

HPC Delivers a 3-D View

NVIDIA accelerates engineering 3-D glasses with a GPU and ANSYS HPC.

By Stan Posey, HPC Industry Market Development, NVIDIA, Santa Clara, U.S.A.



image © istockphoto.com/adventr

The semiconductor industry faces growing competitive challenges to reduce design cycle time and costs, decrease thermal loads and power consumption, and respond to customers who demand high-quality, well-designed products. Because of these drivers, NVIDIA Corporation, based in Santa Clara, U.S.A., deploys ANSYS software to predict phenomena such as large deflection and bending of printed circuit boards (PCBs), thermal creep of solder joints, PCB shock and vibration, and even mechanical performance of 3-D glasses. NVIDIA uses ANSYS technology for a range of applications (including structural, fluids and electromagnetic simulations) in a high-performance computing (HPC) environment that includes workstations and servers. Recently, NVIDIA examined how it exploited HPC in conjunction with the ANSYS Mechanical product; the team found that incremental changes in hardware configuration, including the use of graphics processing units (GPUs) to

accelerate the solver, could increase the number of simulations that could be performed within a given time frame.

Optimizing Performance

Recently, NVIDIA engineers needed to better understand deflection and bending in 3-D glasses, intending to ensure design robustness related to user handling in combination with desirable qualities such as comfort and fit. Engineers planned to apply Mechanical to examine alternative design ideas to sort out the design trade-offs. However, the engineering team found that running the complex models — incorporating nonlinear physics and nearly 1 million degrees of freedom (DOF) — took over 60 hours to complete a single simulation, leaving little or no time to study alternative ideas.

“It’s a virtual cycle: NVIDIA is an industry leader in the design of GPUs; we use ANSYS simulations running on current-generation GPUs to design the next generation of GPUs.”

— Andrew Cresci, General Manager of Strategic Alliances, NVIDIA



NVIDIA® 3D Vision™ glasses

“By optimizing our solver selection and workstation configuration, and including GPU acceleration, we’ve been able to dramatically reduce turnaround time — from over two days to just an hour. This enables the use of simulation to examine multiple design ideas and gain more value out of our investment in simulation.”

— Berhanu Zerayohannes, Senior Mechanical Engineer, NVIDIA

NVIDIA therefore set about optimizing the computing infrastructure and its use of Mechanical to see if it could reduce turnaround time — to pursue more design studies. First, the team considered the use of an alternate solver methodology, moving from the iterative PCG solver to the direct sparse solver, which seemed to have advantages for the nonlinear models in question. This alone generated significant speedup, but the sparse solver used more memory than was available on the workstations being used for simulation. NVIDIA upgraded to newer workstations with increased memory capacity and found that the simulation now executed in-core, using memory instead of slower I/O to manage data during the computation. Finally, NVIDIA added parallel processing in conjunction with acceleration by employing an NVIDIA® Quadro® 6000 GPU, which enhanced processing power. The end result was astounding — the simulations that had previously required 60 hours were completed in just 47 minutes. This 77-fold improvement in time to solution meant that NVIDIA could deploy simulation to help make design decisions to ensure robust performance of the 3-D glasses.

ANSYS Mechanical and GPU Computing

Contemporary CPUs pack multiple processing cores on the same chip to achieve speedup through the use of parallel processing in software like that developed by ANSYS. GPUs take this to an extreme, with hundreds of cores available for computational tasks. The devices function as massively parallel coprocessors to conventional x86 CPUs. NVIDIA, an industry leader in designing computational GPUs, works with ANSYS to enable GPU acceleration of engineering simulation software.

Starting in late 2009, ANSYS and NVIDIA began a technical collaboration to develop GPU acceleration for the ANSYS Mechanical solvers, essentially boosting overall performance by off-loading heavy matrix computations to the GPU. These operations are usually dominated by matrix–matrix multiply for the direct solver and sparse-vector multiply for the iterative solver, operations that have been highly optimized in CUDA™ libraries offered by NVIDIA and used by Mechanical. ANSYS Mechanical 13.0 offers GPU acceleration for shared-memory solvers over a broad range of modeling conditions. ANSYS Mechanical 14.0 will extend GPU acceleration to the distributed memory solvers.



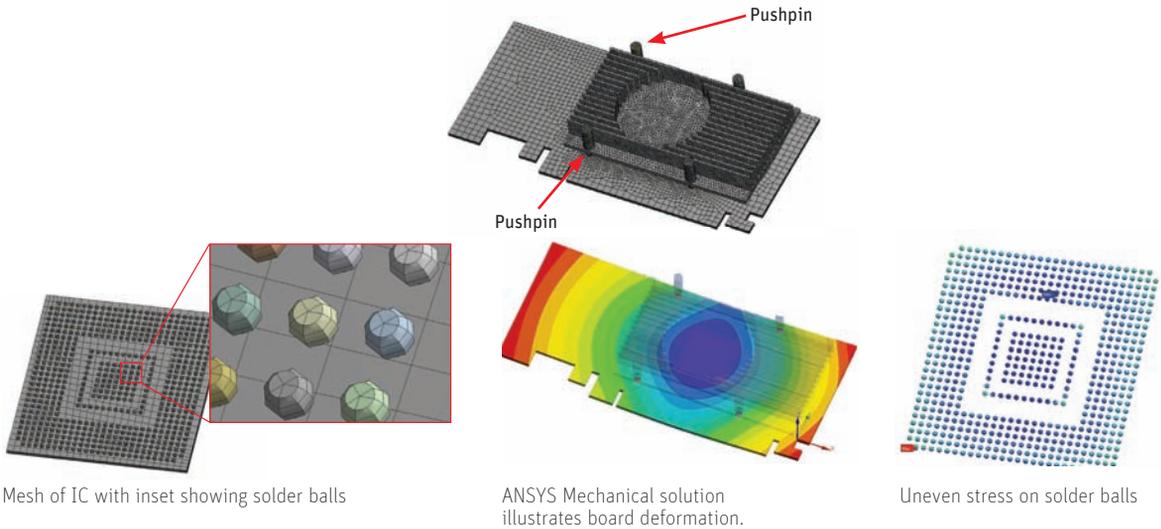
3-D glasses model



Deformation that occurs when worn by 95th percentile of males



Von Mises stresses



Mesh of IC with inset showing solder balls

ANSYS Mechanical solution illustrates board deformation.

Uneven stress on solder balls

Pushpins on the PCB hold the integrated circuit (IC) with a passive heat sink to the board. These exert forces, causing the board to bend. The uneven stress on the solder balls could cause failure if subjected to certain levels of shock and vibration. The NVIDIA team uses ANSYS Mechanical to determine the deflection of the board and stress on the balls to avoid failure.

Because of these HPC changes, the engineering team is exploring the use of parametric modeling for optimization of parameters related to comfort and fit for a range of subjects, and not just the 50th percentile. Based on initial turnaround time for a single simulation, more than 70 variations of the original model can be simulated in the same amount of time.

The results demonstrate that using GPUs to accelerate ANSYS Mechanical software delivers a valuable reduction in wall clock time compared to non-accelerated runs. NVIDIA found that it was equally important to ensure that the simulation first exhausted all the benefits from ANSYS, both

solver options and conventional CPU-only system resources. NVIDIA also examined the impact of using GPU acceleration on other Mechanical workloads, varying from 250K DOF to roughly 1M DOF, and observed performance benefits across many applications. ■

References

Beisheim, J. *Accelerate FEA Simulations with a GPU*, NAFEMS World Congress, Boston, Mass., U.S.A., May 23-26, 2011.

Beisheim, J. *Speed Up Simulations with a GPU*. *ANSYS Advantage*, 2010, V4, I2, web exclusive.

HP recommends Windows® 7.

HP Powers ANSYS
from Desktop to
Data Center
www.hp.com/go/solver

HP Z800 Workstation

SL6500 chassis with
8 x 1U SL390

Everybody On