

GPU Accelerated MoM Matrix Inversion

The appearance of CUDA technology enabled usage of GPUs for general-purpose calculations. Large number of cores and significantly faster memory access allow GPUs significant advantage in the number crunching compared to CPU.

GPU Solver, available in WIPL-D Pro since version 9.0, uses the advantages of GPU computing to significantly decrease EM simulation time. It provides GPU acceleration of three phases in EM analysis: matrix fill-in, matrix solution and near-field calculations.

The most important effect of this modern technology is acceleration of the most time-demanding part of EM simulation - matrix solution. It is possible to use multiple GPUs in parallel to further accelerate GPU matrix solution. Comparing to CPU matrix solver, significant acceleration is achieved, without compromising the accuracy of EM results.

The GPU technology with WIPL-D GPU Solver significantly increase size of problem that can be solved within a reasonable time on a desktop computer. When analyzing electrically large EM problems, where system matrix does not fit available RAM, WIPL-D uses out-of-core (OoC) approach (matrix sub-parts are alternately read from hard disk, processed and re-written to disk). Beside the fact that all calculations are done on GPUs, the GPU Solver has several other improvements in OoC algorithm. First, more than one hard disk can be used in parallel, which significantly increases speed of hard disk I/O operations. Second, I/O operations are done in parallel with GPU calculations, and hence hard disk I/O time will (almost) not appear in the overall solving time. Finally, GPU accelerated OoC reduced algorithm allows matrix inversion time to be significantly reduced when solving problems which system matrix is symmetrical.

Examples

We will illustrate the advantage of using GPUs for matrix solution with two examples. The first example is spherical Luneburg lens, modeled with ten concentric spherical layers with equal thickness and constant relative permittivity, with $\lambda/2$ dipole placed $\lambda/4$ away from the lens (Figure 1). The operating frequency is varied so that the electrical length of the lens is different. We simulated the lens at 9 GHz, 12 GHz and 15 GHz. The corresponding electrical sizes of the lens diameter are 12λ , 16λ and 20λ .

The second example is Cassegrain reflector antenna with dual-mode conical feeder (shown in Figure 2). Operating frequency of the antenna is 25.5 GHz. We simulated antennas with reflector diameter of 140λ , 200λ and 240λ . The model of Cassegrain antenna is fully metallic, and contains no wire-to-plate junctions. Therefore, MoM matrix of the model is symmetrical, and reduced algorithm for matrix inversion can be utilized.

Usage of new GPU OoC reduced algorithm for solving symmetrical systems will be demonstrated by comparing simulations of the two aforementioned examples with similar number of unknowns. The target number of unknowns are 200,000 then 300,000, 400,000 and 500,000. The desired numbers of unknowns are achieved by simulating dramatically larger examples: Luneburg lens with diameter 24λ , 30λ , 35λ and 39λ ; Cassegrain antenna with 240λ , 280λ , 390λ and 440λ reflector.

The simulations were performed on the following workstation:

Intel® Xeon® CPU E5-2650 v4 @ 2.20 GHz (2 processors) with 256 GB RAM and up to four GPU cards NVIDIA GeForce GTX 1080 Ti, 6 SATA HDDs configured in RAID-0.

Owing to the efficient usage of multi-core workstations, on the given computer configuration, WIPL-D Pro can solve system with approximately 150,000 unknowns on CPU (in-core solution) in around 2h. For problems of larger size, time required to solve a system by using the CPU can be estimated assuming that the overall time is proportional to N^3 (where N is the number of unknowns). However, the simulation time is increased if the OoC solver is used due to lack of available RAM and the simulation times are impracticable from the engineering point of view.

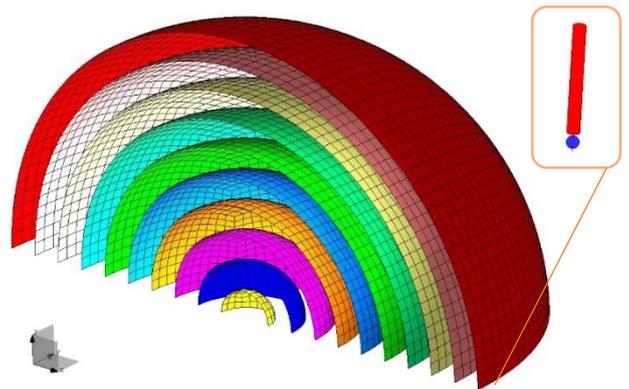


Figure 1. 20λ diameter Luneburg lens excited with a dipole.

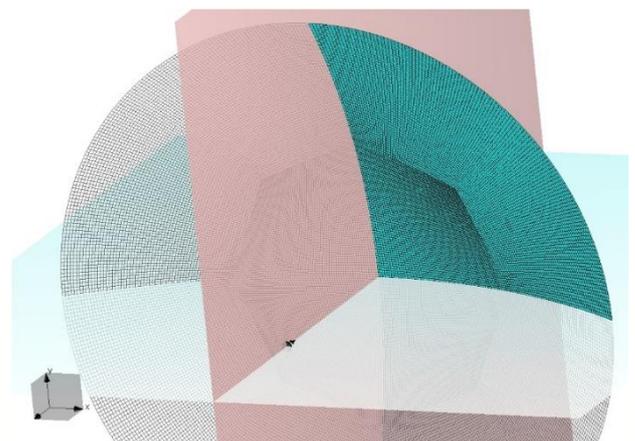


Figure 2. 240λ diameter Cassegrain reflector antenna.

Table 1 shows the comparison of the matrix inversion times for Luneburg lens by using GPU Solver (1 and 4 GPUs) and by CPU (2 CPUs with 12 cores each). Table 2 shows the same comparison for the Cassegrain antenna. Notice that system for 20λ Luneburg lens ($\sim 139,000$ unknowns) is solved in 10 minutes, while the

MoM matrix for 240λ Cassegrain ($\sim 152,000$ unknowns) is solved in less than 15 minutes. As the number of unknowns increases, the improvement achieved by using four GPUs compared to single GPU is improved and it is above 2 times for problems requiring around 150,000 unknowns.

Table 1. Comparison of the measured times for matrix solution for Luneburg lens by using GPU solver (1 and 4 GPUs), and CPU for the same problem.

N	D/λ	t [s] CPU	t [s] 1 GPU	t [s] 4 GPU	Acceleration		
					4/1 GPU	1 GPU/CPU	4 GPU/CPU
52,262	12	377.0	126.6	90.3	1.40	2.98	4.17
85,346	16	1,633.9	379.1	233.0	1.63	4.31	7.01
138,674	20	7,045.0	1,384.8	661.0	2.10	5.09	10.66

Table 2. Comparison of the measured times for matrix solution for Cassegrain antenna by using GPU solver (1 and 4 GPUs), and CPU (2 CPUs with 12 cores each) for the same problem.

N	Antenna diameter	t [s] CPU	t [s] 1 GPU	t [s] 4 GPU	Acceleration		
					4/1 GPU	1 GPU/CPU	4 GPU/CPU
52,223	140	396.9	121.5	82.9	1.47	3.27	4.79
105,728	200	2,460.8	709.4	334.3	2.12	3.47	7.36
151,898	240	7,313.0	2,034.5	875.1	2.33	3.59	8.36

Table 3 shows the matrix solution times for the problems with more than 200,000 unknowns, when normal or reduced matrix inversion is applied. The Luneburg lens was simulated with normal matrix inversion, while the Cassegrain problem was simulated with reduced.

In the case when reduced matrix inversion is applied, matrix with 400,000 unknowns is inverted in less than 3 h, on an affordable 4 GPU workstation. For the problem with for 500,000 unknowns coefficients, which corresponds to Cassegrain antenna which diameter is equal to 440λ , matrix solution time is around 6h!

Table 3. Simulation times for matrix solution using OoC normal and reduced solver on 4 GPUs.

Model	t [s] Normal	Model	t [s] Reduced
Luneburg 24λ 207,254 unknowns	2,671	Cassegrain 280λ 206,468 unknowns	2,256
Luneburg 30λ 302,594 unknowns	7,186	Cassegrain 340λ 304,073 unknowns	5,206
Luneburg 35λ 400,418 unknowns	15,228	Cassegrain 390λ 403,952 unknowns	9,905
Luneburg 39λ 493,058 unknowns	25,991	Cassegrain 440λ 508,748 unknowns	21,945

Calculated radiation pattern in H-plane of the Luneburg lens with 39λ diameter is shown in Figure 3. Radiation pattern of 440λ diameter Cassegrain antenna in E-plane is shown in Figure 4.

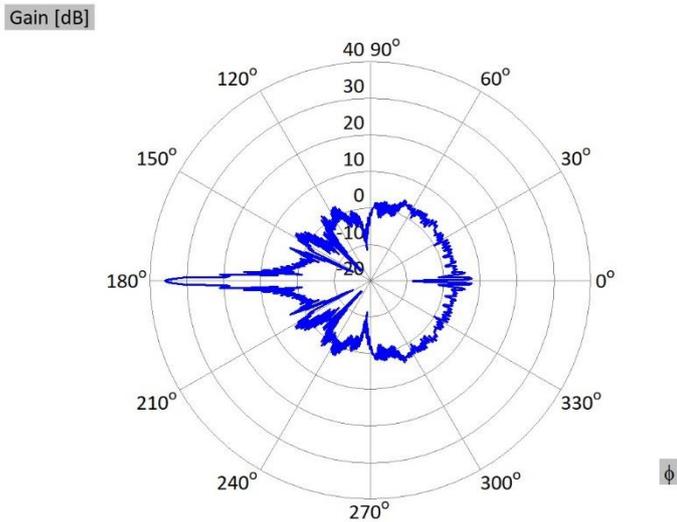


Figure 3. Radiation pattern of the Luneburg lens with 39λ diameter.

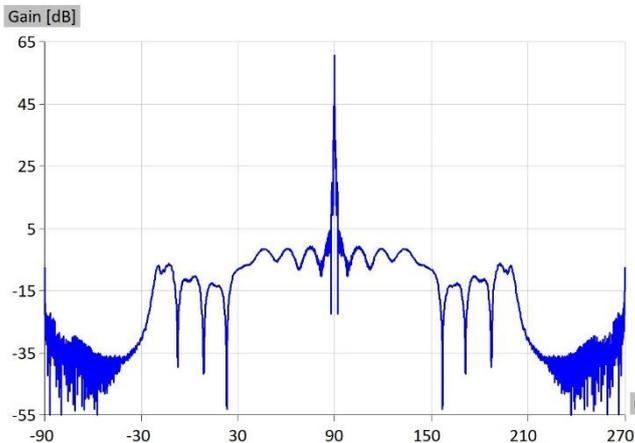


Figure 4. Radiation pattern of Cassegrain antenna with the reflector diameter of 440λ .

Conclusion

The GPU Solver enables WIPL-D Pro to significantly decrease the simulation time, particularly for the dominant component – MoM Matrix inversion. This applies especially in the case of very complex and electrically large structures with dramatically large number of unknowns. In such cases, the suggested simulation approach is using GPUs to perform computing operations instead of CPU. When the MoM matrix does not fit the available RAM, CPU out-of-core solver is additionally slower and the GPU OoC solver is the recommended solver.

The greatest acceleration was achieved in the most time-demanding part of EM simulation - matrix solution for electrically largest problems. Significant acceleration is demonstrated even for the problems with 50,000 - 150,000 unknowns compared to a CPU workstation with 2 CPUs, each with 12 cores. Such problems are within reach of the CPU solution, especially if there is enough available RAM.

For electrically larger problems, practically out of reach of the CPU solution, we show the simulation times for an affordable workstation with 4 low-cost GPU cards and several hard discs supporting the GPU solver. All times are rather short and acceptable from the engineering point of view.

Finally, if the system matrix is symmetrical, as in the cases of pure metallic structures, the reduced matrix inversion can be applied. We demonstrate that the matrix inversion time has been significantly shortened when we compare matrix inversion for two problems. The Luneburg lens requires normal matrix inversion, while the Cassegrain reflector antenna has symmetrical MoM matrix and reduced matrix inversion can be applied.

Electrically large problems, such as 440λ Cassegrain, normally out of reach for any full wave solver, can be solved in reasonable time on a very affordable workstation (2 CPUs each with 12 cores, 4 modest GPU cards, 256 GB RAM and several hard discs supporting the GPU solver). The matrix inversion time is only 6 hours for the 500,000 MoM unknowns problem.