Gemma

Sandia National Laboratories

SAND2021-1899PE

An Electromagnetic Code for Heterogeneous Computer Architectures

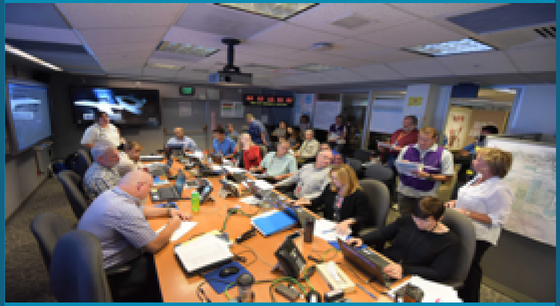

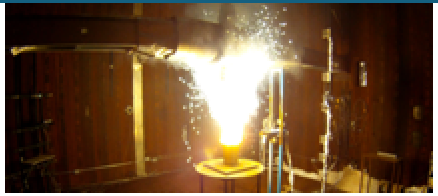

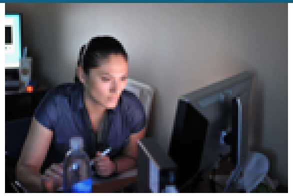

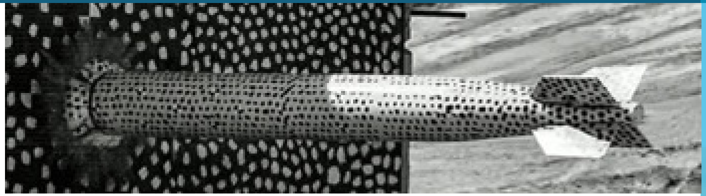

BY

Brian Zinser, Samuel Blake, Robert Pfeiffer, Andy Huang, John Himbele, Brian Freno, Vinh Dang, Joseph Kotulski, Sivasankaran Rajamanickam, William Johnson, Salvatore Campione, and William Langston

U.S. DEPARTMENT OF ENERGY NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Method of Moments and the Problem of Interest

Solving on the Surface instead of in a Volume

Maxwell's Equations:

$$\text{Faraday}: \quad \nabla \times \mathbf{E} = -j\omega\mathbf{B}$$

$$\text{Ampere} - \text{Maxwell}: \quad \nabla \times \mathbf{H} = \mathbf{J} + j\omega\mathbf{D}$$

$$\text{Electric Gauss}: \quad \nabla \cdot \mathbf{D} = \rho$$

$$\text{Magnetic Gauss}: \quad \nabla \cdot \mathbf{B} = 0$$

Wave Equations:

$$\nabla^2\mathbf{A} + \omega^2\mu\epsilon\mathbf{A} = -\mu\mathbf{J}$$

$$\nabla^2\Phi + \omega^2\mu\epsilon\Phi = \rho/\epsilon$$

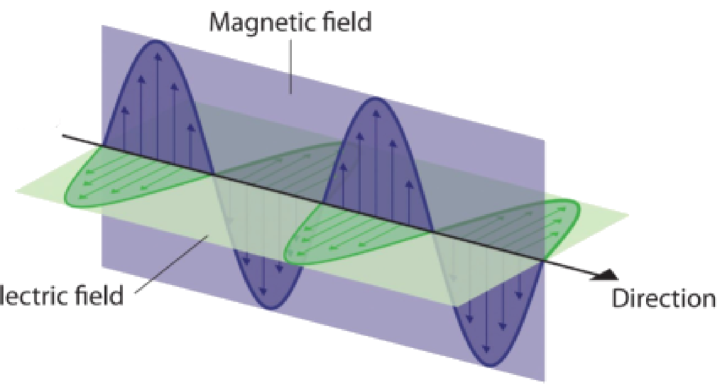

Instead of solving Maxwell's equations in 3D space via the wave equations, we solve them on the boundary between regions.

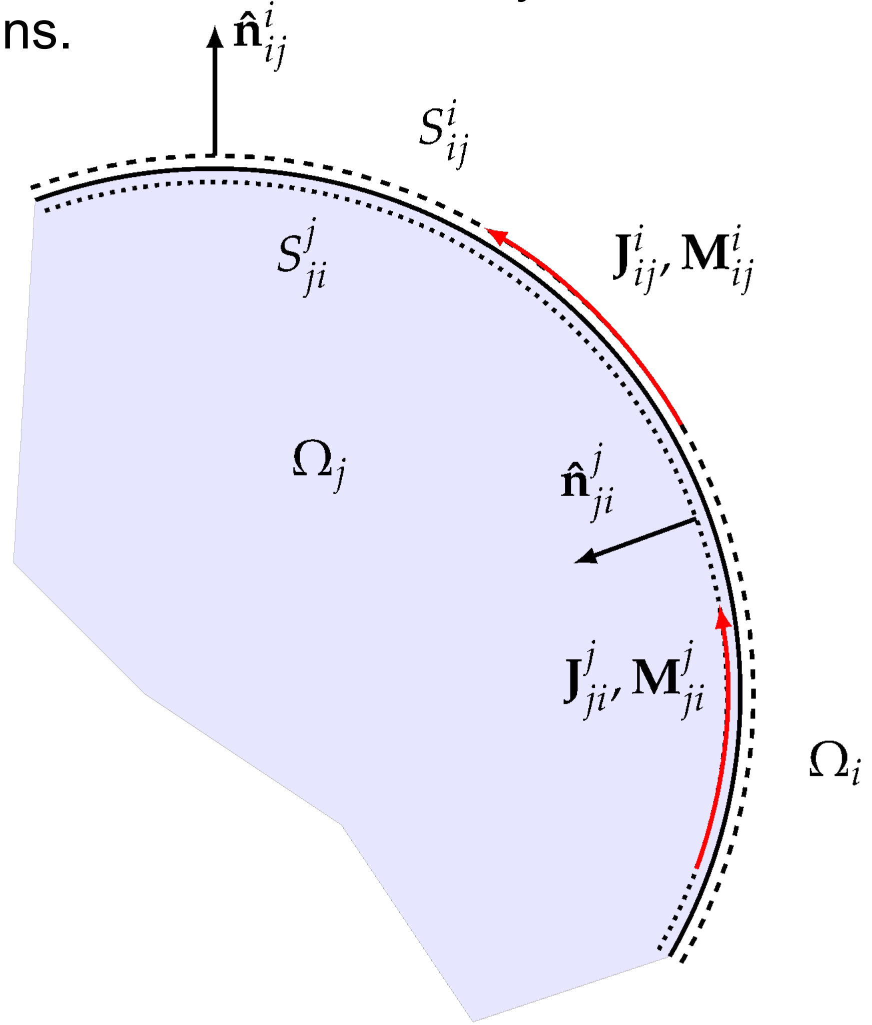

Method of Moments (MoM) Brief Overview

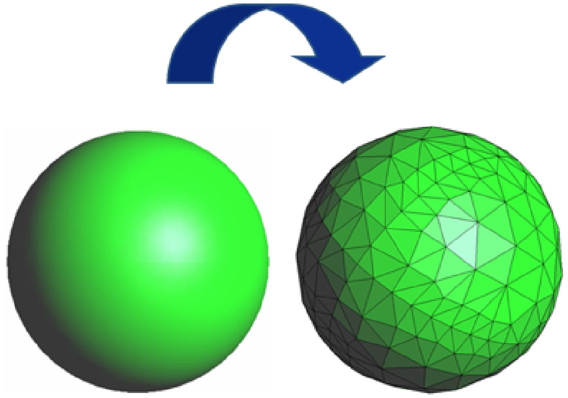

Through the equivalence principle, we consider the current on an objects boundary instead of the field around and inside the object. For electric field $\boldsymbol{E}$, magnetic field $\boldsymbol{H}$, electric current $\boldsymbol{J}$, and magnetic current $\boldsymbol{M}$,

$$\boldsymbol{E} = -i\omega\mu(\mathcal{L}\boldsymbol{J}) - (\mathcal{K}\boldsymbol{M})$$

$$\boldsymbol{H} = -i\omega\epsilon(\mathcal{L}\boldsymbol{M}) - (\mathcal{K}\boldsymbol{J})$$

$$\mathcal{L}\boldsymbol{X} = [1 + \frac{1}{k^2}\nabla\nabla\cdot]\int G(\boldsymbol{r},\boldsymbol{r}')\boldsymbol{X}(\boldsymbol{r}')d\boldsymbol{r}'$$

$$\mathcal{K}\boldsymbol{X} = \nabla\times\int G(\boldsymbol{r},\boldsymbol{r}')\boldsymbol{X}(\boldsymbol{r}')d\boldsymbol{r}'$$

$$G(r) = \frac{e^{-ikr}}{4\pi r}, \mathrm{r} = |\boldsymbol{r} - \boldsymbol{r}'|$$

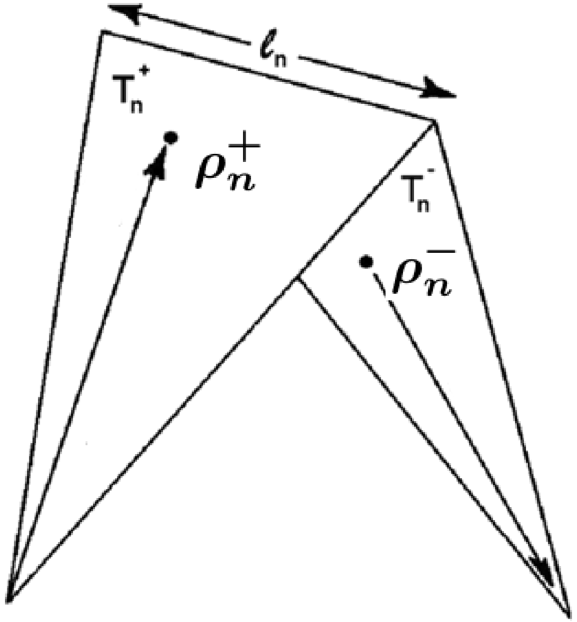

Taking the first equation, but leaving off $\boldsymbol{M}$, gives the electric field integral equation (EFIE). Representing $\boldsymbol{J}$ with a basis $\boldsymbol{f}_n$, testing with a function $\boldsymbol{f}_m$ from the set of basis functions, and moving the derivatives off $\boldsymbol{G}$, its discrete form $\boldsymbol{Z}$ is

$$Z_{m,n} = \int_{f_m}\int_{f_n}\left[i\omega\mu_l \boldsymbol{f}_m \cdot \boldsymbol{f}_n - \frac{i}{\omega\epsilon_l}\nabla\cdot\boldsymbol{f}_m\nabla'\cdot\boldsymbol{f}_n\right]\frac{e^{-ikr}}{4\pi r}$$

$$\mathbf{f_n}(\mathbf{r}) = \begin{cases} \frac{\ell_n}{2A_n^+}\boldsymbol{\rho}_n^+ & \mathbf{r} \in T_n^+ \\ \frac{\ell_n}{2A_n^-}\boldsymbol{\rho}_n^- & \mathbf{r} \in T_n^- \\ \mathbf{0} & \text{otherwise} \end{cases}$$

PEC Sphere Near Field Scattering

For the EFIE, the near field is computed without the test integral:

$$E^{total} = E^{inc} + E^{scattered}$$

where

$$E^{scattered}(\boldsymbol{r}) = \sum_n \int_{f_n} \left[i\omega\mu_l \boldsymbol{f}_n(\boldsymbol{r}') - \frac{i}{\omega\epsilon_l} \nabla' \cdot \boldsymbol{f}_n(\boldsymbol{r}') \right] \frac{e^{-ik|\boldsymbol{r}-\boldsymbol{r}'|}}{4\pi|\boldsymbol{r}-\boldsymbol{r}'|}$$

For a 1 m PEC sphere illuminated by a 377 V/m excitation at 4.77 MHz from above, the scattered near field is given on the right.

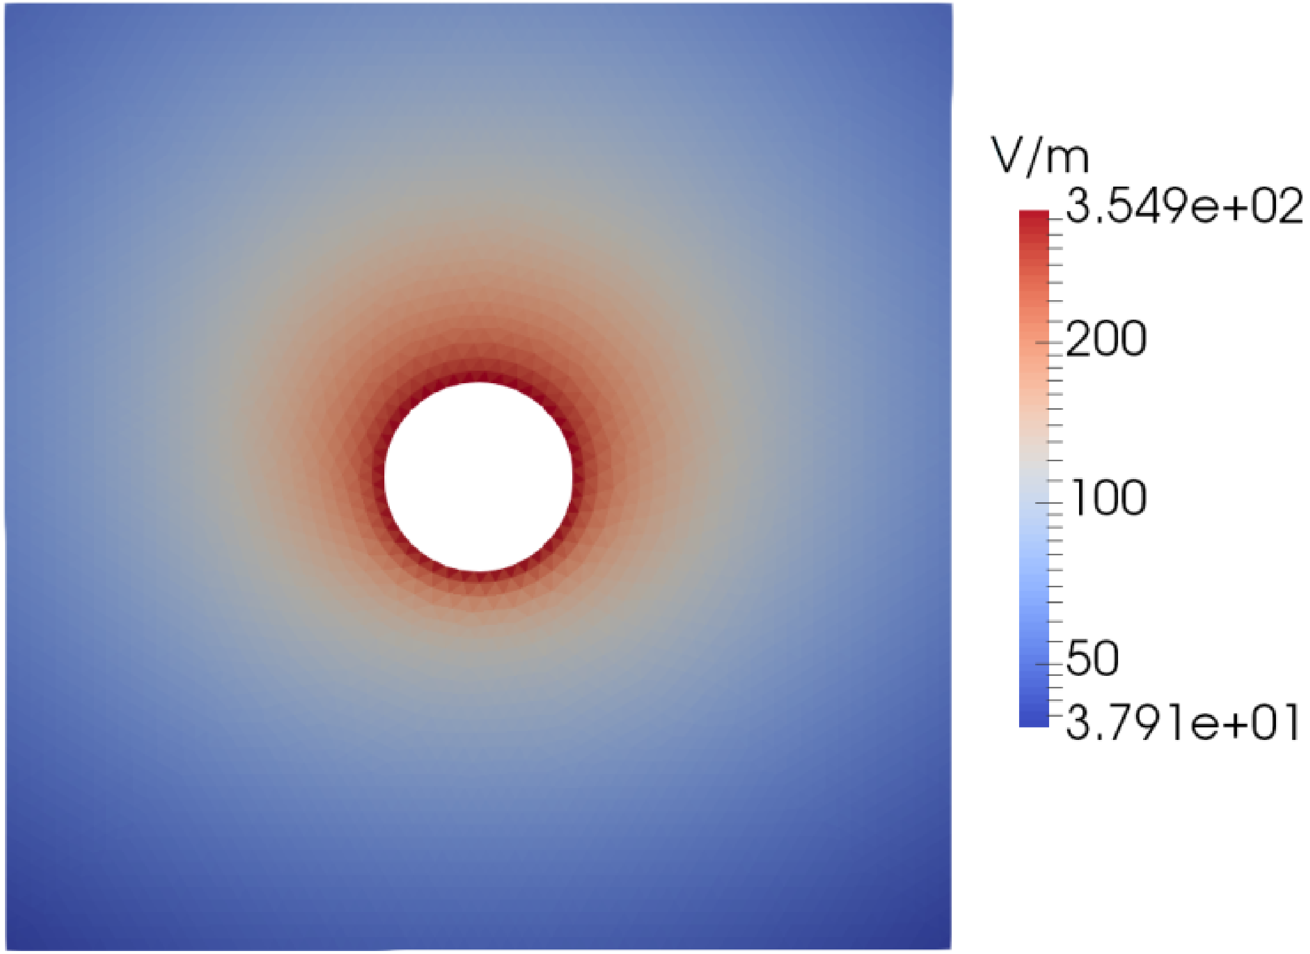

Thin PEC Hollow Sphere

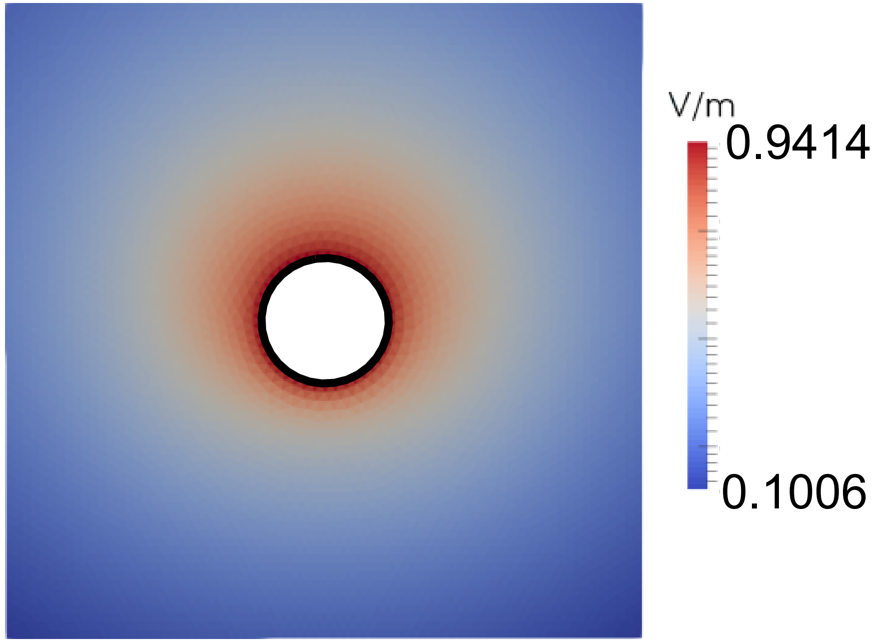

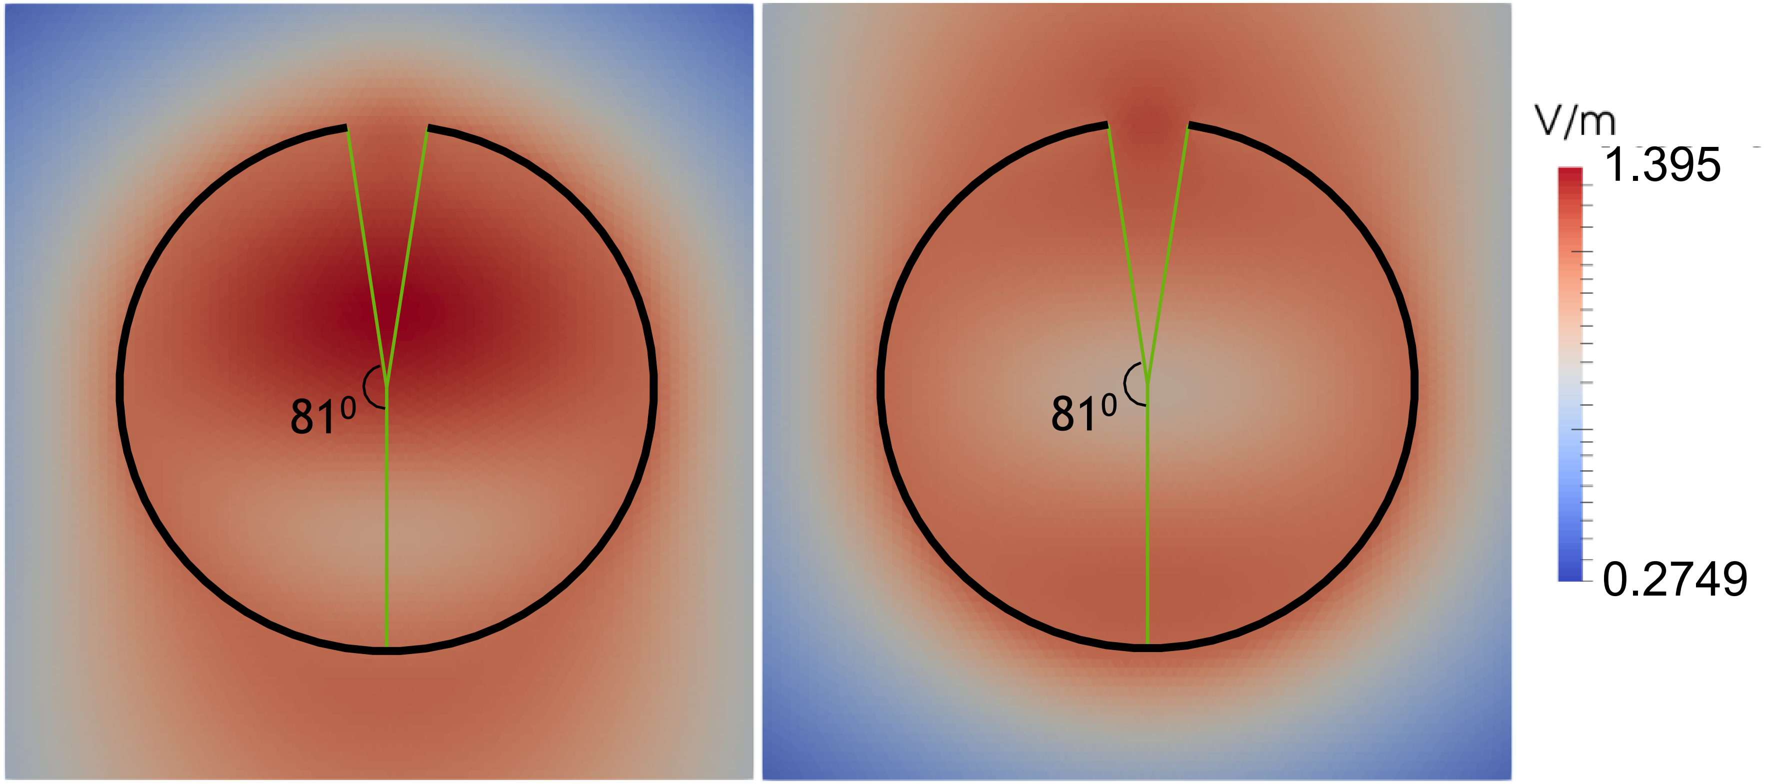

Top:

- Solid sphere
- 4.77 MHz excitation from above, magnitude 1 V/m

Bottom left:

- Hollow sphere
- 130 MHz excitation from above, magnitude 1 V/m

Bottom right:

- Hollow sphere
- 130 MHz excitation from below, magnitude 1 V/m

Performance on Target Architectures

Convergence and Scaling for a Sphere

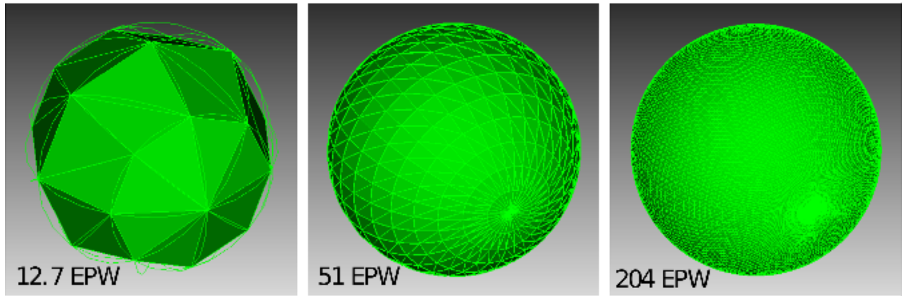

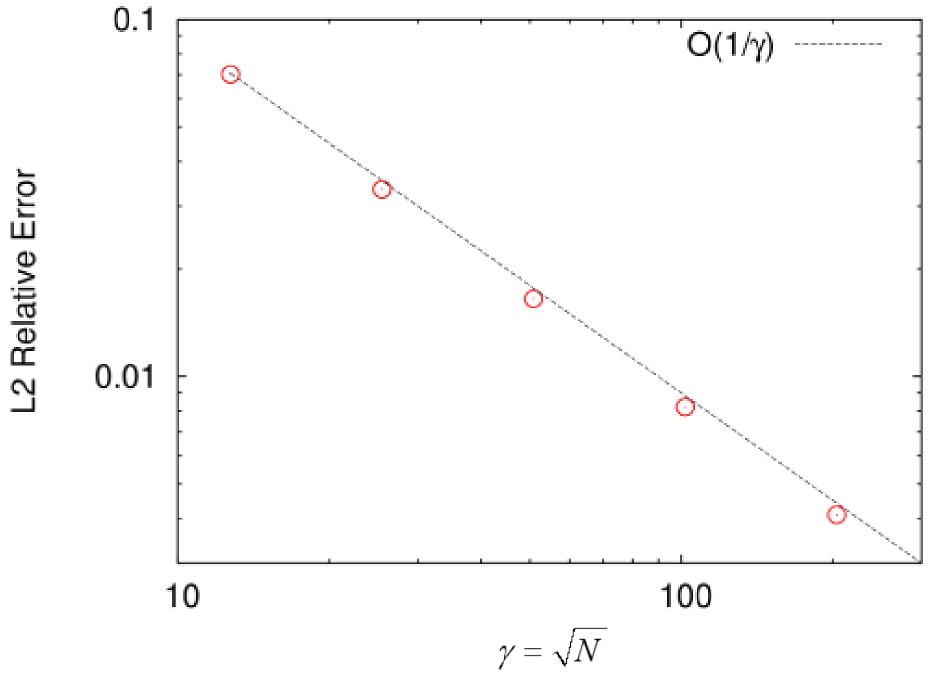

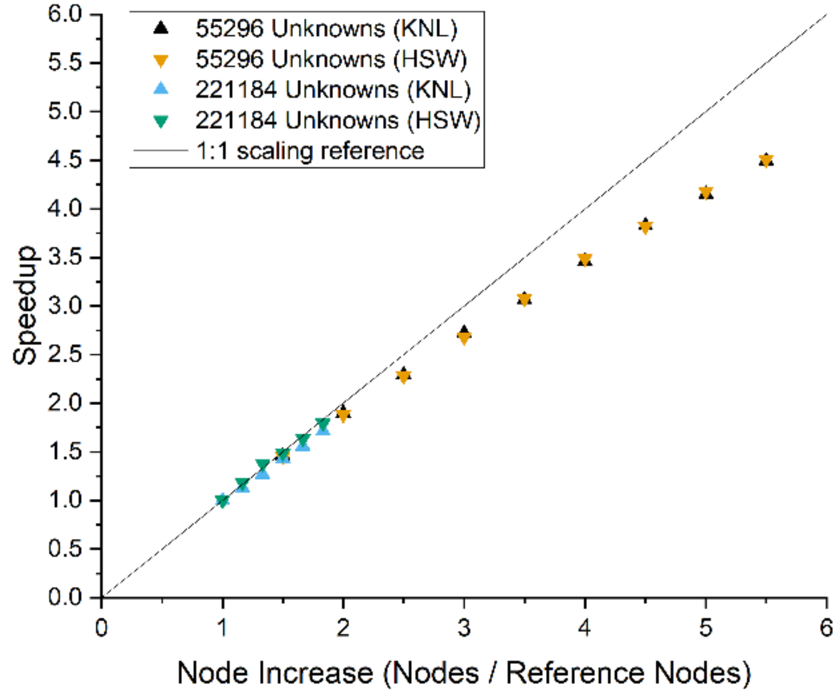

Analytic reference solution for surface current given by Mie scattering solution to Maxwell's equations.

- Slope reduction: communication bound as the amount of computational work per node decreases
- Matrix requires O(N^2) memory to store, calculating its entries is memory bound at the cache level, and O(N^3) computation to solve via LU factorization.

R. Harrington, *Time-Harmonic Electromagnetic Fields*, McGraw-Hill, New York, NY, 1961.

A Deeper Look

HSW CPU Separated Times

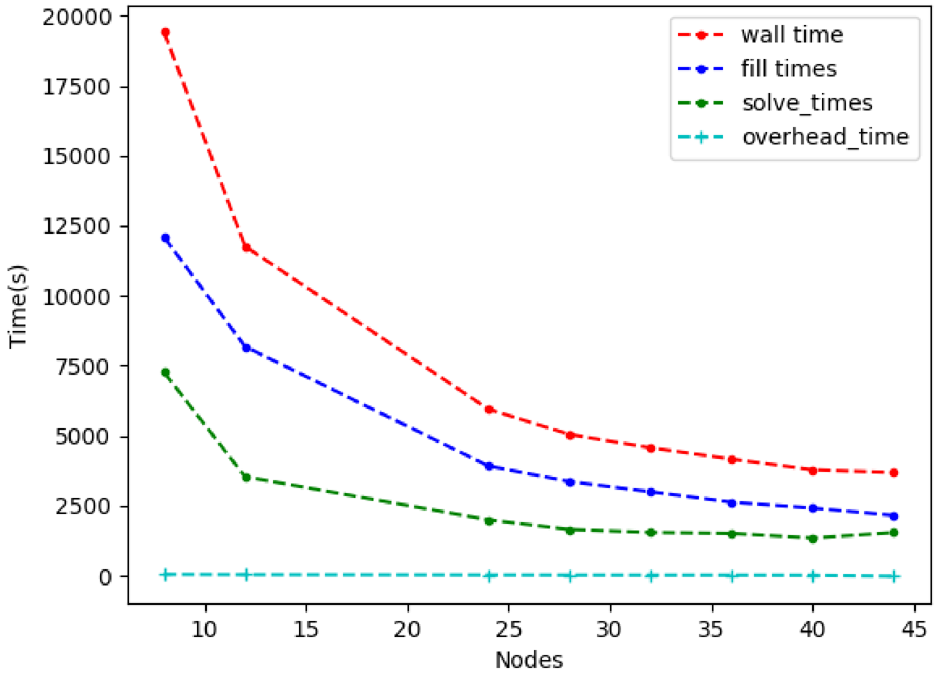

KNL MIC Separated Times

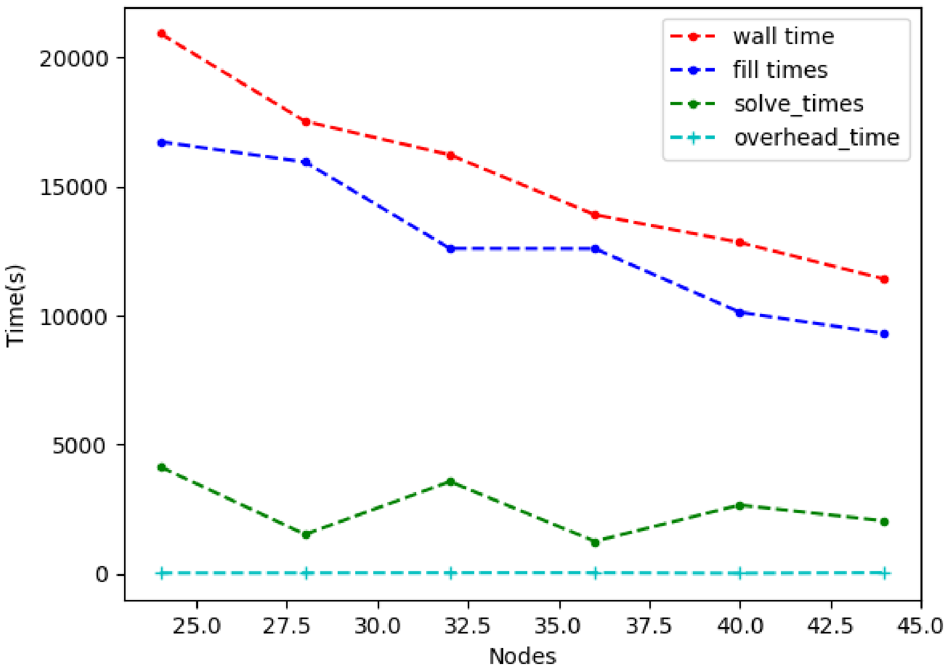

Notes:

- For small problems Gemma is fill dominated. Solve should dominate for larger problems.
- However, Gemma is memory bound at “global” and register levels.

Generally speaking, Gemma is scaling well at tested problem sizes (slope of strong scaling plots near 1 with drop off occurring only as problem size becomes trivial compared to node size).

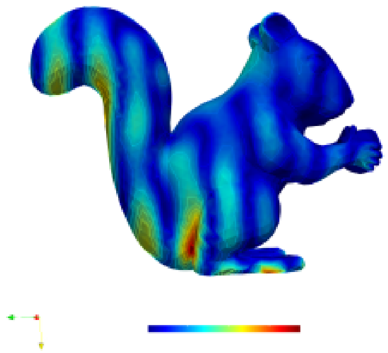

Improving Gemma's Runtime Performance

Gemma

- Efficiency performance: After inspection of the code, identified a large number of calls to the norm function to compute the height of the triangles. Modified code to store the height instead of computing it.
- Tested on CPU (ascicgpu21), GPU (ascicgpu21) and KNL (mutrino)
  - CPU asymptotic reduction: 45%
  - KNL asymptotic reduction: 34%
  - GPU asymptotic reduction: 0.5%
- GPUs are memory bound, not compute bound. In particular, branching and memory access pattern changes are necessary for improvements in performance.
- Future work will include larger problems.

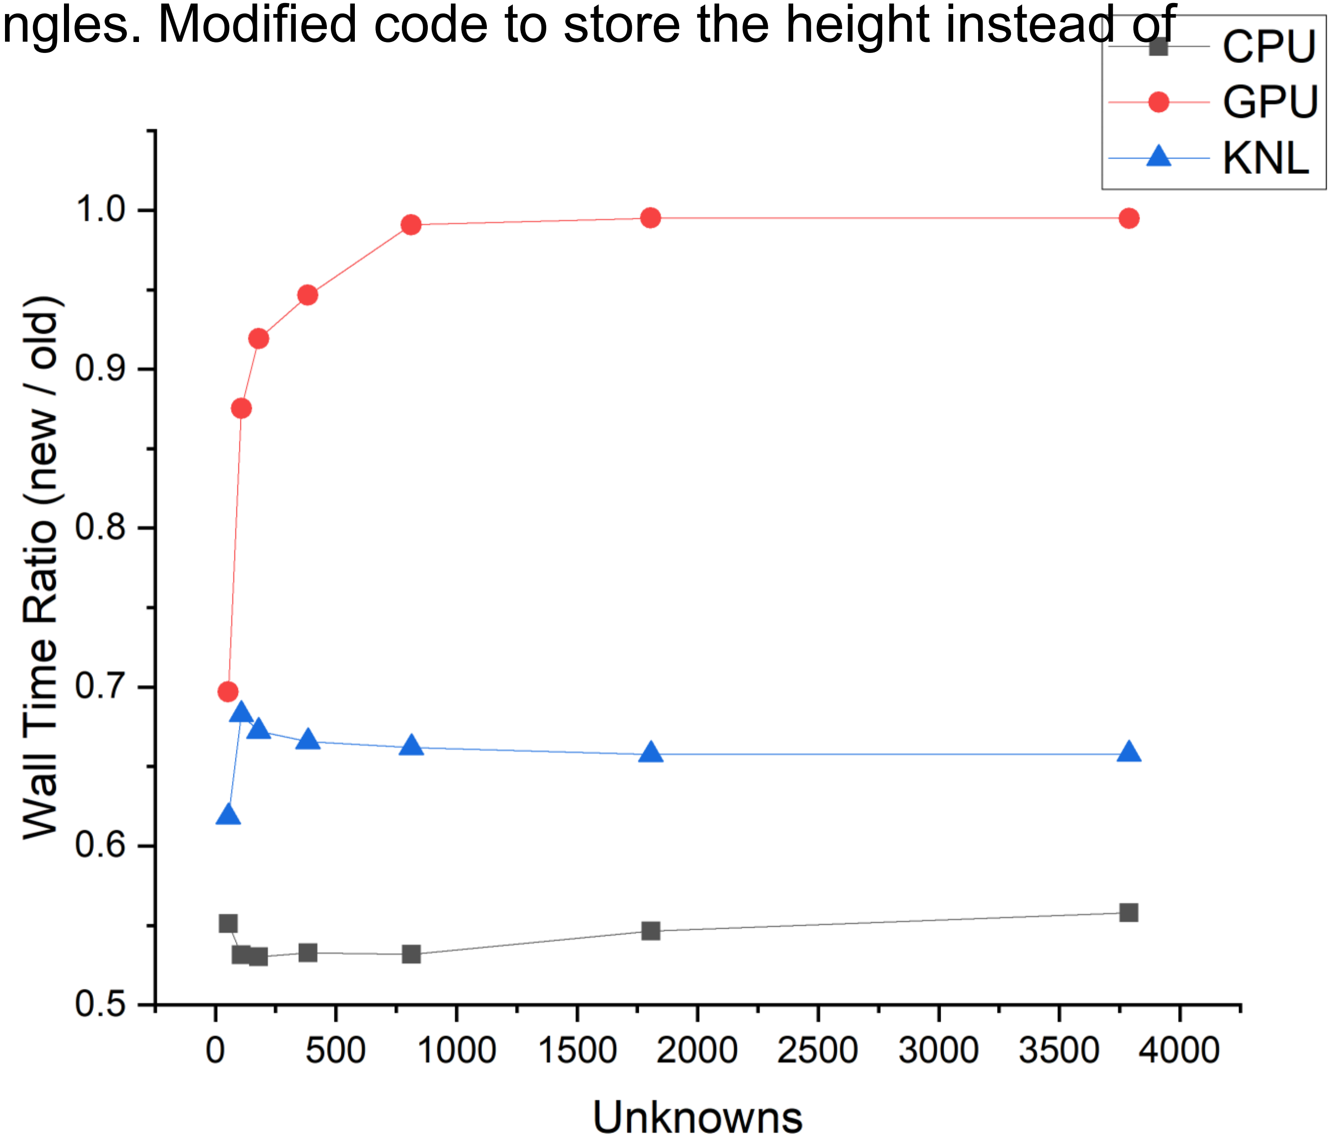

Precomputing Some Information

- Consider filling a single matrix entry (T, S) for the electric field integral equation (EFIE).
- Then compute the 4D integral $\iint G\,(\Lambda^T \cdot \Lambda^S + \nabla \cdot \Lambda^T * \nabla \cdot \Lambda^S)$ over 4 triangle pairs, 2 of which support basis T and 2 of which support basis S.
- Using 3-point and 7-point integration for the test and source elements, respectively, the contribution from a single triangle pair is

$$\begin{bmatrix} \Lambda_1^{Ti} & \cdots & \Lambda_3^{Ti} \end{bmatrix} \begin{bmatrix} G_{1,1}^{Ti,Sj} & \cdots & G_{1,7}^{Ti,Sj} \\ \vdots & & \vdots \\ G_{3,1} & \cdots & G_{3,7}^{Ti,Sj} \end{bmatrix} \begin{bmatrix} \Lambda_1^{Sj} \\ \vdots \\ \Lambda_7^{Sj} \end{bmatrix}$$

- While G changes for every element pair, Λ depends only on the element that supports it.
- At minimum, we would like to precompute Λ.

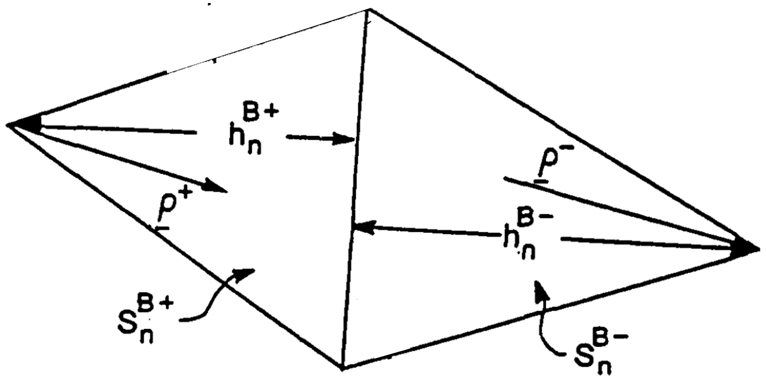

Test Basis

$\Lambda^{T+} = \rho^+/h^+$ and $\Lambda^{T-} = \rho^-/h^-$

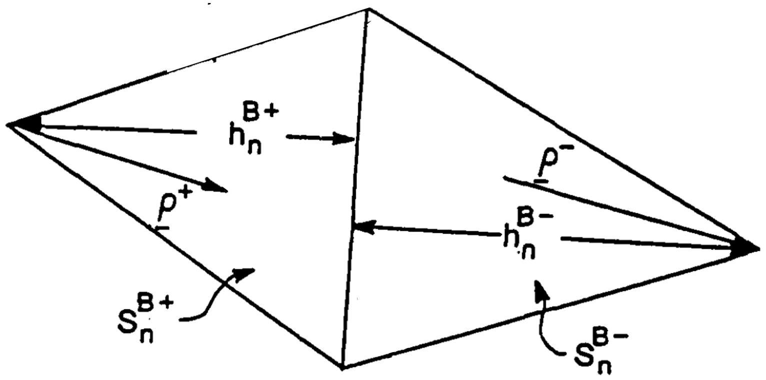

Source Basis

$\Lambda^{S+} = \rho^+/h^+$ and $\Lambda^{S-} = \rho^-/h^-$

Reducing Sample Points

Note the EFIE's $\mathcal{L}$ operator has a weak $O(1/r)$ singularity.

$$\mathcal{L}\boldsymbol{X} = [1 + \frac{1}{k^2}\nabla\nabla\cdot]\int G(\boldsymbol{r},\boldsymbol{r}')\boldsymbol{X}(\boldsymbol{r}')d\boldsymbol{r}'$$

$$\mathcal{K}\boldsymbol{X} = \nabla\times\int G(\boldsymbol{r},\boldsymbol{r}')\boldsymbol{X}(\boldsymbol{r}')d\boldsymbol{r}'$$

$$G(r) = \frac{e^{-ikr}}{4\pi r}, \mathrm{r} = |\boldsymbol{r} - \boldsymbol{r}'|$$

For $\mathcal{L}$, we use a radial angular transformation where the source triangle variables of integration become similar to polar coordinates (ρ, ϕ):

$$\int G(\boldsymbol{r},\boldsymbol{r}')\boldsymbol{X}(\boldsymbol{r}')d\boldsymbol{r}' = \int \frac{e^{-ikr}}{4\pi r}\boldsymbol{X}(r,u)\frac{r}{cosh\, u}drdu,$$

$$u = \ln\tan^{-1}(\phi/2)$$

Instead of selecting the minimal number of samples required for a specific level of accuracy, we take a brute force approach to integration.

While the increased accuracy has helped us with other issues, determining exactly how many points are needed would speed up the fill.

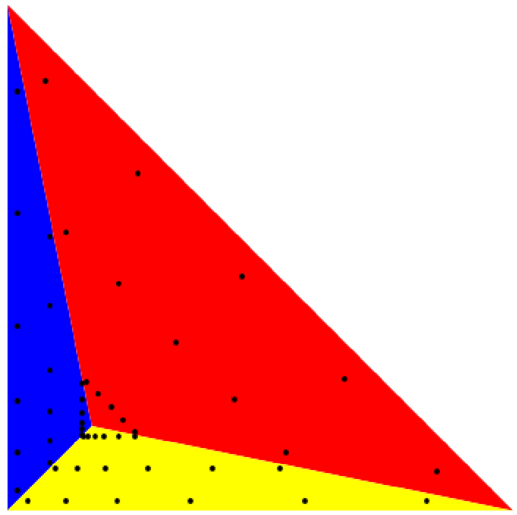

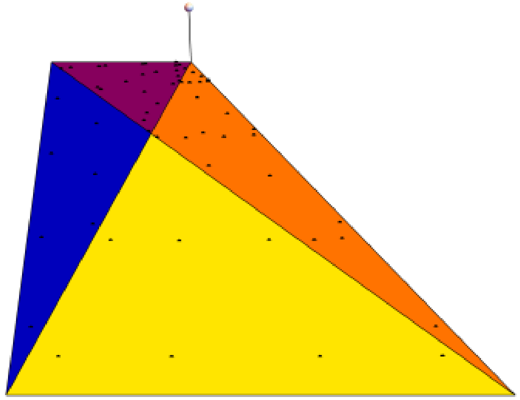

The 3 subtriangles used in the radial angular transformation when the singularity is interior and exterior to the triangle.

M. A. Khayat and R. D. Wilton, "An Improved Tansformation and Optimized Sampling Scheme for the Numerical Evaluation of Singular and Near-Singular Potentials," *IEEE Antennas Wirel. Propag. Lett.*, 7 (2008), pp. 377 - 380.

Weak scalability for Adelus, **A D**ense **LU S**olver Package

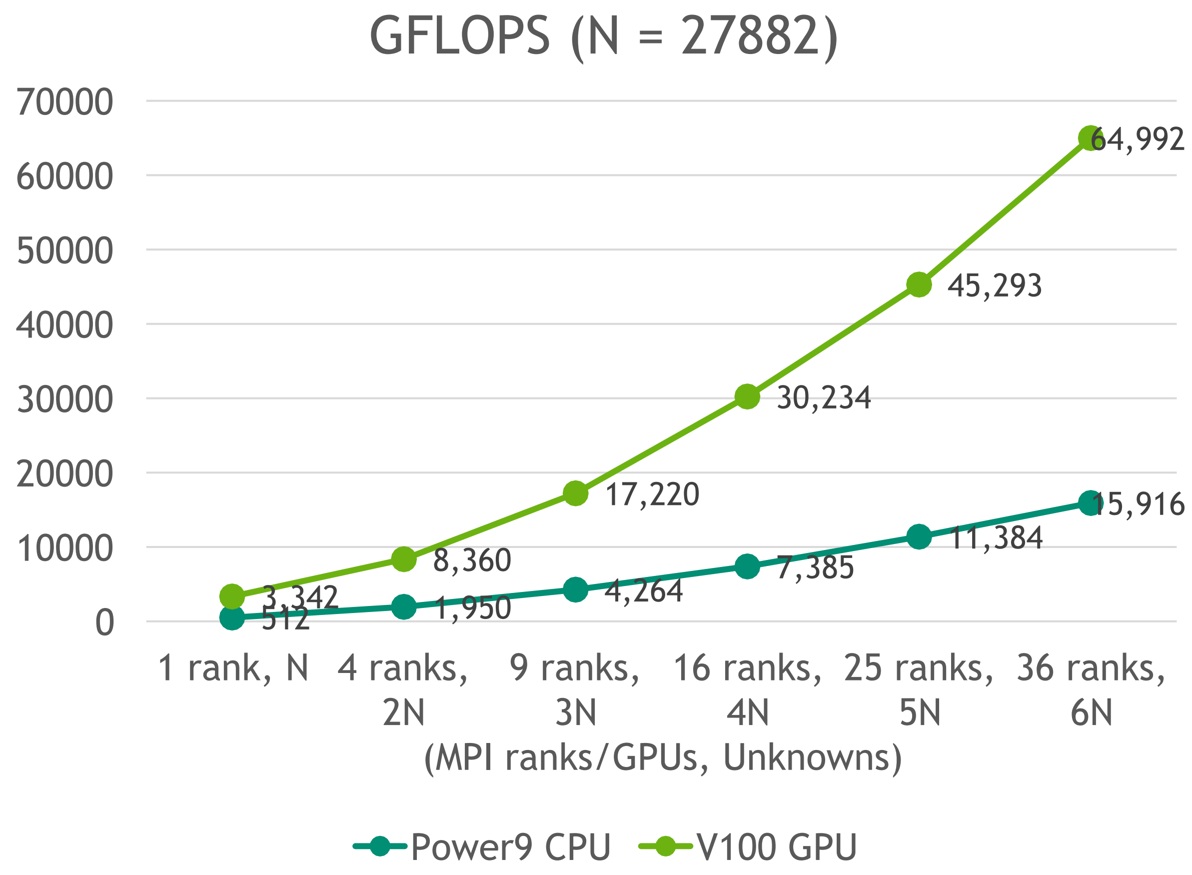

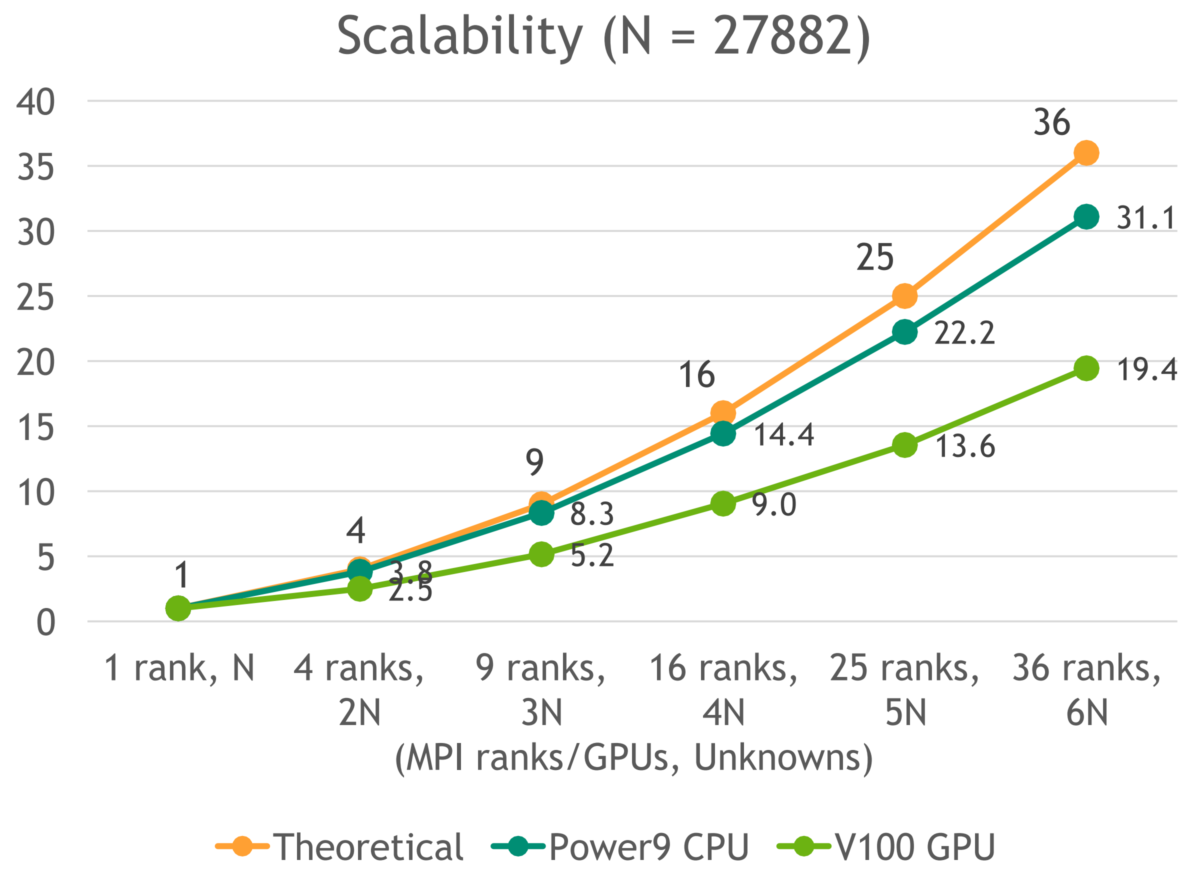

$$S = \frac{\text{GFLOPS(ranks/GPUs,unknowns)}}{\text{GFLOPS(1,1N)}}$$

where ranks/GPUs = 1, 4, 9, 16, 25, 36
and unknowns = 1N, 2N, 3N, 4N, 5N, 6N

Timing on Vortex

Gemma has been successfully run using the GPU's on Vortex

- Each node has 4 GPU's with 16 GByte memory per GPU
- Protocol is to run a MPI rank on each GPU
- A code change was made to GEMMA invoking the Kokkos::Initialize()

Gemma will be used on LLNL's Sierra for more formal testing in the near future

Problem	Unknowns	Nodes	GPUs / Node	Fill Time(s)	Solve Time(s)
Sphere	27882	1	4	1459.6	25.2
Sphere	27882	2	2	1455.0	25.3
Vfy 218	58383	2	4	3391.0	63.3
Almond	112272	4	4	5916.8	141.5

Conclusion

Gemma's fill algorithm and Adelus solver are scaling well, but there is room for improvement and testing on other architectures.

Bibliography

R. Harrington, *Time-Harmonic Electromagnetic Fields*, McGraw-Hill, New York, NY, 1961.

H. C. Edwards, C. R. Trott, and D. Sunderland, “Kokkos: Enabling manycore performance portability through polymorphic memory access patterns,” *Journal of Parallel and Distributed Computing*, 74 (2014), pp. 3202-3216.

M. A. Khayat and R. D. Wilton, ”An Improved Tansformation and Optimized Sampling Scheme for the Numerical Evaluation of Singular and Near-Singular Potentials,” *IEEE Antennas Wirel. Propag. Lett.*, 7 (2008), pp. 377 - 380.

S. Campione et al., *Preliminary Survey on the Effectiveness of an Electromagnetic Dampener to Improve System Shielding Effectiveness*, Sandia Technical Report SAND2018-10548, 2018.

S. Campione et al., *Penetration through Slots in Cylindrical Cavities Operating at Fundamental Cavity Modes*, in review.

Additional slides

Slot subcell model for capturing coupling into a cavity accurately

In free space, the thin slot equation is:

$$H_z^{>}(a,z)+\frac{1}{4}\left(\Delta Y_C \frac{d^2}{dz^2} I_m - \Delta Y_L I_m\right) = -H_z^{inc}(z), I_m = -2V$$

H_z = magnetic field
I_m = current
V = voltage
a = equivalent radius
Y_C, Y_L capture gaskets and wall loss

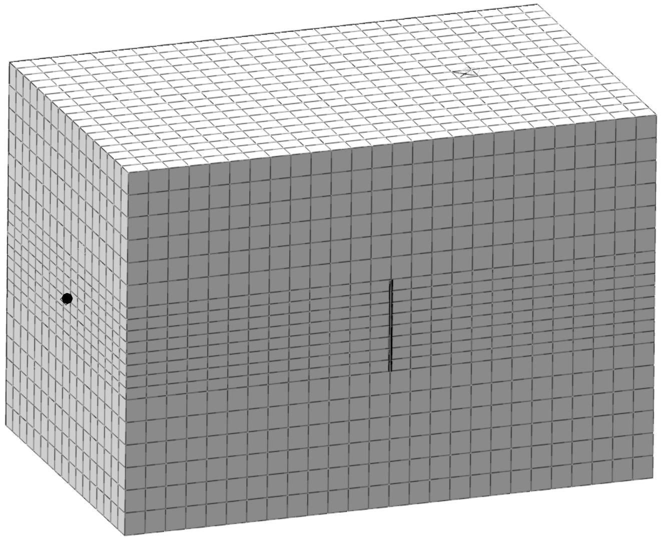

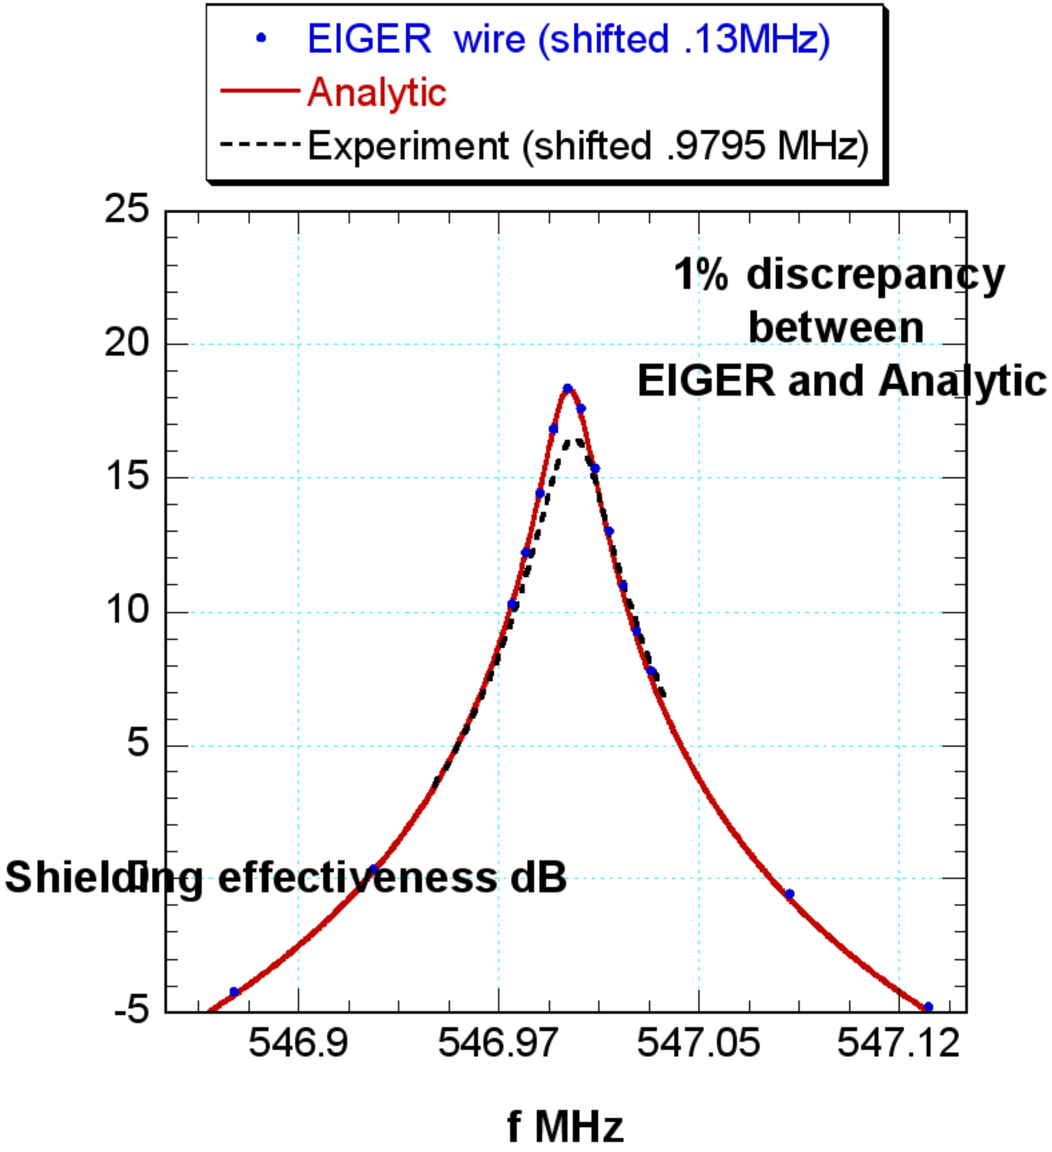

Cavity comparison with analytic and experiment

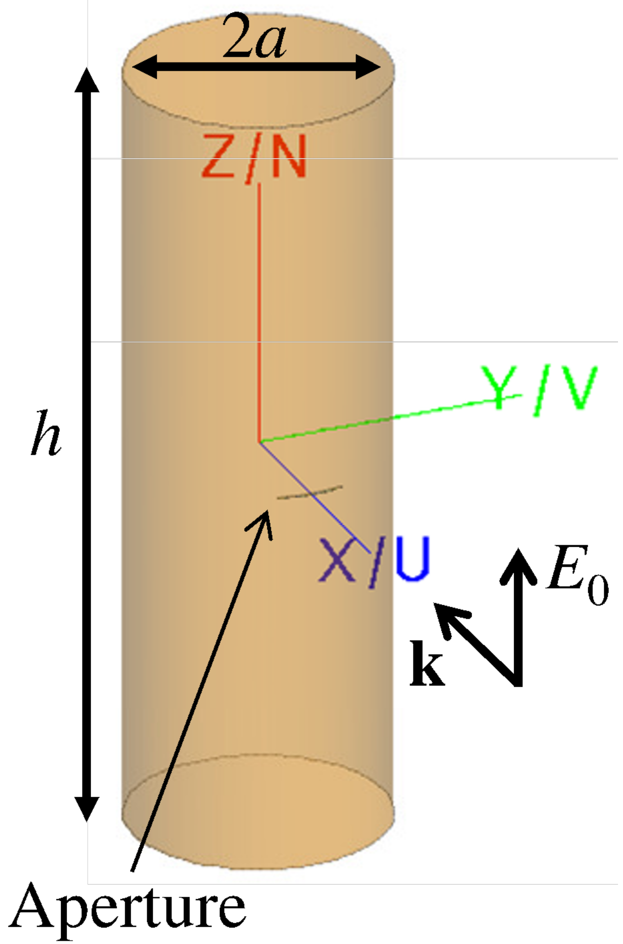

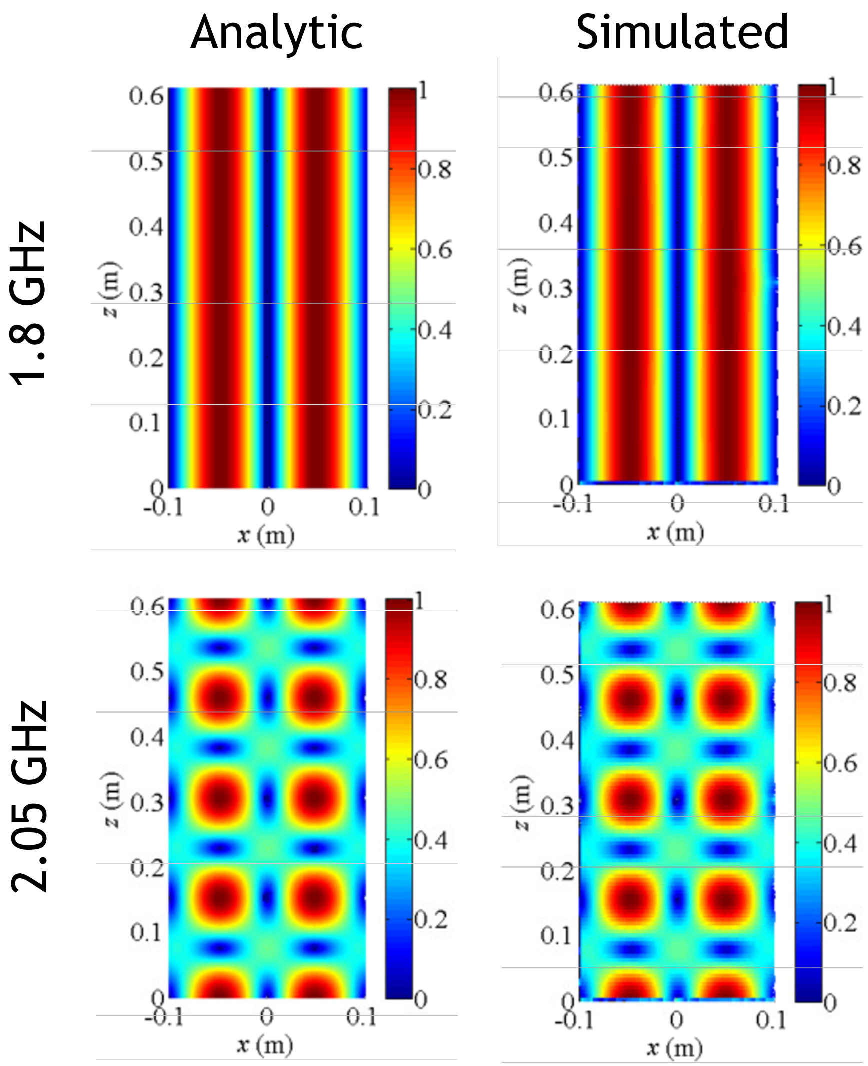

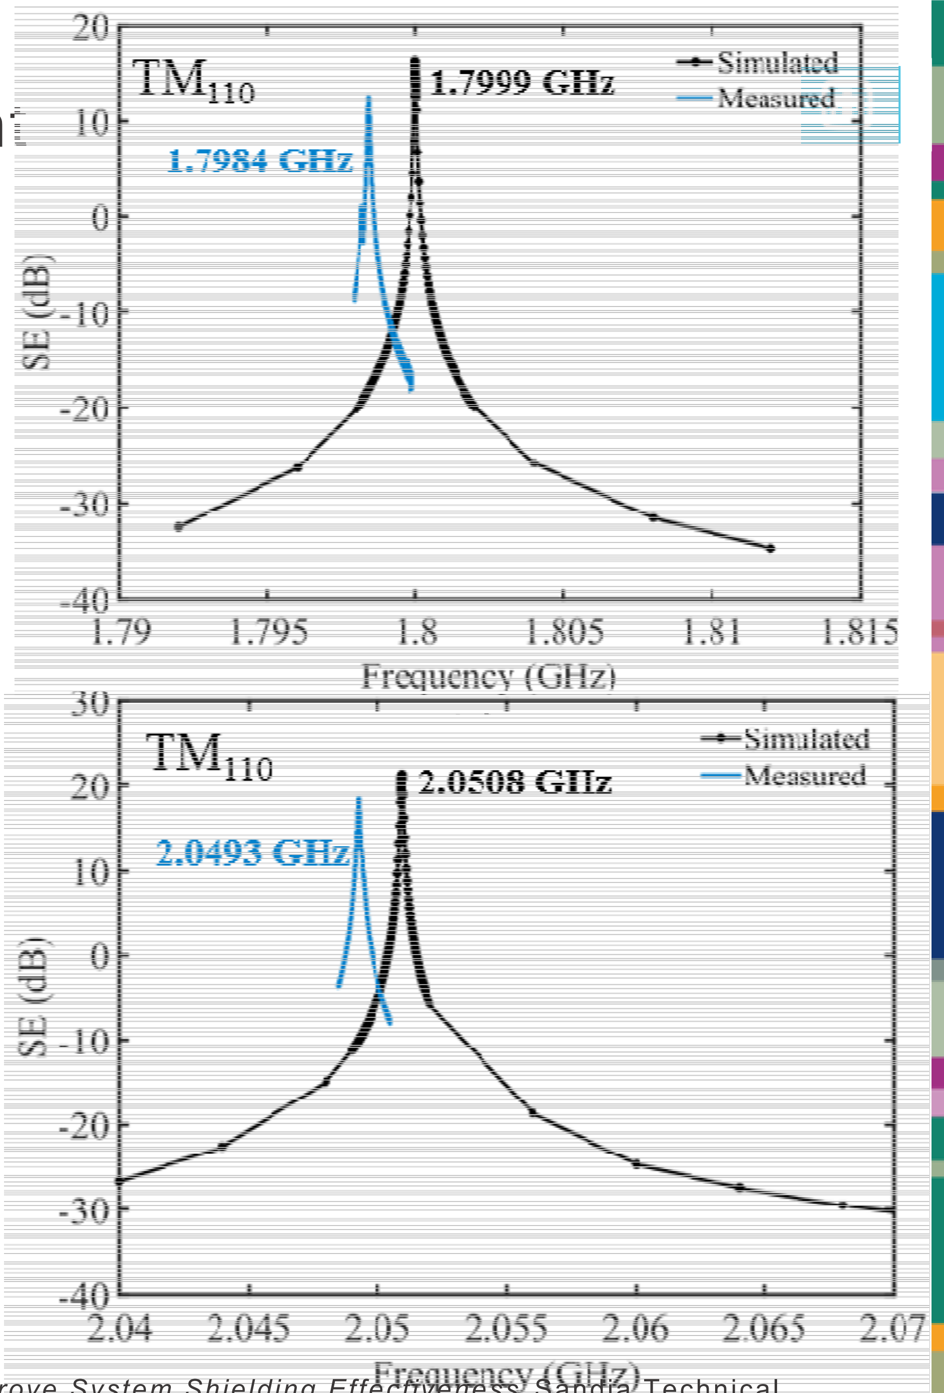

S. Campione et al., *Preliminary Survey on the Effectiveness of an Electromagnetic Dampener to Improve System Shielding Effectiveness*, Sandia Technical Report SAND2018-10548, 2018.