MSC Nastran Sparse Direct Solvers for Tesla GPUs

Cheng Liao

May 16, 2012

MSC Software

Background Information

- MSC Nastran is one of the oldest and most widely used FEA solver codes for Aerospace, Automotive and manufacturing in general.
- First GPU feature for real symmetric sparse direct solver is released in 2012.1, 2^{nd} half of 2011.
- Additional GPU features for complex and unsymmetric sparse direct solvers will be in a 2012 point release. (subject to change)
- GPU rank update kernels are developed by Nvidia Engineers for MSCLDL and MSCLU in house solvers. (Thank you Nvidia!)
- Unsymmetric, 3M and CAS extensions are done by MSC.
- On-going work.

Sparse Direct Factorization

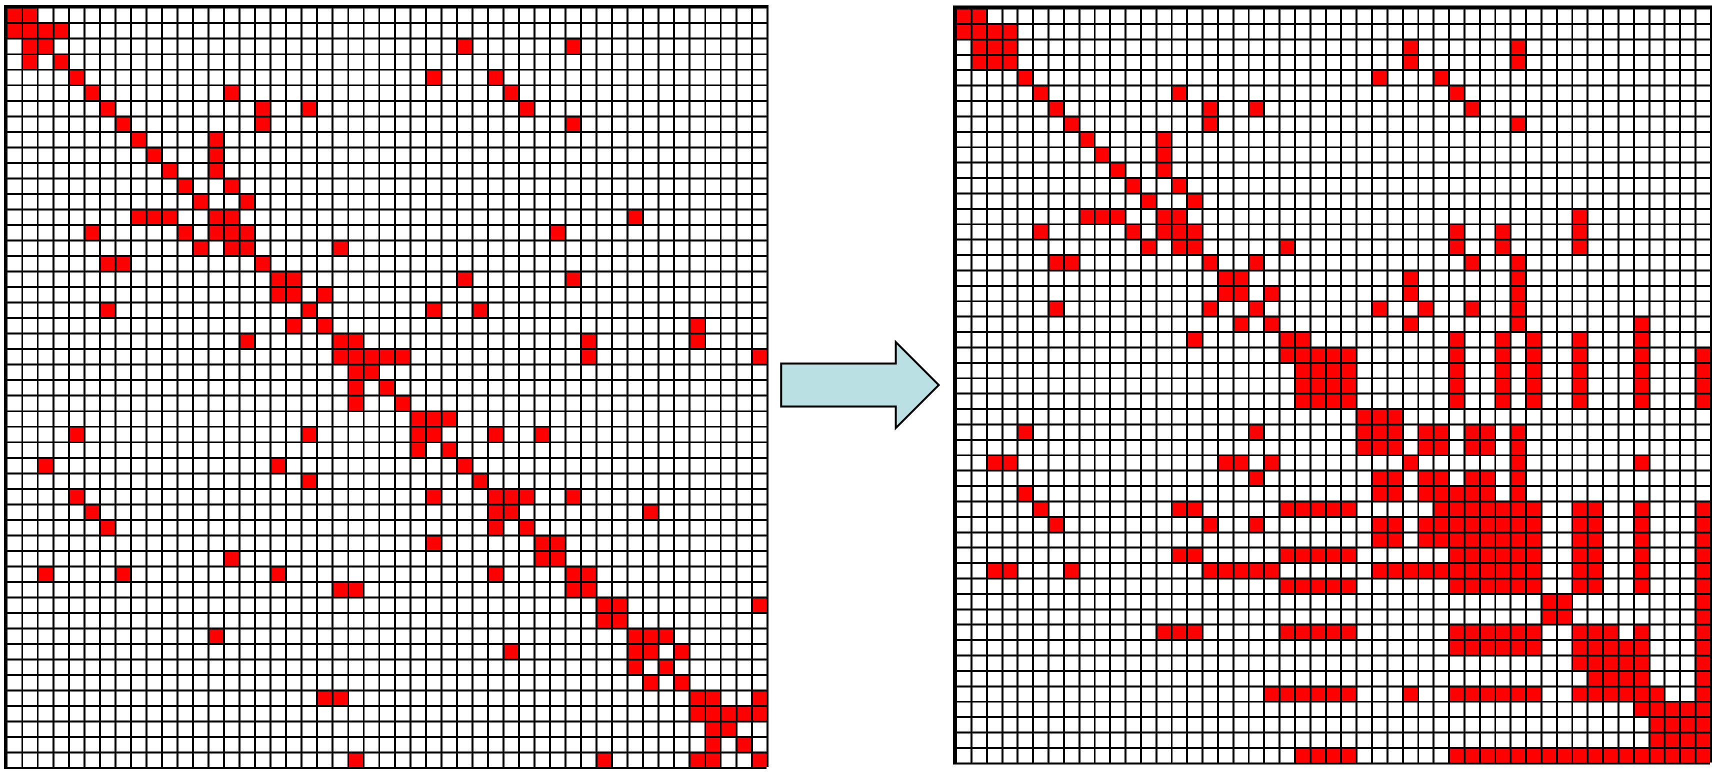

Original Matrix [A]

Factor Matrix [L][D][L]^t

Multi-frontal Algorithm

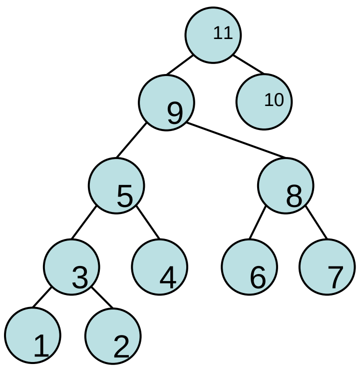

In a proper order, do the following at each node.

Assembly

Pivoting

Block Gaussian Elimination:

From Global Stiffness & contribution blocks

or From Element Stiffness & contribution blocks

$$A = \begin{vmatrix} A_{11} & 0 \\ A_{21} & A_{22} \end{vmatrix} \Rightarrow \begin{vmatrix} L_{11} & 0 \\ A_{21} L_{11}^{-t} D_{11}^{-1} & A_{22} - L_{21} D_{11} L_{21}^{t} \end{vmatrix}$$

Solve

Factorization

Update (Contribution Block)

Three Prong Approach for FP Intensive Kernels

- Sys653=0 (default): segmented rank1 factorization and solve, segmented s/d/c/zgemm update, thresholding 1x1 and 2x2 pivoting, medium segment size, least memory usage.
- Sys653=1: segmented rank-1 factorization, segmented s/d/c/ztrsm solve, segmented s/d/c/zgemm update, no pivoting, larger segment size, slightly more memory usage.
- Sys653=3: s/d/c/zsytrf and s/d/c/zgetrf factorization, s/d/c/ztrsm solve, segmented s/d/c/zgemm update, supernodal 1x1 and 2x2 pivoting, largest memory usage.

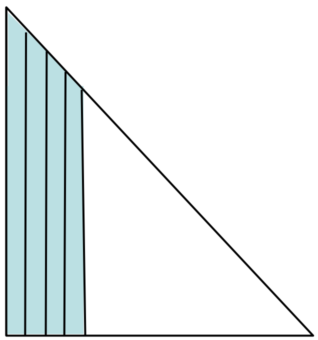

sys653=0

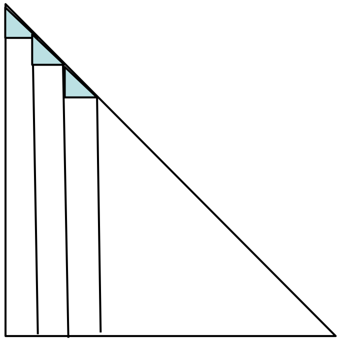

sys653=1

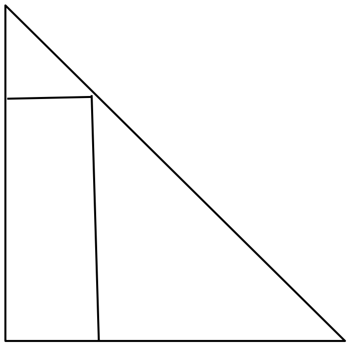

sys653=3

Superelement Matrix Domain Solver

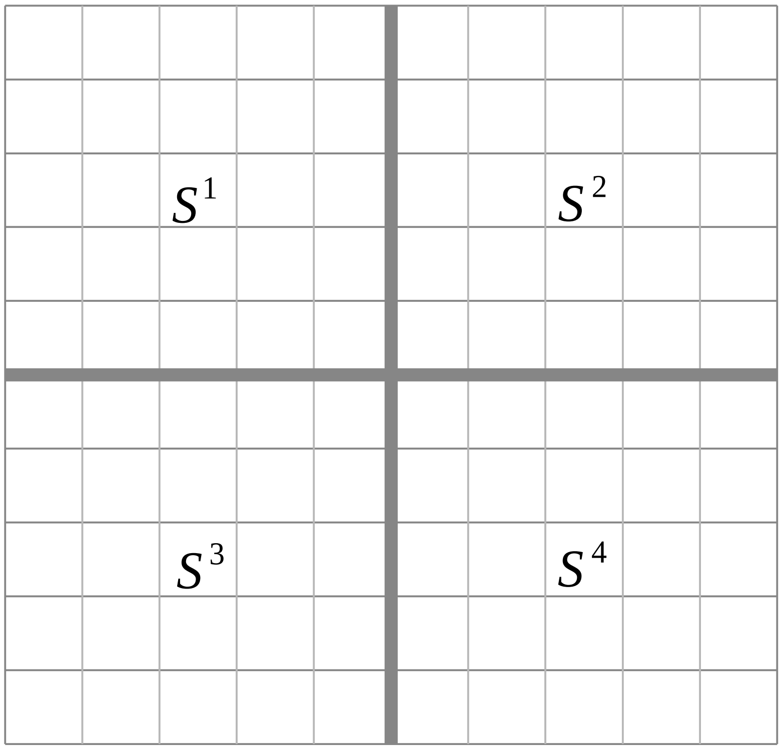

Superelement stiffness:

$$K^j = \begin{vmatrix} K_{ii}^j & K_{ib}^j \\ K_{bi}^j & K_{bb}^j \end{vmatrix}$$

Superelement boundary stiffness:

$$K_b^j = K_{bb}^j - K_{bi}^j \; K_{ii}^{-j} \; K_{ib}^j$$

Global stiffness assembly:

$$K_b = \sum_{j=1}^{n} K_b^j$$

Global equilibrium:

$$K_b X_b = F_b$$

Segmented (rankN) Update with Staged DGEMM

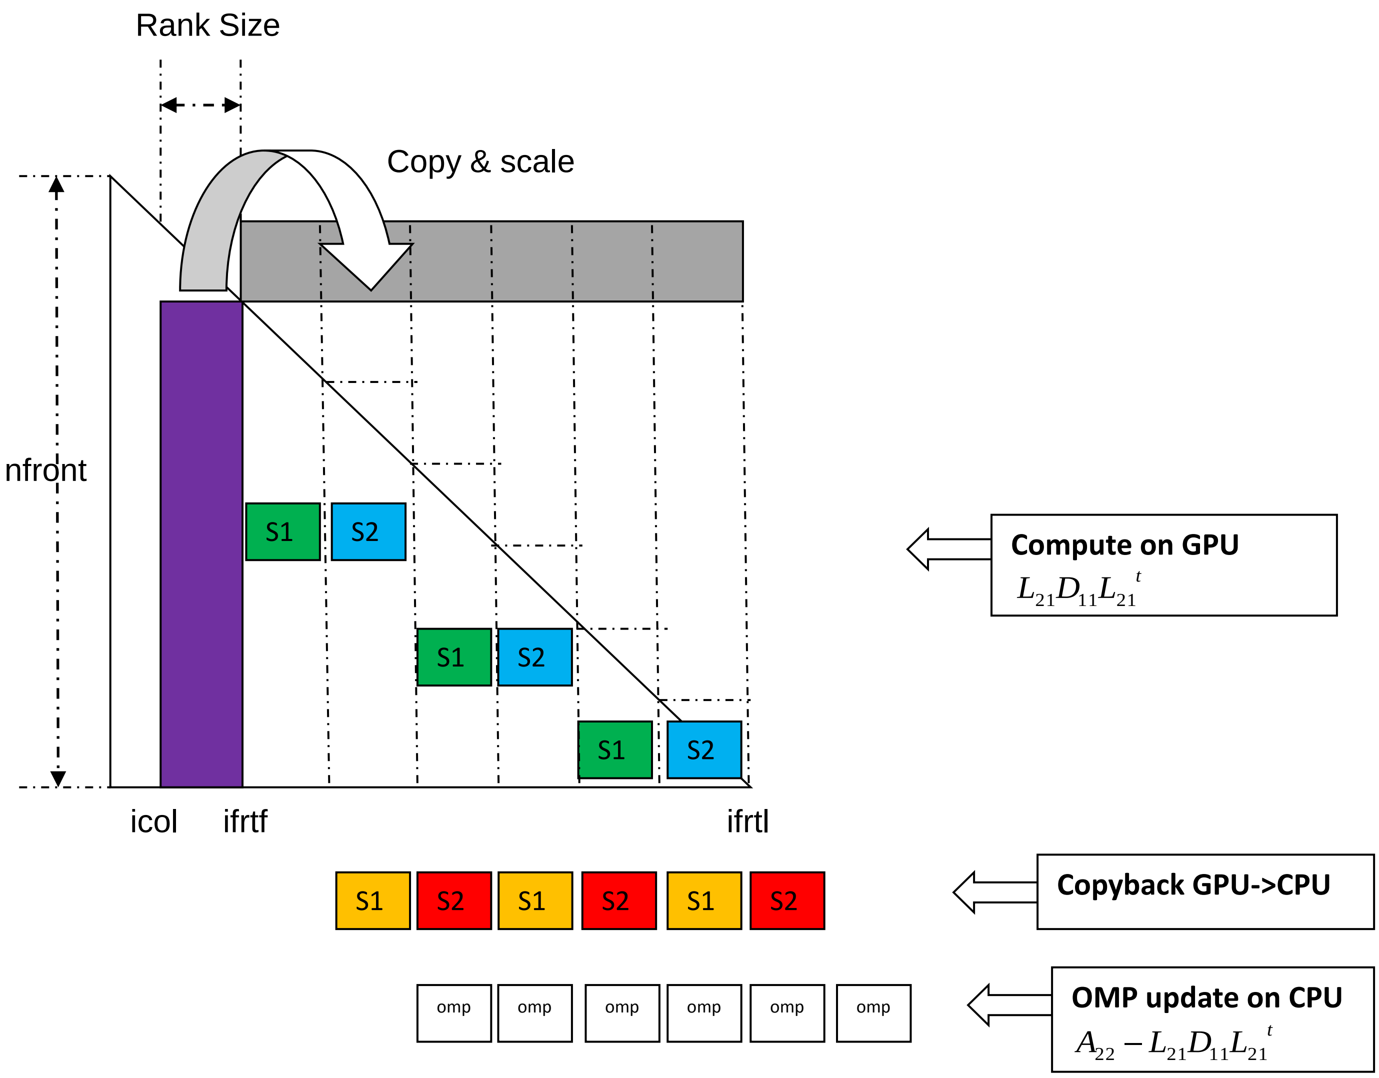

C2070/Westmere-EP Update Gflops vs Rank Size

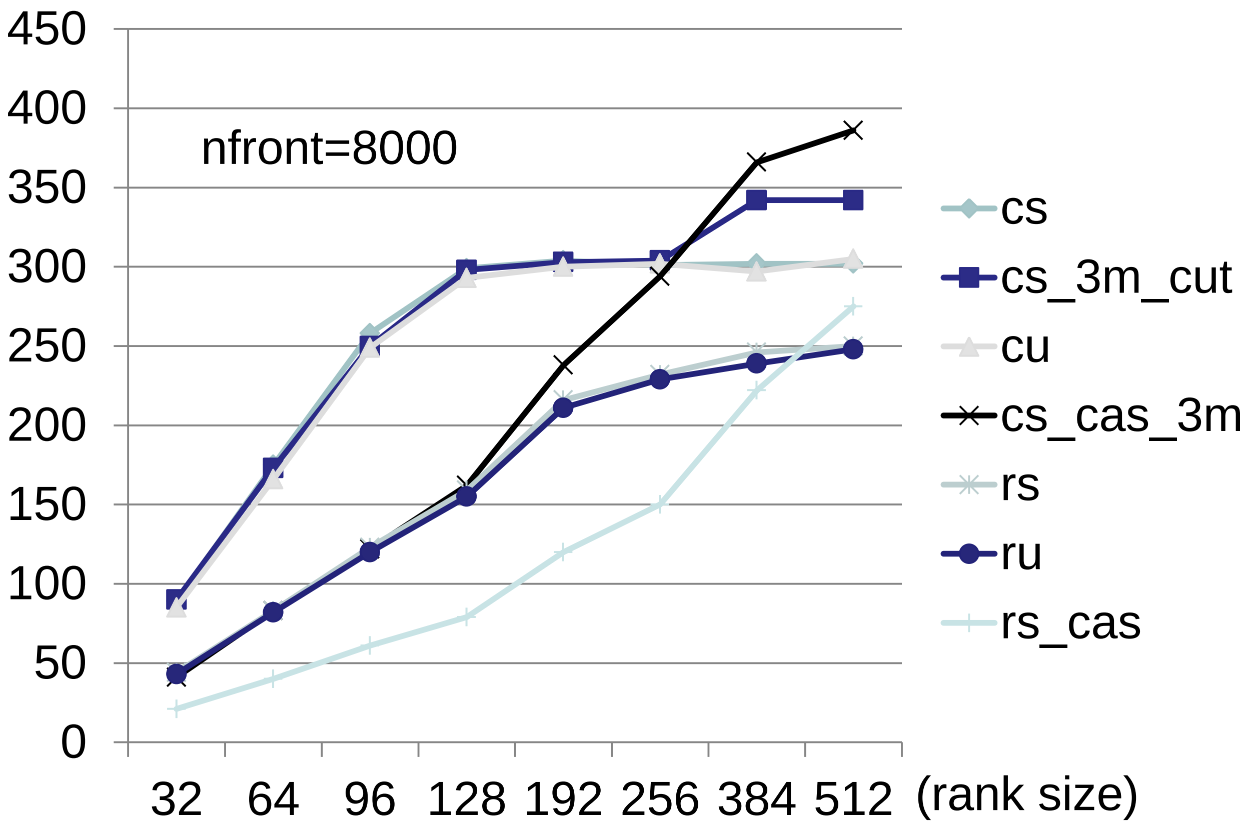

C2070/Westmere-EP Update Gflops vs Front Size

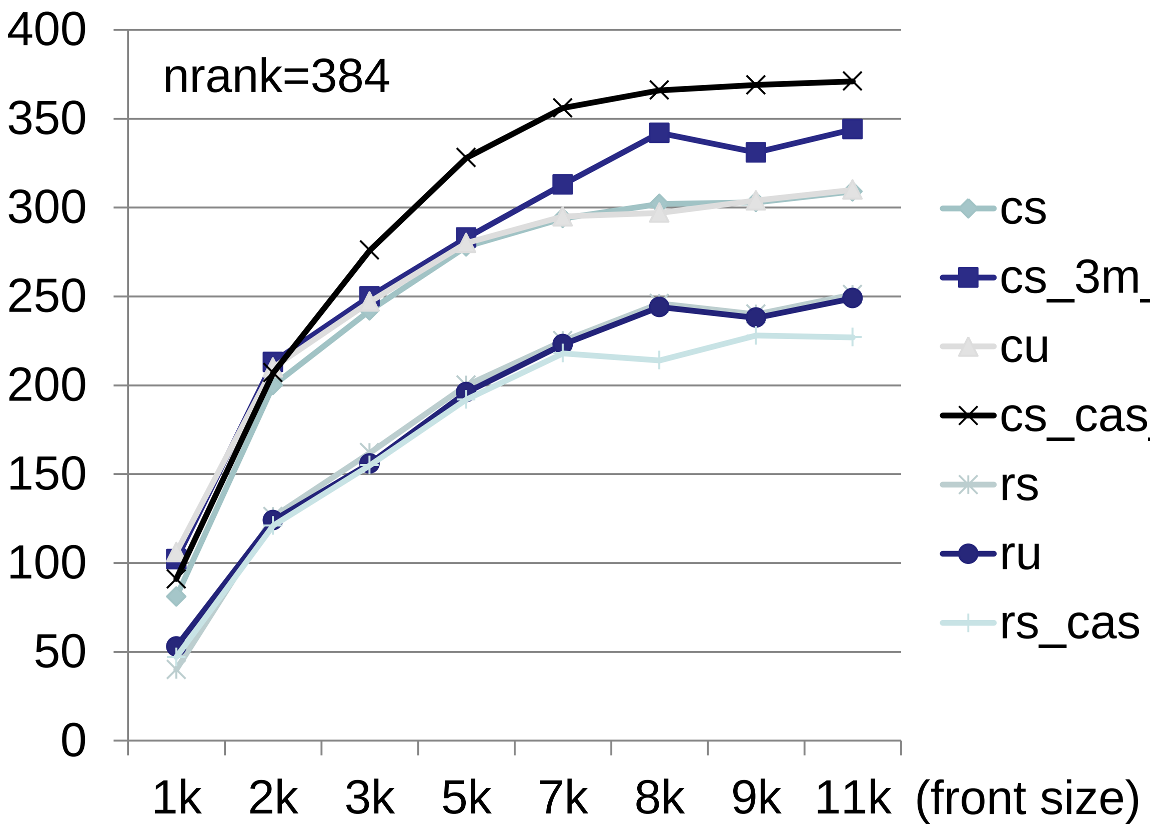

DMP+SMP+GPU (943K dof Crankshaft SOL101)

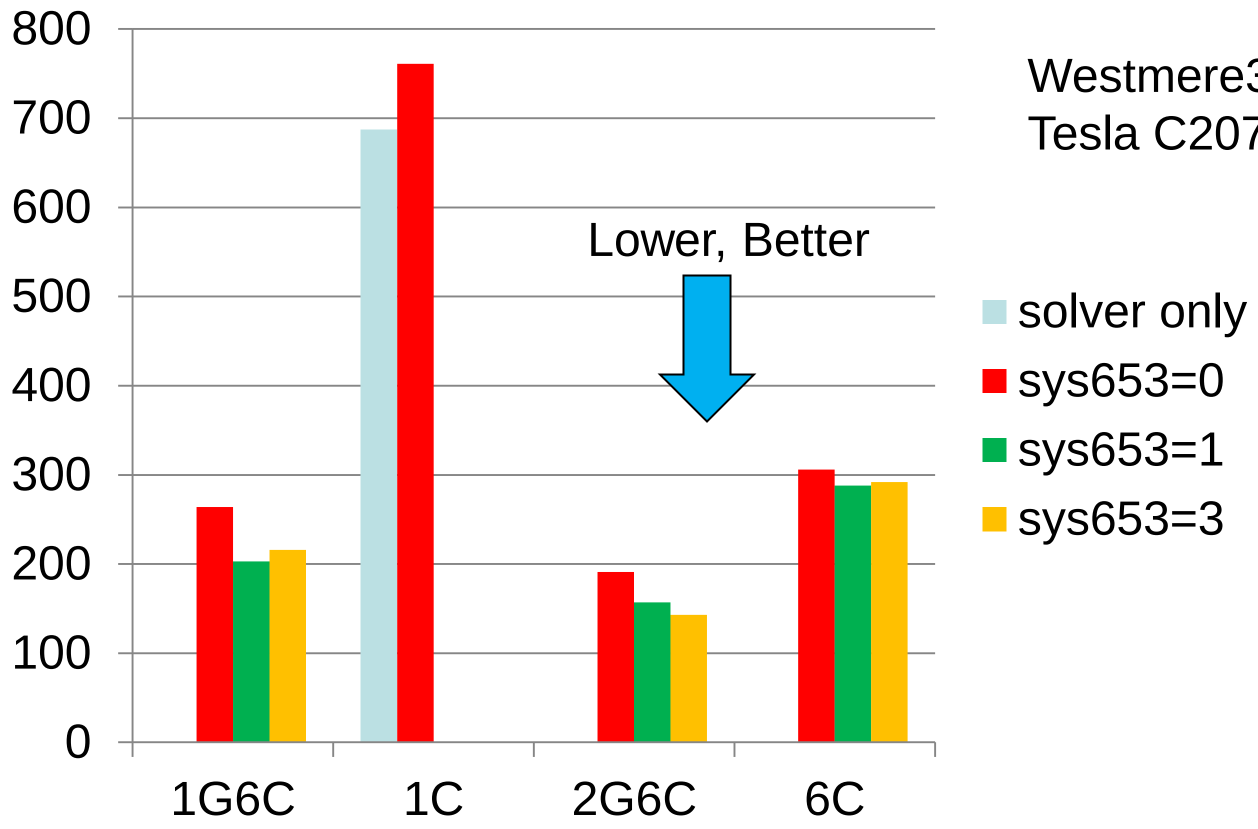

Westmere3.33ghz 2x6
Tesla C2070/2075 48GB

Large 2.4M dof Linear Static SOL101

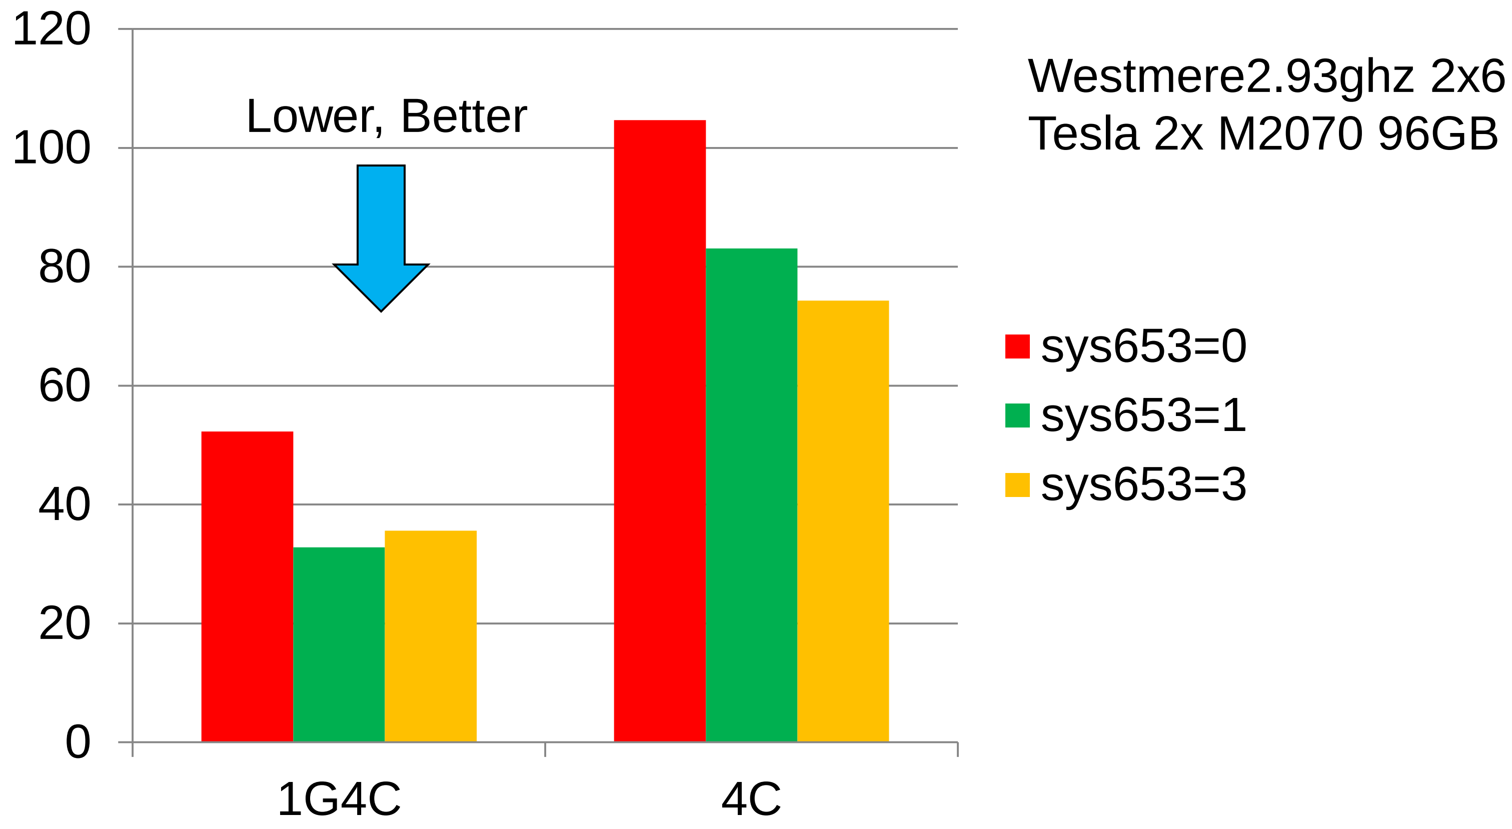

SOL108, NVH
Exterior Acoustics simulation

Exterior Acoustics (3 FREQ1 increments)
Speedup over Serial run

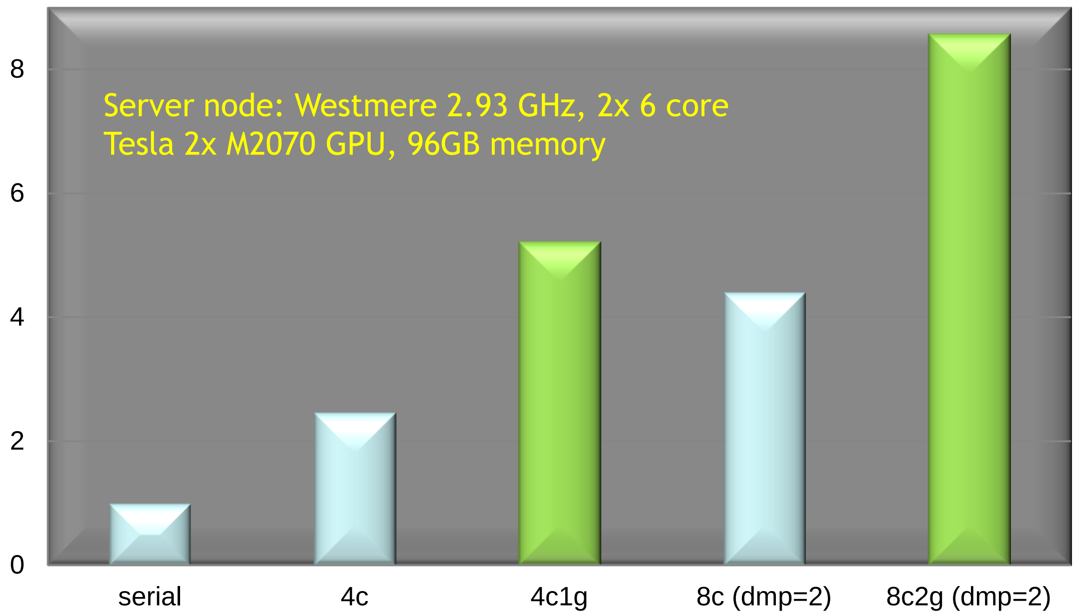

EMEA Customer Model

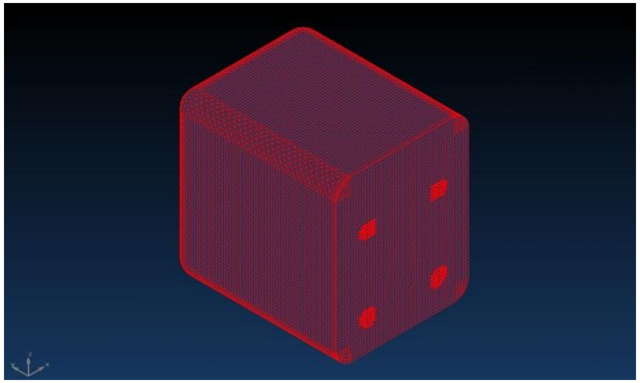

Structural & Acoustic mesh
with Infinite elements
DOF: 0.165M

>5X Speedup with 1 GPU over Serial run
~2x Speedup with 2 GPUs over 8 cores

SOL108, NVH
Trimmed Car Body Frequency Response

Trimmed Car Body NVH
(3 FREQ1 increments; 5 load cases)
Elapsed Time in minutes

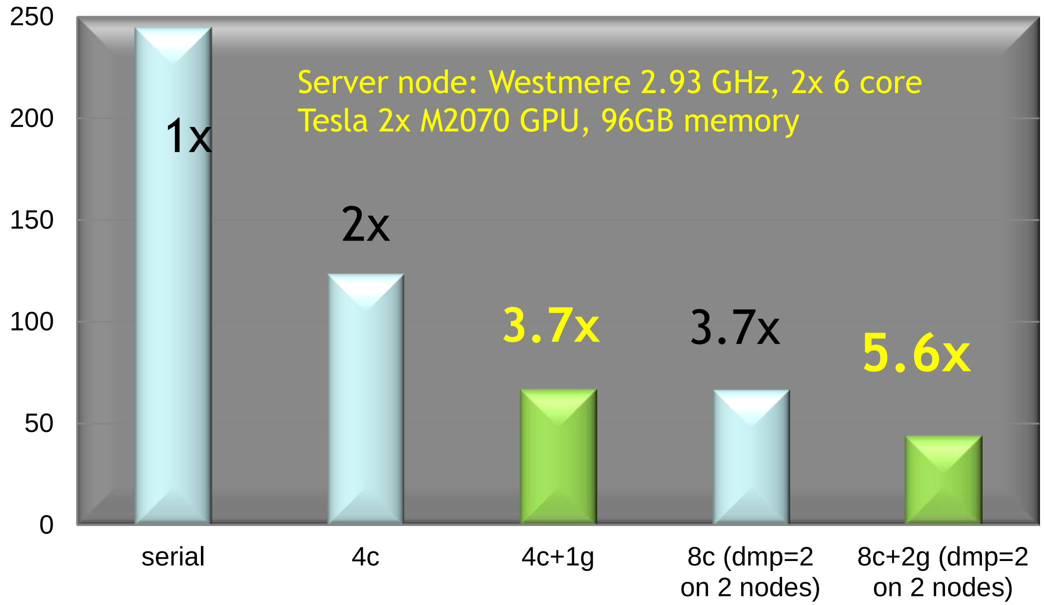

USA Automotive OEM Model

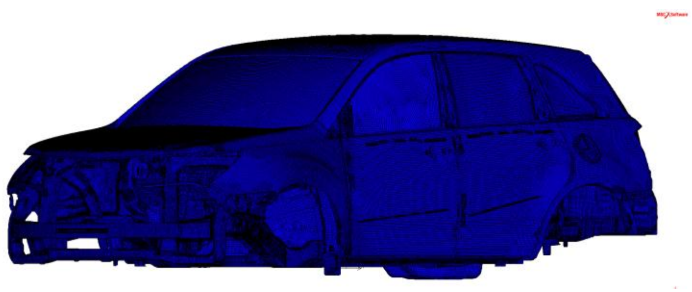

Grids: 1.2M
Shells (CQUAD4): 1.04M
Solids (CTETRA): 0.1M
DOF: 7.47M

~4X Speedup with 1 GPU over Serial run
1.55x Speedup with 2 GPUs over 8 cores

SOL108, NVH
Coupled Structural-Acoustics simulation

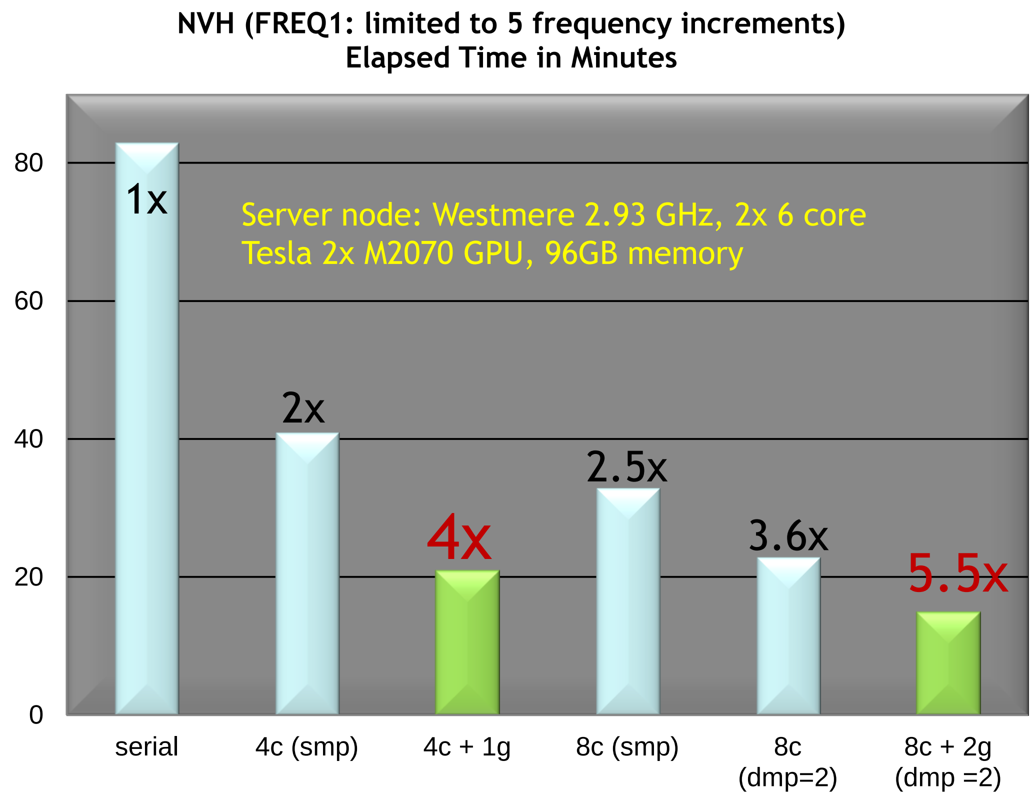

EMEA Automotive OEM Model

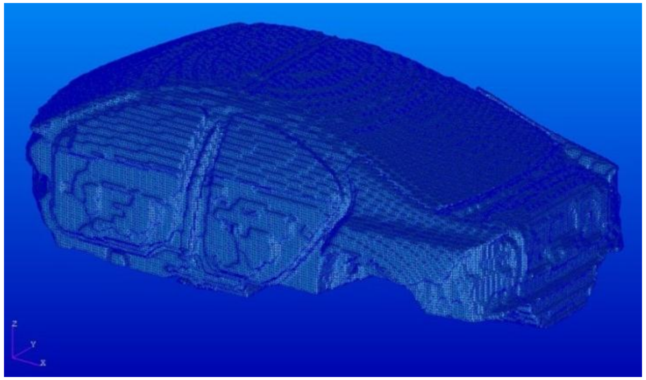

Grids: 0.71M
CTETRAs: 3.8M
DOF: 0.71M

4X Speedup with 1 GPU over Serial run

~1.55x Speedup with 2 GPUs over 8 cores

Conclusions

- GPGPU can give very significant performance boost for direct solver intensive large jobs, ie. max front > 10000 for real data and > 5000 for complex data models.
- Multiple GPGPU performance is delivered by SOL108 (embarrassingly parallel) and matrix domain solver in SOL101.
- Nvidia and MSC continue to work together to tune BLAS and LAPACK kernels for MSCLDL and MSCLU.
- Tree-level parallelism, using GPU memory to cache large sub-tree computing, etc. would provide future road maps to even greater performance.
- A number of other Nastran functional areas are candidates for GPU acceleration.

Acknowledgement

NVIDIA: Steve Harpster, Srinivas Kodiyalam, Stan Posey, Thomas Reed, Steve Rennich, Peng Wang, Shucai Xiao.

MSC: Sunny Gogar, Joe Griffin, Dan'l Pierce, Peter Schartz, Alex TenEyck, Paul Vanderwalt, Ted Wertheimer, Gaowen Ye, members of MSC India QA team.