A WHITE PAPER FROM ANSYS, INC.

Smart Strategies for Large Structural Simulations

One way to gauge the progress of engineering simulation software is through the lens of size. Engineers today routinely run structural simulations with a few million degrees of freedom, and the largest commercial structural simulations have topped even 100 million degrees of freedom. These simulations would have been astonishing just a decade ago.

So what's changed? For one thing, we now live in a world of brute-force computing. More engineers have day-to-day access to multicore engineering workstations and parallel computing clusters. For another, engineers have become increasingly comfortable with large, complex simulations — in part because modern FEA software can automatically generate finely detailed meshes. As they work with larger and larger assemblies consisting of ever finer meshes, today's engineers are giving an entirely new meaning to the idea of a "plus-size model."

There's also a dimension to simulation size that goes beyond the number of elements and degrees of freedom. Add thermal, modal or harmonic transient runs to the mix (which is increasingly the case), and simulations can require a level of computing muscle that's out of proportion to the element count.

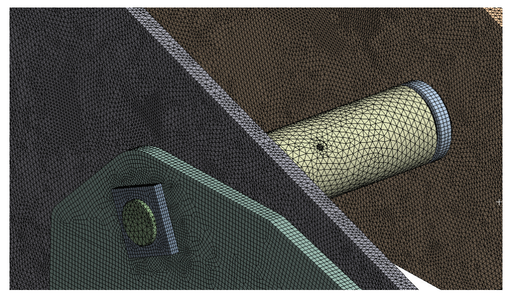

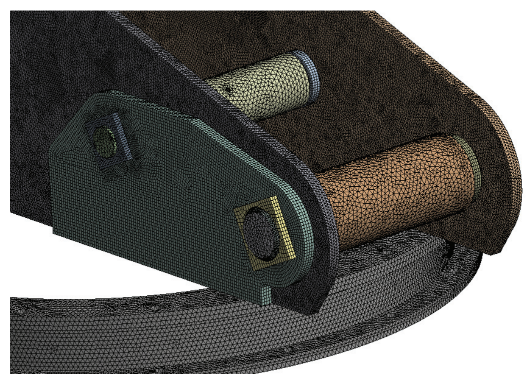

Increased accuracy of the solution requires finer mesh densities.

With computing power and engineering demands continuing to grow, structural simulations will likely get even larger in coming years. Consider what has happened in the world of computational fluid dynamics (CFD), which tends to foreshadow the size trends for structural analysis. Two years ago, a 1 billion-cell CFD problem was solved on commercial computing cluster for the first time. Also, newer hardware techniques such as GPU-based solutions provide further performance for solving structural models.

Structural analysis may not be all that far behind, which is why it's more important than ever to understand the best strategies for solving large simulations efficiently. For many engineers, the first and only impulse will be to throw large problems onto the biggest computer cluster they can get their hands on. Yet even when it's readily available, high-performance computing represents just one of several ways to tackle large problems. But in fact, it's not always the best way. When dealing with parametric changes, for example, an engineer can use reduction techniques on the parts that do not change, as there is no need to entirely recompute them. Software-based accelerators (such as variational technology) provide extra speed, and the only cost to the user is selecting the right option. Here's a closer look at when clusters make the most sense and when alternative techniques — such as problem reduction and software optimization — should be favored.

PARALLEL COMPUTING

When confronted with large simulation problems, engineers want to respond with computing force. And this response has become increasingly viable, given the widespread availability of high-performance computing options. Nowadays, most engineers need to look no further than their own desktops for the closest multicore computing environment. More and more engineering organizations, not just the largest ones, have access to powerful parallel computing clusters.

However, just because parallel computing is readily available, it should not be automatically applied to every large simulation. For example, engineers should consider the scope of design iterations when deciding whether a problem needs to run on multiple cores. Parallel computing tends to be best suited to "single-shot" simulations — those in which the design iterations are so sweeping that each pass through the software needs an entirely new setup and computation. In applications in which the design variations are numerous and more limited in scope, parallel computing's inherent speed advantages will be less noticeable.

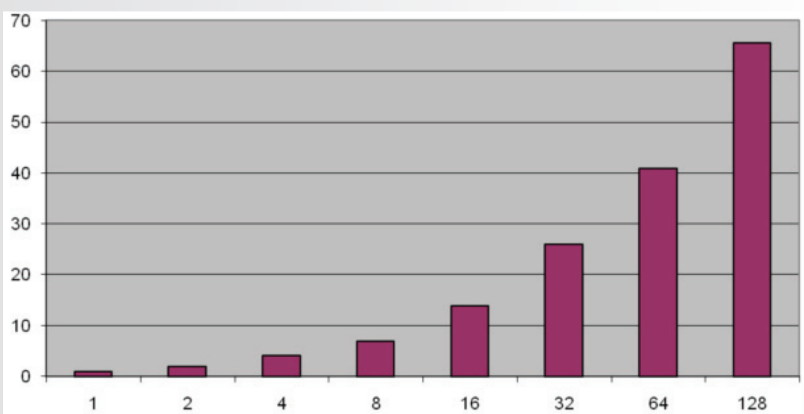

Number of Cores

Solves speedup with increasing number of computational cores

The notion of acceptable accuracy comes into play too. The more stringent a simulation's accuracy requirements, the finer its mesh will be. And the finer the mesh, the more likely the model will have too many degrees of freedom to solve quickly on a single core. Engineers wanting to limit the size of their models might consider lowering accuracy requirements when possible.

There are applications in which the meshes must be fine and every bit of accuracy counts. In these cases, models may grow so large that parallel computing truly is the best way to solve them in a reasonable amount of time.

A side benefit of parallel computing is that it requires less setup work than other methods of dealing with large models. With parallel resources, engineers do not need to worry about creating a mesh that limits degrees of freedom or to look for other ways to condense models. Users can simply click on the defaults for a fine mesh and let the computer do the heavy lifting. Furthermore, the difficult technical aspects of parallel computing have been made largely transparent to users of modern simulation software. For example, most software automatically handles the load balancing among the multiple processors.

One aspect of parallel computing for consideration involves the type of hardware architecture best suited to FEA problems. Multicore systems can employ shared memory processing (SMP), in which CPUs share blocks of memory. Or they can employ distributed processing, in which each CPU has its own dedicated block of memory. Today's FEA software runs on either architecture, but the distributed processing route is much more efficient. Tests show that simulations run on a two-core distributed system are as fast as simulations run on a four-core SMP system.

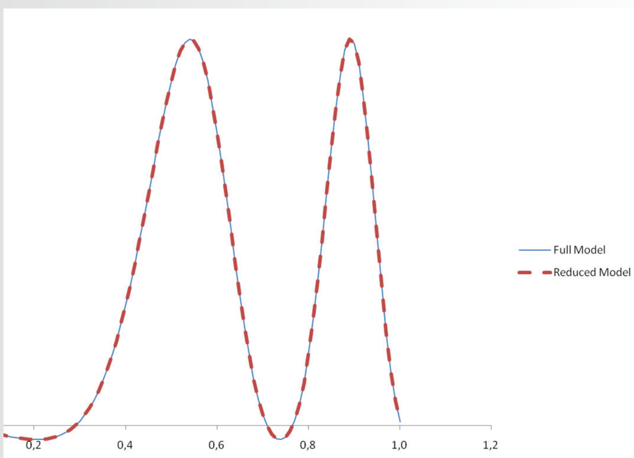

Reduced-order models lead to similar accuracy as full models, but in a shorter amount of time.

REDUCTION TECHNIQUES

Reduction techniques do exactly what their name suggests. With a little extra preparation on the user's part, these techniques can condense models significantly, making them far easier to compute efficiently, even in single-core environments. Model size, in terms of nodes and degrees of freedom, can often be reduced by up to 90 percent compared to the full model. At the same time, properly applied reduction techniques can maintain nearly all the accuracy of a full simulation. Usually, the accuracy penalty is less than 2 percent.

Submodeling and *substructuring* are two reduction techniques that have the broadest applicability, and both techniques have efficiency advantages even in this age of brute-force computing.

Submodeling reduces model size by subjecting only critical areas of the model to a detailed analysis. These areas, which form the submodels, are typically those with highest stresses. The noncritical remainder of the model is meshed more roughly, reducing the model size accordingly. To apply submodeling, a user first computes the deformations on the entire model using a mesh just dense enough to accurately capture global stiffness but not so fine as to predict localized stresses. This initial pass identifies the critical submodels that will be more finely meshed for a detailed analysis of stresses.

The computing improvements associated with submodeling depend on the ratio of critical to noncritical areas, but reductions of 50 percent are common, making submodeling a good choice when a user has limited computing power at his or her disposal.

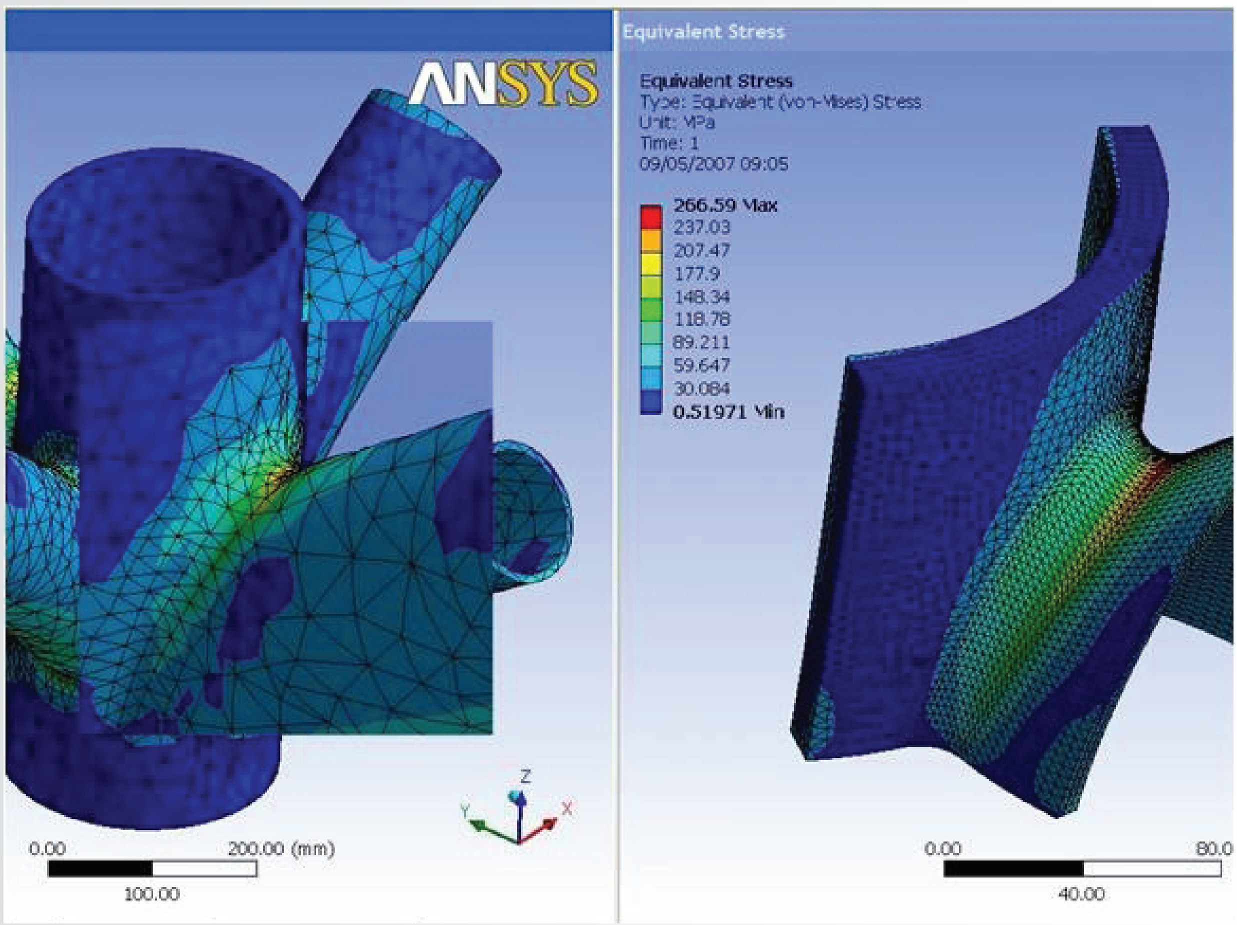

Submodeling is the right solution when only a portion of the model is critical. Start with a coarse (entire) model to capture deformations, then focus on the portions of interest in the model for fast and accurate results.

Even when computing power isn't an issue, submodeling can be a good choice. It comes into its own when a user wants to evaluate a series of circumscribed design iterations that have little effect on the global stiffness of the part. An example is optimizing fillet geometry to decrease local stress concentrations.

The most advanced simulation software can automatically extract the boundary conditions for submodels from the deformations computed for the global model. Still, submodeling does require some extra work because the user has to perform that upfront deformation prediction before getting to the detailed stress analysis. For this reason, submodeling is not typically the best choice for large simulations that can be run in a single pass.

Substructuring, like submodeling, condenses large models and allows them to be solved with limited computing resources. The idea behind this technique is to identify recurring geometric patterns that consist of many elements. These patterns, or substructures, are condensed into superelements that can be computed as a single element; they also can stand in for multiple instances of a recurring pattern. To apply substructuring to a four-legged table, for example, the user can create one superelement that represents a leg and then use that leg superelement four times. Condensing elements in this manner can drastically reduce the degrees of freedom, making models much faster to compute. Typical savings in computation time are on the order of 80 percent to 90 percent for long-transient or frequency-dependent runs.

Of course, most engineering models are not as simple as a four-legged table, and there are advanced forms of substructuring that build on its basic premise of element reduction. The most useful is known as **component mode synthesis (CMS).**

CMS represents large models as a set of interconnected components. Each one acts as a superelement that aggregates many individual elements from the initial model, thereby reducing the degrees of freedom.

Far from providing a skeletal approximation, CMS can actually capture the dynamic behavior of the full model with acceptable levels of accuracy. CMS simulations expand the results from the reduced-model simulation back to the full model before post-processing. So users ultimately evaluate all the CMS results in the context of the full model, not the simplified components. To make this expansion of the reduced simulation results possible, engineering software relies on sophisticated algorithms that can reconstruct the dynamic behavior of the full system from the behavior of the individual components.

Useful for both linear and nonlinear analysis, CMS excels in some applications but should be avoided in others. It is, for example, a particularly good choice for simulations that require long transient runs, especially those that would be difficult to compute when using the full model. Harmonic computations with a large number of frequencies are a good example. CMS also makes it easier to quickly vary the design of components in large assemblies. Rather than taking the time to solve the full model with each design iteration, a user can evaluate each iteration very quickly using the reduced model.

Because CMS requires extra simulation passes to generate the reduced model and expand result back to the full model, it is not the best fit for a one-off static simulation. In some cases, the generation pass alone might take longer than simply solving the full model in a multicore environment. Furthermore, CMS can trade off some accuracy as it creates its summary of the full model. With proper setup of the reduced model, the loss of accuracy can be limited to 2 percent or less. That's not enough to matter in some applications, but can be a deal breaker in others.

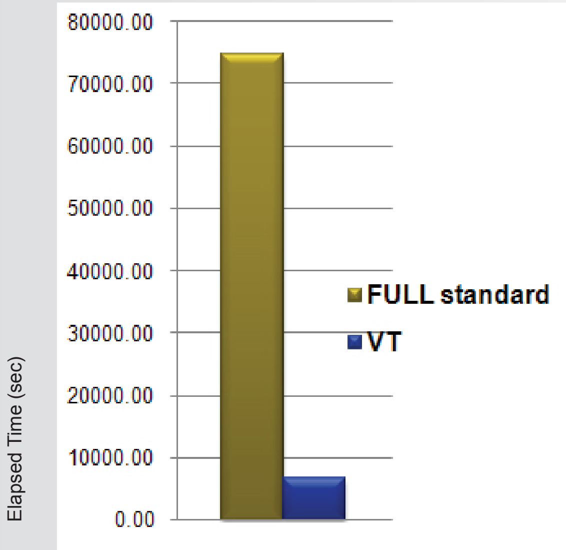

Speedup achieved by variational technology applied to cyclic modal analysis (80 sectors)

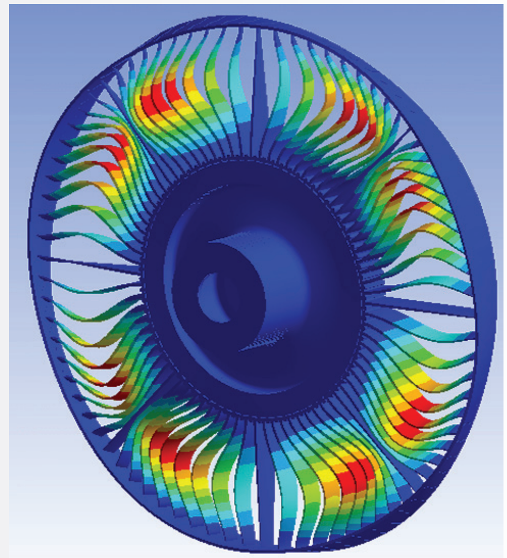

Cyclic symmetry functionality allows engineers to compute patterned geometries using only a sector of the model, leading to faster results.

SOFTWARE OPTIMIZATION

Simulation software increasingly contains optimization algorithms designed specifically to speed up large simulations. These algorithms tend to be proprietary, and they vary from vendor to vendor, so it's worth asking about them when selecting FEA software.

For example, ANSYS has introduced a technology that can dramatically speed up many types of transient simulations as well as simulations in which there are parametric variations of an initial design.

Called variational technology (VT), this capability leverages the results from each time step or design iteration to speed up subsequent computations. Every time step or design iteration will be computed faster than the previous ones. The time savings add up, with VT typically speeding up simulations by five to 10 times. The more time steps or iterations involved in the simulation, the greater the computational savings. In the appropriate applications, running simulations with VT on a single core can be faster than running it without VT on multiple cores. Combining VT with multicore processing can lead to even faster results. The only trade-off involves a slight loss of accuracy — usually about 1 to 2 percent.

THE BEST LARGE SIMULATION STRATEGY

When considering any of the large simulation strategies discussed here, it's important to realize that they are not mutually exclusive. Simulations that run quickly in a multicore environment can run even more quickly if reduced and run in a multicore environment. The same goes for VT and parallel computing.

The trick to solving large problems with the most efficiency, then, is not necessarily picking one way to solve every problem but picking the best way to solve the problem at hand. Rather than assuming the brute-force route will be the fastest in every case, an engineer should consider the nature of the large problem carefully.

Problems with many small-scale design iterations or transient runs may derive more benefit from reduction techniques or variational technology than they will derive from a few extra processors. Other problems may need all the computing muscle they can get. Still other problems can benefit from a combination of large simulation strategies. Some FEA software, however, does not support this multi-faceted approach to large problems. So as always, it's important for engineers to choose software tools wisely, as not all simulation tools are equal.

To help ensure that your product designs succeed in the market, engineering and design teams must accurately predict how complex products will behave in a real-world environment. The best simulation software spans the design continuum to fuel open communication between diverse engineering teams — from electrical and mechanical to thermal and fluid dynamics. Software that incorporates deep and broad multiphysics tools built on a single, integrated platform helps to ensure that no detail is missed, no potential risk factor is overlooked, and no product is released before it is ready.

For more information, visit www.ansys.com/Products/Simulation+Technology/Structural+Mechanics

www.ansys.com

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

Toll Free U.S.A./Canada:
1.866.267.9724
Toll Free Mexico:
001.866.267.9724
Europe:
44.870.010.4456
eu.sales@ansys.com

ABOUT ANSYS, INC. SOLUTIONS

ANSYS designs, develops, markets and globally supports engineering simulation solutions that are used to evaluate and support the product development process. Its integrated, open portfolio of tools is flexible enough to fit into any customer's product development process yet powerful enough to drive it. The ANSYS Simulation Driven Product Development vision is to enable customers to optimize designs throughout the product development process, especially in the early stages when changes can be efficiently and cost-effectively implemented. The process can help a business reduce development time, prototype testing and time-to-market — and ultimately help win the race in product innovation. The solutions ANSYS provides in the areas of structural, fluids, chemical, electromagnetic and coupled simulation are adaptable to customer-specific needs.

ANSYS tools provide customers strategic advantages by:

- Offering time- and cost-saving alternatives to expensive prototype development and experimental efforts
- Providing the opportunity to examine and optimize more design alternatives in the product definition and design stages of development
- Reducing product development time
- Enabling the customer to use simulation in order to optimize and support a product through its entire lifecycle

As one of the fastest-growing simulation companies worldwide, ANSYS remains dedicated to offering best-in-class simulation solutions that repeatedly demonstrate success.

Visit www.ansys.com for more information.