

Reducing Parallel Communication Costs in Algebraic Multigrid

Amanda Bienz
William D. Gropp
Luke Olson


University of Illinois at Urbana-Champaign

What is Algebraic Multigrid (AMG)?

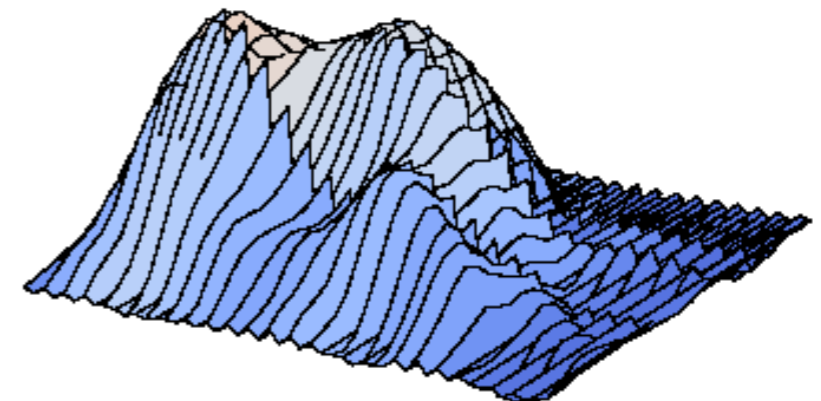
Iterative method for sparse linear systems



→
Relaxation
(Jacobi, Gauss-Seidel)



↓
Restrict to coarser grid



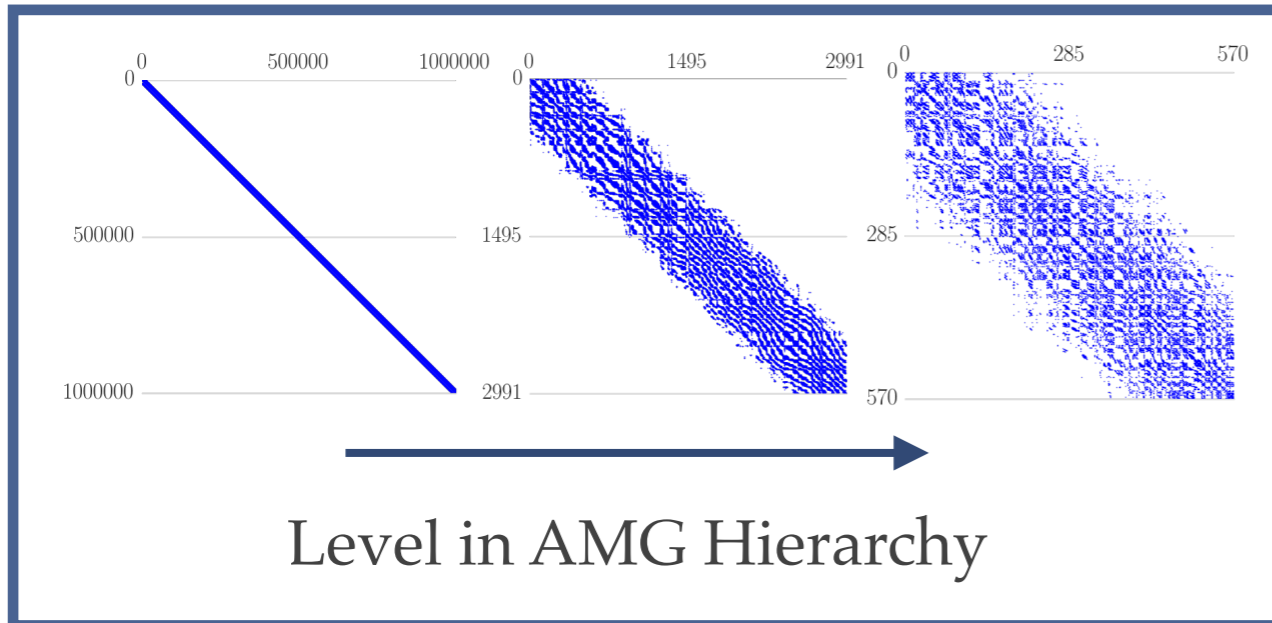
Main Operations:

Matrix-Matrix Multiplication (Both Sparse)

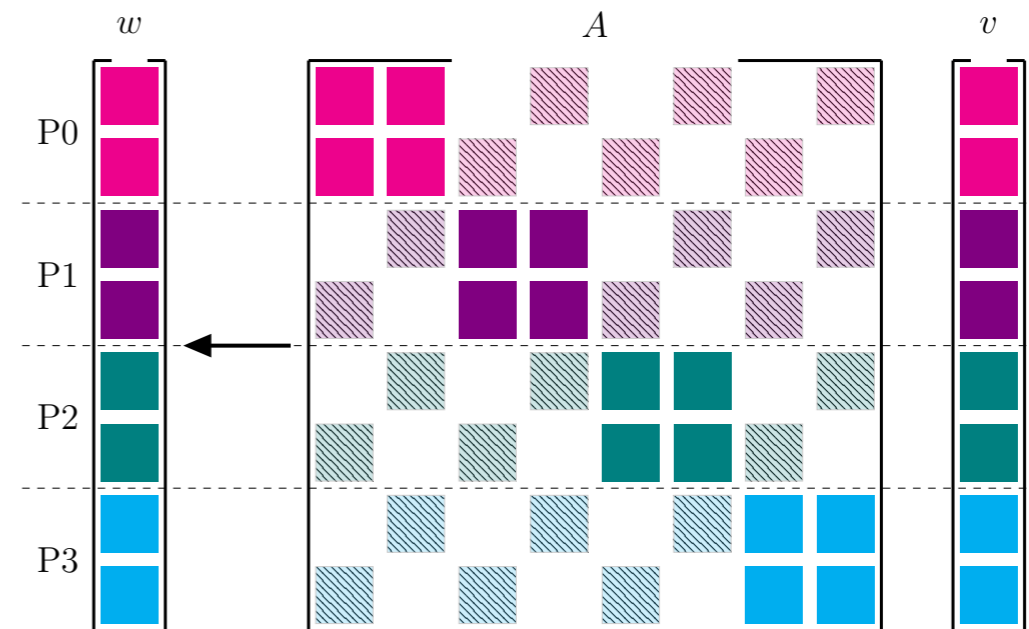
Sparse Matrix-Vector Multiplication

What's the Problem?

Increased density on coarse levels



- Linear system split across four processors
- **Solid blocks:** fully local portion of SpMV
 - **Patterned blocks:** require communication



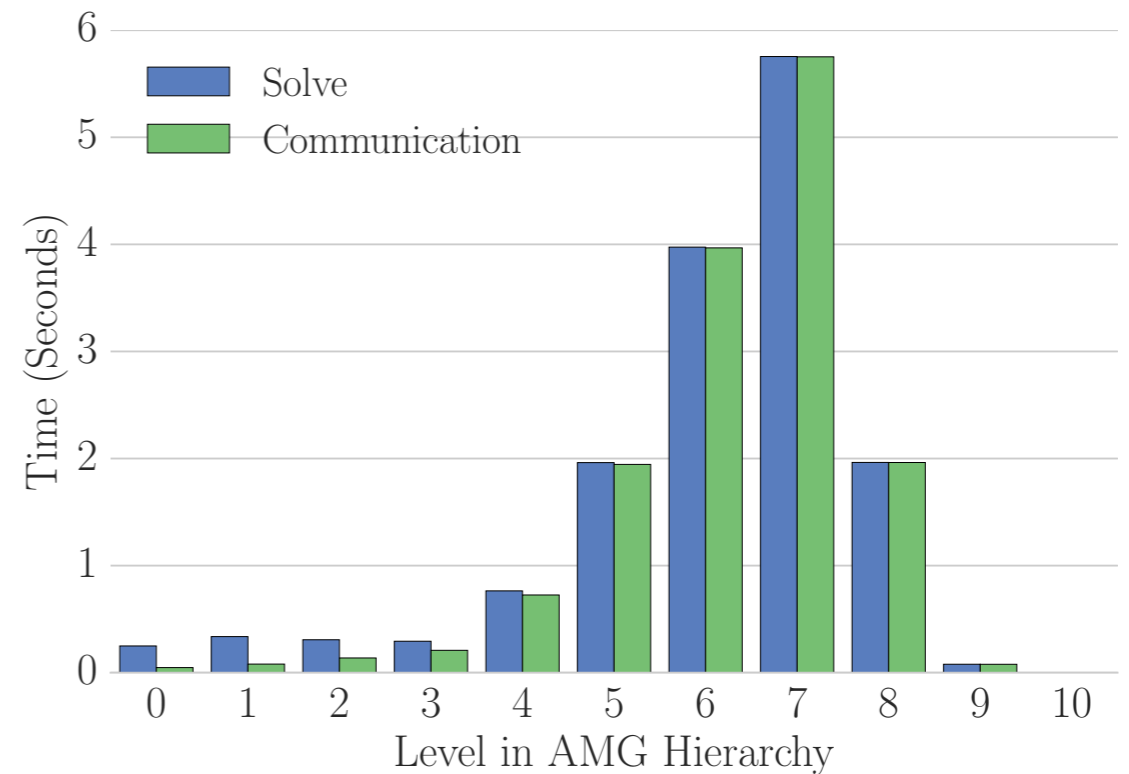
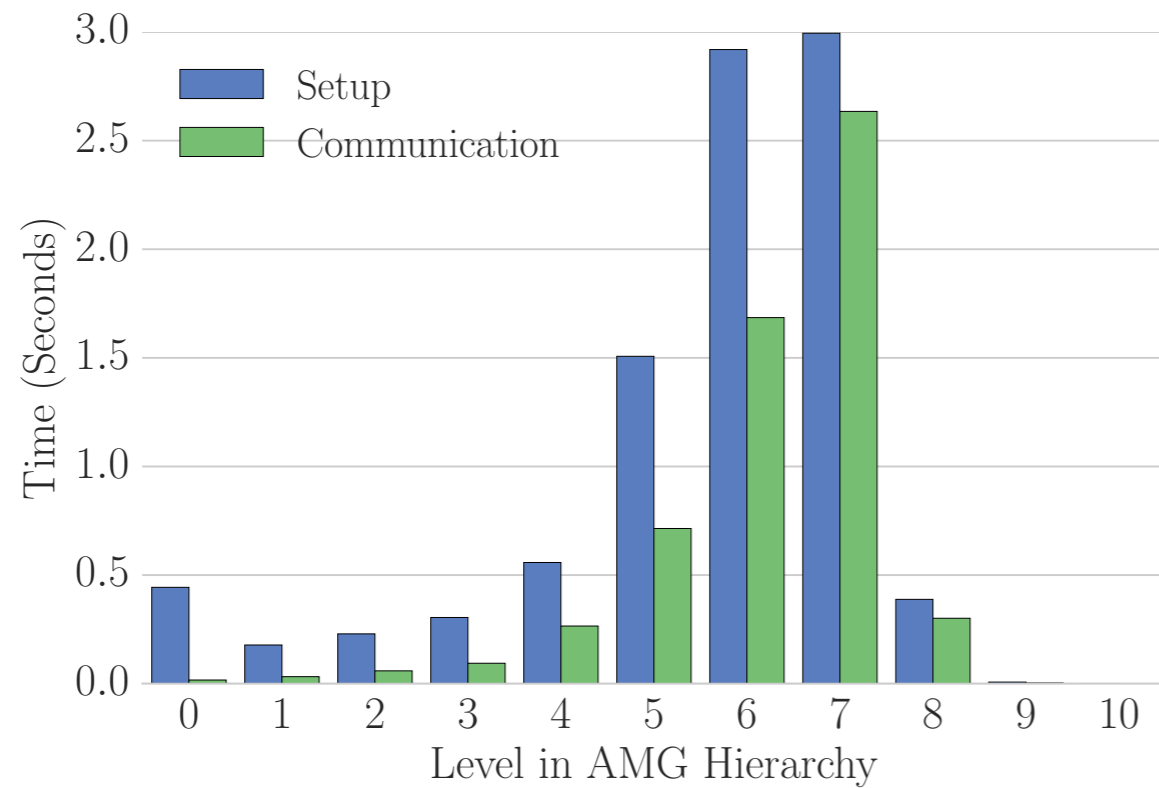
More non-zeros per row



More Communication

What's the Problem?

Large increase in time on coarse levels



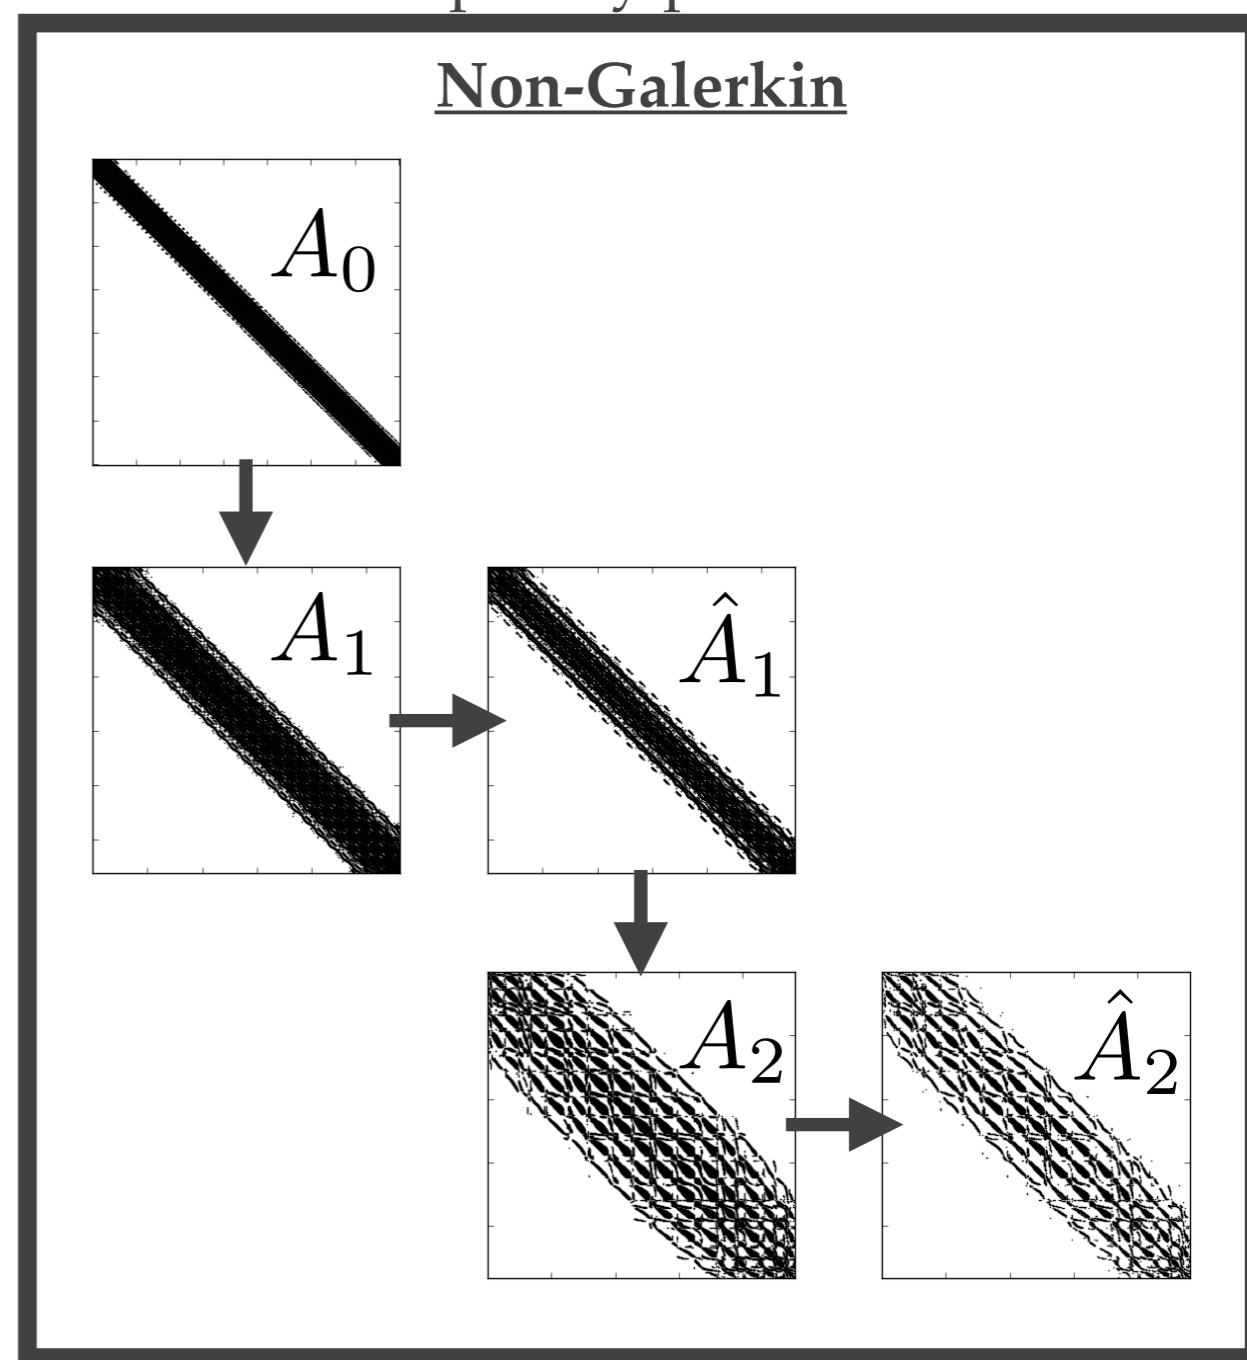

Want to reduce communication costs

Reducing Communication Costs

- ❖ **Alter Algorithm:**
 - ❖ Reduce amount of communication
 - ❖ Possible increase in convergence factor
- ❖ **Parallel Implementation:**
 - ❖ Leave algorithm unchanged
 - ❖ Reduce cost associated with communication



Altering Algorithm to Reduce NNZ/Row

- Based on method of non-Galerkin coarse grids
- Systematically remove **small** entries outside of minimal sparsity pattern



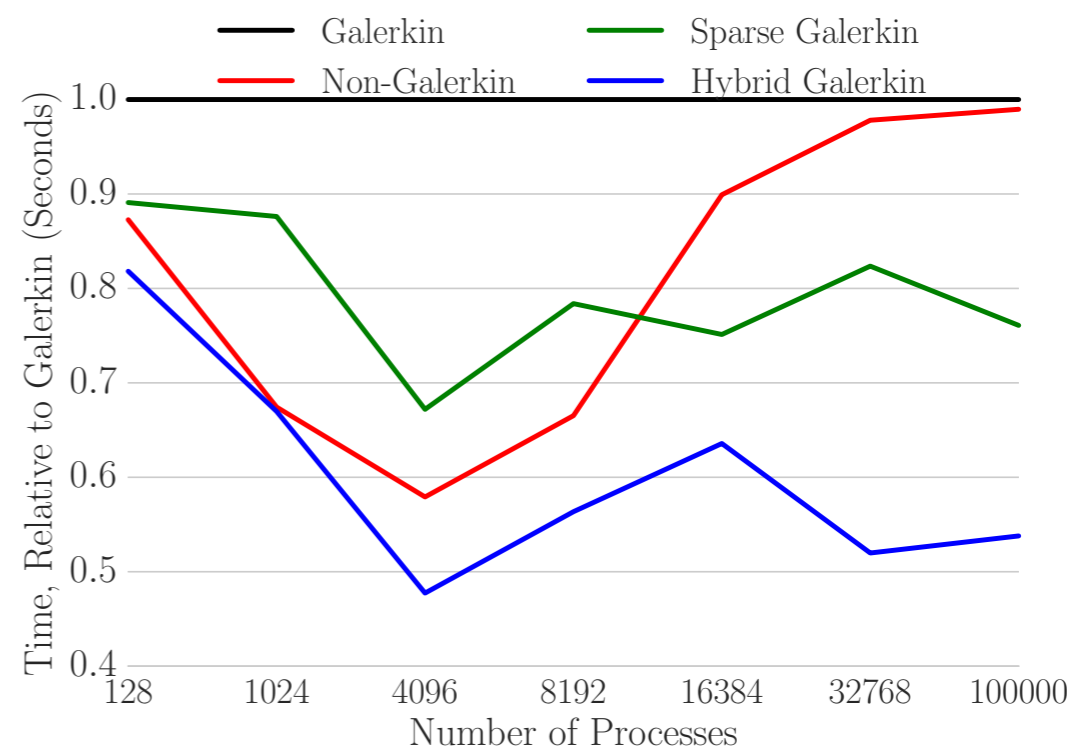
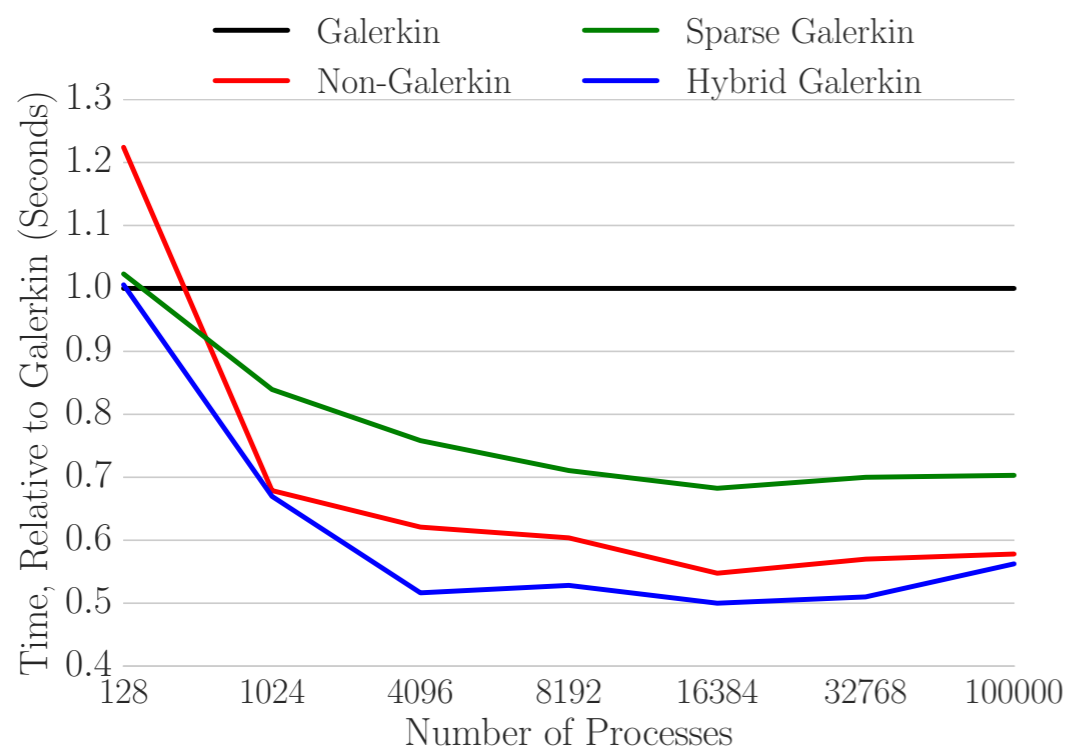
Altering Algorithm to Reduce NNZ/Row

- Based on method of non-Galerkin coarse grids
- Systematically remove **small** entries outside of minimal sparsity pattern



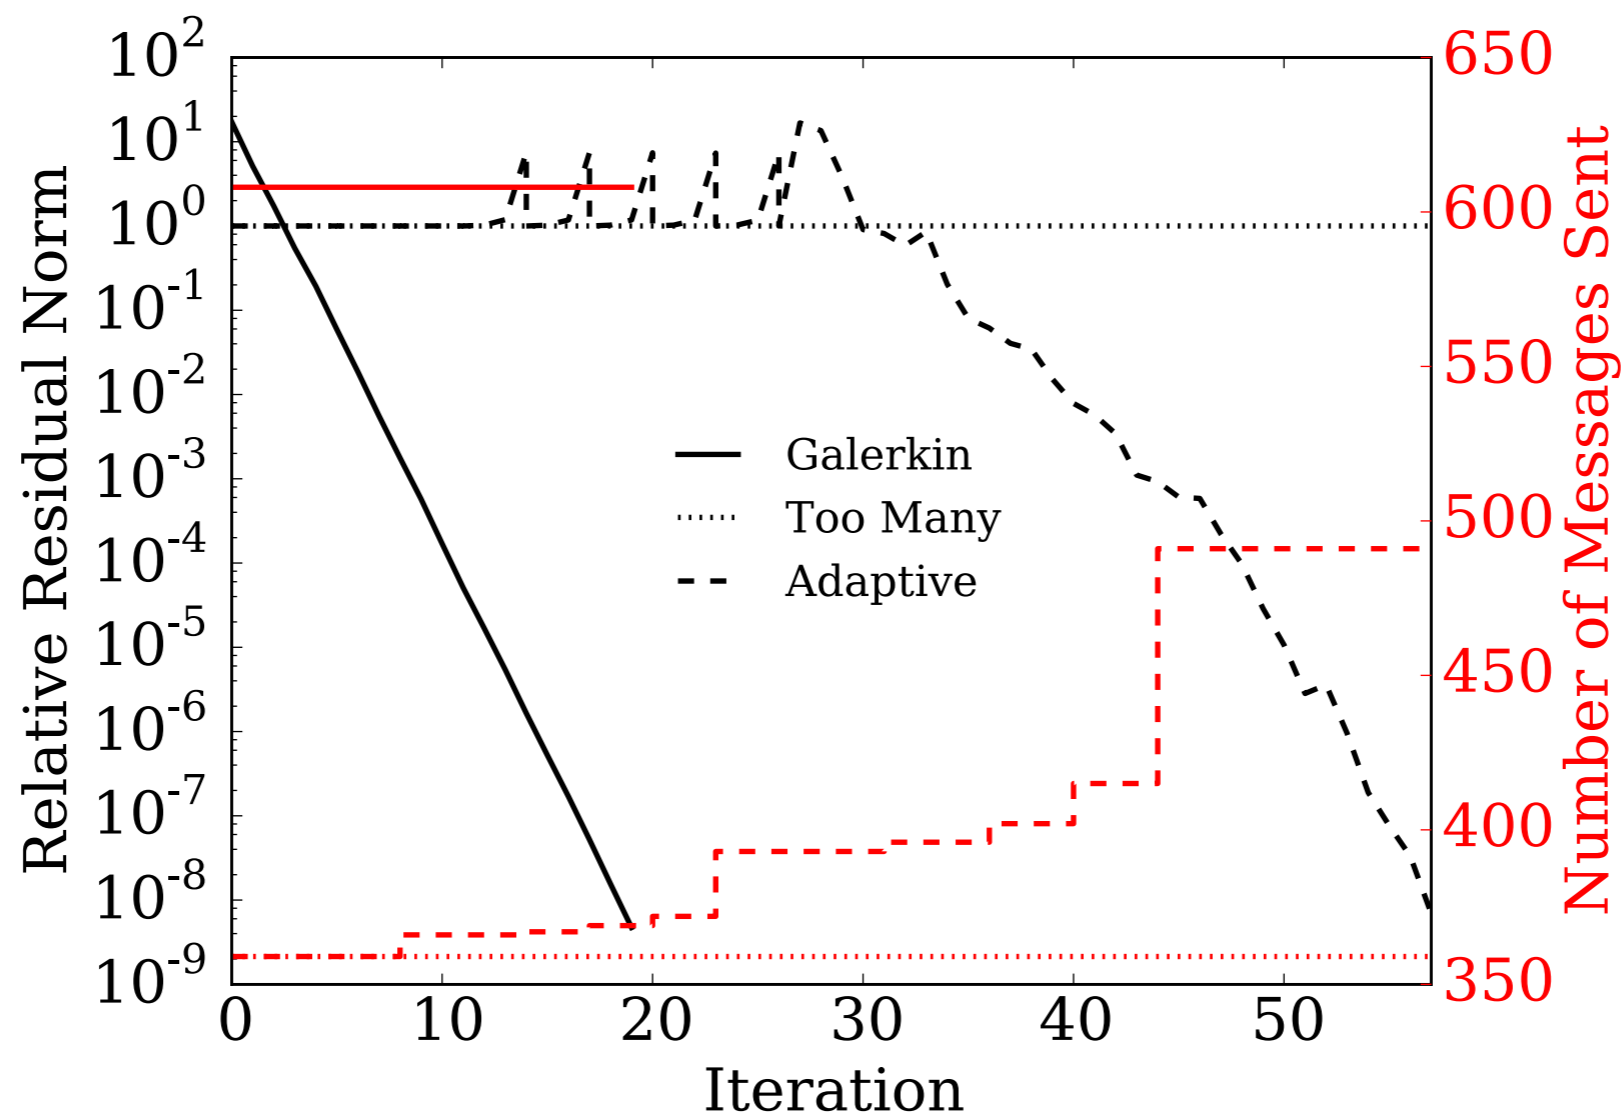
Altering Algorithm to Reduce NNZ/Row

Time required to solve 2D rotated anisotropic diffusion
(relative to standard Galerkin AMG) on Blue Waters



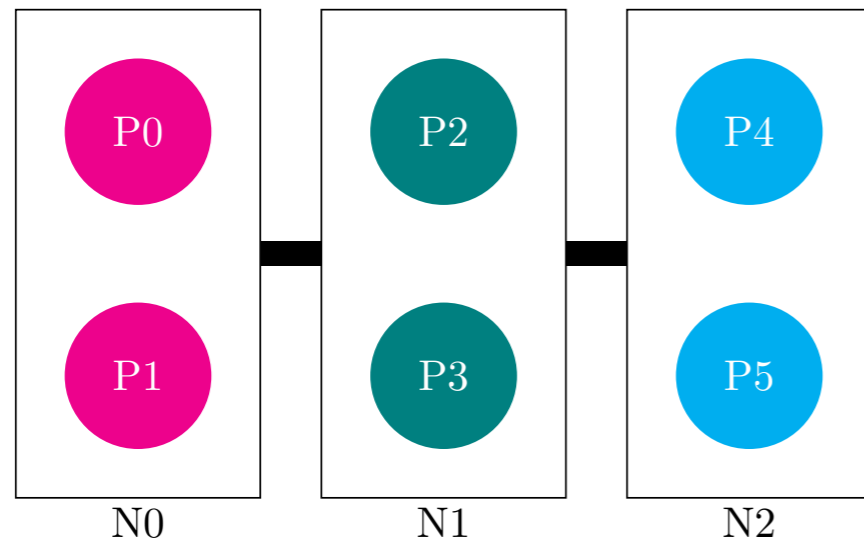
Altering Algorithm to Reduce NNZ/Row

Adaptively update sparsity to control convergence




A. Bienz, R.D. Falgout, W. Gropp, L.N. Olson and J.B. Schroder, Reducing Parallel Communication in Algebraic Multigrid Through Sparsification. SIAM J. Sci. Comput., pp. 332-357. 38 (2016).

Parallel Implementation: Where is the Cost?



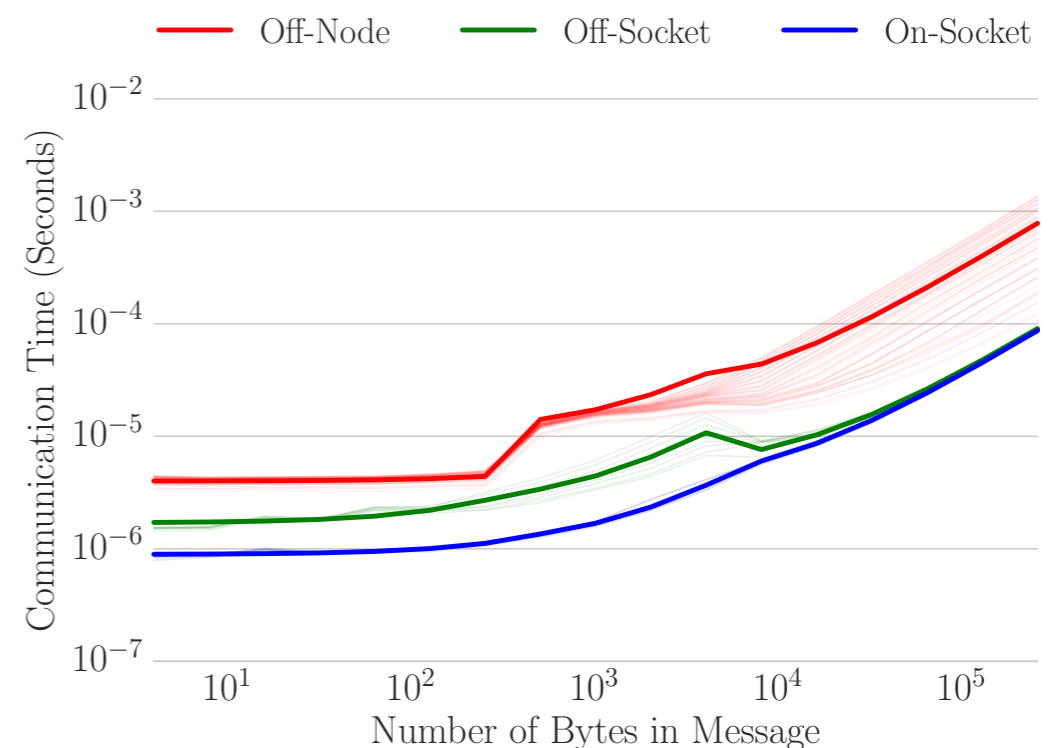
Six processes distributed across three nodes



Linear system distributed across the processes

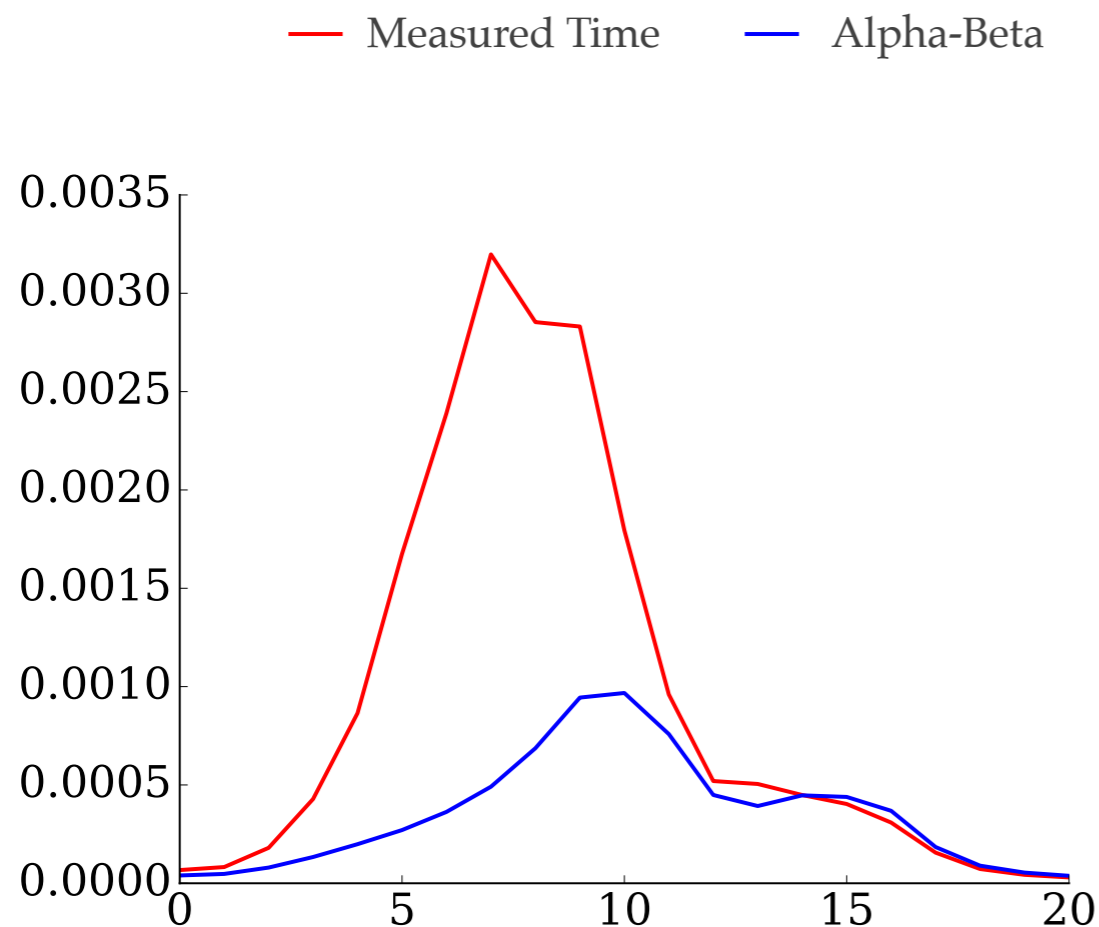
Many different ways to map rank to node:

1. Set environment variable on Cray, BGQ
2. Topology-aware methods
 - Nodecomm (William D. Gropp)
 - Topomgr (Ahbinav Bhatele)

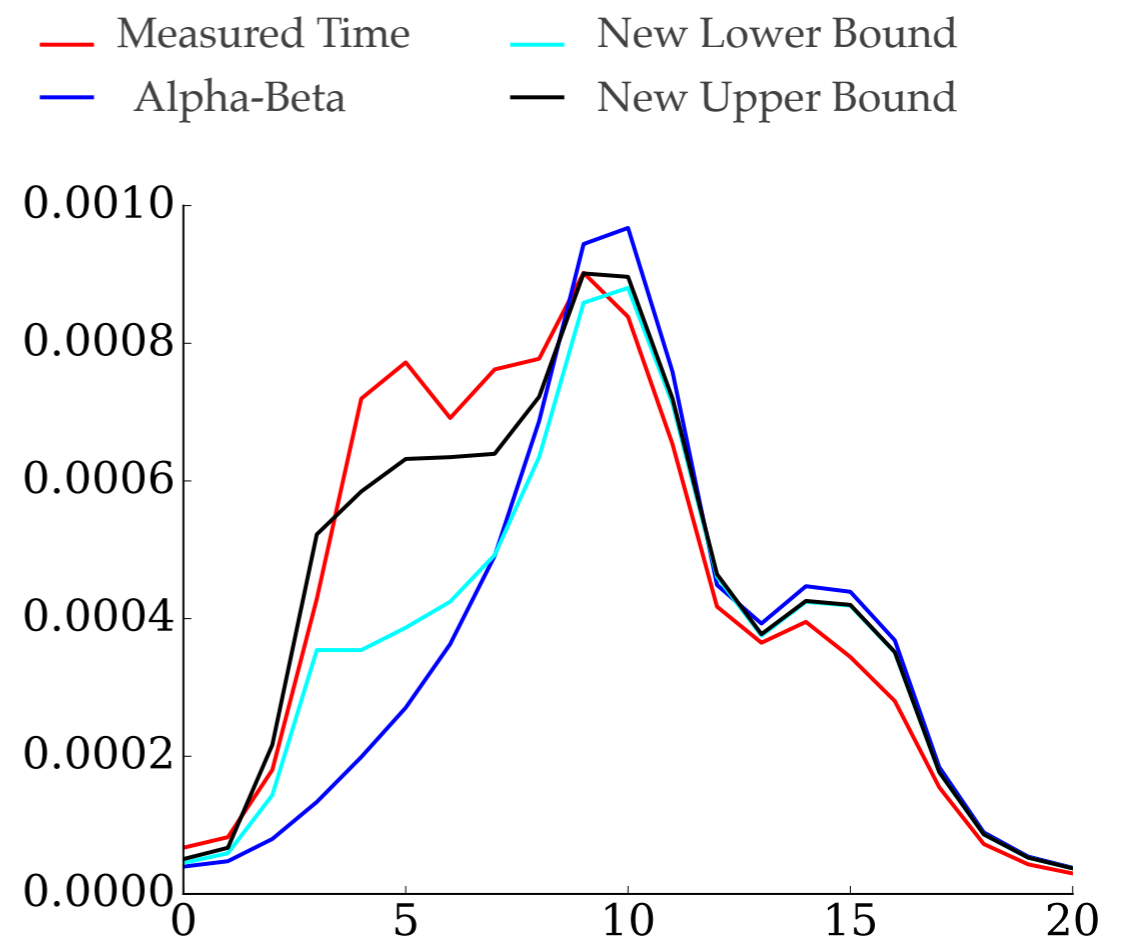


Parallel Implementation: Where is the Cost?

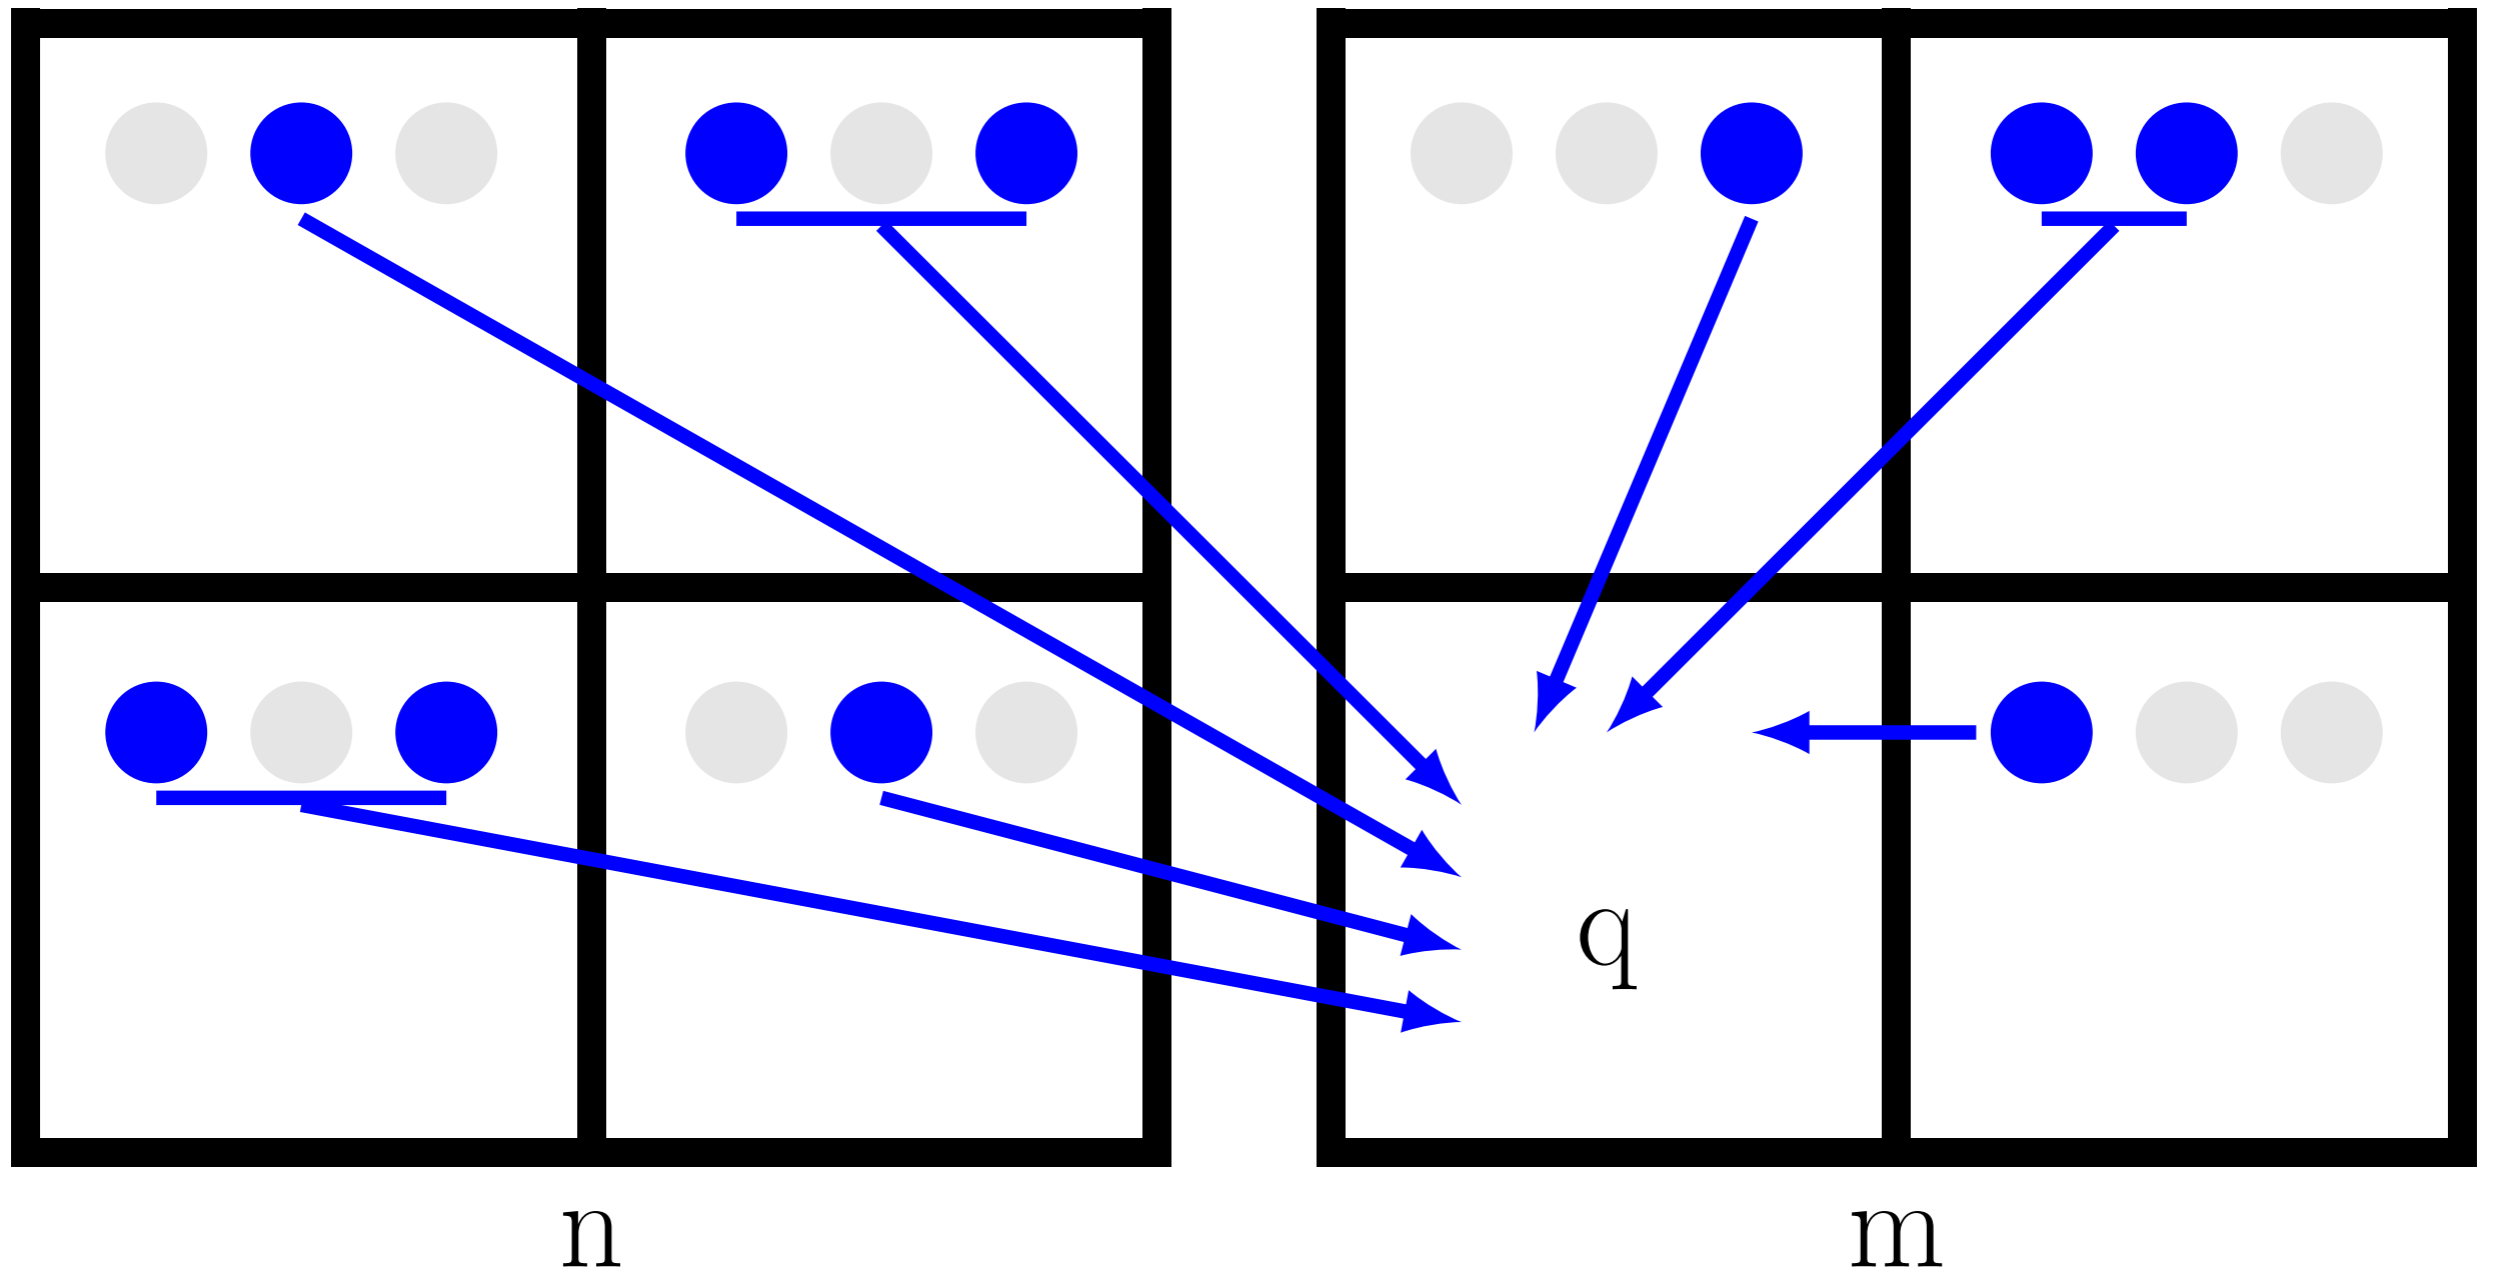
Current alpha-beta models do not accurately capture cost of coarse levels




Can accurately model cost only if network contention is removed



Standard Communication

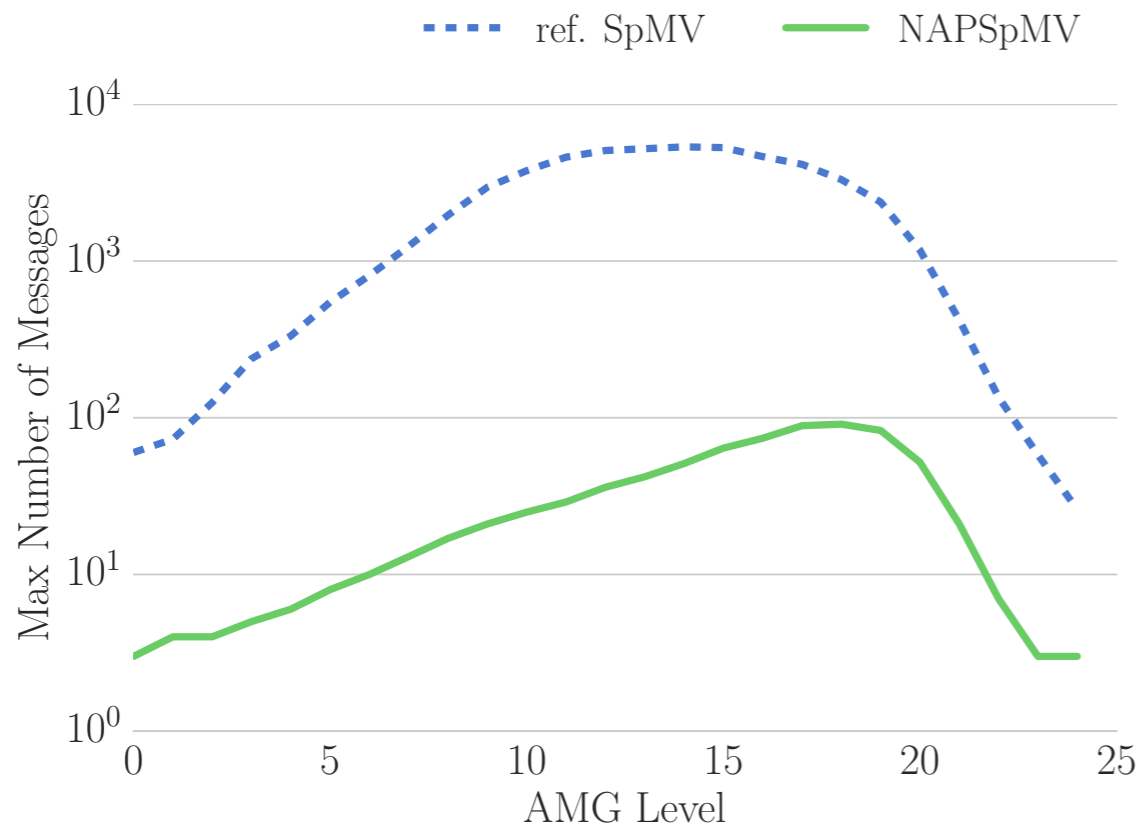


Node Aware Parallel (NAP) Communication




Linear Elasticity Hierarchy

Maximum number of messages sent from any process on 32,768 processes




Maximum number of
off-node messages




Maximum number of
on-node messages

Linear Elasticity Hierarchy

Maximum size of all messages sent from any process on 32,768 processes



Maximum size of
off-node messages



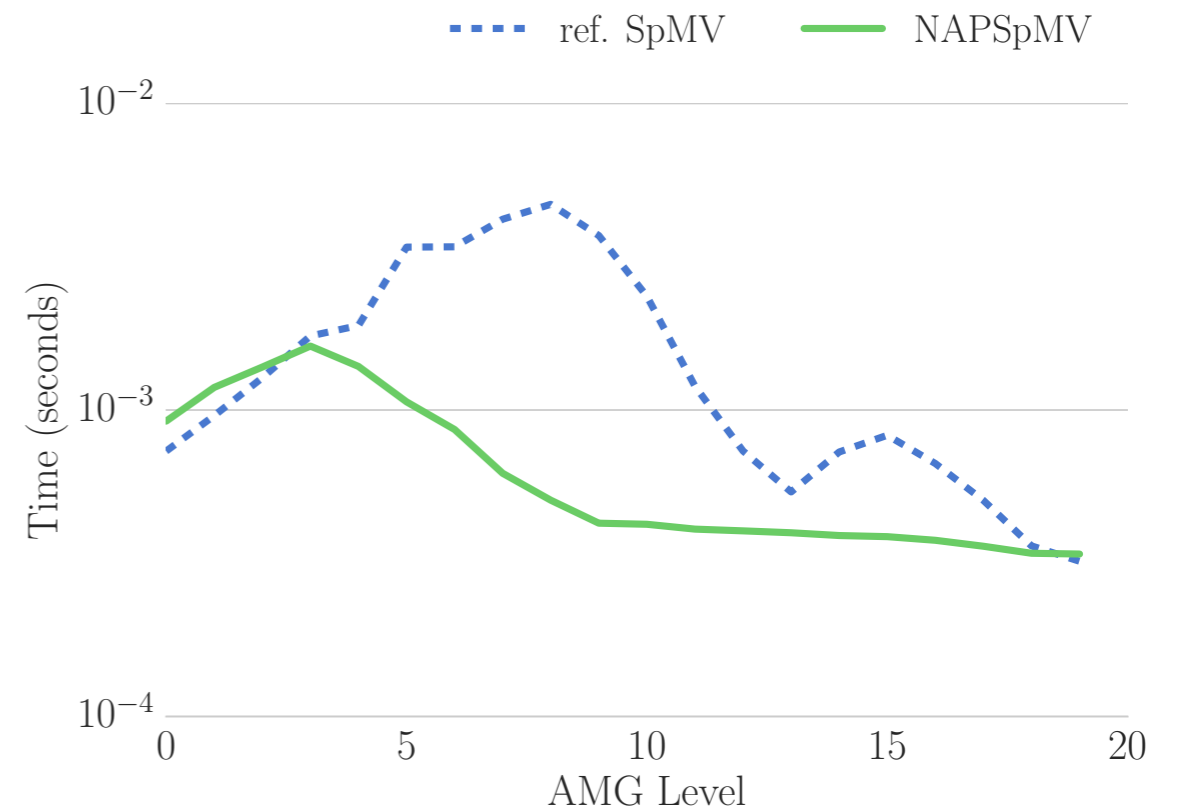
Maximum size of
on-node messages

AMG Hierarchy SpMV Times

Time spent performing a SpMV on each level of the hierarchy with 32,768 processes



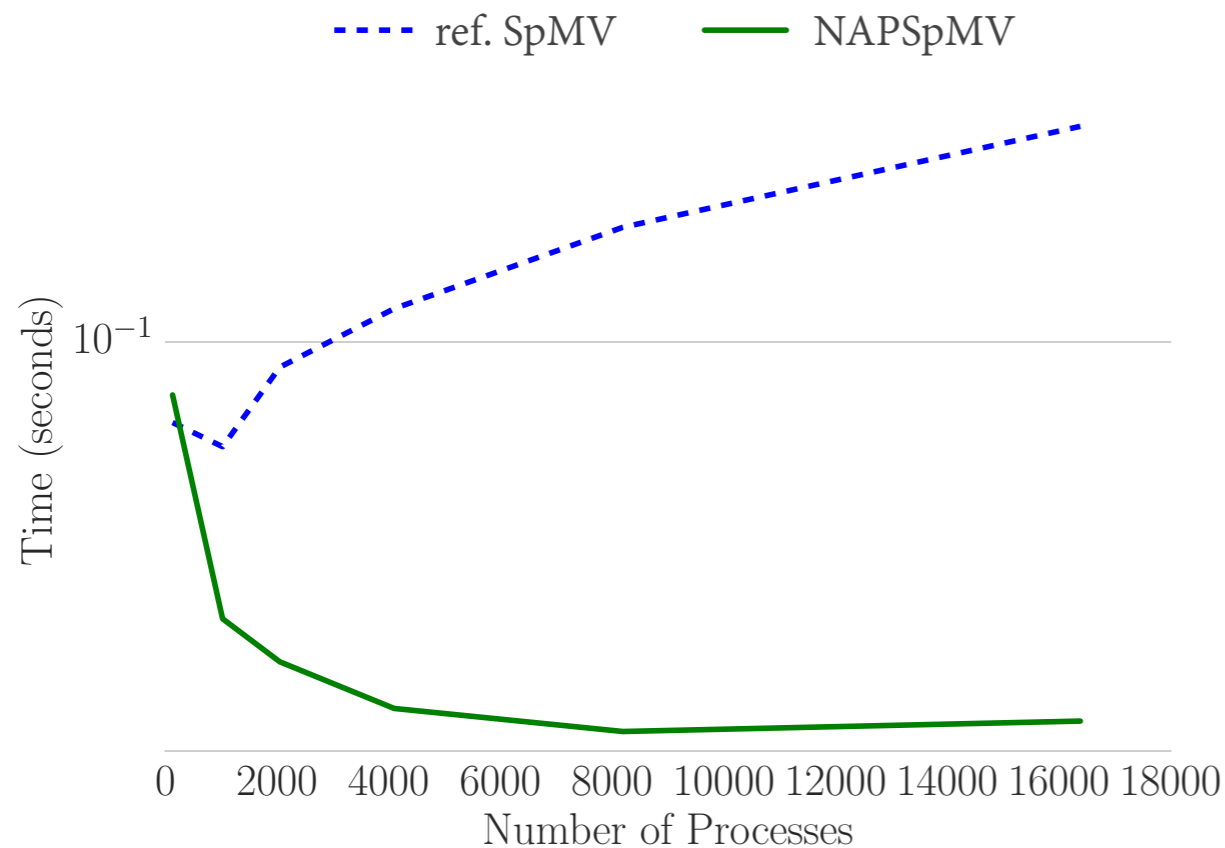
Linear Elasticity (MFEM)




2D Rotated Anisotropic Diffusion

AMG SpMV Times: Strong Scale

Time spent performing a SpMV on all levels of the hierarchy at various scales



Linear Elasticity Hierarchy
Fine Level: 839,619 rows

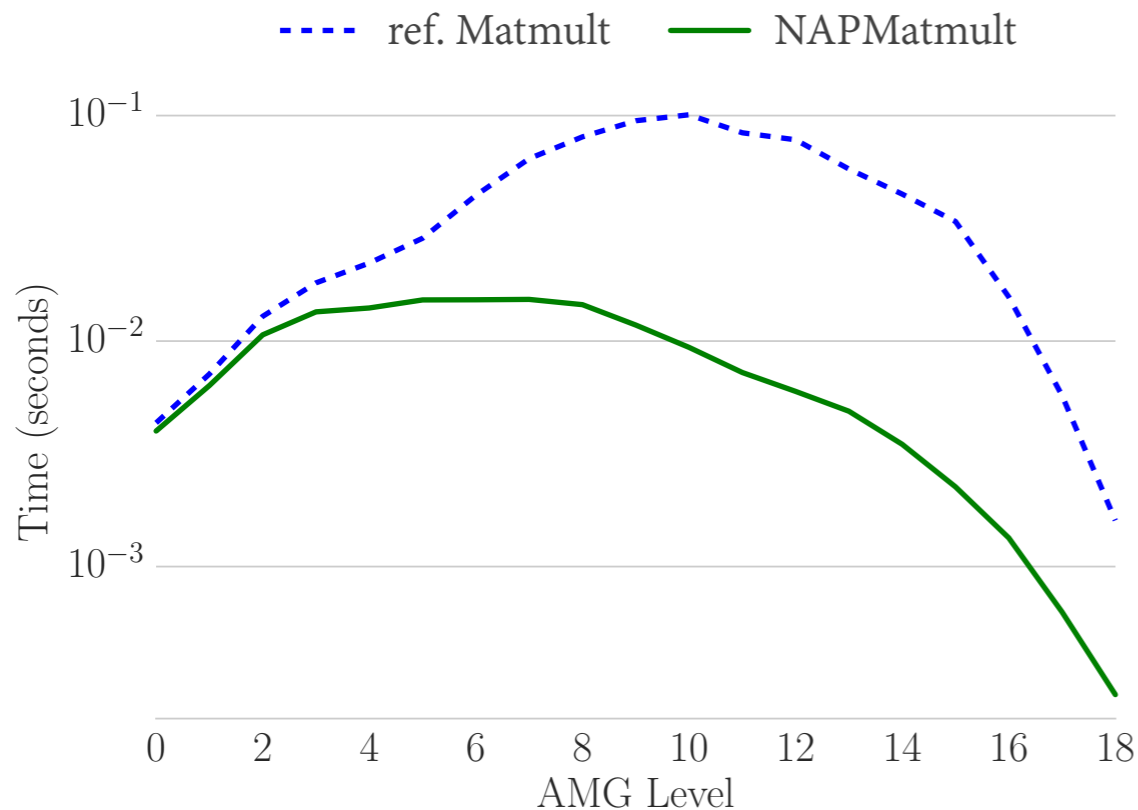


Rotated Anisotropic Diffusion Hierarchy
Fine Level: 10,240,000 rows


On Arxiv: [TAPSpMV: Topology-Aware Parallel Sparse Matrix Vector Multiplication](#)

AMG Hierarchy Mat-Mat Times

Time spent performing matrix-matrix multiplication on each level of the hierarchy with 16,384 processes



Linear Elasticity (MFEM)



2D Rotated Anisotropic Diffusion

Conclusions & Future Work

- ❖ Able to improve cost of AMG through changes to algorithm and parallel implementation
- ❖ Codebase Raptor: parallel AMG with node aware communication
- ❖ Future Work:
 - ❖ Modeling Network Contention in SpMV, Matrix-Matrix
 - ❖ Extending node aware communication to other parts of AMG and other solvers
 - ❖ Combination of sparse / hybrid Galerkin algorithm with node aware communication

Questions?