Microwave Journal 50 Year Retrospective Series:

A Few Personal Remarks on the Evolution of Computational Electromagnetics

Zoltan J. Cendes
Ansoft Corporation
Pittsburgh, PA, USA

Introduction

I remember reading the Microwave Journal as a graduate student in the Microwave Lab at McGill University in Montreal in 1969. Then, as now, the Microwave Journal was filled with articles, observations and product information on microwave and wireless hardware, although today there are also articles on microwave software. The graduate students around me were also building hardware – remote sensing devices, communication devices, antennas, and the like – but my project was different: I was writing a computer program to model microwave fields.

Computers in 1969 were still primitive. This is an odd statement since the electronic computer was over two decades old by that time. Yet the basic procedure to program a computer had been unchanged for years: in 1969, one still used punched paper cards to write a computer program and submitted this card deck to a central computer facility. The memory available on the mainframe computer was miniscule: the IBM 360 mainframe computer, figure 1 at McGill had 128 Kbytes of available main memory. I remember limiting my computations to solve about a 100 X 100 matrix; anything larger was likely to cause a program overflow.

Figure 1. Researchers (not the author) with the IBM 360 Mainframe Computer

Since memory was small, the early work in computational electromagnetics focused on simple algorithms. Chief among these was the finite difference method that was simple to program and avoided explicitly storing the coefficient matrix. K. Lee had published the finite difference time domain algorithm in 1966. This was all changed by the two pioneering giants of computational electromagnetics, Roger Harrington and Peet Silvester. Professor Harrington developed the method of moments procedure for solving the integral form of Maxwell's equations, while Professor Silvester developed the finite element method for solving the differential form of Maxwell's equations.

I was fortunate to join Peet Silvester's research group in 1969 as his third graduate student and the first one to work on finite element methods for microwave engineering. Finite element methods had originated as an efficient procedure to solve structural engineering problems; Professor Silvester was the first to apply this method in electrical engineering. In particular, he had just published the finite element solution of homogeneous waveguides; my research project was to extend this work to inhomogeneously loaded dielectric waveguides.

The excitement of reaching beyond the simple finite difference algorithm and developing new and more efficient computational algorithms propelled me forward. Professor Silvester was a fountain of inspiration and genius and developed several basic techniques still used in finite element analysis today. Nevertheless, my project failed; the procedures I developed at that time are now known to generate unphysical solutions of Maxwell's' equations called spurious modes. While some of the solutions produced by the procedure were correct, others were not, and since there was no easy way to tell the difference, the method was useless. I moved on and did my PhD on a different topic.

After graduation, I went to work at the General Electric Corporate Research and Development Center in Schenectady, NY. My initial projects were to develop finite element computer programs to simulate the behavior of large steam-turbine generators and power transformers. The finite element method worked well in these cases, but one day I was assigned to develop a program to simulate microwave ovens. GE had just come out with the Space Saver^{TM} microwave oven to fit above the range in the space where the ventilation fan used to be. The problem was that the thin vertical dimension of this microwave oven caused highly non-uniform heating; hence the need for better understanding through simulation.

While simulating the fields in a 3D microwave oven was significantly different from 2D waveguide analysis, spurious modes were generated once again. I was intrigued: while I didn't understand the reason for the breakdown of the application of finite element methods to microwave engineering at the time, I knew the answer was important. After all, mathematicians had proved that the finite element method had a higher rate of convergence than did the finite difference method and the method had already become the method of choice for solving structural engineering problems.

The breakthrough came after I joined Carnegie Mellon University in the 1980. J. C. Nedelec published a tersely worded paper in 1980, which examined the continuity conditions on finite element approximation function spaces and showed that the electromagnetic approximation space was smaller than had been previously thought. To generate a practical finite element procedure, at CMU, my graduate student Mike Barton and I developed new types of finite element within the Nedelec spaces called edge elements that interpolate to the tangential components of the field along the element edges as shown in Figure 2 (a). We applied these new elements to the solution of 3D magnetostatic field problems in 1986 and showed that they provide correct solutions. My graduate student Jin-Fa Lee and I also created the higher-order edge elements illustrated in Figure 2 (b).

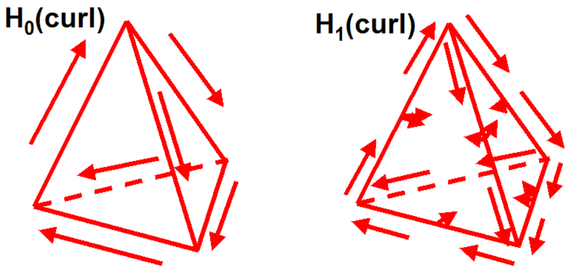

Figure 2 (a) Edge elements (b) High-order tangential vector finite elements

It remained to show that edge elements solve the problem of spurious modes that had plagued the earlier attempts to solve microwave field problems with finite element methods. This was done by my graduate student Steve Wong; in 1988 we published the key result that edge elements satisfied a property now called discrete compactness which correctly maps the nullspace and the range space of the finite element system. This proved that spurious modes were eliminated by using edge elements. Although it took many years, with edge elements we are able to model microwave ovens reliably and accurately.

The 1980's were also a period of entrepreneurial opportunity. The new numerical methods and the introduction of the IBM PC in 1981 led a number of my fellow professors and researchers to start electromagnetic field simulation software companies: Sonnet, MAFIA (now CST), Vector Fields, IES, Infolytica, Magsoft/Cedrat. I started Ansoft in 1984; our first products were for the solution of 2D low frequency field problems. Memory was still small by today's standards; the original PC was limited to 640 Kbytes. By 1988, we had also developed 2D microwave field simulation software based on the new edge elements; we introduced this software at a booth at the IEEE IMS event in Las Vegas in 1988. Shortly thereafter I was called to Santa Rosa by the Hewlett Packard Company to present our approach to microwave field simulation software. This led eventually to an OEM agreement whereby Ansoft Corporation developed the High-Frequency Structure Simulator (HFSS) for sale by HP.

Our initial excitement in signing the OEM agreement with HP turned into to a crisis for the young Ansoft. One of the requirements in the agreement was to create software that could be used by engineers having little or no experience with finite element analysis. Since the accuracy of all electromagnetic field simulation methods depend upon the discretization of the problem, the key to making finite element analysis easy to use is to automate the mesh generation process. At Carnegie Mellon, we had developed algorithms for adaptive mesh refinement that not only created the mesh automatically but also concentrated the mesh refinement according to the distribution of the field. To provide an optimal allocation of resources, and hence high accuracy with a minimal mesh, adaptive mesh refinement generates small elements in locations with high field gradients but allows large elements where the field is smooth. Figure 3 shows an animation of the adaptive mesh refinement process for a patch antenna. Not only is this process automatic and refines the mesh near the input trace and the edges of the patch as needed, it also provides a measure of the error in the solution. As seen in Figure 3, the fast frequency sweep of the input admittance of the antenna converges to the exact solution as the mesh is refined.

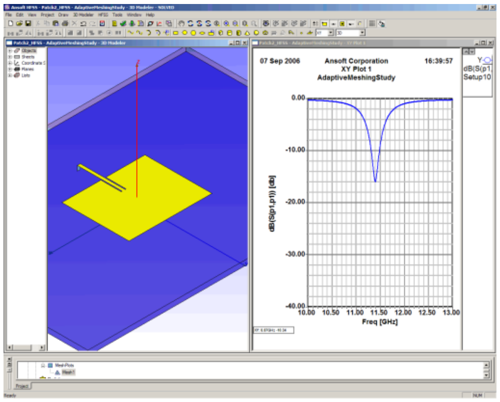

Figure 3. Adaptive mesh refinement for a patch antenna showing convergence of the input admittance as the mesh is refined.

While the adaptive mesh refinement algorithm generates outstanding meshes and provides a way to measure solution convergence, the resulting varying element sizes generate poorly conditioned finite element matrices. To make HFSS efficient for shipment of HFSS by HP, we had counted on using an iterative matrix solution algorithm. Yet these iterative matrix solution algorithms were unreliable with the combination of edge elements and adaptive refinement procedures we had devised. By the time we discovered the unreliability of the iterative matrix solution procedure, we had essentially bet the company on HFSS. Almost all of our resources were focused on developing this product and little revenue was coming in from other sources. Finally, after a delay of many months, we managed to develop a direct solver that was efficient enough to ship with HFSS, although it was not as efficient as we would have liked. HFSS v1 shipped in October 1990.

HFSS was introduced to the world in the Microwave Journal in February 1990 (Figure 4). I remember being thrilled in seeing the work that I had participated in being featured on the cover of the leading trade journal I had read since my graduate student days. The discerning HFSS user will notice that the mesh displayed in this cover picture is a regular mesh from our pre-direct solver days and not the adaptive mesh shipped with version 1 of the software. Many of today's users would also be surprised to learn that the simple coax-to-waveguide adapter featured on the Microwave Journal cover required 16 hours to solve for a single frequency point on a 1990 HP computer. Using HFSS v11 and a modern computer, the same adapter solves for 10 frequency points in 3 seconds today. That's a near 200,000 to 1 speed-up!

Figure 4. The February 1990 Microwave Journal cover announcing HFSS.

How did HFSS get to be so much faster? Part of the answer lies in faster hardware – Moore's Law states that the number of transistors on an IC and consequently computer hardware performance doubles every 18 months – but an even larger part of the answer lies in improvements in software – we may formulate an analogous law to Moore's Law for computational algorithms, namely the speed of electromagnetic software algorithms doubles every 12 months. HFSS has been made faster in every release. Significant speed improvements have been made in dozens of algorithms within HFSS including better adaptive mesh refinement criteria, improved meshing algorithms, superior port solutions, and the ability to perform broadband fast frequency sweeps. Fast frequency sweeps first appeared in version 3 of HFSS and provide hundreds of frequency points for little more than the cost of a single frequency solution. Our latest HFSS speed-up is a result of 17 years of research. Ever since our failure to develop a reliable iterative solver for HFSS v1, we have pursued ideas for achieving this goal. Finally, HFSS v11 shipped last year with an iterative matrix solver. As Figure 5 shows, the growth rate in both computer time and memory grows nearly linearly with problem size. This means that as you double the size of a problem, computer time and memory also roughly doubles. Since storing twice as many numbers requires double the memory, this is near optimal performance.

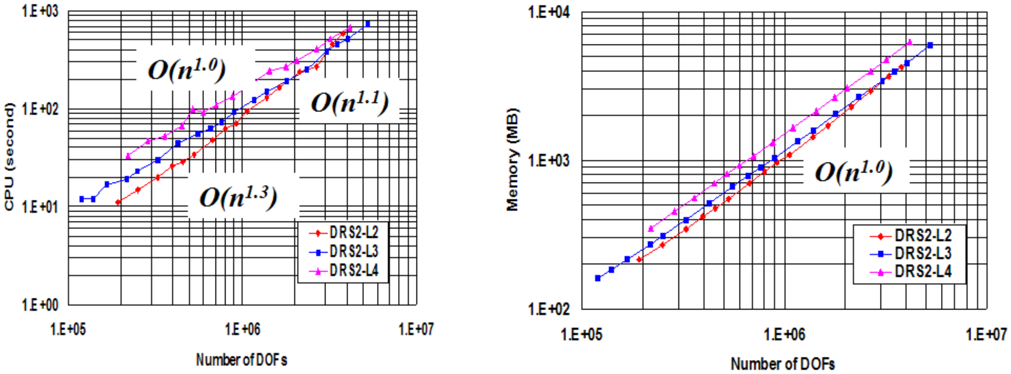

Figure 5 (a). FEM Time Scales as N1.0 to N1.3. (b) FEM Memory Scales as N1.0.

The new iterative solver allows massive problems to be solved. Figure 6 shows the magnitude of the electric field in a dish antenna while Figure 7 shows the field and currents generated by a 4 GHz antenna on a drone aircraft. Solving such large problems is made possible by all of the speed improvements including the iterative solver in HFSS v11.

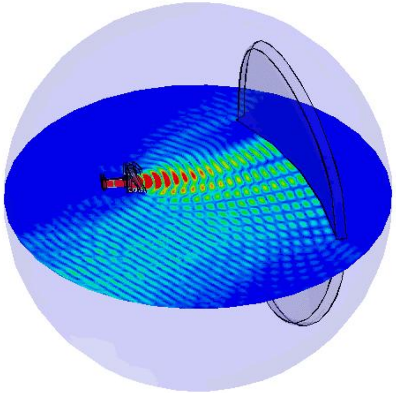

Figure 6. The magnitude of the electric field in a dish antenna.

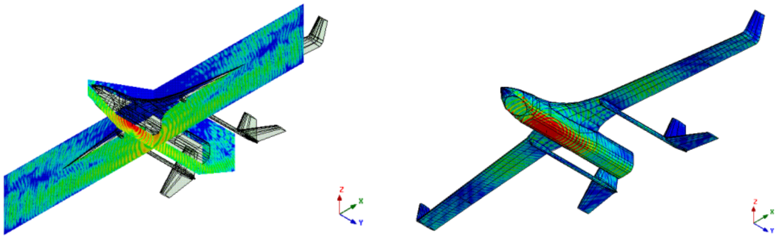

Figure 7. (a) Field and (b) current on an airplane from a 4 GHz Antenna

Electromagnetic field computation continues to evolve. The early days of computational electromagnetics were a fertile time of innovation that set the stage for today's innovations. While the basic algorithms in use today were developed twenty, thirty or even forty years ago, numerous improvements have been made to these algorithms to make them faster and more robust. Today's engineers are used to solving for the electromagnetic fields from entire airplanes or complex antenna structures. Such solutions would have been impossible just 5 years ago. Microwave engineers now have the ability to solve 3D microwave field problems easily and efficiently. The world of high frequency design is forever changed and the Microwave Journal was there from the beginning and continues to communicate the evolution of computational electromagnetics.