Large and Massively-Parallel Image Reconstruction Accelerated with the Multilevel Fast Multipole Algorithm

Mert Hidayetoğlu, Carl Pearson, Izzat El Hajj, Weng Cho Chew, Levent Gürel, and Wen-Mei Hwu
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

ILLINOIS 150 1867-2017

ECE ILLINOIS
Department of Electrical and Computer Engineering

BLUE WATERS

Inverse-Scattering Problems

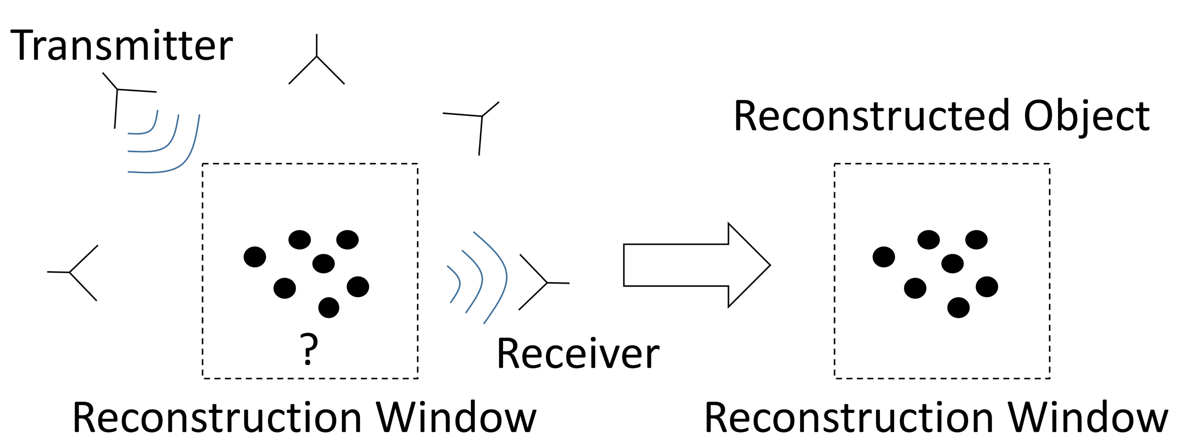

Full-wave methods do not impose any approximation by solving the wave equation with its full glory.

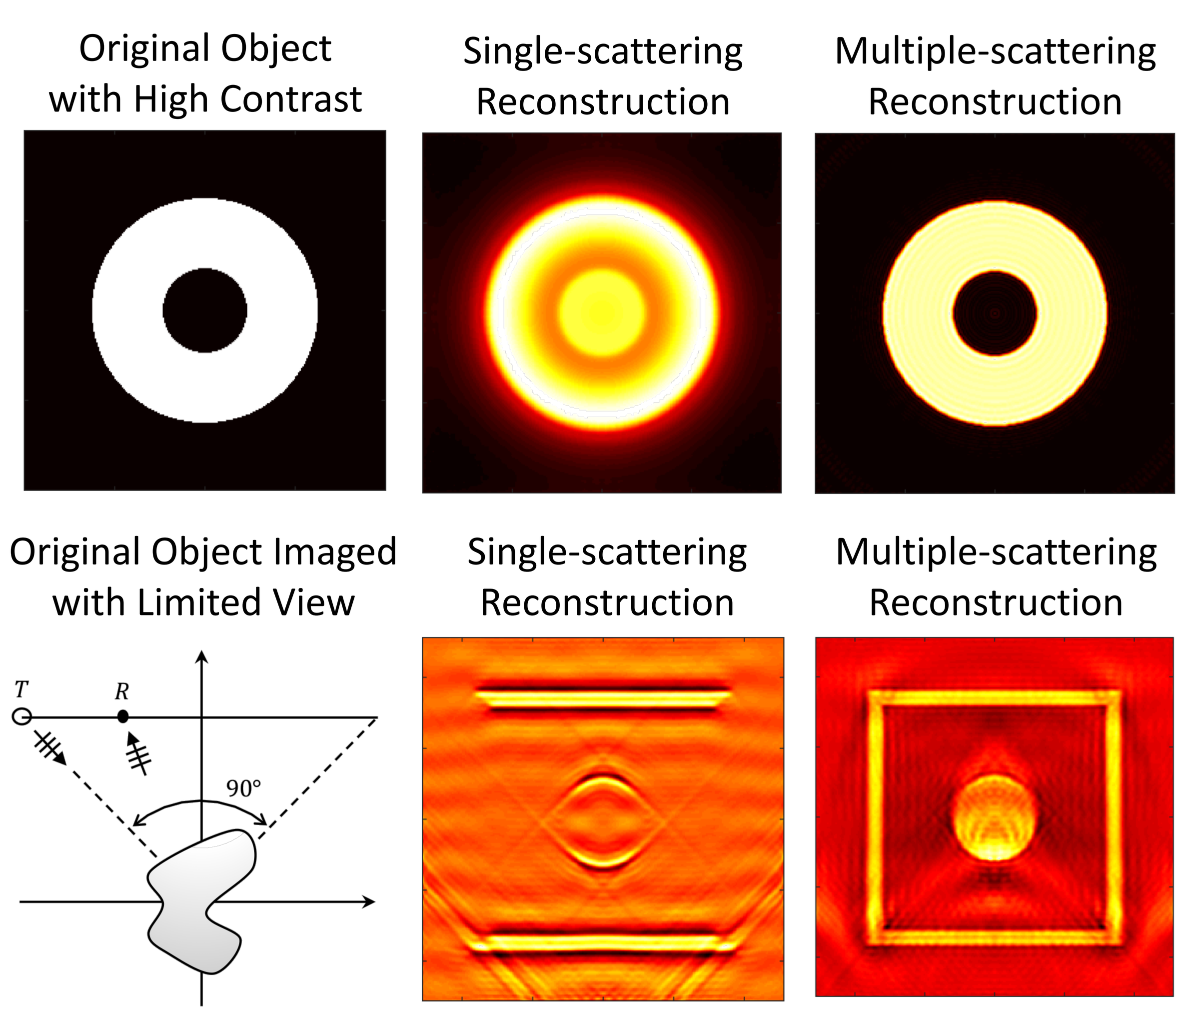

Nonlinear Optimization

There is not a single way to do this, however, we found out nonlinear conjugate-gradient method is efficient for this algorithm.

- Gradient-Descent
  $-\nabla\Phi = -\boldsymbol{F}^H\boldsymbol{b}$
- Conjugate-Gradient

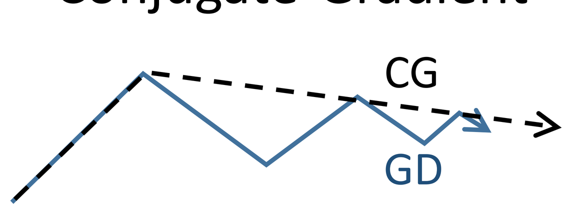

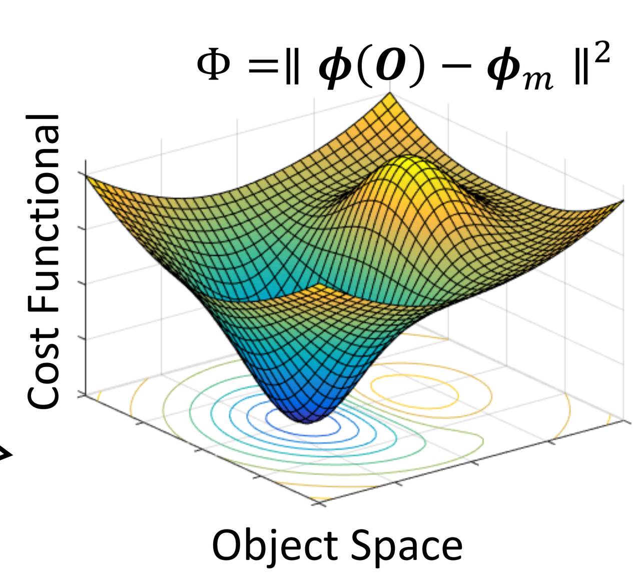

- Newton-Type Methods
  $[\mu^2\boldsymbol{I} + \boldsymbol{F}^H\boldsymbol{F}]\delta\boldsymbol{O} = \nabla\Phi$

$\boldsymbol{F}$: Functional Derivative Operator

Multilevel Fast Multipole Algorithm (MLFMA) Schematic

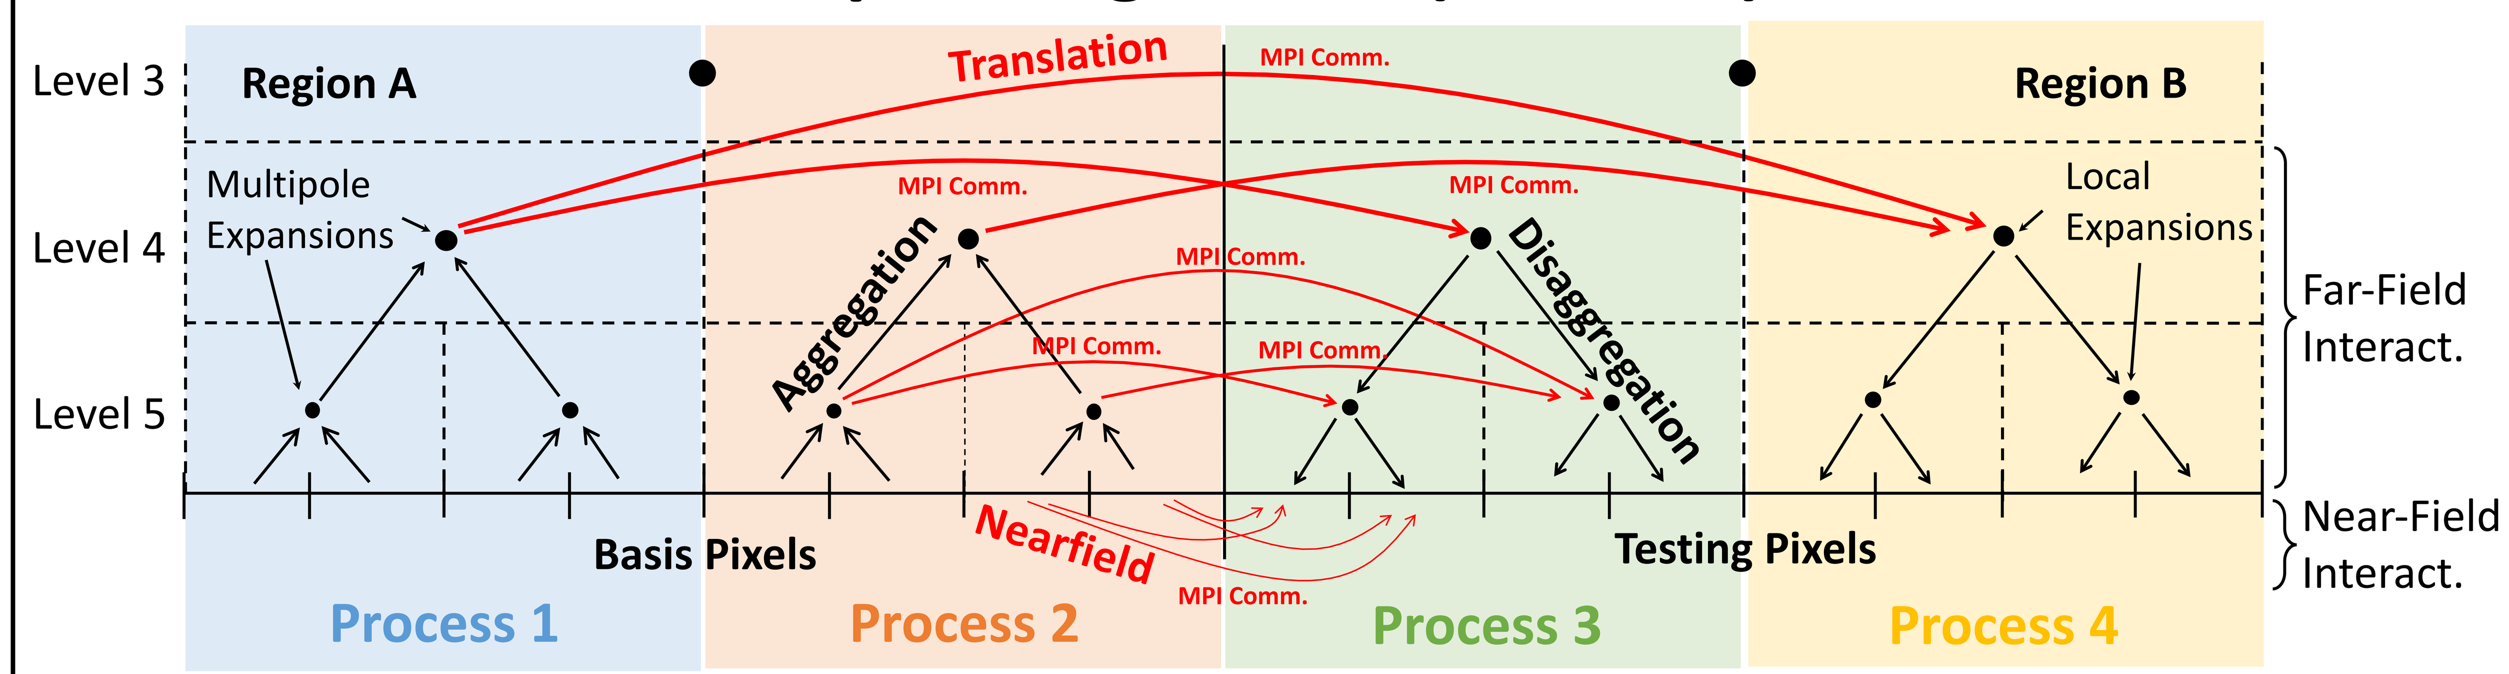

MLFMA Provides fast solutions of forward-scattering problems with O(N) computational complexity.

Distorted-Born Approximation

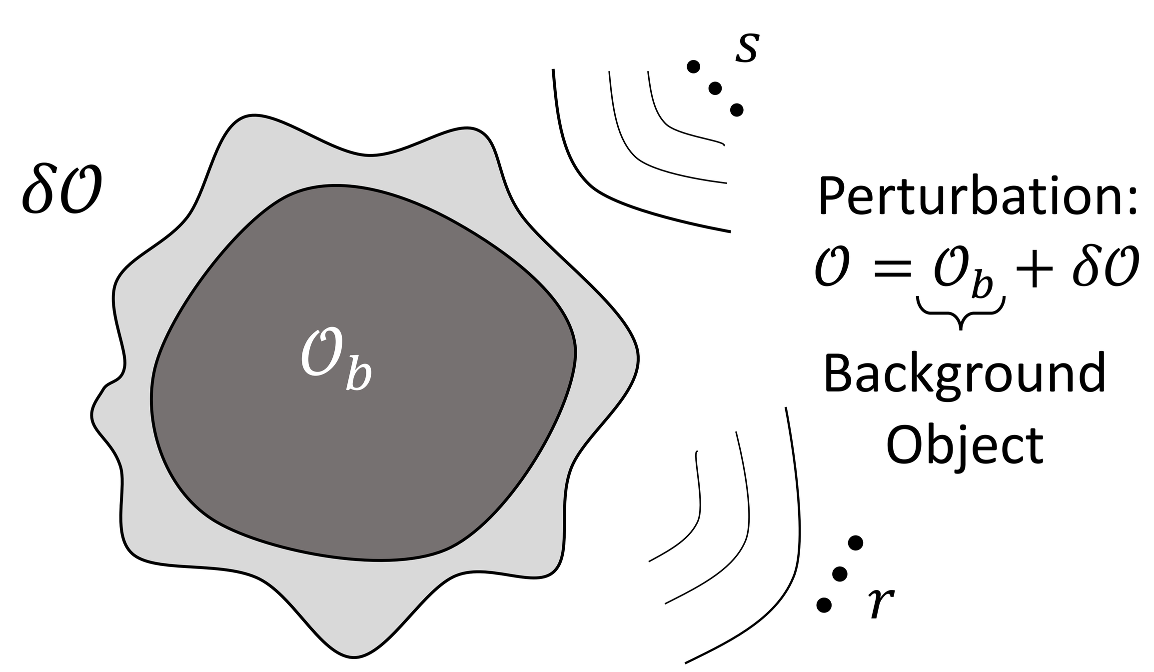

Scattering Equation: $\phi = \phi_0 + \mathcal{G}_0 \mathcal{O} \phi$

Variational Equation: $\delta\phi = \mathcal{G}_b \delta\mathcal{O} \phi_b + \underbrace{\mathcal{G}_b \delta\mathcal{O} \delta\phi}$

Distorted-Born approximation provides a semi-analytical way to find functional derivatives.

Higher-order Variations (neglected)

$\delta\phi \approx \delta\phi^{(1)} = \mathcal{G}_b \delta\mathcal{O} \phi_b = \mathcal{G}_b \phi_b \delta\mathcal{O}$

$\delta\phi^{(1)} = \mathcal{F}\delta\mathcal{O}$

$$\boldsymbol{F} = \boldsymbol{G}_R\{\boldsymbol{I} + \boldsymbol{O}_b\overbrace{[\boldsymbol{I} - \boldsymbol{G}_0\boldsymbol{O}_b]^{-1}\boldsymbol{G}_0}^{\text{MLFMA}}\} \underbrace{\mathrm{diag}\{[\boldsymbol{I} - \boldsymbol{G}_0\boldsymbol{O}_b]^{-1}\boldsymbol{G}_T\boldsymbol{s}\}}_{\text{MLFMA}}$$

Discretization with a subspace projection method.

$\overline{\boldsymbol{G}}_0$: Dense, $N\times N$
$\overline{\boldsymbol{G}}_T$: Dense, $N\times T$
$\overline{\boldsymbol{G}}_R$: Dense, $R\times N$
$\overline{\boldsymbol{O}}_b$: Diagonal, $N\times N$
$\overline{\boldsymbol{I}}$: Diagonal, $N\times N$

Massive Parallelization

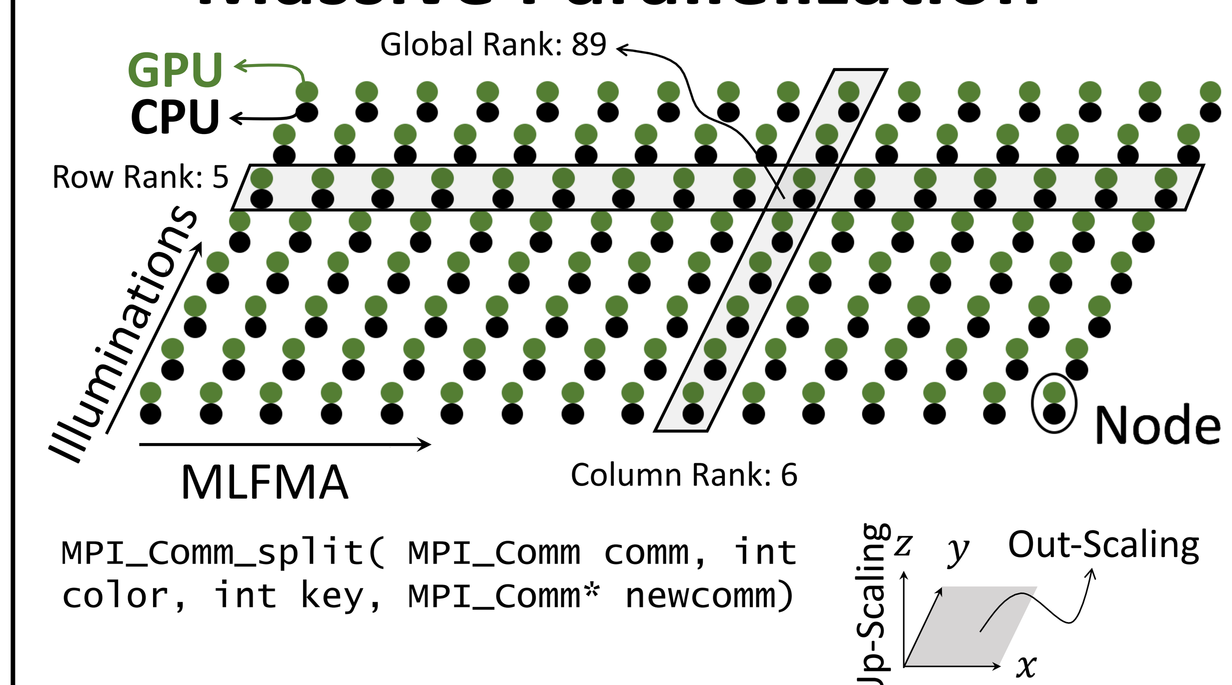


```
MPI_Comm_split( MPI_Comm comm, int
color, int key, MPI_Comm* newcomm)
```

- **Hierarchical parallelization provides high granularity for out-scaling.**
- **Each MPI process employs OpenMP and CUDA threads to employ multi-code CPUs and GPU for up-scaling.**

Large Reconstruction

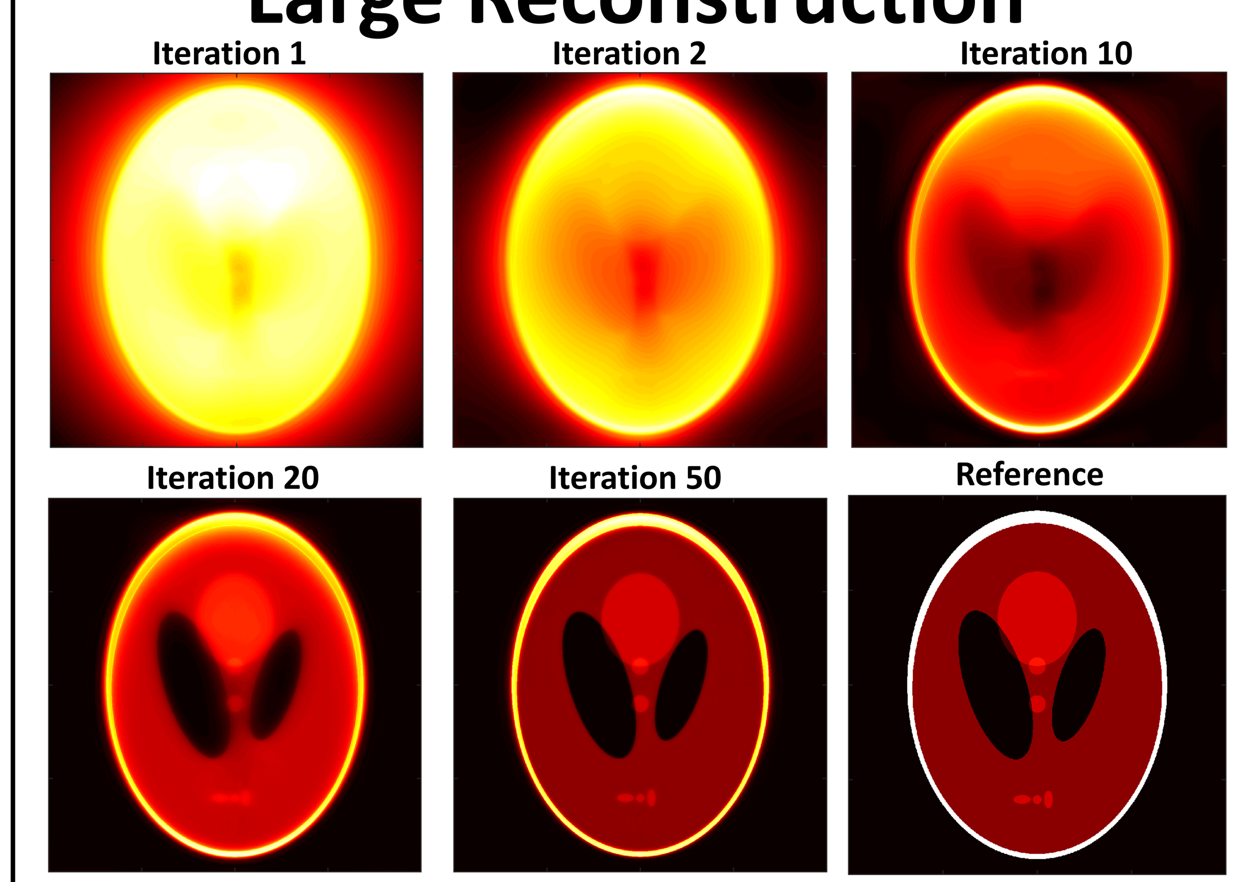

Strong Scaling

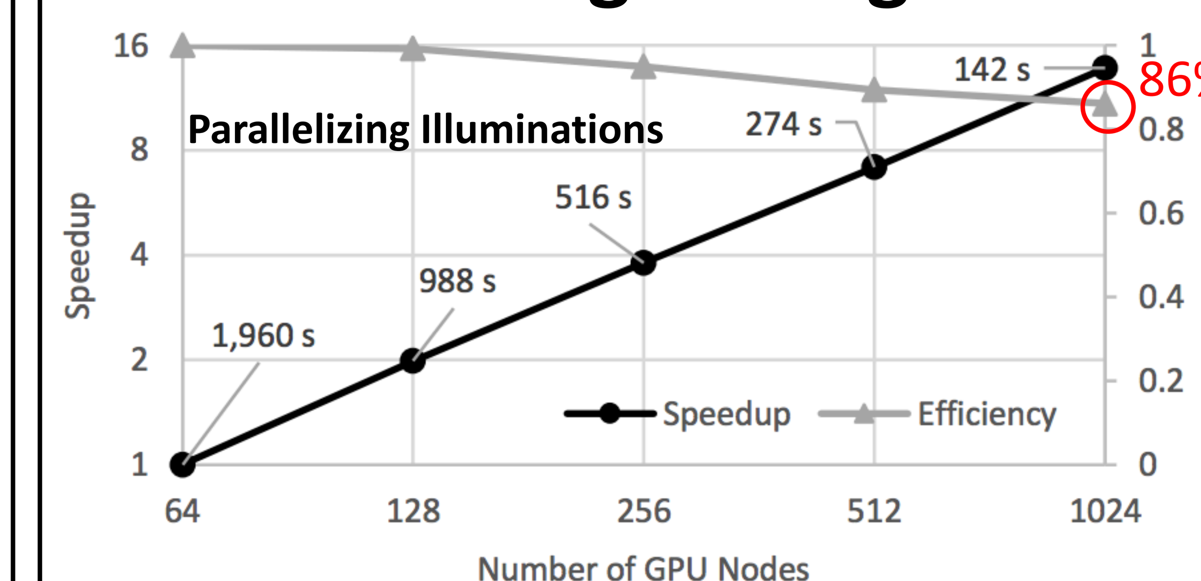

Parallelizing illuminations is almost perfectly efficient due to the independent nature of forward solutions.

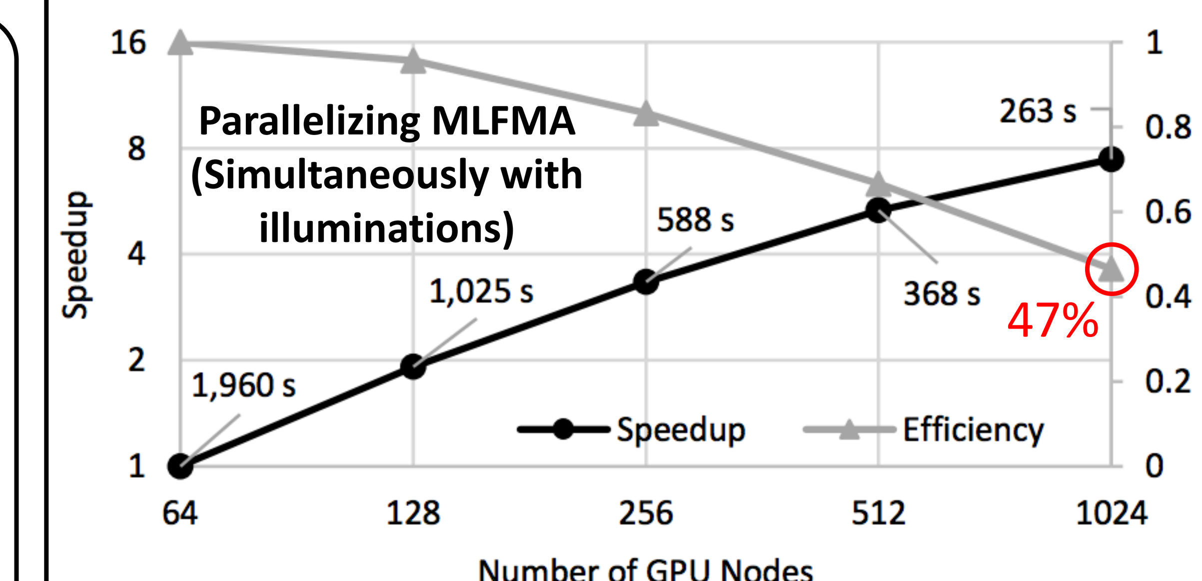

MLFMA scaling is not perfect due to the MPI communications and underutilization of GPUs.

4,096 GPU Nodes: 40.2 sec.	1 GPU Node: 1.5 day
4,096 CPU Nodes: 2.4 min.	1 CPU Node: 6 days
64 CPU Nodes: 2.4 hours	Sequential: 85 days!

Conclusions

- Exploiting multiple-scattering is effective, but has a huge computational burden.
- MLFMA provides algorithmic speedup.
- A hierarchical parallelization strategy improves the scalability on large supercomputers.
- GPUs and multi-core CPUs provides massively-parallel reconstructions.

See More Results & Animations:

Future Plans

- **Further algorithmic improvements like compressive-sensing.**
- **2.5-D and 3-D extensions (not trivial because of computational requirements).**
- **Real measurement data will be used for imaging. Not trivial because of noise, calibration, etc.**