From February 2010 *High Frequency Electronics*

High-Frequency Algorithmic Advances in EM Tools for Signal Integrity—Part 2

By John Dunn
AWR Corporation

Part 2 of this two-part series on new advances in the algorithms underlying EM simulation techniques examines the use of time-saving "fast frequency sweep" techniques

The first article in this series looked in detail at 2.5D and 3D method-of-moments (MoM) EM solvers as they relate to signal integrity designers and illustrated how, along with the power of multi-core processors and computer memory, new mathematical algorithms and techniques such as state-of-art-preconditioners, compression techniques, and multipole formulations have dramatically increased the capabiliteis of EM software.

This article examines a second class of mathematical algorithms that have dramatically increased the speed of EM simulators. It describes ways to reduce solution times by using fewer frequencies, while ensuring the frequency resolution of the resulting dataset. The methods are known generically as "fast-frequency-sweep" techniques, and they attempt to reduce the number of frequencies required to obtain the simulation response of the problem while maintaining accuracy. The methods are also useful for solving the problem of using the EM simulation results in time-domain simulators.

In an ideal world, every designer would receive his or her results instantaneously with no errors, while solving it on the computing equivalent of a digital watch over gigahertz bandwidths and 1-Hz increments. Although this is obviously not possible, improvements in EM codes offer many different ways to tackle some or all of these challenges. As in all things electronic, trade-offs in accuracy, speed, problem complexity (or computing capacity), and frequency span and granularity, are required.

Advanced Frequency Selection Versus Convolution

Advanced frequency selection (AFS) allows interconnects to be studied at many frequencies by simulating only at a few frequencies. The computer rather than the designer automatically chooses the optimal simulation frequencies. The benefits are significant. The computer has the ability to choose fewer frequencies (which reduces the total simulation time), and from these selected frequencies provide coarser spacing, interpolating the intervening ones with a high degree of accuracy without actually solving the EM problem at these finer frequencies. In addition, the methods are normally iterative. That is, a few frequencies are initially selected and more frequencies are then solved for in subsequent iterations based on the initial frequency results. This gives the designer some idea of error or convergence rate.

These algorithms are not really providing a true error, which is only possible if the exact results are known and the problem is solved with very high accuracy at each frequency generated by the AFS algorithm. They instead provide a glimpse of how much the answer is changing with further refinement, which is often useful as it provides a practical measure of how well the problem has converged. Unfortunately, the obvious technique of simply choosing a few points and drawing straight lines between them is not very useful, as a resonance in the frequency response for a high Q circuit can be completely missed. Fortunately, mathematicians have developed a number of clever ways of estimating the response that attempt to include the basic physics of the underlying system.

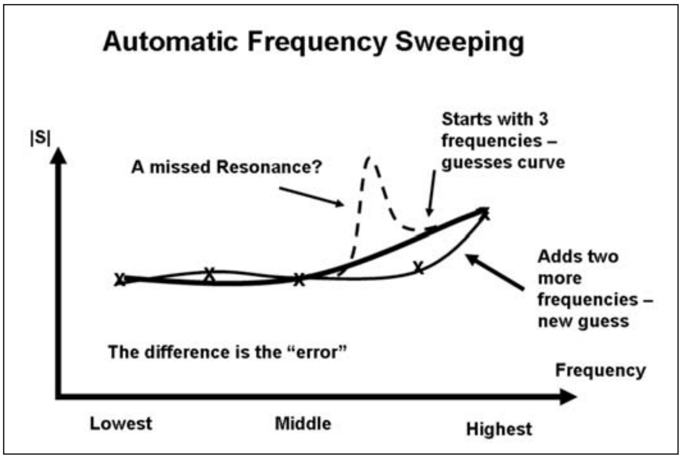

Figure 1 · Estimating error in approximate frequency sweeps.

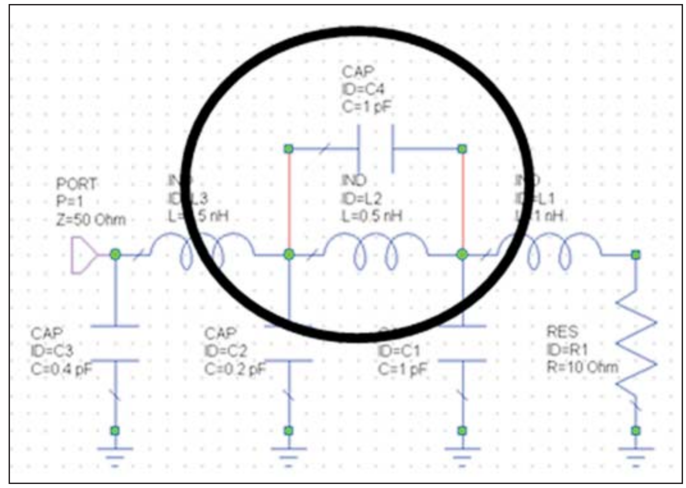

Figure 2 · An EM response fitted to circuit elements. As an LC pair can represent a resonance and it is derived from Maxwell's Equations, it can be used to map a finite set of discrete frequency results for an EM analysis onto a continuous frequency range.

These methods are also useful in solving another outstanding problem for EM simulations. One way to classify EM simulation technology is in terms of the frequency and time domains, and the results are interrelated by the frequency-time duality inherent in circuits. A designer can switch between the time-domain and frequency-domain response using Fourier transform techniques, but a problem arises in using frequency-domain data in a time-domain simulator like SPICE or its variants. This technique, called convolution, is extremely slow, but the advanced frequency sweeping algorithms can also be used to efficiently model the frequency-domain data, allowing the time-domain simulator to quickly and efficiently use the resulting reduced data set.

Speed-Increasing Sweeping Techniques

Frequency-domain solvers, both 2D and 3D, must solve EM problems at multiple frequencies; a single frequency is rarely useful to a designer. For expediency, designers normally choose a fixed step size spanning the range of frequencies needed for the simulation. So it's reasonable to wonder if the number of frequencies required to simulate the problem can be reduced while also delivering a reasonable answer over the frequency range of interest. If care is taken when developing the algorithm, the answer is yes. There are two goals: the first to reduce the required simulation time, and the second to give the designer an estimate of simulation error. A generic approach for estimating this error is shown in Figure 1. The details vary with the specific algorithm being used.

The problem is first simulated at three frequencies: the minimum frequency (F_{min}), the maximum frequency (F_{max}), and the midpoint (F_{mid}). The results for the *S*-parameters can be approximated for all frequencies by drawing an interpolating curve, in this case a quadratic. Two more frequencies are then simulated and are shown in Figure 1 as the midpoints of the lower and upper halves of the frequency range. A new curve is drawn, and while the curves will not agree at all frequencies, the difference between them is an estimate of the error, with the largest difference used as the error criterion. If this error is less than a specified goal, the problem can be considered converged. If it is too great, more points are added. Just where in the frequency range the points are added depends on the algorithm used, but eventually either the error criteria are met or the maximum number of simulations is reached.

The trick in this approach is to determine the estimating curve. The straightforward method of using some form of polynomial fit is not a good choice for several reasons. First, high order polynomial interpolations are notoriously unstable. This problem can be overcome by patching together lower-order polynomials over subsets of the entire frequency range. Spline fits, for example, could be made to produce smooth curves that are numerically stable. A bigger problem is that the interpolation can completely miss a resonance. This is shown in Figure 1, where the actual structure has a resonance, and a naïve interpolation approach completely misses it. This can be especially problematic for high Q structures, but sophisticated techniques have been developed to mitigate it. The results of this continuing work are available in the technical literature.

It doesn't seem possible upon first inspection that a few frequencies can approximate all of the frequency points needed to accurately represent a resonance, especially when the sample points are not near the resonance in question. However, mathematical representations and

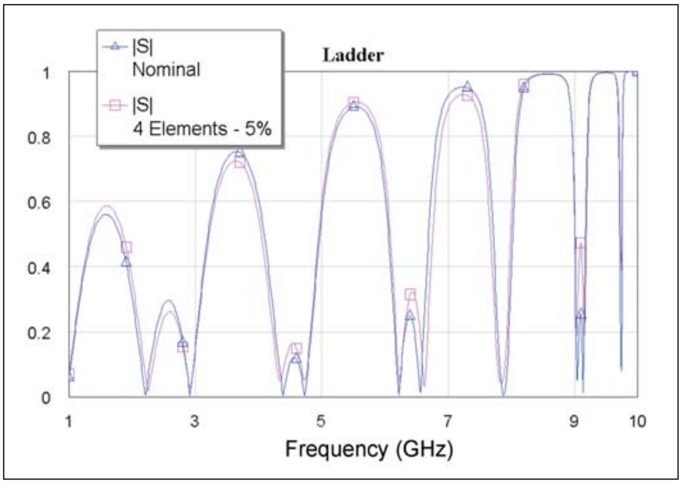

Figure 3 · A drawback of the circuit-fitting approach is that more resonances require more elements, which leads to higher order polynomical representations with greater sensitivities. In this plot, a four-element LC ladder demonstrates a 2x error at 9 GHz with just a 5% LC deviation.

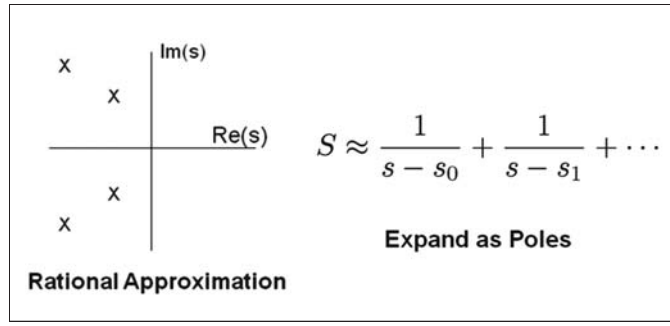

Figure 4 · Rational approximation fitting maps an S-plane representation of the EM results at a small, finite number of points to a rational function, *S*, defined by fitting a set of poles to the *s*-plane representation.

methods to parameterize them have been developed to capture the underlying physics of the problem.

Consequently, using the notion that the EM simulations themselves obey Maxwell's equations, and that they also are the basis of circuit theory, it should be possible to make a reasonable circuit model that can reproduce the S-parameters generated by the EM simulator as shown in Figure 2. This approach has merit and is extremely powerful, but can require many elements if the circuit of interest is complicated, as are very dense SI interconnects. The large number of elements isn't necessarily a problem if the fit is exact everywhere. If one resonance is missed, the designer must add more elements to create a larger order polynomial. Higher order polynomials are notoriously unstable and can introduce greater and greater error. This approach was intensively investigated in the 1990s and found wanting. For example, Figure 3 shows the response of the circuit drawn in Figure 2. It also shows the change in the response when the values are changed by 5% for the circuit elements. Note that the prediction of the resonance at about 9 GHz is off by 100%. This is because the high-order polynomial underlying the circuit response is poorly behaved numerically.

New methods were therefore investigated starting in the mid 1990s. Researchers realized that it is not necessary to use actual circuit elements to approximate the response. Mathematically, the circuit elements result in equations in the frequency domain that can be represented by polynomials in the complex plane. So it is logical to think of poles and zeros in the complex plane as a good way to represent the S-parameters. The goal is to choose a few poles in the complex plane that reasonably represent the complex response, where "reasonably" is the error previously discussed for AFS. There is no attempt to start with an actual circuit model. Rather, the entire approach is to model the system response as an abstract model of poles.

The motivation for this approach is a compact way to find and represent the resonances. As resonances normally occur because the response is being dominated by a single pole in that frequency range, the entire response might be represented by a rational function defined by them. If the designer or the frequency selection algorithm can guess the dominant poles, the result should be a frequency response that works well over the entire range (Figure 4).

The remaining question is how to choose the poles, a problem the EM community has been trying to answer for 20 years. A large portion of the underlying mathematics is based on advances in approximation theory and control theory, and the "trick" is to choose the poles so as not to violate the basic underlying physics, ensuring it is causal, stable, and passive. It helps to look at each one of these requirements individually.

Causality states that a response cannot be obtained before the excitation, so a reflected signal cannot be generated before the incident wave arrives. In other words, the signal cannot show up at the load before it leaves the source. This condition is not that difficult to meet for pole-zero models. Incidentally, it is easy to violate this condition with traditional circuit models when improper values are used. For example, microwave engineers often create models with negative inductance. This works in the frequency domain but is a disaster when tried in time domain simulators.

Passivity simply implies that the structure cannot create energy unless there is an energy source such as a power supply. For example, referring to Figure 5, S-parameters generated in EM simulators should theoretically be passive (the exact solution in Figure 5). Passivity is the hardest criterion for the algorithms to meet (the spline solution). The details of implementation are quite extensive, but essentially the algorithm must guess pairs of poles and then test for passivity (non-passivity vector fitting). Passivity has been the most difficult physical constraint to satisfy in the current algorithms. Methods that ensure passivity are only practical on small problems. The problem becomes more difficult for multi-dimensional S-parameters with many ports. For a large numbers of ports, there is no practical way to ensure passivity, and algorithms rely on reasonable checks that work in most cases. When using these frequency-reduction methods, it is important to perform a passivity check on the raw dataset, especially if a SPICE simulation is to be performed.

It was mentioned earlier that these methods can also be used for time-domain simulations. EM simulators generate S-parameters and many SI engineers want to use time-domain simulators to look at eye diagrams, use time domain models, and study hysteresis effects and switching issues. S-parameters must therefore be placed into a time-domain circuit simulator, which presents a big problem. The straightforward way to include S-parameters in the time domain (in SPICE for example) is to invoke convolution, which is extremely slow. The method requires that for each time step taken in Spice, integrations be performed over all previous time steps. A much more computationally efficient method is for the S-parameter file to be approximated by poles in the complex plane, making it straightforward to represent these poles in SPICE by voltage-controlled current sources. The simulation is much faster and S-parameter effects are included. The same caveats mentioned above apply when choosing the poles: the approximation must be constrained by keeping the results causal, stable, and passive. The very same technique employed to speed up EM simulation is also used to produce a model for the time-domain simulation that speeds up the time-domain simulation itself.

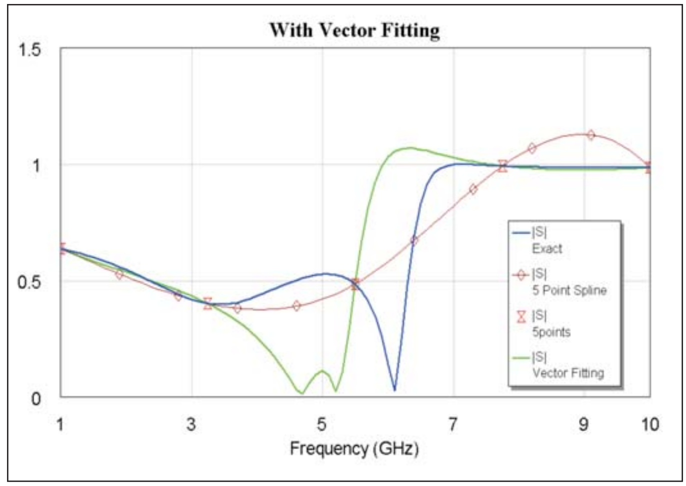

Figure 5 · Passivity analysis. Comparison of different AFS approaches for five EM analysis points: spline fitting of five points versus pole-zero vector fitting. The exact solution created by the EM solver is shown in blue.

Summary

Advances in EM theory, applied mathematics, and computing are making EM simulators of more practical value to SI engineers. Most designers appreciate that computers are becoming ever more powerful. By discussing two important examples, we have attempted to show that that is not the only reason for the increased power of EM tools. New mathematical techniques that have been developed over the past 20 years are being incorporated into commercial simulators. Compressed, iterative solvers are solving large problems that could not be imagined even a few years ago. Advanced frequency sweeping methods are giving simulators the capability of predicting circuit performance over large bandwidths in a fraction of the time necessary with discrete frequency stepping methods.

The obvious question at this point is what is next for new solution techniques, and what kinds of problems can be solved? Although the future is obviously hard to predict, parallel computing algorithms are an obvious candidate. Researchers are hard at work trying to cleverly use the immense power now readily available with inexpensive clusters of machines. Nevertheless, two things are certain: dedicated researchers are hard at work coming up with the next great EM algorithm and SI engineers will eagerly exploit the predictive power these new algorithms bring.

Author Information

Dr. John Dunn is a senior application engineer with AWR whose area of expertise is electromagnetic simulation and modeling. He was a principal engineer at Tektronix for four years before joining AWR and was a professor of electrical engineering at the University of Colorado for 15 years. He received his BS degree in physics from Carleton College, and his MS and PhD degrees in applied physics from Harvard University.

Readers should address questions or comments to: info@awrcorp.com