HPC Options for HFSS

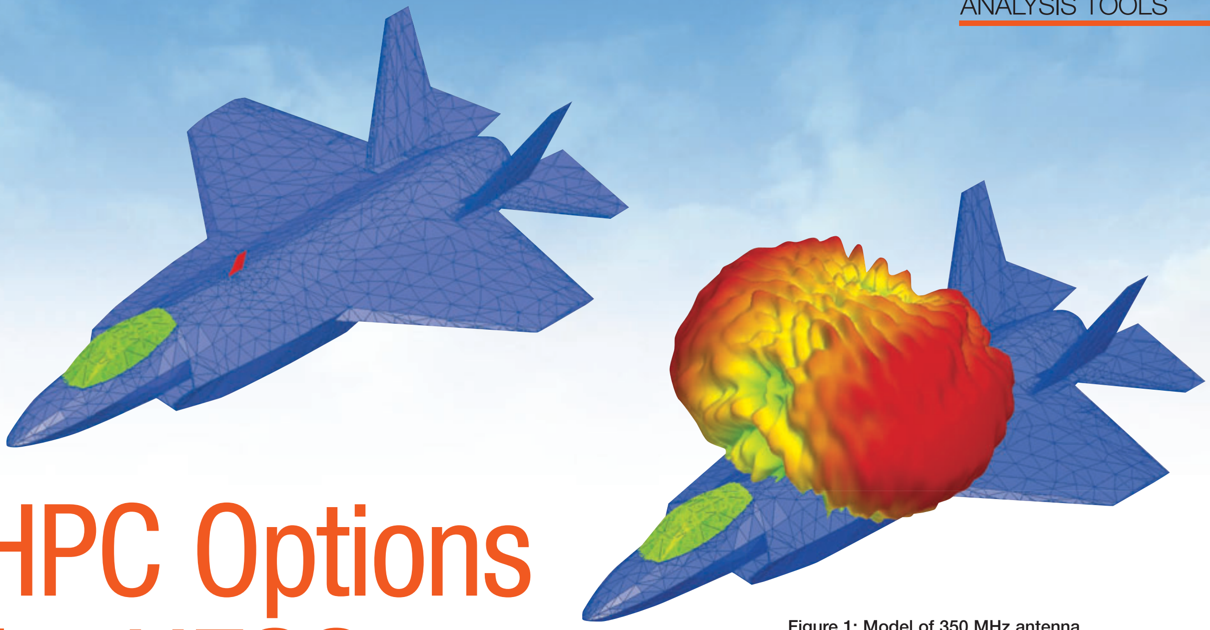

Figure 1: Model of 350 MHz antenna mounted on an F-35 Joint Strike Fighter (left) and resulting radiation pattern (right)

HFSS offers options for high-performance computing that deliver fast and efficient simulation for complex problems.

By John Silvestro, Senior Member of Technical Staff, and Matt Commens, Lead Product Manager for HFSS, ANSYS, Inc.

The latest HFSS release contains a wide array of powerful new features and enhancements. Included are some dramatic additions to the product's high-performance computing (HPC) capability. These improvements allow efficient simulation of large, complex problems and enhanced insight that would be difficult or impossible to obtain any other way.

Engineers today use multicore computers that are often connected to other machines via high-speed networks. In many cases, organizations have access to an enormous amount of computer power. To assist in leveraging these resources, HFSS offers three HPC solution options: multiprocessing (MP) and the domain decomposition method (DDM), both used for exploiting multiple cores for a single simulation, and the distributed solve option (DSO), which enables consideration of multiple design points using distributed processors. All three options are designed to make efficient use of multicore and networked processing power. This article illustrates how the options can be applied and the potential of each.

Multiprocessing

The MP option takes advantage of multiple cores on a single computer. The most time-consuming element of a simulation is the matrix solve. With the MP option, that part of the solution procedure can be run in parallel on multiple cores, reducing solution times.

For example, the F-35 shown in Figure 1 has a VHF antenna attached to its airframe. Engineers performed simulation at 350 MHz with eight adaptive passes, and the resulting patterns are shown in the simulation image. To illustrate how the MP option improves solution times, the team ran the final adaptive pass five times using a different setting for MP each time. The speed improvement in solution times is shown in Figure 2.

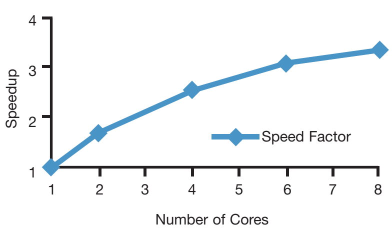

Figure 2: Solution speedup relative to single-core execution for antenna model using MP option

Other overhead, such as meshing operations, that are performed during a standard solve were not included in the data shown here, so the overall speedup factor is not as great as the times shown in Figure 2. Improvement in solve times is highly problem dependent, but there is potential for considerable improvement in solve time when using the MP option.

Domain Decomposition Method

The most exciting new feature in HFSS is DDM, which allows the model to be split into smaller domains. The solution for each domain is executed on a separate core and/or node and can be distributed over a network. Once simulations are complete, an iterative procedure combines the separate results into a single solution that gives the complete response for the entire model.

Here's how DDM works: When the solution region is broken into smaller pieces in a DDM simulation, each piece is simulated on a separate core

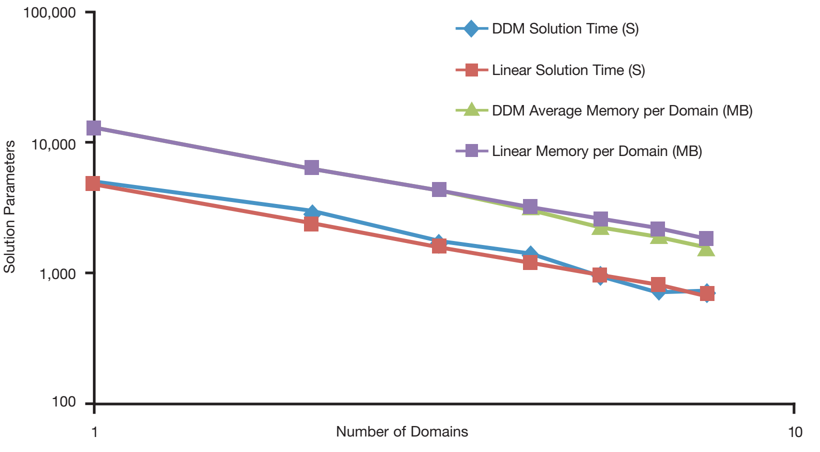

Figure 3: DDM results for F-35 aircraft model running a single pass on a well converged mesh; plot is log–log scale

or set of shared cores. These separate cores can reside on a local machine or be shared across a network. By using a distributed solution, the memory available for a single simulation is no longer limited to a single machine, and thus the size of models that can be simulated grows dramatically. In addition, speed and memory improvement may be realized. Solution times and memory for the direct solver in HFSS grow roughly as $n^{1.4}$ (where n is the number of unknowns). Thus on an eight-core machine with one core reserved for the head node, the domain solution time can speed up by as much as $7^{1.4}$ (or 15) times relative to the solve time of the entire mesh. There is, of course, additional overhead involved in the DDM calculation such as mesh partitioning and the iterative solution process. But, there is potential to realize better-than-linear speed and/or memory improvement when using the DDM.

Using the same F-35 model shown in Figure 1, the engineering team studied the solution speed and memory required to solve the final pass of this simulation using different numbers of domains. For this case, the team used two quad-core machines. Because one core is always a master process, this model can be simulated using from one to seven domains. The solution speed and memory improvement are shown in Figure 3; it also demonstrates time and memory as scaled linearly. The results show a close-to-linear improvement in solution time and a better-than-linear improvement in memory use for the solutions using five or more domains.

Distributed Solve Option

The distributed solve option allows simultaneous consideration of multiple discrete design points by distributing the individual solutions over available processors. This is a useful feature when an engineer is solving a parametric or frequency sweep. Industry-standard schedulers can be used to distribute the simulations, including Platform LSF® from Platform Computing and the Windows® HPC Server 2008 Job Scheduler from Microsoft. If sufficient hardware is available, DSO and MP may be used together for further performance improvement.

An example of the DSO option is the signal integrity (SI) model shown in Figure 4. This model is a set of coupled microstrip lines that transition to a set of coupled striplines through a differential via pair. To characterize the behavior of this transition, a designer typically needs to know its behavior over a variety of model variations — a process that can be time consuming if run serially. The engineering team created a parametric sweep with 32 variations of the model dimensions. The solution for each variation required a full adaptive simulation and an interpolating frequency sweep. The results from this analysis can be utilized for design optimization or yield studies.

The DSO simulation was demonstrated by ANSYS partner Hewlett-Packard® on an HP® ProLiant™ BL 465c series server. These machines use six core processors, and Platform LSF was used to schedule the simulations. Data was compiled for simulations run with one, two, four, eight, 16 and 32 cores. Figure 5 shows that significant improvement in performance can be realized with DSO. For example, with eight cores, the speed improvement was approximately seven times.

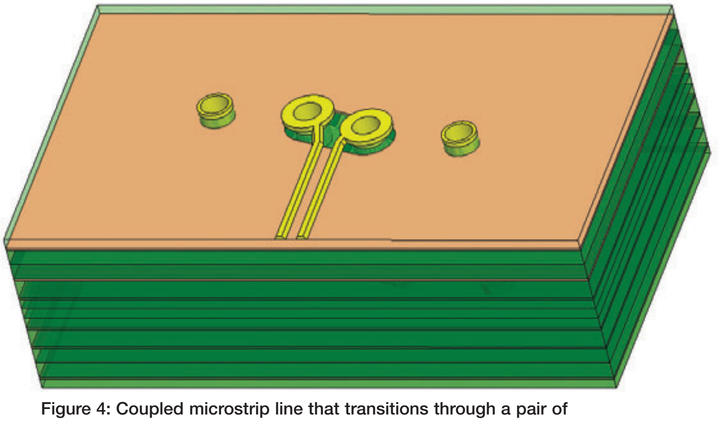

Figure 4: Coupled microstrip line that transitions through a pair of vias to a coupled stripline on a lower layer

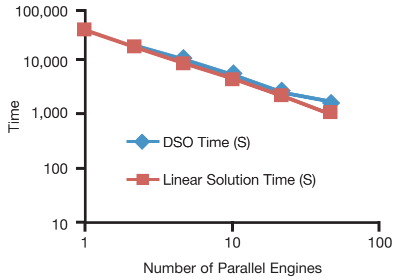

Figure 5: Log–log plot of the solution time for a parametric sweep of the model shown in Figure 4 when solved using DSO
Data courtesy Hewlett-Packard.

As this is a parametric sweep with a new geometry for each variation, the solution times for each variation will likely differ due to overhead, but the resulting speed improvements are exceptional.

HFSS software has three very powerful capabilities that take advantage of multicore and/or networked computers. These new and enhanced features allow organizations to leverage multicore/multimachine environments and compute clusters. This ability to simulate large, time-consuming problems in a highly efficient manner allows for further, higher-fidelity insight into a company's design. ■

Acknowledgment
The authors acknowledge the contribution of Donald Mize at Hewlett-Packard, who performed the DSO simulations reported here.

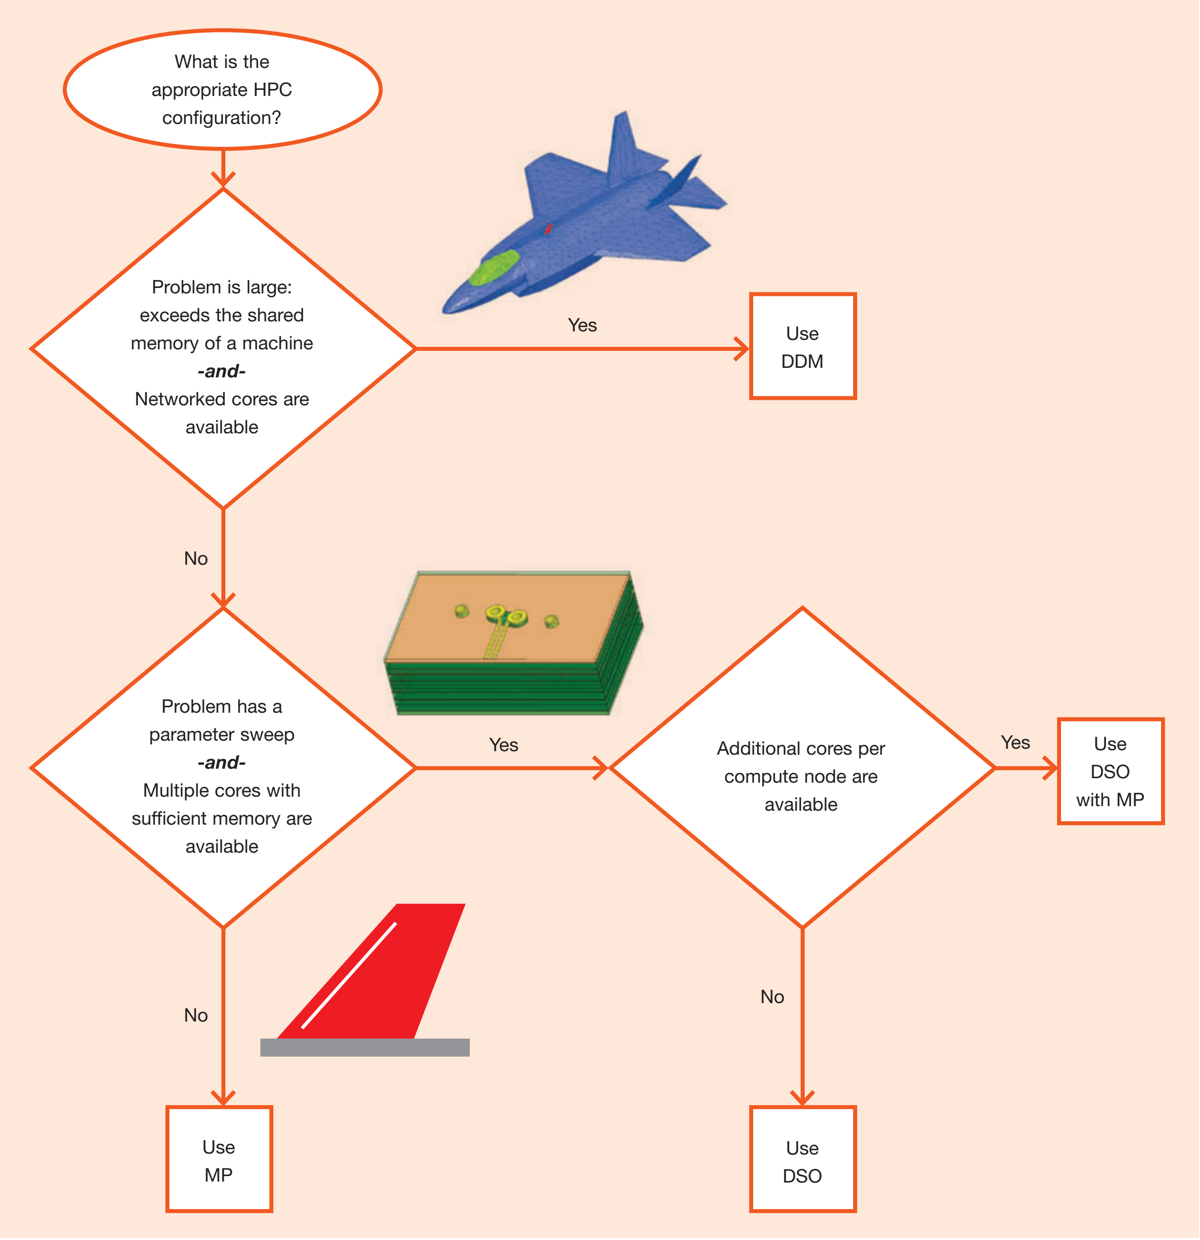

Figure 6: HPC decision tree