GPUs SPEED THE SOLUTION OF COMPLEX ELECTROMAGNETIC SIMULATION

The ANSYS HFSS transient solver leverages NVIDIA's leadership in GPU computing to enable quick solutions for transient electromagnetic simulation.

By Hsueh-Yung (Robert) Chao, Lead R&D Engineer, and
Stylianos Dosopoulos, Senior R&D Engineer, ANSYS

Transient electromagnetic (EM) analysis is indispensable for modern electronic design. Since the ANSYS HFSS transient solver's debut, it has been applied to a variety of design challenges in the electronics, semiconductor, energy, automotive, aerospace and defense industries. The product provides robust solutions for applications related to electromagnetic interference and compatibility (EMI/EMC), signal integrity, time-domain radar cross section (RCS), time-domain reflection/transmission (TDR/TDT), lightning strike and ground penetration radar (GPR). The HFSS transient solver is most beneficial for applications that require time-domain intuition and field visualization. It can also perform frequency-domain analyses such as S-parameters and frequency-domain far fields.

There is a growing interest in studying the electrostatic discharge (ESD) on touch screens of mobile handheld devices. ESD is often considered the top reason for post-shipment failure of solid-state electronics. This phenomenon is inherently transient and well suited for analysis by a transient EM solver for field visualization. By simulating field strengths on sensor pads, engineers can determine if thin-film oxide will potentially be damaged by dielectric breakdown at hot spots. Because of the miniature scale of the structures and extremely short duration of discharge, it is difficult to obtain reliable prediction through measurement.

In addition to analyzing structures at millimeter scales, the HFSS transient solver can solve such large-scale problems as the spilled fields induced by switching extra/ultra-high voltage (EHV/UHV) bus charge currents in a power substation. The transient spilled fields pose a potential safety hazard to personnel and equipment and cannot be predicted by steady-state

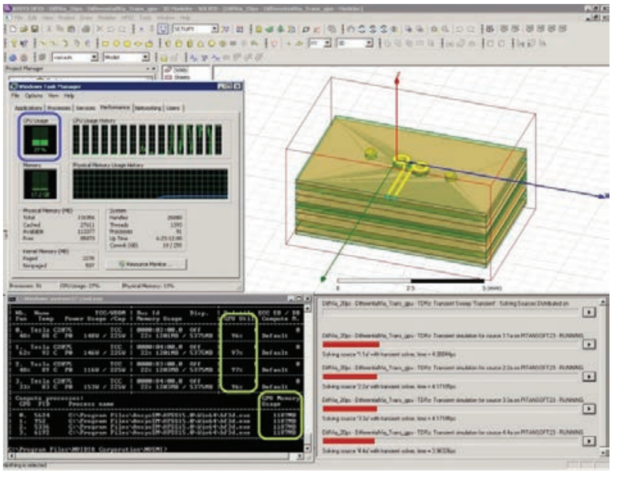

▲ Network analysis of differential stacked vias on four GPUs. Simulation jobs are run in parallel as shown by the progress bars and NVIDIA System Management Interface (nvidia-smi). The model was discretized into 68,401 tetrahedrons.

Transient EM analysis is indispensable for modern electronic design.

Expanded GPU Support Across Multiple Physics

To leverage cutting-edge hardware and deliver faster engineering simulation technology to users, ANSYS has teamed up with NVIDIA® to develop and release a GPU-accelerated computational fluid dynamics (CFD) solver. The result of a multi-year strategic partnership, this new solver addresses customer demand for increased speed and the ability to handle larger, more complex CFD simulation models. Available in ANSYS 15.0, the solver — and a new HPC licensing that enables all HPC users to take full advantage of GPU technology — broadens support for GPU acceleration within the ANSYS portfolio. GPUs can now speed up fluids, structural and electromagnetic simulations to increase the value of ANSYS HPC capabilities.

— **Wim Slagter,** Lead Product Manager, ANSYS

ACCELERATING ANSYS FLUENT 15.0 USING NVIDIA GPUs
ansys.com/83GPU

electromagnetic analysis. It is therefore highly advantageous to simulate different options of bus layout and equipment placement during the early stage of plant planning. By extracting the resistance, inductance, capacitance and conductance (RLCG) equivalent circuits of buses and switches through ANSYS Q3D Extractor and performing broadband SPICE circuit simulation using ANSYS Designer, engineers can obtain time-varying voltage sources for an HFSS transient simulation.

GPU ACCELERATION

The ANSYS HFSS transient solver's engine is based on the discontinuous Galerkin time-domain (DGTD) method, which is especially compatible with parallel execution on general-purpose graphics processing units (GPUs) with thousands of cores. With advances in GPU acceleration, more than four teraflops of computing power can be achieved on a single GPU that consumes less than 240 watts of electric power. Release 15 of the ANSYS electromagnetics suite enables engineers to leverage NVIDIA® CUDA™ technology for GPU computing to accelerate the HFSS transient solver. The GPU-accelerated solver can typically achieve two-times speedup on one NVIDIA Tesla K20 versus eight cores of Intel® Xeon® X5675. As a general rule, problems that require intense computational effort tend to result in higher speedup factors. In benchmarks of 15 examples, a maximum speedup of 5.2 times was achieved. Moreover, the solver can detect cases in which GPUs may not provide speedup and automatically fall back to CPUs. In those cases, there is no significant performance decrease since CPUs using OpenMP multi-threading are employed.

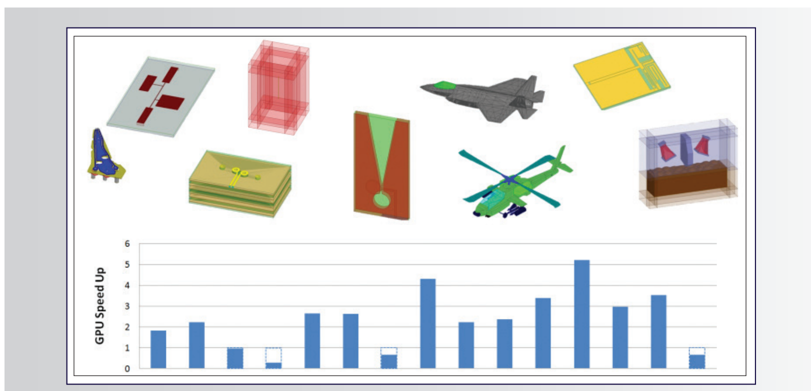

▲ ANSYS HFSS transient benchmarks on one NVIDIA Tesla K20 compared with eight cores of Intel Xeon X5675. Only nine of the 15 benchmark structures are shown in the figure. Dotted lines indicate cases in which the GPU did not provide speedup and CPUs were automatically used.

[redacted]

The speedup with GPU acceleration scales linearly with respect to the number of GPUs when the simulations are run with multiple HPC tasks for parametric sweeps or network analyses with multiple excitations. For example, if the transient analysis of a four-port network takes 40 minutes in serial on one GPU, the simulation time is reduced to around 10 minutes with four GPUs running in parallel. When solving four excitations consecutively, the speedup of one NVIDIA Tesla C2075 versus eight cores of Intel Xeon E5-2650 is 7.2 times. Therefore, the overall speedup is 28.9 times when all four NVIDIA Tesla C2075s are used. The assignment of multiple GPU jobs is fully automatic and requires no user intervention. The HFSS transient solver uses NVIDIA's exclusive process compute mode to assign one process to one GPU. Therefore, engineers will encounter no issues of load balancing or multiple HFSS processes competing for hardware resources on a single GPU.

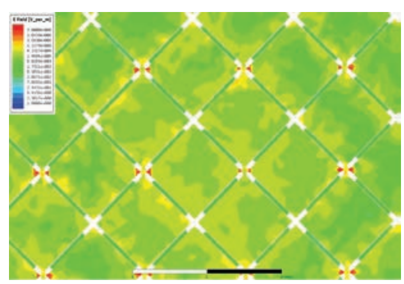

▲ Strong electric fields (red) on the bridges between sensor pads of the touch screen of a handheld device cause dielectric breakdown of indium-tin-oxide (ITO) thin films.

[redacted]

For example, simulation speedup occurs when performing power-surge analysis of a smartphone using the HFSS

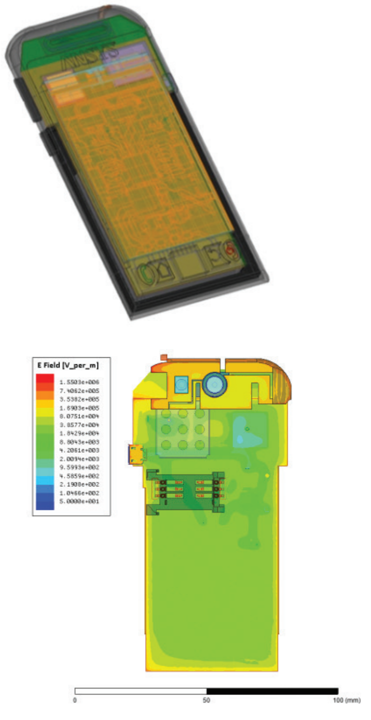

▲ A smartphone for transient field analysis on CPU, memory, GPS and Bluetooth ports due to power surge during battery charging. The model was discretized into 1,093,376 tetrahedrons.

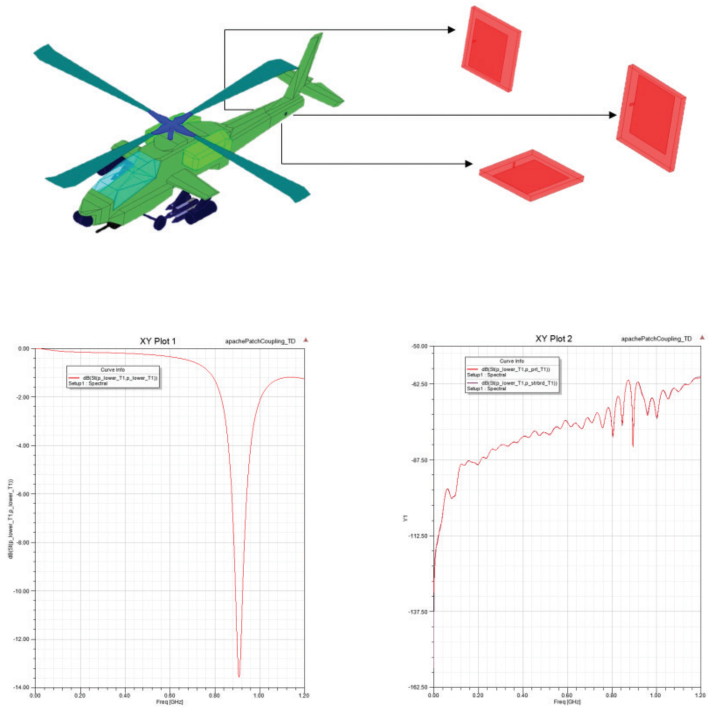

▲ Three patch antennas placed on the tail of a helicopter (top) and their S-parameters from DC to 1.2 GHz (S11 bottom left, S12 and S13 bottom right). The length, height and wingspan of the helicopter are 17.73 meters, 4.05 meters and 5.23 meters, respectively. The model was discretized into 549,247 tetrahedrons.

transient solver on a GPU. One NVIDIA Tesla K40, using 5.99 GB GPU RAM, speeds up the simulation 4.8 times when compared with eight cores of Intel Xeon E5-2687W. For signal integrity analysis of the printed circuit board, one NVIDIA Tesla C2075, using 3.98 GB GPU RAM, is twice as fast as eight cores of Intel E5-2650.

Moreover, when applying the transient solver to analyze the antenna coupling on a helicopter, the speedup of one NVIDIA Tesla K40 versus eight cores of Intel Xeon E5-2687W is 4.5 times, and the GPU RAM requirement is 4.35 GB. In all cases, the GPU solver requires less memory than its CPU counterpart. The frequency-domain S-parameters are calculated dynamically during the transient simulation. The simulation shows the resonant frequency of the patch antenna at 0.91 GHz and low mutual coupling (below –60 dB) between the antennas.

INSTALLATION AND SETUP

To access GPU acceleration, you must have NVIDIA GPUs and drivers installed on your computers and clusters. GPU acceleration in the ANSYS HFSS transient solver is officially supported with the Tesla and high-end Quadro series cards. For optimal performance, GPUs used for running simulation jobs should not be simultaneously used for visualization jobs. Only GPU cards with CUDA compute compatibility 2.0 and above should be used. To improve the speedup of transient field visualization, the GPU cards should be installed on a system with PCI-E 3.0 slots. A mixture of interface cards with lower PCI-E versions may result in the data not being transferred from GPU to CPU at the highest speed.

Before running HFSS transient simulations, it is important to ensure that GPUs are set with error correction code (ECC) disabled for performance, Tesla compute cluster (TCC) enabled for remote execution, and exclusive process enabled for GPU distributed computing. The GPU processor and memory usages are monitored through NVIDIA's utility program nvidia-smi. ▲

Thanks to Rickard Petersson and Matt Commens for their valuable input to this article. Application examples courtesy Jack Wu, Ally Liu and Sara Louie.

The ANSYS electromagnetics suite allows you to leverage NVIDIA CUDA technology for GPU computing to accelerate the HFSS transient solver.