White Paper

Turning Signal Integrity Simulation Inside Out

DesignCon is the premier conference on high-speed electronics at the chip, package and system level. Held annually in the heart of Silicon Valley, it brings together top design engineers, researchers and industry executives who drive the frontier of electronics. At the conference Dr. Zoltan Cendes, the founder of Ansoft Corporation (acquired by ANSYS in 2008) and a Life Fellow of IEEE, delivered a keynote presentation on how signal integrity (SI) simulation has been turned inside out with electromagnetic (EM) field simulation.

Inside Out

So what does it mean to turn signal integrity simulation inside out? Modern simulation has physics-based solvers in the foreground supported by circuit and system simulation rather than the other way around. All electronic design is fundamentally based on Maxwell's equations, so naturally the most rigorous way to deliver accurate simulation of high-performance systems is to solve those equations directly. In the past, electronics were not dense enough or fast enough to require solutions of the fully coupled Maxwell equations (what we call "full-wave" field solutions), nor did we have the computers able to solve these giant problems. Traditional electronic design validation relied mainly on circuit simulators like SPICE. Only isolated portions of the design were modeled using electromagnetics. Electrical parasitics were typically extracted using approximate, quasi-static table look-up methods. We worked with circuit simulation, quasi-static solutions and hybrid methods for signal integrity design. Electromagnetic simulation has now evolved to the point where entire electronic systems can be simulated from a layout. Advanced numerical methods, high-performance computing and new technologies that automate the handling of massive EDA data sets have enabled this transformation. It is even possible to couple the electromagnetic analysis with multiphysics simulations, so that designers can evaluate the impact of thermal and mechanical stress effects.

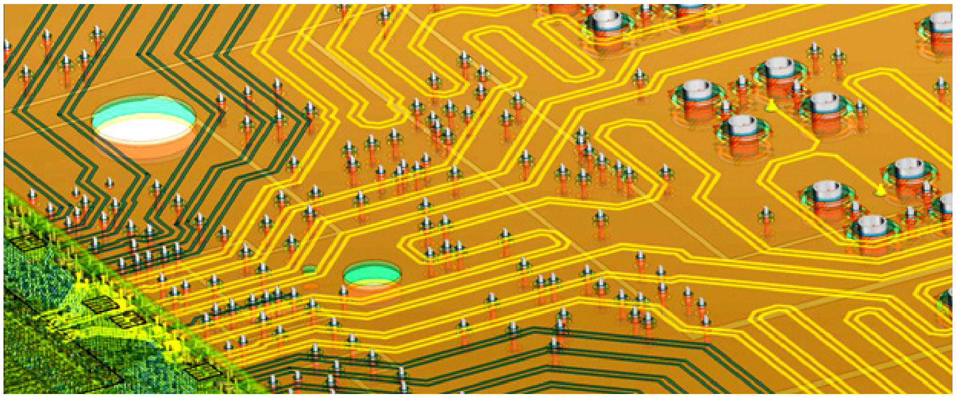

Electromagnetic field simulation is a critical part of electronics design. Power and signal integrity problems are directly due to EM phenomena. Building prototypes is impractical and the cost of fixing problems after a device is built is exorbitant.

ANSYS is pursuing a strategy in which electromagnetic simulation is primary and circuit analysis supports the electromagnetic solution, rather than the other way around. To achieve this goal, we have added new software features that enable engineers to perform transient circuit analysis directly from the layout in ANSYS HFSS and ANSYS SIwave. The idea is to “assemble” an electronic system with IC packages, sockets, printed circuit boards, connectors and cables just as you would in the real world, then perform analysis of that system using appropriate technology. The result is I/O waveforms and eye diagrams just as you might measure in the lab. Because the electromagnetic behavior of the system has been solved rigorously and efficiently, the signal integrity design engineer gets accurate results to discover how the design performs. The ANSYS strategy is to solve “Big EM” systems using ANSYS HFSS and ANSYS SIwave full-wave electromagnetic field solvers. The isolated nonlinear driver and receiver circuits are solved by circuit simulation with accurate, full-wave models of the interconnect. Thus the physics is primary, and circuits are secondary. That is what is meant by “inside-out.” This new era permits circuit and system analysis to be part of the broader physics-based assembly solution. By relying on our layout solution we can automate much of the process.

Why it Matters

Engineers have relied on circuit-level simulation for years. Often they use Synopsys HSPICE® for signal integrity transient solutions and RF simulation methods like harmonic balance for wireless and microwave. Circuit analysis is the foundation of electrical engineering design and it is essential to our craft. We draw a schematic containing models for passive and active components then netlist and run simulation for DC, linear, transient or frequency-domain behavior. All components are coupled together at circuit nodes by applying Kirchhoff’s laws so that interaction among components at those nodes can be analyzed. There are some components that require more specialized models, such as transmission lines, nets on a PCB, connectors or microwave filters. HSPICE uses the W-element, for instance, for transmission lines and nets, and can accept S-parameter models for other components, such as connectors. A challenge occurs as signaling speeds and frequencies increase to the point where we exceed the applicable range of a model, or when the models don’t take into account coupling among components. Consider a complex PCB with thousands of nets and hundreds of components. Can we partition the system into the thousands of individual models for all the nets and know where coupling occurs among them? Which coupling is important and which can we neglect? If we knew that, then maybe we really didn’t need to perform simulation in the first place! We would simply correct any coupling challenge in the layout before simulating.

Starting in the 1980s, electromagnetic simulators were introduced to create models of that complex coupling behavior. Rather than measuring the performance of a connector or the escape routing beneath, we would run an electromagnetic simulation to extract an S-parameter model to be used in circuit simulation. This method works well for isolated EM issues, and has been used to design countless high-speed networking and wireless systems. As switching speeds, frequency and layout density increased, more and more of the system required EM simulation. It also became apparent that system resonances are important for predicting power integrity and EMI/EMC. For those effects, circuit simulation does not offer any remedy; the behavior is dictated by the physical size and shape of the PCB. Larger PCBs have more resonances across a wide frequency band. How do you know when one of those resonances is being excited by the layout and then causing the system to fail in the EMC chamber? With circuit simulation alone you cannot.

We began to address the issues of full system electromagnetics in the late 1990s. It would have been terrific to simply pass a complex PCB to ANSYS HFSS to extract all the EM effects, but at that time computers and algorithms were not able to handle the massive size of the simulation. It would have required tens of gigabytes to solve a single board. Instead ANSYS (Ansoft) and others developed specialized “hybrid” EM solvers (ANSYS SIwave) that automatically combine circuit models for transmission lines, layer-to-layer vias and board resonance models to capture the behavior of full PCB systems. This was a significant breakthrough, enabling engineers to model signal- and power-integrity plus EMI quickly and efficiently. ANSYS HFSS and ANSYS SIwave coexist in the toolbox of the design engineer, enabling extremely detailed full-wave analysis for structures that require that rigor, and full system analysis for large layout-based systems, respectively. HFSS and SIwave are routinely used to extract the passive structures found in electronic and microwave circuits, with results connected to circuit- and system-level simulations to handle the nonlinear and system behavioral effects.

New Modeling and Algorithms

Fast forward to today. What if modern EM simulation techniques and large-scale, high-performance computing (HPC) could handle simulation of an entire PCB with HFSS full-wave accuracy? If you could do it in a reasonable time, why wouldn’t you? There would be no reason to attempt to model a net using a transmission line model or a via with a behavioral model. You simply present the entire system to the EM solver and run a simulation. No concerns over which coupling to include; no worries about board carve-outs that may place artificial boundary conditions on the simulation model, altering the results. New algorithms, automation and HPC are making it possible today.

For instance, several years ago ANSYS introduced a new geometric modeling paradigm in HFSS. Since its inception, HFSS had always been based on a 3-D mechanical CAD (MCAD) interface. Models were created using 3-D primitives like rectangular prisms, cylinders and spheres, along with sheet surfaces for planes. That type of modeling works fine for waveguides, housings, antennas and vehicles. Any PCB had to be recreated in terms of the 3-D primitives, creating a very large MCAD model to support thousands of nets. To rectify this, ANSYS created a layout-based modeler built upon electrical CAD (ECAD) constructs such as layer stackups, nets, vias and padstacks. Indeed, there is an underlying electronics database (EDB) built into the system to support efficient and parametric modeling of ECAD data. Linkage to EDA layout tools from Cadence, Mentor, Zuken and others becomes natural using this data model and layout paradigm. So HFSS inherently has two geometric modelers: MCAD and ECAD. This enables engineers to combine MCAD and ECAD in an electronic assembly to support, for instance, a 3-D MCAD connector placed upon an ECAD PCB layout.

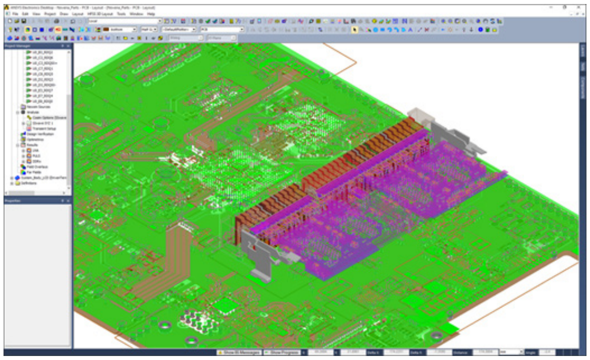
ANSYS HFSS enables engineers to combine a 3-D connector with a PCB layout and solve the connector attached to the board.

Modeling electronics with ECAD has some very significant benefits for simulation. Not only do we have greater capacity to handle the thousands of traces in the design by virtue of lighter weight ECAD, we also “know” the difference among nets, power/ground planes, vias and padstacks. We can pass this information to the solver to be handled in smart ways using new technology. One of those new technologies is our Phi meshing engine. The Phi mesher was specially designed to handle electronics layout data. MCAD models that might take an hour to mesh can now produce the same result in minutes. We then run HFSS using its traditional adaptive process to converge on the right solution.

Knowing more about the ECAD design also leads to significant opportunities to automate the analysis and speed up the simulation. For example, the simulation time in HFSS increases whenever a port is added. Each two-terminal SMT device, like a capacitor in the system, requires two ports to be added to the HFSS model. Complex systems may contain hundreds of these components and thus have traditionally required hundreds of ports. Recent advances in HFSS have dramatically improved the speed of handling models with hundreds of ports; in typical cases simulation can be accelerated by a factor of 20 or more. Another approach is to incorporate the device’s S-parameter representation into the HFSS system matrix, eliminating the need for the ports. As an added benefit, the schematic representation of the EM model is greatly simplified. By preserving the ECAD representation, we can place SMTs, connectors, flex circuits and other components into an assembly. We know where to place them and which pins need to connect to particular pads. This information was lost when we flattened a design to a traditional MCAD model.

Imagine the case of adding a 32-pin connector to a PCB. In a purely schematic-based flow, we would have to extract the S-parameters for the connector in HFSS, and then place that S-parameter black-box into a schematic. We would then have to make 64 connections in the schematic; each conductor in the connector has an input and an output, so 64 individual nodes would have to be connected. Now consider an ECAD plus MCAD assembly method. One simply picks the MCAD model of the connector, orients it correctly, then places it onto the board. Sophisticated algorithms in the background automatically make the electrical connections without further user input on a schematic. This is a huge time savings and also prevents inadvertent errors. What's more, we can script that functionality in the layout and automate the process to match any organization's workflow to make it even easier. Simulation can proceed by extracting the PCB using, for example, SIwave; the cascading of the connector plus PCB is then accomplished using linear circuit simulation. A simple menu-pick will tell the system to instead use HFSS to mesh and solve the entire assembly (no linear circuit cascade) if that is desired for final full-wave verification.

High-Performance Computing

A major enabler for turning SI simulation inside out is advanced high-performance computing. With the right numerical procedures and algorithms, it is possible to take advantage of large compute clusters to accelerate solutions, solve bigger problems, and sweep frequencies and parameters. HPC opens up the possibility of optimizing a design and improving its reliability by making 3-D full-wave simulations of many design variations both practical and efficient.

One of the newest technologies is distributed adaptive meshing. This is the most fundamental change to the adaptive process ever since HFSS was initially released. HFSS has always used a central, user-defined frequency for the adaptive solution, followed by various frequency sweep methods to get a broadband response. The frequency sweep methods used the adapted and converged mesh from that single frequency to compute the system response at all other frequencies. For many years it has been possible for users to converge at one frequency, and then use a "dependent mesh" setting to further adapt at another frequency, but this was a manual process. The new technology automatically adapts the mesh at multiple frequencies. Frequency selection can be done fully automatically with the "Broadband Adaptive Meshing," or at user-defined frequencies using "Multi-Frequency Adaptive Meshing."

These frequency sweeps can be greatly accelerated using HPC. We recently added two new capabilities: distributed frequency sweeps and the S-Matrix-only solver. The first leverages a message passing interface (MPI) to distribute multiple frequency points across nodes in a cluster. Having multiple machines solve individual frequency points allows linear scaling of the speed with the number of nodes. The second, S-Matrix-only solver, is useful when the S-parameters are the only quantity of interest. If the EM fields are not required, we can realize substantial computational gains by not solving for them. For problems in which you don't need to explicitly examine the fields at all frequencies, this greatly reduces the RAM used and speeds the solution due to efficient hardware utilization.

The new Distributed Direct Solver increases capacity by distributing the finite element matrix solution across multiple physical compute nodes. This makes it possible to solve very large and complex designs that cannot be solved on a single machine. It also makes it possible to solve large problems with four moderately priced machines with, say, 256 GB of RAM, instead of having to buy an expensive machine with 1 TB of RAM. Engineers can combine these different distributed simulation technologies. Multilevel distribution allows two- and three-level HPC distribution for parameters, frequency sweeps and the solver. For example, you can do two-level distribution, with distributed design parameter sweeps plus distributed frequency sweeps. You can also do three-level distribution, with distributed design parameters, distributed frequency sweeps, plus the Distributed Direct Solver. With large compute clusters, multilevel distribution makes it possible to simulate large systems in record times.

Breaking Down Barriers and Silos

Signal integrity designers care about time-domain (transient) signals. ANSYS provides the fine details of the EM fields they need to design a chip-package-board system. Recent advances in our layout-based automation and underlying electronics database representation allow us to link to popular EDA tools, solve highly complex electronic assemblies and automatically produce transient signal plots for time domain reflectometry (TDR), eye diagrams and compliance reports.

Several top companies leverage this new level of automation to perform rapid simulation and diagnostics on their designs. For example, Design Rule Check (DRC) capabilities in traditional EDA tools might flag numerous problems in a package or board. Now what? If DRC says there are 25 faults in a design, an experienced engineer will need to examine every one by setting up 25 models. This is hardly scalable. With new automation methods based on layout we can have the software create 25 submodels, solve them across an HPC cluster and provide full reports based on HFSS. Now the knowledge and experience of an organization’s best engineers are embedded in an automated flow that can be used by all SI engineers. This breaks down a significant barrier that has already been implemented by some of our top customers.

The ability to couple full-wave field solvers, which operate most naturally in the frequency domain, with time-domain circuit analysis is a key requirement for signal integrity. Reduced-order modeling enables this coupling. To solve signal integrity problems involving possibly hundreds of ports, it must be possible to efficiently extract the reduced-order model, while at the same time preserving its fidelity to the original S-parameters and ensuring that the model is both stable and causal. ANSYS’ reduced-order modeling incorporates several advanced technologies: the Loewner matrix-based TWA algorithm, fast Loewner SVD, and passivity enforcement through iterated fitting of passivity violations (IFPV). These technologies enable engineers to freely use full-wave models of large boards and packages with transistor-level models of drivers and receivers.

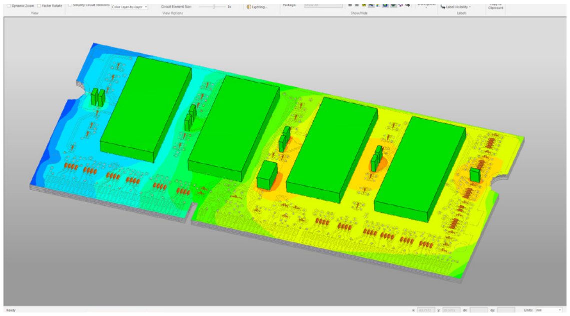

PCBs are densely populated with components and have high currents resulting in temperature increases. Joule heating can affect copper traces and result in delamination and failure. A design flow that spans electromagnetic, thermal and mechanical analysis must be deployed to avoid costly failed prototype and warranty costs.

We're also breaking the silo between electrical and mechanical domains. While we are primarily turning SI (electrical) simulation inside out, we can apply the same idea to other domains. Reliable electronics must meet requirements in both electrical and mechanical domains. Recent advances have enabled electrothermal analysis of layout-based designs. Previously, full PCB geometries were far too complex for mechanical FEA. We combine ANSYS RedHawk to extract chip electrical models, SIwave to model the DC IR drops in the board and package, ANSYS SpaceClaim to create the 3-D MCAD model, and a unique metal fraction mapping algorithm to automatically set up package/board designs for ANSYS Mechanical. For the first time, we provide a design flow that can perform DC analysis, map Joule heating to a mechanical solver, and then produce temperature profiles and associated mechanical deformation and stress. It's a chip- package-board solution that allows engineers to evaluate electrical, thermal and structural behavior.

Transforming Business with Pervasive Engineering Simulation

Advanced simulation is a critical driver for industry-leading companies. We've always argued that leaders adopt simulation to create their competitive advantage, and today we are seeing that coming to full fruition. The automation methods mentioned above break down silos and make simulation more accessible to the design engineer. The future of simulation is that businesses will build it into their processes even more deeply, capturing the know-how of their best engineers in an automated system that becomes available to entire design teams.

Strong executive buy-in makes the change possible. The importance of having senior management engaged as the agent of change cannot be understated. The CTO of an RF module company, for example, came into an ongoing effort for streamlining their design flow. He recalibrated the entire team by setting the goal to improve their design turnaround time by 10x. This goal became the marching orders for the entire team, from design to fabrication to test. Automation of simulation became a huge opportunity for them to meet that challenge.

We are seeing engineering simulation becoming more and more pervasive across the entire product lifecycle, as well as with all types of engineers. While simulation was once the sole domain of experts and used mainly for verification, it is now moving up front in the development process to quickly evaluate changes in design. At the same time, it's also moving downstream of the product lifecycle process to analyze real-time operational data from connected machines in the industrial internet. We call this expansion of our tools "pervasive engineering simulation."

Simulation's pervasiveness isn't surprising because we're growing the applications for simulation technology: We are building into our software the ability to simulate the behavior of thousands of different materials; to understand temperature and warping of a PC board; and to understand how components are wired and what power runs through them. We can simulate not only the electromagnetic behavior of an antenna, but we can digitally place it in an operating environment to determine how it will perform in the real world given the interference with other systems.

Companies large and small are turning signal integrity inside out – and realizing fantastic results.

ANSYS, Inc.
Southpointe
2600 ANSYS Drive
Canonsburg, PA 15317
U.S.A.
724.746.3304
ansysinfo@ansys.com

If you've ever seen a rocket launch, flown on an airplane, driven a car, used a computer, touched a mobile device, crossed a bridge or put on wearable technology, chances are you've used a product where ANSYS software played a critical role in its creation. ANSYS is the global leader in engineering simulation. We help the world's most innovative companies deliver radically better products to their customers. By offering the best and broadest portfolio of engineering simulation software, we help them solve the most complex design challenges and engineer products limited only by imagination.
Visit www.ansys.com for more information.