Thank you!

Thank you for downloading one of our ANSYS whitepapers – we hope you enjoy it.

Have questions? Need more information?

Please don't hesitate to contact us! We have plenty more where this came from.

Sales Support Training Emag FEA CFD

Ozen Engineering Inc. – ANSYS Channel Partner & Distributor

1210 East Arques Ave. #207, Sunnyvale, CA 94085

Telephone: (408) 732-4665

E-mail: info@ozeninc.com

Web: www.ozeninc.com

Channel Partner & Distributor

White Paper

16X Speedup in ANSYS Maxwell DSO on 32-Core High-Performance Compute Farm Doubles Traction Motor Design Productivity at General Motors

By Dr. Bradley Smith
HPC Architect
General Motors Corporation

Optimizing the design of traction motors used to drive electric vehicles (EVs) and hybrid electric vehicles (HEVs) is challenging because automakers have little experience in this area. Electromagnetic simulation plays a critical role by evaluating the performance of design concepts, such as by computing the torque profile of the machine. In the past, it took hundreds of hours to perform electromagnetic simulation on a single design iteration, which hurt productivity by leaving engineers waiting for results. General Motors (GM) developed a high-performance computing (HPC) environment using surplus hardware with a 32-core compute farm that provided a 16 timesx speedup on electromagnetic simulation, reducing solution time to less than five hours.

The company estimates that the HPC system doubled the productivity of engineers working on traction motor design. In addition, it provided the ability to evaluate more design alternatives, reduced time to market, reduced manufacturing costs, and improved the quality of traction motors. The team has since made an order-of-magnitude further improvement by implementing an enterprise-scale HPC environment, whose details are confidential. This paper presents information on the predecessor HPC platform that served as a stepping stone to the current approach.

"The GM team achieved major gains with a stepping-stone system using low-cost castaway hardware. For a small organization, this can be used to implement a no-hardware-cost entry into the HPC supercomputing arena."

Traction Motor Design Challenge

HEVs, such as the Chevy Volt, and EVs have been rapidly gaining in popularity because they provide large gains in fuel efficiency and near-zero emissions at increasingly affordable prices. Electric motors provide highly efficient conversion of the energy contained in their power sources — batteries, ultra-capacitors and fuel cells — into tractive power for the wheels. The traction motor is a special-purpose electric motor that propels the HEV or EV. It takes the electrical energy given to it and converts it into motion. Likewise, during braking, it takes motion and convert it into electrical energy. Thus, it is responsible for bidirectional energy conversion. This impacts available power and efficiency of the vehicle. Minimizing electromagnetic losses is an important factor in maximizing the efficiency of the motor.

Engineers must consider a wide range of design alternatives to optimize these and other aspects of traction motor design. Possible motor architectures to consider are interior permanent-magnet (IPM) synchronous machines, switched-reluctance machines (SRMs) and induction machines (IMs). In most cases, IPM motors are used; then the question arises of where to position the magnets and how many magnets are needed per pole. Engineers must decide whether the stator should have solid or stranded

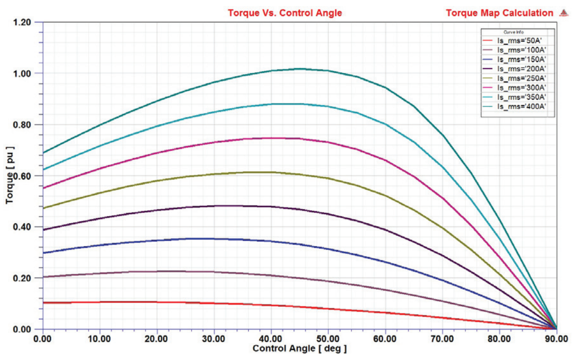

Figure 1: Torque plot provides typical example of simulation results.

winding design. They must determine whether the windings are concentrated or distributed. Making these and dozens of other design decisions required to optimize the design of a traction motor often requires evaluating thousands or even hundreds of thousands of design iterations.

The engineers responsible for the motor design first work on the electromagnetics of the electrical machine by defining the main elements of the motor, including magnet materials, coil configurations, number of turns and air gaps. When designing a new internal combustion engine, the team can rely upon experience to determine the bore and stroke of the engine. The automotive industry has much less experience with electric traction motors than with internal combustion engines. There is almost no historic knowledge in traction motor design. The engineering team cannot take an educated guess, then build and test the outcome. It is necessary to use design techniques based on simulation, with a single build for confirmation.

Critical Role of Simulation

As with many HEV and EV OEMs, GM uses ANSYS® Maxwell® electromagnetic field simulation software to compute the motor's torque profile, how the torque ramps up over time in motor mode, and electrical resistance in stopping the vehicle in regenerative brake mode. GM also applies ANSYS HFSS™ and ANSYS Q3D Extractor® as part of the electromagnetic design process, to help address any compatibility issues that may arise. Based on the simulation results, the team modifies the design by changing basic design parameters to balance traction motor system performance against its size, weight, cost, and interactions with other systems.

Later in the development process, the computed torque profile is used in additional downstream analysis, such as for computing mechanical stresses, loads, deformations and vibrations of the powertrain's physical parts, including the driveshaft and gearing. Vibration analysis is important because traction motors can be a prominent source of noise in EVs. A fluid dynamics solver can help in studying thermal management issues, mapping energy losses, and determining heat distributions in the motor/generator assembly.

When GM first began work on traction motor design, the company ran electromagnetic simulation on a powerful single-core desktop computer, which took about 72 hours to solve each design iteration. As simulation complexity increased, analyses requiring these single simulations took in excess of 720 hours, placing a strain on the corporate infrastructure due to uptime requirements of a Microsoft® Windows®-based network. This reduced engineering productivity while the team waited for the crucial simulation results to determine which direction to take. Members of the GM IT team and the Traction Motor Design Team experimented by running the ANSYS electromagnetics distributed solve option (DSO), which enables Maxwell, HFSS and Q3D Extractor to distribute parametric studies across available hardware to expedite electromagnetic model extraction, characterization and optimization. The team used a computer with eight cores and a large disk and memory space. This platform speeded up the solution by a factor of 3.5, a considerable improvement but less than the target.

Up until the HPC corporate system was installed and functional, engineers employed multiple multicore desktop units to perform processing. As such, there are literally thousands of runs on multiple desktops, thousands of runs on the stepping-stone HPC, and, now, tens of thousands of runs on the large-scale corporate system.

Configuring a Stepping-Stone System

Once it became clear to the organization that a desktop solution was inadequate, payback of a large-scale HPC system was not proven. So, it became GM's goal to build an inexpensive stepping-stone system. This system is a great model for people who are new to HPC to get significant performance increases with little or no hardware expense.

HPC does not require the latest or fastest computing hardware to deliver high levels of performance. So to save money, the team collected outdated rag-tag Hewlett-Packard® desktop computers from the scrap bin as the heart of new HPC system. Microsoft Windows 2008 Server Edition HPC, a stripped-down version of Windows that is designed to consume fewer hardware resources, was used as the operating system. LSF by Platform Computer (now IBM) was used as the parallel task manager and queuing software. LSF is configurable and can comprehend differences in processing speeds, memory sizes, and number of cores of the scavenged desktops used in the stepping-stone system. The corporate FLEXLM server used to serve production applications was used to service the licenses. Licenses are pulled from one central location, which simplifies administration and maintenance. Due to the relatively low frequency of license pulls, a single quorum can serve both desktop and HPC license support at the same time with minimal impact on either environment. GM still leverages this approach on the production server.

The whole stepping-stone concept was intended to be low cost, supporting a small, tightly connected user community. For a global IT environment, this microscale distributed HPC is difficult to sustain; but for the first-time or small organization, it represents a low-budget, high-return undertaking.

For the stepping-stone approach, it was assumed that no more than two concurrent users would be interfacing with the HPC system at one time. Their jobs would continue after they logged off, and they would need to log on only to review results.

Two interface node computers, each supporting one user, comprised the external interface to the HPC. Additional client LSF licenses are required to perform submits from nodes other than the two interface nodes. Although not the case with the stepping-stone system, these two interface nodes can be used for desktop applications with direct access via the keyboard, or via Remote Desktop in Windows, or with one of the many Linux® applications. To properly calculate speedup, the interface nodes on the stepping-stone system were used solely for submission and retrieval of simulation results. In either situation (other tasks on interface nodes or dedicated), the interface nodes should not be used as compute nodes because they are busy with scheduling the HPC, user applications, antivirus, and data sharing with mounted systems in the outside world. Later on in the study, the team installed high-end graphics cards in the interface nodes to assist with graphical rendering. Doing rendering on the interface node computers saved a lot of strife. The parametric variation result files are immense and time/bandwidth consuming. So, keeping them internal to the HPC eliminated a number of bottlenecks and kept everyone more content.

Rather than having a client directly submit jobs to the HPC, files were loaded onto an HPC storage area and queued by LSF. There are a number of reasons for this approach. First, some client machines are laptops, which can be disconnected at any time. If this disconnection happens mid-simulation, all mapped links are lost, which can lock the simulation and eventually create a denial-of-service situation. Second, the amount of traffic required to generate the large output files can swamp the link from the HPC to the client. Third, post-processing data on the HPC would require retransmission of the intermediate data files from the client to the HPC.

The network was not connected to the outside world, except at the interface hosts, which are dual NIC card computers. This structure provides a firewall for the HPC environment. It also gives a separate and less congested infrastructure for the HPC environment and exposed external segments alike. Firewalling the HPC network also addresses security concerns while avoiding the need for antivirus software that consumes valuable CPU cycles. As the test environment was very new to GM, each node had full antivirus security on it, requiring a large overhead. Eliminating antivirus software avoided the need to reboot the system for virus file updates. Rebooting the server would make it necessary to stop jobs and have a brief “reset to 0” time that would reduce HPC throughput.

A mix of disk- and network-attached storage was utilized to run the operating system, LSF and Maxwell application software as well as to perform user home directory tasks and store solution files. All the disk space on the HPC was treated as temporary, because old data can cause disks to fill up and create a denial-of-service situation. Windows Active Directory allows users to mount a virtual drive to the active directory area, then read and write files as if they were connected to their machine. This avoids the need to send files via FTP to a queuing area, and connections to the spool area are maintained whether the user is logged on or not.

Each task ran under a user ID, so a common sign-on was required on all machines. Once a user logged in, all directories and files were identical for each user. The author once managed a small Windows HPC without LDAP. Changes had to be done independently on all machines, which took hours per upgrade. For a 16-host HPC without LDAP, each password update had to be done 16 times. If one machine was out of synch, jobs would fail.

16X Speedup in ANSYS Maxwell DSO on 32-Core High-Performance Compute Farm Doubles Traction Motor Design Productivity at General Motors

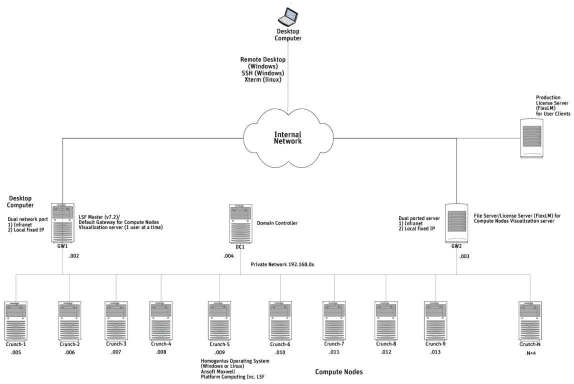

Figure 2: Stepping-stone system that provided a 16 times speedup and two times increase in engineering productivity

Controlling Software Expenses

Software is the primary expense of an HPC undertaking, but it can be minimized with careful planning. For a larger system, with N computers, each with M cores, N(M)-1 licenses of ANSYS Optimetrics™ software are needed. Optimetrics automates the design-optimization process for high-performance electronic devices by quickly identifying optimal values for design parameters that satisfy user-specified constraints. N operating systems and N LSF licenses are also required.

For the HPC, one must think in terms of both cores per box and boxes. For N cores, there should be approximately N queues. More queues per core can be used, but one is a good starting point. This means that for an N core system, you should have N DSO licenses (available in 10 packs at a reduced rate). For an M box HPC, you will need M LSF server licenses and one client license per submitting box.

In a typical environment with served licenses, a safe assumption is that 40 percent of users will be employing the tool at any one time. Engineers can and should learn to execute multiple jobs at a time trying to increase project bandwidth, but this is a situation that is grown into, rather than started out from. So for P users, there should be about 0.6 P licenses of Maxwell and Optimetrics. In general, HFSS and Q3D Extractor have a lower frequency of usage and can run at about the 0.3 P level. Maxwell 3-D requires larger memory and processing capacities than Maxwell 2-D. You should also use up to 1 HPC/MP license per box if you are using Maxwell 3-D. As an alternative, you can limit the size of your jobs to the number of CPUs that can be handled by your HPC/MP license.

The author utilized the configuration shown in Figure 2 with 16 compute nodes in addition to the LSF master gateway, domain controller and license server. The domain controller served as the file server and had 2 TB disk storage for all the user files. It also served as the server for the Maxwell files, so it was not necessary to install Maxwell independently on each box. The stepping-stone system required 32 DSO, six Maxwell 2-D, six Optimetrics and 16 LSF server licenses. Since users were limited to submitting from the two client computers, only two client licenses were required. If users had been allowed to submit from their PCs, one client LSF license would have been needed for each user.

Further improvements are on the way. When ANSYS Maxwell supports graphic processor unit (GPU)-computing by offloading highly parallel number-crunching algorithms from the central processor unit (CPU) cores onto GPUs, it will result in a substantial increase in speed. A GPU such as the NVIDIA® M2070 Tesla™ Accelerator offers immense floating point speed and array manipulation features, typically providing a two times speedup compared to an HPC system without GPUs.

16X Speedup in ANSYS Maxwell DSO on 32-Core High-Performance Compute Farm Doubles Traction Motor Design Productivity at General Motors

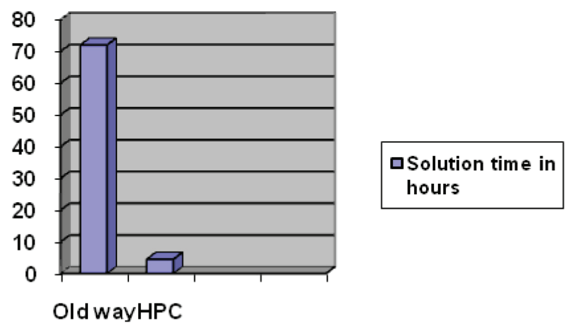

Performing simulations running on 16 computers with 32 cores reduced solution time by a factor of 16 to 4.5 hours. Let's assume that a fully trained engineer costs about $100,000 per year. To perform the same work as two engineers in the stepping-stone environment in a non-HPC environment requires at least four engineers. So the stepping-stone environment saves about $200,000 per year in personnel costs. The cost of the HPC software is roughly $30,000 per year. With software costs factored in, the stepping-stone system saved $170,000 per year with virtually zero upfront investment. The cost of the HPC system would have been slightly higher if the team had purchased new hardware rather than recycling old hardware, but the numbers would not have changed very much. A key value of this mini- or stepping-stone HPC is that it is scalable downward to a few boxes or upward to 20. The only hard costs are the software.

An interesting twist to the stepping stone system is to look at the nodes as individual computers attached to a small company LAN. LSF offers the capability to start jobs and end jobs at given times. So, for a company that works 12 hours a day normally, jobs can be queued during the day and executed on idle machines at night. This approach allows even desktop computers with adequate memory to be used 24 hours per day — 12 hours for daily work and 12 hours for HPC work, allowing a small company to perform world-class Maxwell parametric electromagnetic analysis for only the license cost. With some of the "fair-share" features intrinsic in LSF, jobs can be queued during the business day, when desktop hardware becomes available.

HPC combined with the latest generation of analysis tools has led to dramatic improvements in traction motor design productivity. The details of GM's current enterprise-scale HPC environment are confidential, but this article explains how the team was able to achieve major gains with a stepping-stone system using low-cost castaway hardware. For a small organization, this can be used to implement a no-hardware-cost entry into the HPC supercomputing arena.

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