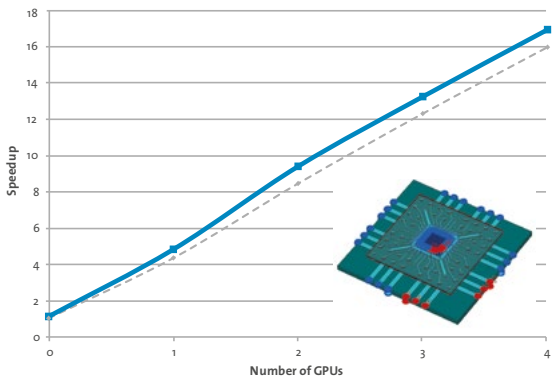
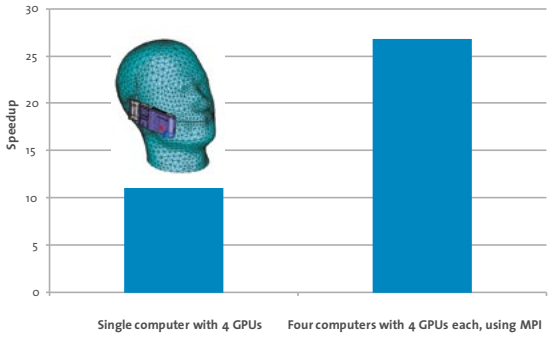
CST STUDIO SUITE

GPU Computing

*The high memory bandwidth and parallel processing abilities of GPU cards mean that GPU computing can provide significant simulation speed advantages over conventional CPU computing. CST STUDIO SUITE® supports GPU computing in the transient solver, integral equation solver and particle-in-cell (PIC) solver, using NVIDIA Tesla® GPU cards. The GPU Computing Guide, available on our website, gives a complete, up-to-date list of supported hardware. With the transient solver and integral equation solver in CST STUDIO SUITE, up to eight GPUs can be used per host system, with additional GPUs conferring further performance improvements. All GPU-capable solvers support the newest generation of Tesla cards, Kepler, which offers a performance improvement of around 70 % over the previous generation, Fermi.**



Speedup of solver loop: IC package (82,542,410 mesh cells), CST STUDIO SUITE 2014 blue, CST STUDIO SUITE 2013 grey.



Speedup of solver loop: mobile phone and phantom (>100,000,000 mesh cells).

Supported solvers

Transient solver

The transient solver is the flagship solver of CST STUDIO SUITE, and is a versatile full-wave electromagnetic solver. GPU computing is supported for both the FIT solver, which is the standard time domain solver, and the transient solver on the TLM mesh, which is well suited to wEMC problems.

Both solvers support the use of up to 8 GPU cards per machine using a domain decomposition approach. In CST STUDIO SUITE 2014 the transient solver uses CPU and GPU resources in parallel to provide even greater speed-up than either technique can provide alone.

"GPU computing has allowed us to perform some complex simulations that were previously impractical"

Matt Fuller, Selex ES

The FIT solver also supports the combination of GPU computing and MPI computing (cluster computing) to allow the simulation of very large models which can't be handled by a single computer. For more information, please see our HPC flyer.**

* Based on a CST comparison of a Tesla K40 to a Tesla C2075.

** Note that the benchmarks and speedup numbers shown in this document are always related to a high-end CPU configuration with best possible memory bandwidth (Dual Xeon E5-2630 v2, quad-channel RAM, memory clock: 1600 MHz), compared to an equivalent system including Tesla K40 GPU cards.

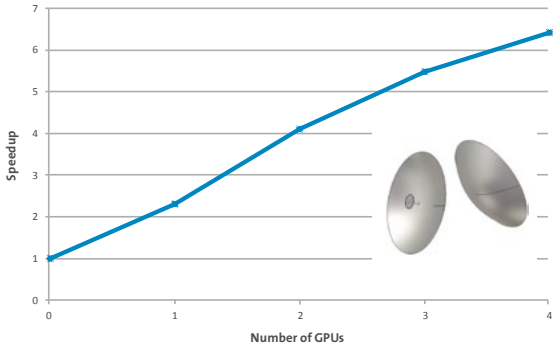
GPU Computing

"We at NVIDIA designed the Tesla family of GPU cards especially for scientific applications. All Tesla cards undergo stress testing with zero error tolerance before release, and they can also take advantage of driver features optimized for performance, making them the optimal choice for GPU computing."

Sumit Gupta, General Manager, Tesla GPU Business Unit, NVIDIA. Learn more at www.nvidia.com/object/why-choose-tesla.html

Integral equation solver

The integral equation solver is a solver especially suitable for electrically large objects, using surface meshing instead of volume meshing. The direct integral equation solver supports GPU computing, with up to 8 GPU cards per machine, and the efficient out-of-core algorithm allows for acceleration of arbitrary model size, as long as the host system has enough memory to hold the data.

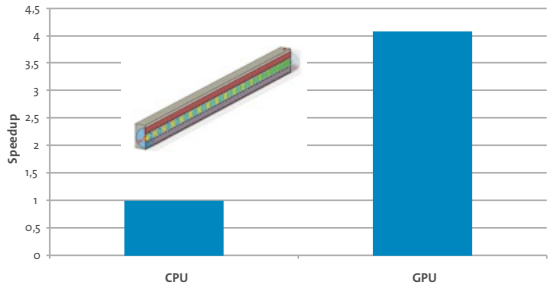


Speedup of matrix factorization: reflector antenna (52,313 surface mesh cells).

Particle-in-cell solver

The PIC solver is a self-consistent charged particle simulation tool. It calculates both the movement of the particles and the behavior of electromagnetic fields in the time domain, making it well suited to devices such as vacuum tubes and magnetrons.

The PIC solver supports one GPU per solver process at the moment, with development of a multi-GPU solver ongoing. New generation Tesla cards are especially beneficial to the PIC solver, since they offer excellent performance for the double-precision calculations needed to perform particle simulations.



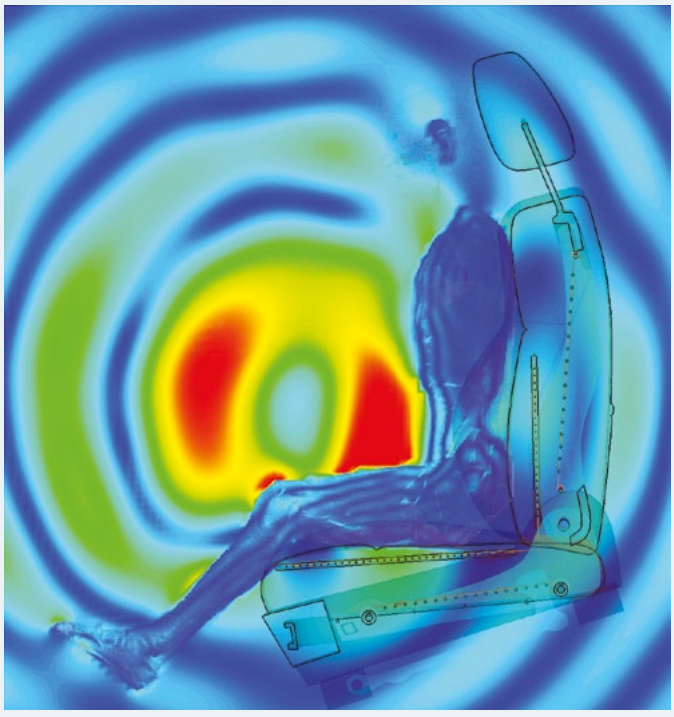
Speedup of solver loop: travelling wave tube (21,157,872 mesh cells).

Case study

The Chair of Electromagnetic Theory of Prof. Clemens at the Bergische Universität Wuppertal uses 16 NVIDIA Tesla M2090 GPUs with up to 8 cards per server for GPU acceleration and has recently added a new designated computer cluster with up to 100,000 CUDA cores (40 Tesla K20 GPU cards distributed on 8 compute nodes).

"In multiple heavy-duty research projects using the CST transient solver, HPC computing featuring multi-GPU acceleration provided us with a powerful tool to remarkably reduce overall simulation time."

Carsten Cimala, Bergische Universität Wuppertal



GPU accelerated simulation of a human model in a car seat.
Model provided by Bergische Universität Wuppertal.