shows the instantaneous distribution (positive and negative values) of the total gap (E_y) electric field component in a two-dimensional observation plane that cuts through the center of the three-dimensional engine inlet. This photograph was derived from a color videotape display of the propagating electric field penetrating the inlet, generated directly by the FD-TD time-stepping. The display was taken late in the time-stepping when the field had settled into a repetitive sinusoidal oscillation (standing wave). It may be possible to use such highly detailed near-field information (very difficult to obtain from measurements) to improve future RCS designs. Comparatively, if MoM were applied to model the same engine inlet,

Figure 2. Geometry of engine inlet embedded in a rectangular metal box coated with commercial radar-absorbing material. The 10-GHz incident wave propagates from right to left.

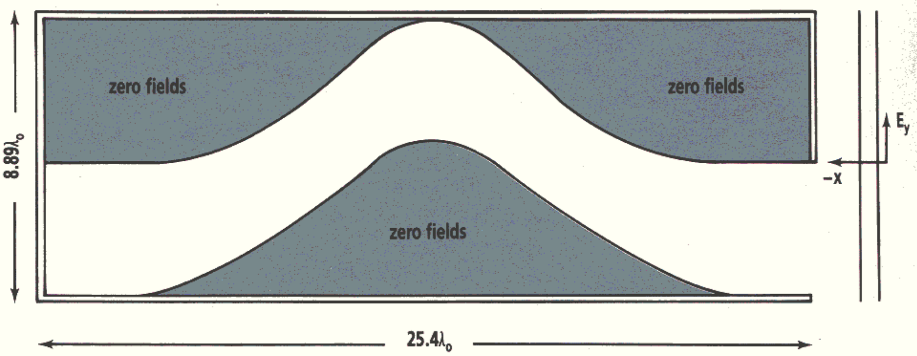

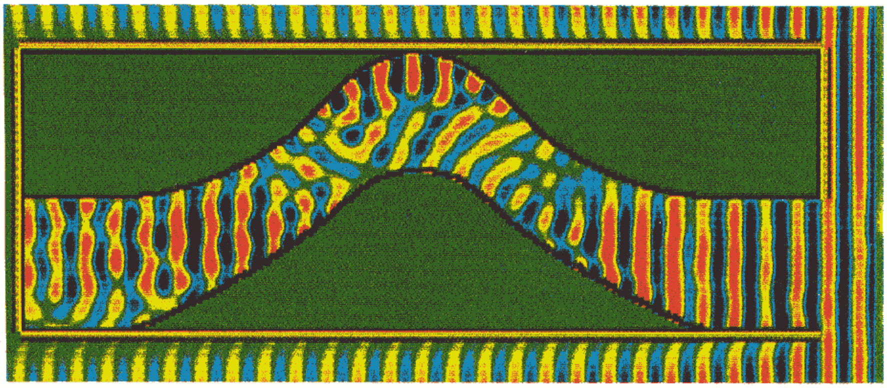

Figure 3. FD-TD computed map of the instantaneous distribution of the total E_y vector field component in a two-dimensional cut through the z-center of the engine inlet geometry (at the sinusoidal steady state).

a complex-valued linear system involving approximately 450,000 equations would have to be set up and solved. This assumes a standard triangular surface patching implementation of the electric field integral equation[1], with the $1500\lambda_0^2$ surface area of the engine inlet discretized at 10 divisions per λ_0. Using the 2.1-GFLOPS out-of-memory subroutine for LU decomposition discussed earlier, the CRAY Y-MP system running time for this matrix would be about 2.6 years for 5000 monostatic angles. This compares to only about 12.7 days for FD-TD for the same number of monostatic angles, a speedup factor of 75. Additional problems involved in error accumulation in the LU decomposition and reliability of the computer system over the multiyear solution time probably would combine to render a traditional MoM solution useless for this target and those of similar or larger electrical sizes. We note also that MoM does not directly provide details of the penetrating near-field distribution.

Present work and future directions

At present, grid-based time-domain CEM models of three-dimensional structures that span more than $30\lambda_0$ are being developed for the eight-processor CRAY Y-MP system. Work at this time addresses several areas:

☐ Automated mesh generation
☐ Multiprocessing out-of-memory software
☐ Subcell models for fine-grained structural features such as coatings
☐ Higher-order algorithms
☐ Application to nontraditional CEM areas, including design of ultra-high-speed electronic computer circuits, electro-optic components, and all-optical switches.

Extrapolating from benchmarks with the eight-processor CRAY Y-MP system, the next-generation CRAY Y-MP/16 system should provide a steady 10 to 13 GFLOPS computation rate for grid-based time-domain CEM codes when using automatic multitasking across 16 processors. The proverbial "billion-unknown" CEM problem (a three-dimensional computational volume of about $150{,}000\lambda_0^3$) could be completed in as little as 40 minutes per monostatic RCS observation. Multiprocessing out-of-memory software should enable even larger volumes to be modeled in their entirety. Using such software, the era of the "entire airplane in the grid" would be opened for a number of important aerospace systems for radar frequencies of 1 to 10 GHz. Automated geometry generation would permit CEM modelers to use structure databases developed by non-electromagnetics engineers, leading to lower design costs and the possibility of innovative design optimizations.

Acknowledgments

Daniel S. Katz and Allen Taflove were supported in part for this work by Cray Research, General Dynamics P.O. 4059045, National Science Foundation Grant ASC-8811273, and Office of Naval Research Contract N00014-88-K-0475. Jeffrey P. Brooks and Evans Harrigan were primarily responsible for the MoM, FD-TD, and FV-TD code optimizations and benchmarks. Thomas G. Jurgens of Fermilab, Batavia, Illinois, implemented the double-sphere FD-TD model.

About the authors

Daniel S. Katz is a graduate student pursuing his Ph.D. degree in the Electrical Engineering and Computer Science Department, McCormick School of Engineering, Northwestern University. He served an internship in Cray Research's Industry, Science & Technology Department during the summers of 1989 and 1990.

Allen Taflove is a professor of electrical engineering and computer science at Northwestern University. He recently was named an IEEE Fellow for his pioneering of FD-TD methods in CEM since 1973. He also is a national lecturer for the IEEE Antennas and Propagation Society and technical program chairman of the 1992 IEEE Antennas and Propagation Society International Symposium in Chicago.

Jeffrey P. Brooks is a manager of benchmark analysts in Cray Research's benchmarking department. He received an M.S. degree in mathematics and an M.B.A. degree from the University of Minnesota in 1985 and 1991, respectively. Before joining Cray Research in 1989, he was an instructor of mathematics at St. Olaf College in Minnesota and a technical representative at the Minnesota Supercomputer Center.

Evans Harrigan is a senior systems engineer at Cray Research, with 23 years of experience in supercomputing. He received a degree in mathematics from Michigan State University.

References

1. Rao, S. M., D. R. Wilton, and A. W. Glisson, "Electromagnetic Scattering by Surfaces of Arbitrary Shape," *IEEE Transactions on Antennas and Propagation,* Vol. 30, pp. 409-418, May 1982.
2. Newman, E. H. and D. M. Pozar, "Electromagnetic Modeling of Composite Wire and Surface Geometries," *IEEE Transactions on Antennas and Propagation,* Vol. 26, pp. 784-789, November 1978.
3. Taflove, A., "Review of the Formulation and Applications of the Finite-Difference Time-Domain Method for Numerical Modeling of Electromagnetic Wave Interactions with Arbitrary Structures," *Wave Motion,* Vol. 10, pp. 547-582, December 1988.
4. Shankar, V., A. H. Mohammadian, and W. F. Hall, "A Time-Domain Finite-Volume Treatment for Maxwell's Equations," *Electromagnetics,* Vol. 10, pp. 127-145, January-June 1990.
5. Jurgens, T. G., "Conformal Finite-Difference Time-Domain Modeling," Ph.D. dissertation, Northwestern University, EECS Dept., 1990.
6. Ludwig, A. C., General Research Corp., Santa Barbara, California, private communication with T. G. Jurgens.