High Performance Computing Techniques for Efficient 3D Full-Wave Simulation of EMC Problems

Ilari Hänninen
Felix Wolfheimer, Andreas Barchanski, and Darryl Kostka
CST AG
Darmstadt, Germany
Email: Ilari.Hanninen@cst.com

***Abstract*—All electronic devices must meet electromagnetic compatibility (EMC) standards to ensure safe and reliable operations. Using simulations to investigate possible EMC issues during the design stage can be used to quickly identify possible issues. Realistic simulations may however result in very large scale problems, especially if one wants to include all possible environment effects. This article describes how to use high performance computing techniques for efficient EMC simulations, introducing different strategies that can be applied to a wide range of problems. Two examples are discussed: A large scale PCB simulation, and an ESD simulation of a cable system in a vehicle. A number of simulation techniques and technologies are shown and discussed, including domain decomposition, GPU computing, combining solver technologies, and the use of model segmentation to study how a large system can be simulated in a piece-wise fashion.**

I. Introduction

The currents flowing inside electrical devices generate electromagnetic fields. When multiple devices operate in a shared environment, these fields can create couplings between the devices, which can affect their performance or even lead to failures. There are regulatory limits that define the emissions devices can produce, in order to reduce the risk of electromagnetic compatibility (EMC) problems. However, the effects that give rise to EMC problems, such as resonances, couplings and field leakage, are complicated and often hard to compute, and to include all possible effects in a realistic operating environment often also leads to very complex and hardware intensive simulations. So traditionally EMC engineering was associated with measurement.

Nowadays it is possible to investigate and check the EMC properties of a design at any stage in the design process, as the hardware and the simulation techniques have advanced. In particular, the results of simulation can influence the design, allowing multiple possible configurations - for example, the alignment of a board or the position of a component - to be tested comparatively quickly and cheaply. In this paper several high performance computing techniques are discussed with the aid of two examples: Analysis of a high speed single-ended bus on a large scale PCB, and an ESD simulation of a cable system in a vehicle chassis. The techniques proposed include MPI computing (domain decomposition method), GPU computing, combining different solver technologies, and the use of model segmentation which can be used to simulate a large model in a piece-wise fashion.

II. Simulation Acceleration

The time spend on the development of a product until it is released to the market can be critical for success or failure. As simulation has become an essential part of the product development process there's a high demand for simulation programs which allow to make best possible use of available hardware resources to optimze the time-to-result. Besides algorithmic ways of speeding up a calculation, e.g. choosing the algorithm best suited to the problem, or using advanced model decomposition approaches as described in section II-C, acceleration techniques based on special HPC hardware can be used to speed up the computation itself and, thus, minimize the time needed to obtain the reqired results. CST STUDIO SUITE® [1] which was used to obtain the results shown in this paper offers several acceleration methods based on the use of special hardware as well as on parallelization on different levels of the simulation process.

A. Multithreading and GPU Computing

The performance of the transient solver which was used for the computations presented in this paper, a very small amount of computational operations are performed on a large amount of data. Therefore, many clock-cycles are spent on waiting until the required data arrives at the computational units and the memory bandwidth is in most cases the factor which determines the performance of the solver. In the current CPU architecture all CPU cores on a socket are competing for the same memory controller which limits the scalability of the algorithm on a multicore CPU (see Fig. 1). The higher memory bandwidth of the recent NVIDIA Kepler GPUs which is 288GB/s for a Kepler K20X or K40 device as compared to the memory bandwidth on a standard dual socket workstation equipped with Intel Xeon E5 (Sandy Bridge EP) processors which is 59.7GB/s per socket, and the data parallel nature of the FIT algorithm allows for a very efficient implementation of the algorithm for GPUs. To obtain the best possible performance a domain decomposition is applied to the model under study to assign the computational workload onto all available devices, i.e. onto the GPUs as well as the CPUs available in the machines. The amount of work assigned to each of the available resources is based on a weight function which estimates the computational workload of each mesh cell as well as a performance evaluation of the computing devices (CPU and GPUs) which is done in advance of the actual simulation.

Copyright 2014 IEICE

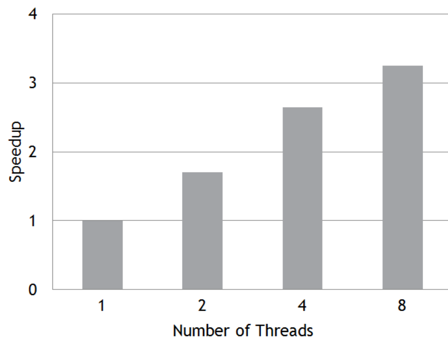

Fig. 1. Typical scaling of the CST MWS® transient solver as a function of the number of CPU threads.

B. MPI Computing

If a single computer has not enough resources (memory or CPU/GPU performance) to handle a large simulation model efficiently, the computational workload and the simulation data can be distributed onto a computer cluster. In contrast to DC described in the previous section MPI [2] uses a domain decomposition to assign the computational workload of a *single* model to the computing resources of the cluster (see Fig. 2). The domain decomposition is again based on the computational workload of each mesh cell to achieve an optimal load balancing on the cluster and, thus, the best performance. In case the cluster computers have GPU hardware attached the computations of the subdomains can be accelerated further by applying the domain decomposition described in section II-A. As the computations of the subdomains are not independent of each other data exchange between the cluster computers is required after each time step and, therefore, a high speed low latency interconnection network such as Infiniband helps to improve the performance of such a parallel simulation.

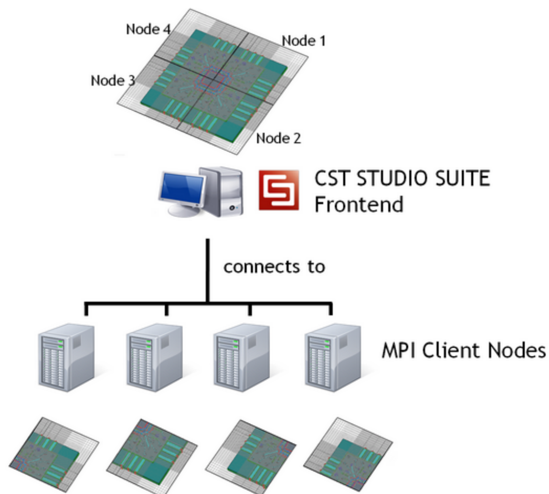

Fig. 2. MPI method working principle.

C. Model Segmentation

Although it is not a hardware based simulation acceleration technique, model segmentation can also be used to reduce the total time required to perform the simulation. That is, instead of using the full model, it is often possible to divide the model into separate independent parts, each of which can be simulated separately. The simulation results from each part can then be combined using e.g. CST DESIGN STUDIO™. This approach has some limitations, namely the 3D fields of the full model cannot be simulated, and it is possible that some coupling and radiation effects between the pieces are not taken into account.

Using the model segmentation technique also allows one to choose the best solver technology for each part of the model. E.g. a full model consisting of a feeding network, filter, and an antenna, could be segmented into three parts. The feeding network could be simulated with a transient solver, the filter with a frequency domain based method such as FEM, and the antenna with an integral equation based solver. Naturally this approach requires being able to use different solver technologies and to easily combine the results from each single simulation afterwards. CST STUDIO SUITE® offers that capability in a single user interface, which makes the use of model segmentation a practical tool for large scale simulations.

III. Analysis of a High Speed Single-Ended Bus on a PCB

Large PCBs are very challenging applications to simulate since although they are not usually electrically large, they include lots of electrically small details and usually operate at very high frequencies. Accurately simulating such structures requires a very fine mesh to resolve the small details, and to be able to correctly compute the conductivity and connectivity information. The total mesh size usually grows very large, which makes the memory demands of the simulation very high. Due to the small mesh cell sizes, the high frequency, and the length of the signal line, the propagation of the signal pulse through the PCB typically takes a long time to simulate using time domain based solvers. Frequency domain based simulation methods are usually not applicable for high frequency applications as the memory requirements and simulation times are impractical, although they can be useful for low frequency studies, e.g for power integrity.

The example we discuss here is a study of a high speed single-ended bus from Intel Corporation [3]. The purpose of the study was to investigate the high speed connection between the chips A and B on the PCB, see Fig. 3. Studying the full 3D link is desirable for checking and verifying the functionality of the board in the post-layout stage of the design process. In addition to calculating the S-parameters, with a 3D simulation one can also easily study the field and current distributions on the board, which provides additional tools for the designer to check that the PCB functions as required. If possible problems are found, implementing routing changes and re-simulating the structure is a quick process compared to making a new prototype board and measuring it.

There are two viable approaches to simulate the full 3D link of the single-ended bus: First is a straightforward brute force approach. The full PCB model requires approximately 2.8 billion mesh cells in CST MICROWAVE STUDIO® with the hexahedral mesh. A simulation of that size is practically impossible to run on a single computer, so MPI computing

Copyright 2014 IEICE

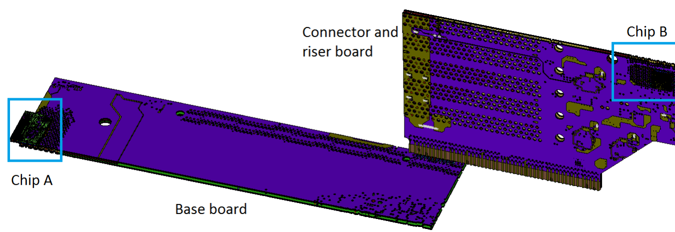

Fig. 3. Picture of the full PCB board indicating the locations of the chips A and B.

is a more viable choice of simulation method. Using a four node computer cluster, we were able to run the simulation with the full model. The results of the simulation are shown in the Figs. 4 and 5 for the insertion and return loss (IL and RL respectively), and in Figs. 6 and 7 for the near-end and far-end crosstalk (NE XTLK and FE XTLK, respectively), as measured from the chip A (all results are shown in dB-scale). As can be seen, the simulation and the measurement results are in good agreement. A full board simulation is thus entirely possible to perform even with relatively affordable hardware investment.

Due to the large size of the full PCB the model segmentation approach was used to compute the full 3D link in [3]. The board was divided into four separate parts: Chip A and the socket, base board, connector and riser board, and chip B. As the resulting models were smaller, it was possible to simulate them on single computer and also with the added benefit of using GPU cards to further speed up the simulation. (The full board was too large to fit into the GPU memory, thus we were not able to use GPUs with the domain decomposition approach.)

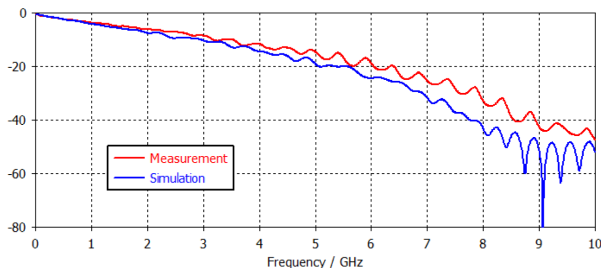

Fig. 4. Insertion loss.

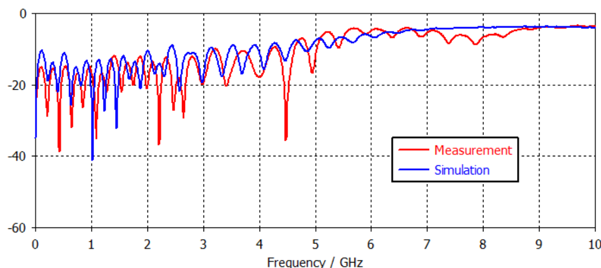

Fig. 5. Return loss.

The results from the model segmentation approach were comparable to the full board simulation with the MPI approach, and are thus not repeated here. However, the model segmentation approach has some limitations. Namely, it is possible that some coupling and radiation effects between the model segments are not taken into account. The full model approach naturally allows us to consider all such effects, and also to study the field and surface current distributions in critical areas between the model segments, thus providing additional insights into possible EMC problems.

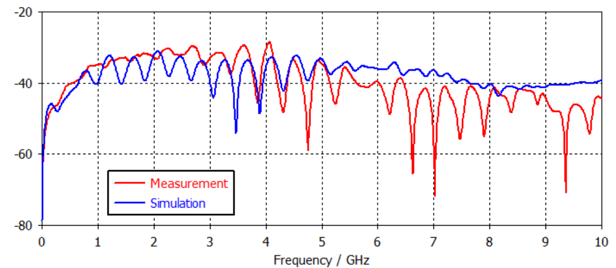

Fig. 6. Near-end crosstalk.

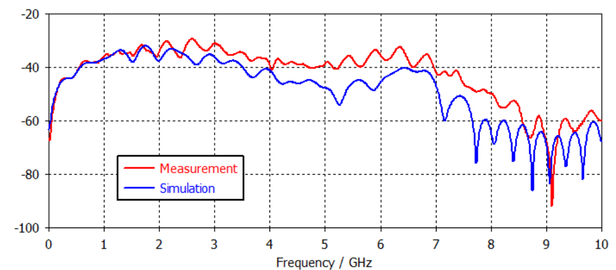

Fig. 7. Far-end crosstalk.

IV. ESD Analysis of a Cable Harness in a Car Chassis

The environment where an electrical device is placed often puts considerable constraints on the simulation. As an example we study the ESD simulation of a cable system including a car chassis. The size of the car chassis and the small cross-section of the cable would result in a very large memory requirement and a long simulation time using a straight forward approach. Instead, we use the combined strengths of CST CABLE STUDIO® and CST MICROWAVE STUDIO® solvers to study the radiaton to and from the cable harness in a realistic environment. The layout of the cable harness inside the car, from the charge point (on the left-hand side of the car at the back) to the battery, and further to the motor (in the right-hand side of the car in the front), is shown in Fig. 8.

The cable harness consists of the power distribution part, which includes three single shielded cables from the charging point to the battery and further to the motor, and in addition of a shielded twisted pair cable for the data connection from the battery to the motor. The diameter of each cable is approximately 5mm, and the length of the cable route is approximately 5m. The frequency range that is simulated is from 0Hz up to 1GHz. At 1GHz the total length of the cable route is several wavelengths. In addition, due the small diameter of the cable the mesh would need to be very fine using a straighforward 3D simulation approach. All these constraints mean that considerable computational resources would be required using a normal 3D full wave simulation.

Copyright 2014 IEICE

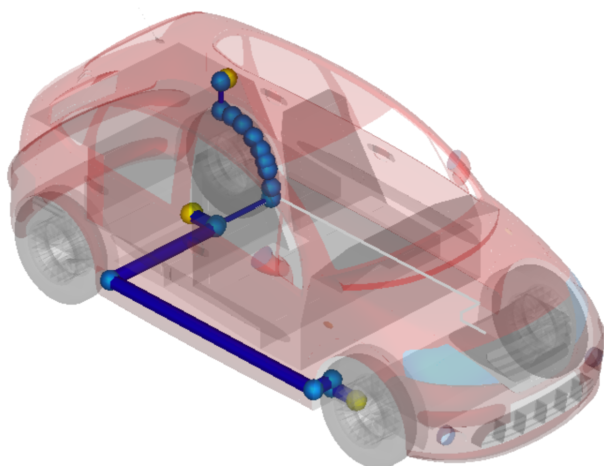

Fig. 8. Cable harness and car chassis.

Instead, we will apply a co-simulation approach and use CST CABLE STUDIO® for the cable harness simulation and CST MICROWAVE STUDIO® TLM-solver for the 3D simulation. The co-simulation allows the use of bi-directional coupling, so that the radiation from the cables to the 3D simulation and the irradation from the 3D simulation into the cables can be taken into account. In addition, we have also used GPU cards to accelerate the simulation. A comparison of the simulation times with up to two NVIDIA Tesla C2075 GPU cards for both CST MICROWAVE STUDIO® TLM- and T-solvers are presented in Table I.

We apply an ESD pulse on the charging point and study the radiated emissions from the cable. A snapshot of the resulting electric field inside the car chassis 4ns after the simulation start can be seen in the Fig. 9. The pulse propagating inside the cable harness has just reached the bottom of the rear left tyre well, as can be seen by the strong field values in the figure. However due to the field escaping outside the cable harness, the induced field actually propagates to the car battery quicker through the air than via the cable bundle itself, and the resulting field on the car battery (the box-shaped structure between the rear tyres) can clearly be seen. Insights like these can prove valuable when researching or resolving possible EMC issues.

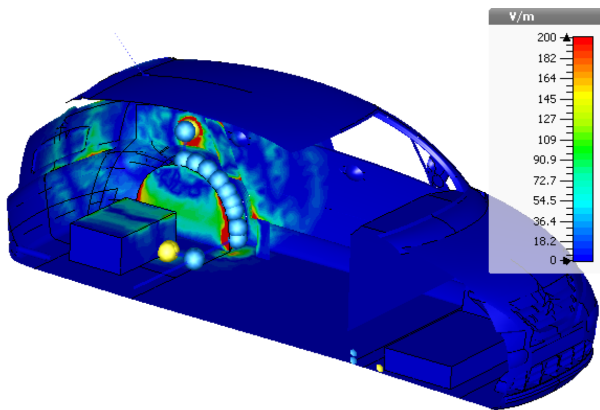

Fig. 9. Electric field induced by the ESD pulse inside the car chassis 4ns after the start of the simulation.

TABLE I. SIMULATION TIMES FOR THE CABLE HARNESS AND CAR CHASSIS MODEL

Number of GPUs	TLM-solver	T-solver
No GPUs	72 min.	136 min.
1 x GPU	25 min.	52 min.
2 x GPU	18 min.	42 min.

V. CONCLUSION

Efficient simulation of 3D full-wave EMC problems can use several different approaches. We have discussed some methods which can be succesfully used for such simulations, and have shown with the aid of two large scale practical examples that it is entirely possible to incorporate simulation tools in the analysis of EMC problems. With the aid of sophisticated simulation methods and acceleration techniques, it is possible to run extremely large scale and detailed models in a reasonable time, and to have access to analysis tools that are not available with measurement based approaches.

ACKNOWLEDGMENT

The authors would like to thank Mauro Lai, Jonathan Casanova, Madhumitha Seshadri of Intel Corporation for the high speed single-ended bus on PCB model and the measurement results.

REFERENCES

[1] CST STUDIO SUITE®, CST AG, Germany, *www.cst.com*

[2] The Message Passing Interface Forum, *A Message Passing Interface Standard, Version 3.0*, 2012

[3] M. Lai, D. Kostka, J. Casanova, M. Seshadhri, *High Speed Single-Ended Bus: Full-Wave Modeling Methodology and Correlation*, 2013 IEEE EMC International Symposium on Electromagnetic Compatibility

Copyright 2014 IEICE