Abstract
Presentation
Back to Agenda
Next Abstract

Parallelization of a Electromagnetic Analysis Tool

Milissa Benincasa
Black River Systems Co.
162 Genesee Street
Utica, NY 13502
(315) 732-7385 phone
(315) 732-5837 fax
benincas@brsc.com
United States

Chris Card
Black River Systems Co.
162 Genesee Street
Utica, NY 13502
(315) 732-7385 phone
(315) 732-5837 fax
card@brsc.com
United States

Alan George
Black River Systems Co.
162 Genesee Street
Utica, NY 13502
(315) 732-7385 phone
(315) 732-5837 fax
george@brsc.com
United States

1. Executive Summary

This paper describes research work currently being conducted under the Common High Performance Software Support Initiative (CHSSI) sponsored by the DoD High Performance Computing Modernization Program (HPCMO). A scalable, portable, parallel electromagnetic modeling tool is being developed that will provide the capability to rapidly generate scenes of radiating and scattering structures (targets and their surrounding environment) in realistically complex electromagnetic environments. This tool allows users to accurately model targets embedded in their environment. It will be able to solve problems 10 to 100 times larger in liner dimension than previous models. The parallel electromagnetic modeling tool is providing exciting new design and research possibilities for electromagnetic analysis. The team assembled to conduct this research effort consists of the Air Force Research Laboratory, Naval Research Lab. (NRL), US Army Space & Missile Defense Command (USASMDC), Black River Systems Company, RADC, Syracuse University, University of Toronto, and SUNY Binghamton.

2. Background

Electromagnetic analysis requires the solution of Maxwell's equations in either the time or frequency domain. In realistic applications, closed form solutions do not exist and numerical solutions to either the differential form or the integral form of Maxwell's equations must be employed. Finite Element Method (FEM), Finite Difference (FD), and Finite-Difference-Time-Domain (FDTD) are examples of differential equation techniques that can be employed. However, integral equation techniques are inherently better suited to the analysis of antennas and thick or thin dielectric/magnetic materials. Of all the integral equation techniques, Method of Moments (MOM) solves Maxwell's equations more directly, leading to more accurate analysis.

The electromagnetic modeling application, WIPL-D (Wires, Plates and Dielectrics), is a well known commercially available analysis tool based on Method of Moments. It has evolved from over ten years of research in numerical electromagnetics. The development of this code has been driven by the strong demand for analysis of composite conducting and dielectric structures. WIPL-D's strength is in its efficient use of a bilinear quadrilateral, entire domain technique, which greatly reduces the number of unknowns needed for large problems. It requires only 10-20 unknowns per square wavelength, as opposed to the 100s needed by subdomain approaches. This efficient methodology is critical for the solution of large problems where the required number of unknowns can grow unmanageably. This processing advantage can be observed by looking at a typical benchmark problem, the analysis of electromagnetic scattering from a 16 wavelength diameter conducting disk. The analysis using WIPL-D requires 16 hours using 5320 unknowns and 144 MB of main memory. This same problem using a triangular patch subdomain approach on the same computer requires 45 hours, 18,000 unknowns and 2 GB of required memory.

3. The Need

While WIPL-D is an efficient code, there is a desire to solver larger and larger problems. As one begins to include large mounting structures and more of the environment, the size of the problem grows geometrically. It is very difficult for the current version of WIPL-D to solve problems of this magnitude. Problems 10 to 100 times larger in each linear dimension will require not only efficient code, but also advanced parallel processing techniques. Our research extends the current WIPL-D tool to develop a parallelized tool that will address these needs.

Abstract Presentation Back to Agenda
Next Abstract

This parallelized WIPL-D tool will be applicable to a wide range of end-user applications. Applications of interest include: Detection of Targets Under Trees (propagation through foliage), Ship Radar Performance (multipath scattering from many structures), Strategic Subsurface Target Detection, and Land Mine Imaging/Detection (propagation through and scattering from the soil) just to name a few. These applications have one unifying feature: the need to accurately model the scattering from a target (or the radiation from an antenna) within a complex scattering and propagation environment. The other feature that these applications all have in common is the need for a parallel implementation of WIPL-D to solve the problem.

4. Progress to Date

This year's parallelization effort has focused on two main areas to speed up WIPL-D's processing time. These include distributing the main frequency loop and the solving of the impedance matrix. The next step will be to focus on distributing the construction of the impedance matrix. Parallelizing the matrix construction will provide additional increased functionality. The goal for this phase of the project is to achieve linear speedup through parallelization on two high performance computing machines (HPC). A demonstration application has been defined for accomplishing this. The selected application is a simulation of a cell phone next to a human head. To solve this problem requires solving 3,549 unknowns. The target HPCMO machines for this phase are Huinalu (Linux cluster) and Tempest (IBM SP3) which reside at the Maui High Performance Computing Center.

Frequency parallelization is currently implemented and functioning on the target HPC machines. The main loop of the program performs iterations for all frequencies requested by the user. Each loop iteration is independent, providing for a high level of parallelism. All frequencies are evenly distributed among the available processors. The root processor enters the program and inputs data from a file created by WIPL-D. It then distributes this data to all other processors. Output data from each processor is stored in separate files that are collected and joined by the root processor. The results to date are consistent with expectations. We will have more concrete numbers by August 2003 when the Alpha Testing is conducted. With a single processor on Huinalu the run time is approximately 6000 (+/- 100) seconds for two frequencies. When run in parallel on two processors the processing time is around 3000 (+/- 100) seconds. For four frequencies the single processor time is over 14,000 seconds, and running on four processors it is still approximately 3000 (+/- 100) seconds.

The solution to the impedance matrix is currently implemented and functioning on a single processor Linux machine. We are currently working to get it running on the parallel machines. The matrix solution is performed using the Scalapack parallel LU decomposition and solution functions for complex data. A new function was created to replace the original matrix solution function in WIPL-D. In it, the processors are grouped into a processing grid as required by Scalapack. The impedance matrix is then divided among the processors in the grid in a 2-D block cyclic distribution. The Scalapack functions pzgetrf() and pzgetrs() are used. This parallelization should have a significant impact on the performance of WIPL-D. Through profiling, it was determined that the solving of the matrix consumed 75% of the total processing time for larger simulations. A new function was inserted into the code in a way such that alternative matrix solution functions can be substituted in place of the Scalapack functions.

The parallelization of the impedance matrix construction is the next step. Once completed it should provide additional functionality and faster execution. A current shortcoming of WIPL-D is that as the simulation becomes very complex, the impedance matrix grows too large to be held in memory. This limitation inhibits the solution of a number of DoD applications. Parallelizing the impedence matrix will allow the matrix to be constructed as separate sub-matrices that are distributed among the processors. Splitting up of the matrix could potentially cause slower execution times due to communication overhead. This overhead may be overcome by exploiting sections of the matrix construction that exhibit levels of parallelism.

Using the current parallelization of WIPL-D, we are able to achieve the required speedup for this phase of the project. As new levels of parallelism are added, the potential exists for greater speedup performance. The eventual goal for this effort is to provide the capability to perform simulations in less time and expand the application base of the models.