LA-UR- 11-06165

Approved for public release; distribution is unlimited.

Title:	Solution of the Within-Group Multidimensional Discrete Ordinates Transport Equations on Massively Parallel Architectures
Author(s):	R. Joseph Zerr Yousry Y. Azmy
Intended for:	American Nuclear Society Winter Meeting 2011 Presentation

Los Alamos
NATIONAL LABORATORY
EST.1943

Form 836 (7/06)

Abstract

The advent of massively parallel computing environments has made the application of efficient schemes to solve the transport equation in parallel a high priority. The parallelization of the transport equation has been researched intensively for over two decades. Distributing work based on energy or angular variable decomposition alone is inadequate when tens of thousands of processing elements (PEs) are available. Alternatively, problems with ever-increasing numbers of spatial cells are highly desired to analyze physically larger regions and/or to employ finer spatial meshes. Spatial domain decompositions (SDD) have been developed to use the greater number of unknowns to achieve high scalability on increasingly large computing clusters.

This work has focused on the development of a novel kernel for handling the local problem within the PBJ framework of the global problem based on the integral transport matrix method (ITMM). The ITMM seeks operators that act directly on the cells' scalar flux and incoming angular fluxes on the boundaries of the sub-domains. Such an approach abandons repetitive mesh sweeps and other expensive iterative routines.

Solution of the Within-Group Multidimensional Discrete Ordinates Transport Equations on Massively Parallel Architectures

R. Joseph Zerr

The Pennsylvania State University

Los Alamos National Laboratory

Yousry Y. Azmy

North Carolina State University

LA-UR-11-?????

0. Outline

1. Introduction: Parallel Solution of the Transport Equation
2. Integral Transport Matrix Method (ITMM)
3. The Parallel Gauss Seidel (PGS) Method
4. Weak Scaling Studies
5. Periodic Heterogeneous Layers (PHL) Studies
6. Summary and Conclusions

1. Parallel Solution of the TE (S_N Form)

- **Domain Decomposition in Energy, Angle, and/or Space**
- **Current focus on Spatial Domain Decomposition (SDD)**
  - KBA/Wavefront
    - Map m-dimensional mesh onto $m-1$ processor topology
    - Compute angular flux for cell center and edges
    - Communicate outward edge flux to neighbors along diagonal wavefront
    - Sweep cells on wavefront concurrently; pipeline ordinates
    - **Synchronous method**
  - **Parallel Block Jacobi (PBJ)**
    - Divide problem spatially into sub-domains
    - Work on transport equation for sub-domains concurrently
    - Communicate angular flux at boundaries
    - **Asynchronous method**

1. PBJ Forms a Local-Global Problem

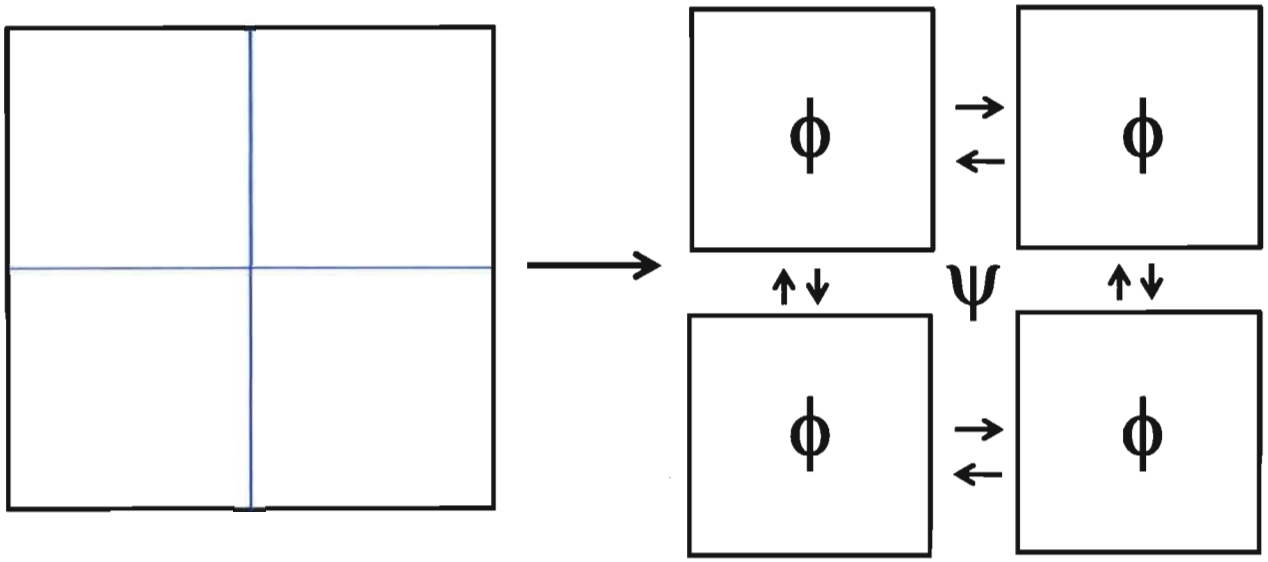

$$\phi = \Phi\left(\mathbf{q}, \Psi_{in}\right)$$

$$\Psi_{out} = \Psi\left(\phi, \mathbf{q}, \Psi_{in}\right)$$

- Local: operators compute the scalar and outward angular flux of each sub-domain → Sweeps, Krylov, other?
- Global iterations on angular flux – pass among sub-domains
  - Sub-domains modify/update with own information
  - **Decoupling**: sub-domains effects more local, less global

2. Integral Transport Matrix Method (ITMM)

- Reconsiders the typical SI scheme as an iterative process to determine a new iterate of scalar flux from previous value of the scalar flux

$$\boldsymbol{\phi}^{(l+1)} = \mathbf{J}_{\phi}\left(\boldsymbol{\phi}^{(l)} + \boldsymbol{\Sigma}_s^{-1}\mathbf{q}\right) + \mathbf{K}_{\phi}\boldsymbol{\Psi}_{in}$$

- $\mathbf{J}_{\phi}$ is the iteration Jacobian, constructed with **differential mesh sweeps** in all directions, per group, to attain full scalar flux coupling
- $\mathbf{K}_{\phi}$ describes the attenuation of the BC flux throughout the system and the contribution to the cell-average scalar fluxes
- At convergence limit:

$$\left(\mathbf{I} - \mathbf{J}_{\phi}\right)\boldsymbol{\phi}^{\infty} = \mathbf{J}_{\phi}\boldsymbol{\Sigma}_s^{-1}\mathbf{q} + \mathbf{K}_{\phi}\boldsymbol{\Psi}_{in} \rightarrow \boldsymbol{\Phi}\left(\mathbf{q}, \boldsymbol{\Psi}_{in}\right)$$

2. Outgoing Angular Flux

- Computing the outgoing angular flux at system boundaries requires additional equation of matrix-vector products

$$\boldsymbol{\psi}_{out} = \mathbf{J}_{\psi}\left(\boldsymbol{\phi}^{\infty} + \boldsymbol{\Sigma}_s^{-1}\mathbf{q}\right) + \mathbf{K}_{\psi}\boldsymbol{\psi}_{in} \rightarrow \boldsymbol{\Psi}\left(\boldsymbol{\phi}, \mathbf{q}, \boldsymbol{\psi}_{in}\right)$$

- $\mathbf{J}_{\psi}$ and $\mathbf{K}_{\psi}$ are straightforward extensions of $\mathbf{J}_{\phi}$ and $\mathbf{K}_{\phi}$, respectively
  - Construction based on the coupling of the system's cells to hypothetical cells neighboring the system's boundary cells
- Reflective BCs: Set $\boldsymbol{\psi}_{out}$ to $\boldsymbol{\psi}_{in}$ and re-solve $\boldsymbol{\phi}^{\infty}$ (iterative)

2. ITMM Operator Sizes

- New algorithm partially limited by memory
- ITMM operators provide full coupling among cells, among cells to boundaries, and among incoming/outgoing boundaries
- Operators are large and grow super-linearly with number of cells, N
  - $\mathbf{J}_{\phi}$: $O(N^{2})$
  - $\mathbf{K}_{\phi}$: $O(N^{5/3}M)$
  - $\mathbf{J}_{\psi}$: $O(N^{5/3}M)$
  - $\mathbf{K}_{\psi}$: $O(N^{4/3}M)$
- Differential mesh sweep time grows linearly with angles, M
- ITMM is applicable to higher order spatial discretizations and anisotropic scattering—sizes of operators become dependent on number of modeled spatial and angular moments of flux

3. Global Solution as an Iterative Problem

- Global system of equations

$$\left(\mathbf{I}-\mathbf{J}_{\phi}\right)\boldsymbol{\phi}^{(l+1)}=\mathbf{J}_{\phi}\boldsymbol{\Sigma}_{s}^{-1}\mathbf{q}+\mathbf{K}_{\phi}\boldsymbol{\psi}_{in}^{(l)}$$

$$-\mathbf{J}_{\psi}\boldsymbol{\phi}^{(l+1)}+\mathbf{I}\boldsymbol{\psi}_{out}^{(l+1)}=\mathbf{J}_{\psi}\boldsymbol{\Sigma}_{s}^{-1}\mathbf{q}+\mathbf{K}_{\psi}\boldsymbol{\psi}_{in}^{(l)}$$

  - Iterating on the outgoing angular flux, NOT scalar flux
  - Not an inner/source iteration; departure from previous PBJ methods
  - Scalar flux is an intermediate value for convenience and efficiency
- #of global iterations driven by the tightness of sub-domain **coupling**
  - Decreasing optical thickness $\rightarrow \sigma_t h$
  - Increasing scattering ratio $\rightarrow c$
- PBJ iterative solution is most straight-forward method: Incoming interfacial angular fluxes lag one iteration

3. Improving PBJ → Parallel Gauss-Seidel (PGS)

- Red/Black Parallel Gauss-Seidel (PGS)
- Split single sub-domain into many smaller sub-domains → **PRO:** Operators' memory requirement and construction time super-linearly decrease with $N = I \times J \times K$
- Each sub-domain is either **red** or **black**
- At each global iteration:
  - Solve the local systems for ϕ, ψ_{out}
  - Copy/Send $\psi_{out} \rightarrow \psi_{in}$
  - Solve the local systems for ϕ, ψ_{out}
  - Copy/Send $\psi_{out} \rightarrow \psi_{in}$
- Competing effects
  - **Faster convergence rate of PGS vs. PBJ**
  - **Increase in number of global iterations**

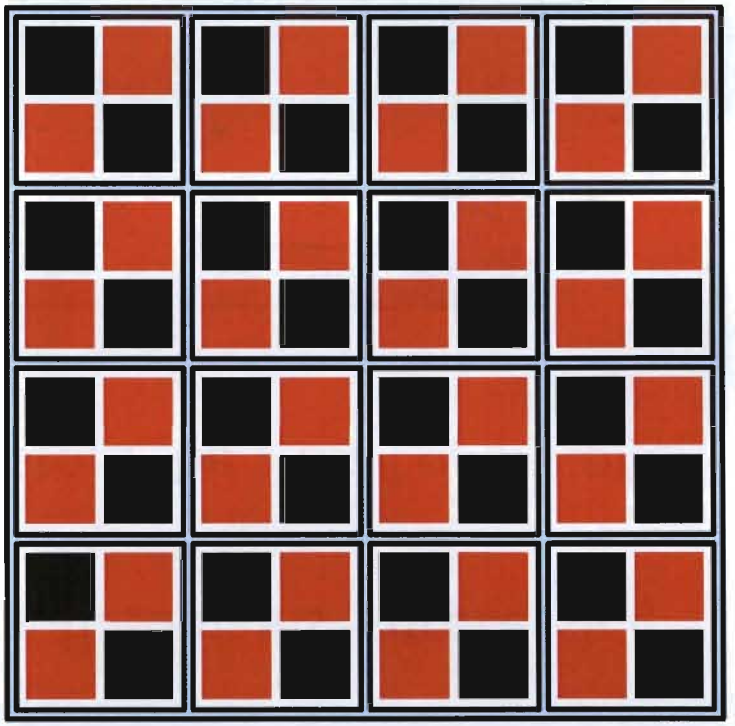

3. Multiple Sub-Domains Per Processor

- Assign each *P* several, still independent, sub-domains – **processing elements (PEs)** do not suffer idleness
- Construct (differential mesh sweep) and store the ITMM operators for each

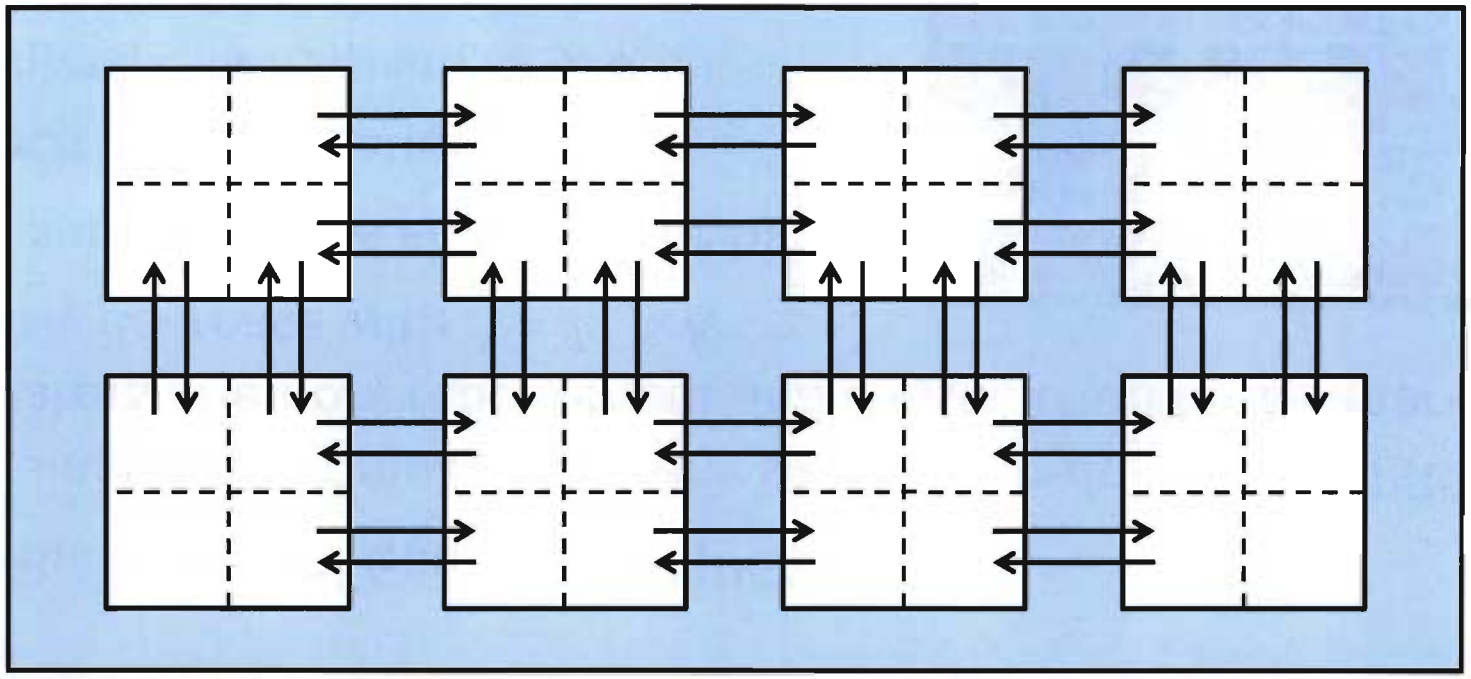

4. Numerical Experiments in Weak Scaling

- Parallel Integral Discrete Ordinates Transport Solver (PIDOTS)
- Testing with one-group, S_N, DD problems with isotropic scattering
- Testing on JaguarPF(JPF) [ORNL] up to P = 32,768
- Evaluate performance as problem size grows with P
- Problem size per P fixed
- Start with $L \times L \times L$ domain
  - 4 materials
  - Vary cell size h (0.1, 1.0, 10.0 cm)
  - c = 0.9, 0.99
- Choose L and S_N together: memory limited

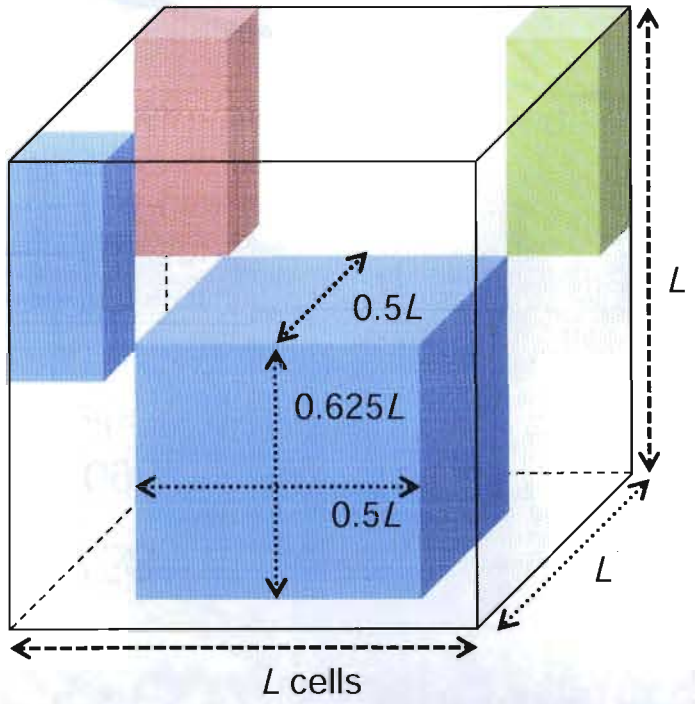

4. PGS-KBA (PARTISN) Comparison on Jaguar

- 16×16×16-cell model per P, S_8, up to P = 32,768
- PGS: 64 4×4×4-cell sub-domains per P
- c = 0.9, h = 0.1 cm results:

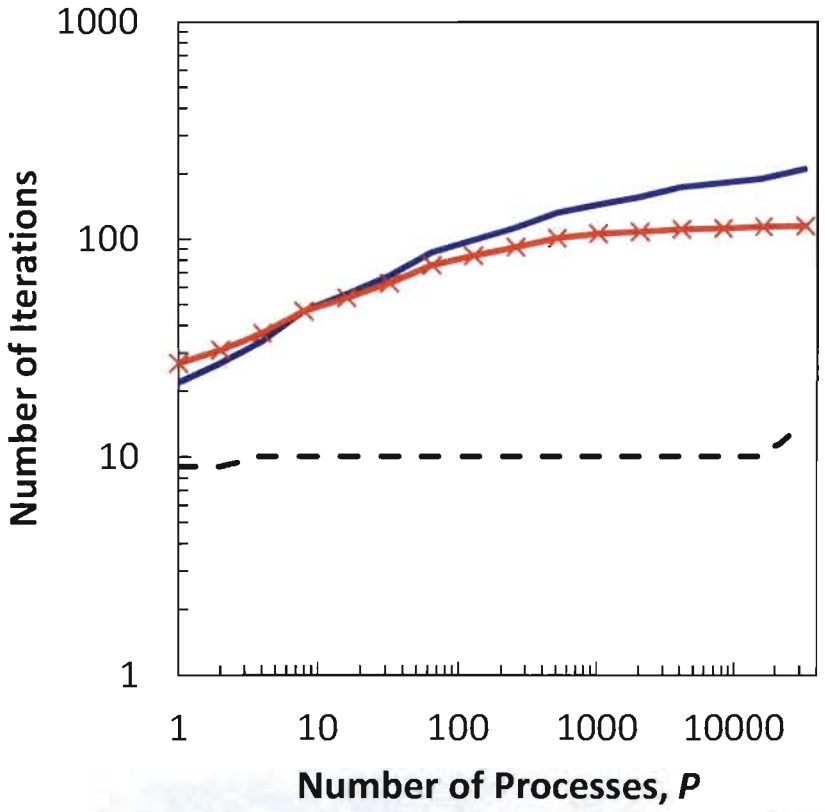

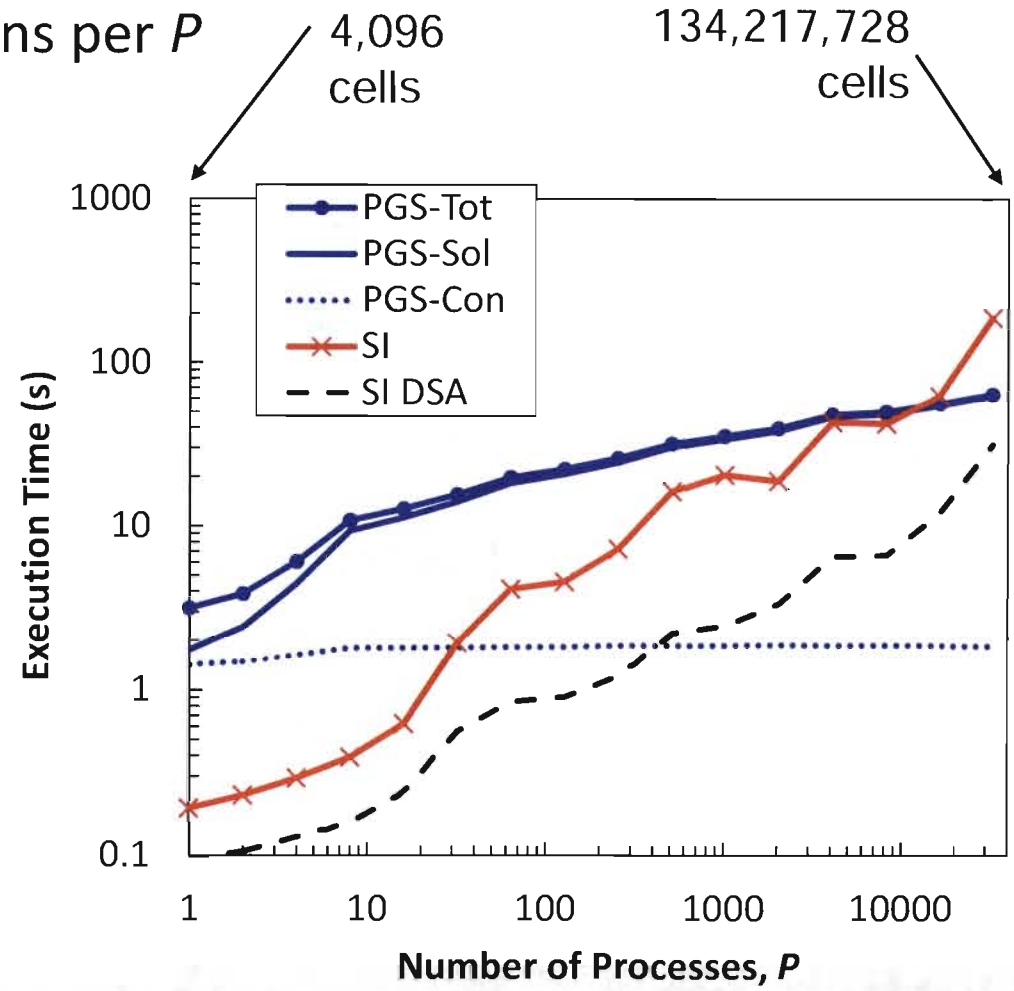

4. PGS-KBA Comparison on Jaguar

- $c = 0.9$, $h = 1.0$ cm results:

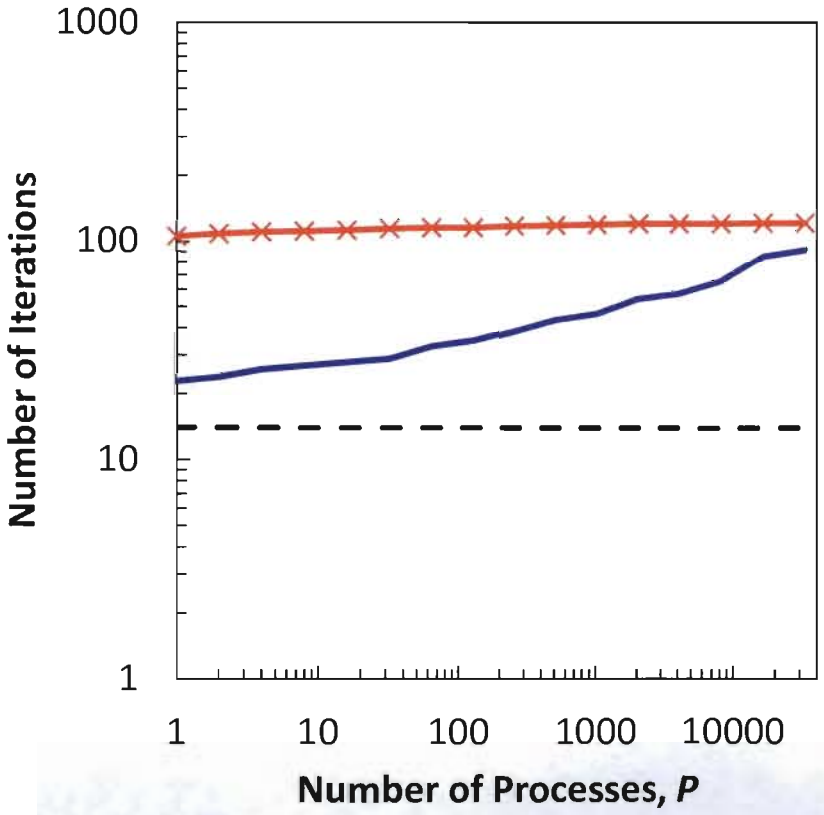

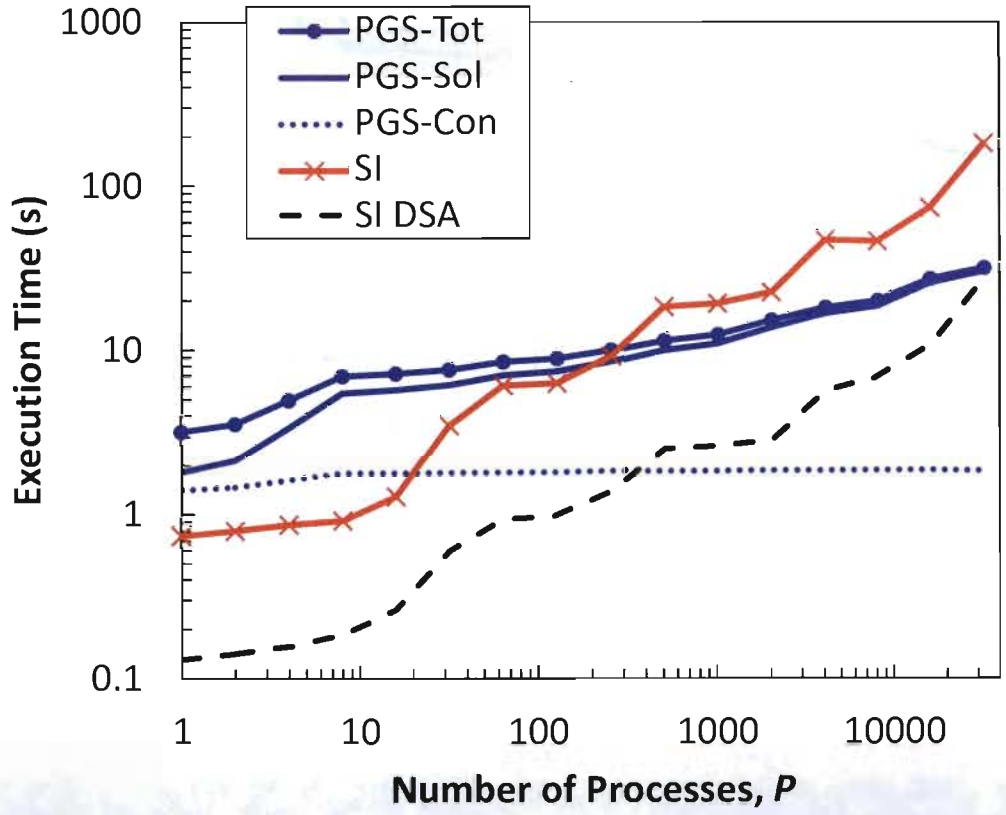

4. PGS-KBA Comparison on Jaguar

- $c = 0.9$, $h = 10.0$ cm results:

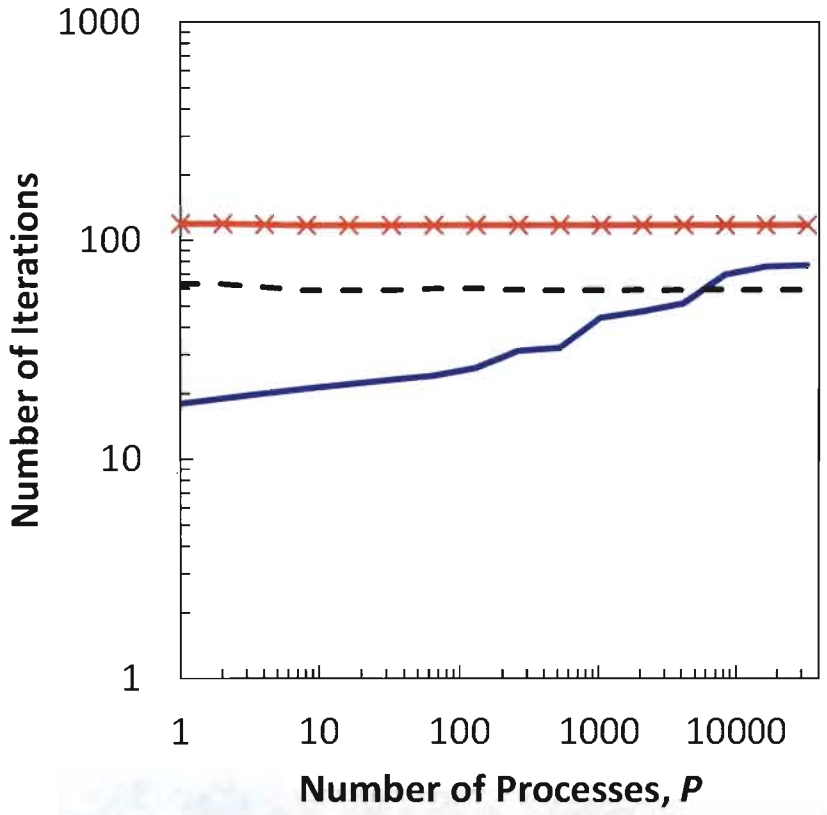

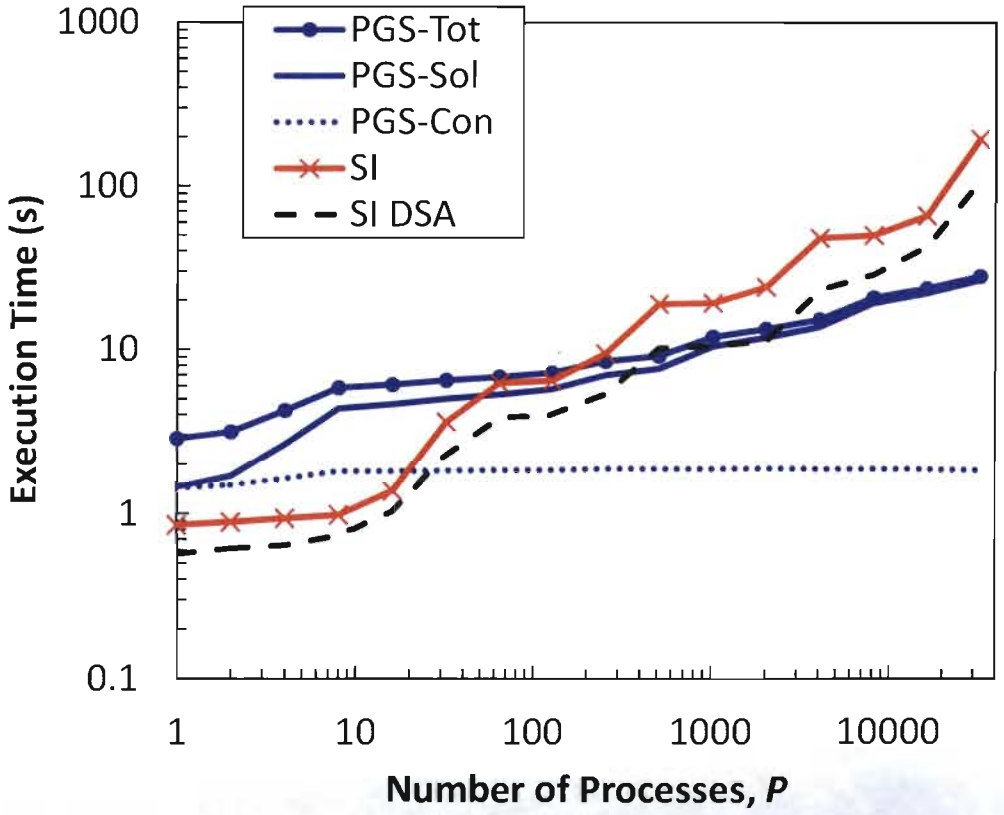

4. PGS-KBA Comparison on Jaguar

- $c = 0.99$, $h = 0.1$ cm results:

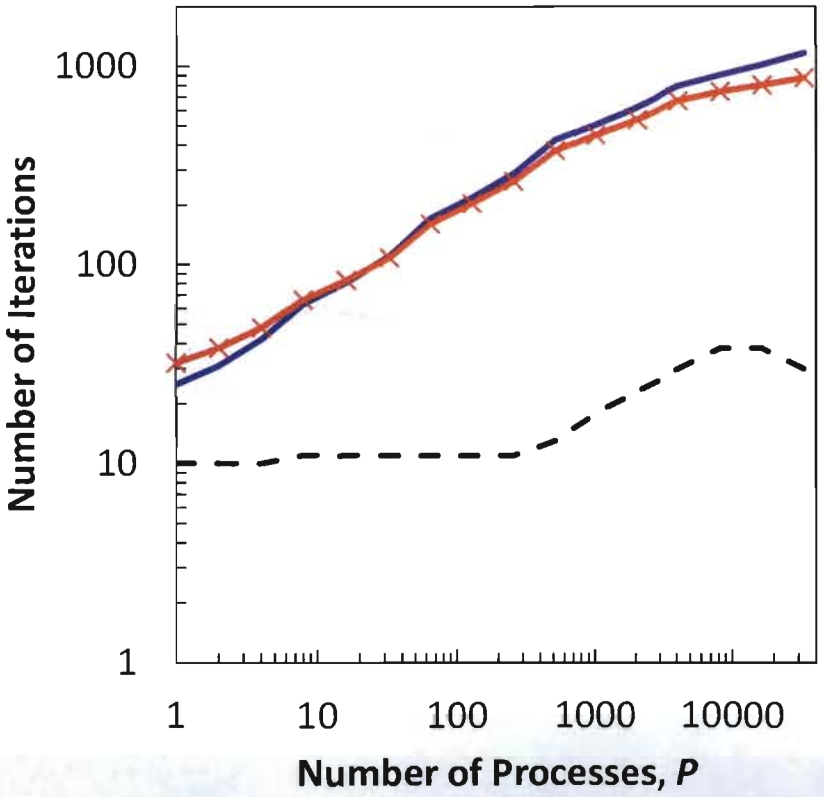

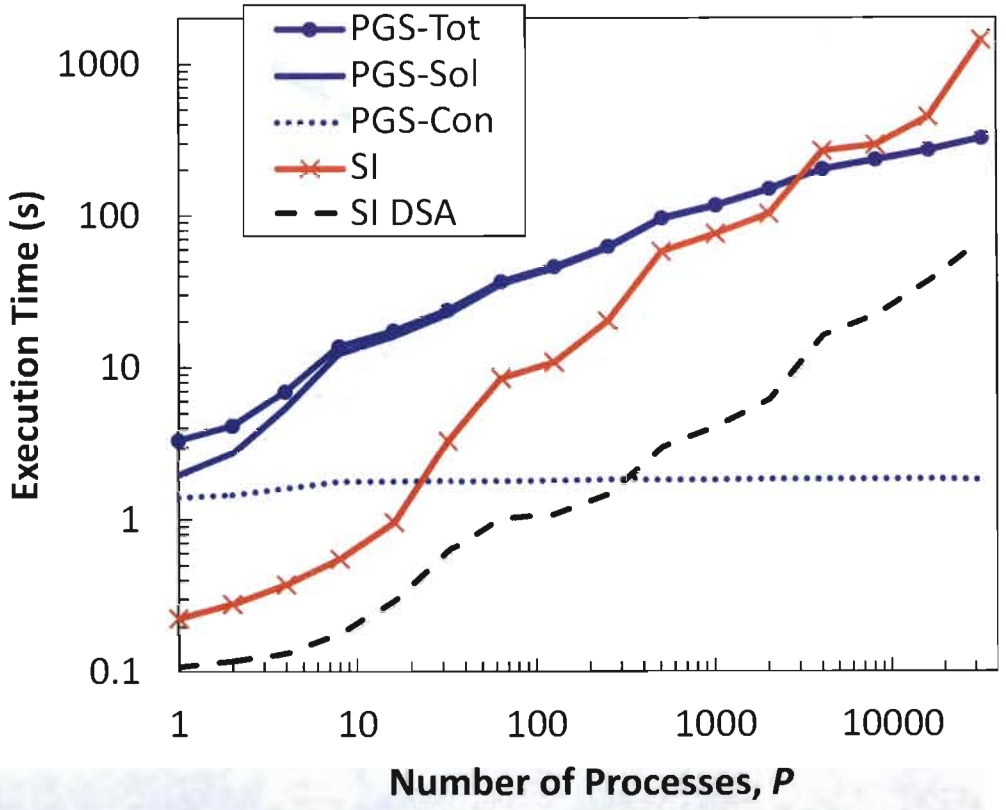

4. PGS-KBA Comparison on Jaguar

- $c = 0.99$, $h = 1.0$ cm results:

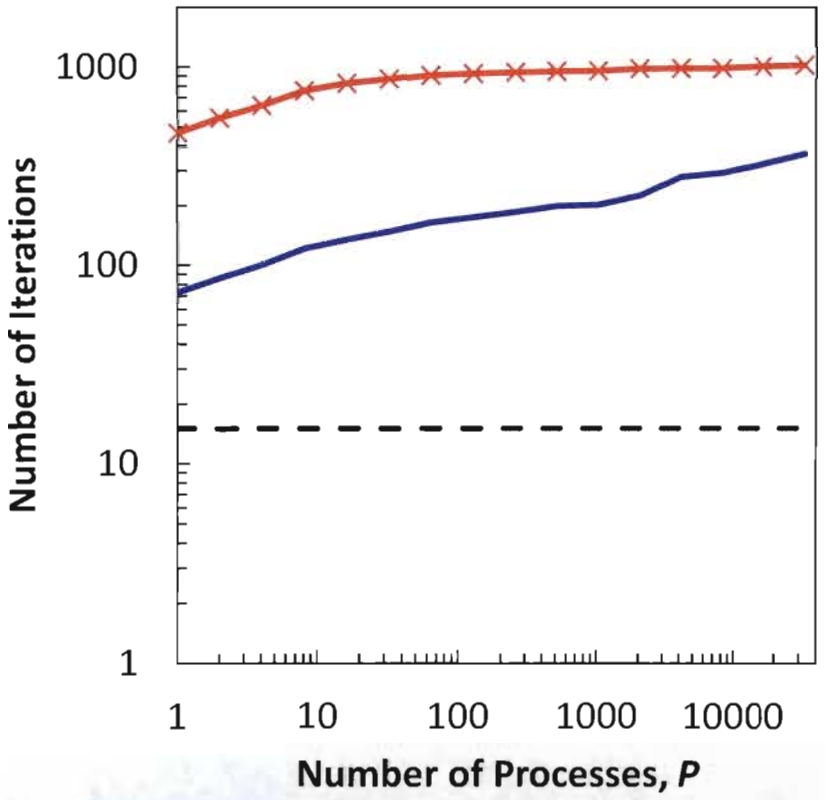

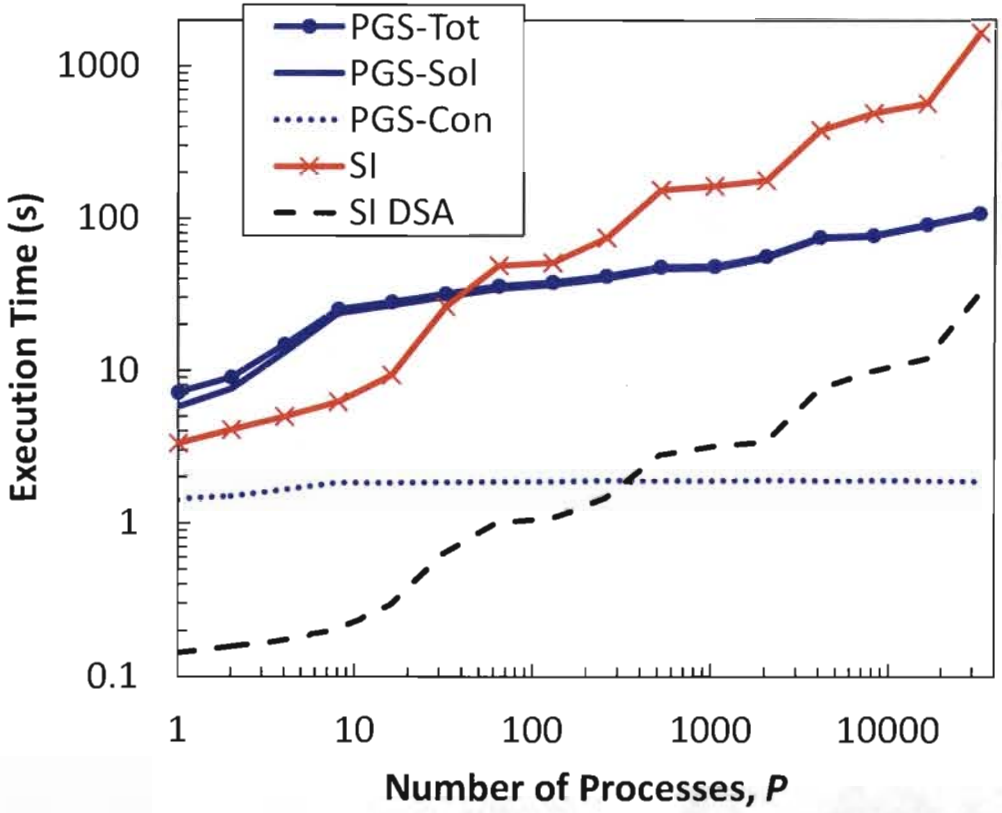

4. PGS-KBA Comparison on Jaguar

- $c = 0.99$, $h = 10.0$ cm results:

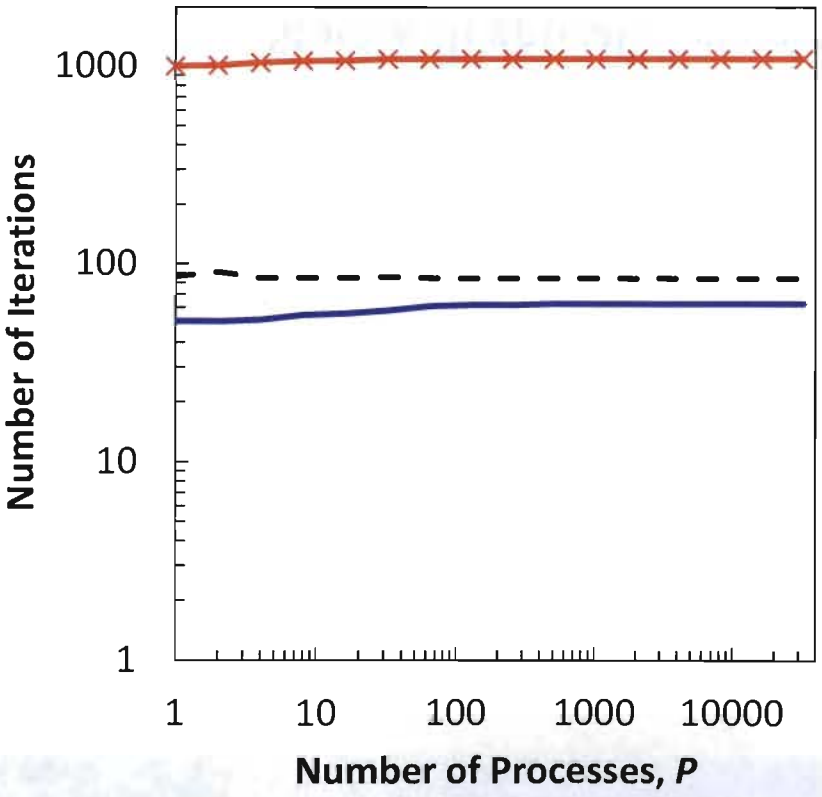

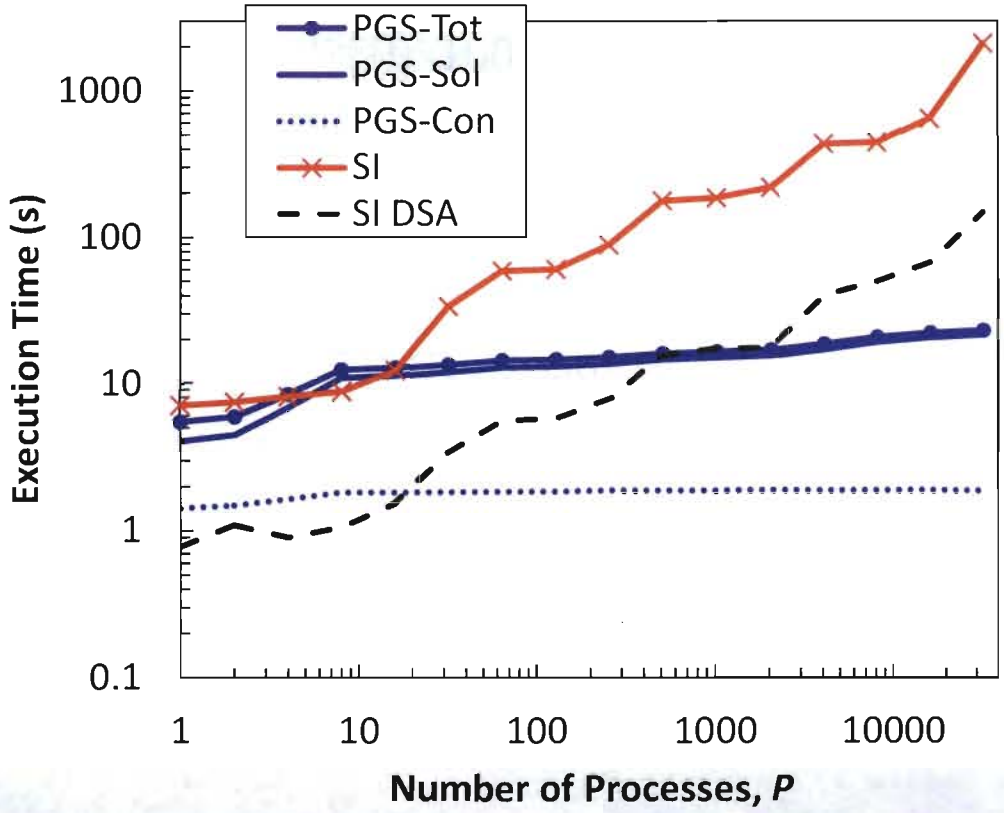

5. Periodic Heterogeneous Layers (PHL)

- Examine the performance of ITMM-PGS compared to SI-KBA in the presence of sharp material discontinuities
- SI DSA ineffectiveness in the presence of such heterogeneity
- Problem Description
  - $8\times8\times8$, S_{16}
  - Periodic x-y planes (z-layers) of optically thin and thick layers
  - Start with optical thickness $\sigma_t h = 1$, scale by a, a=10, 100, 1000
  - Scattering ratio $c = 0.9$
  - Weak scaling, copy the base model as each dimension is increased
  - Run up to $P = 4{,}096$ on JPF

5. PGS-KBA with PHL on Jaguar

- a = 10 results:

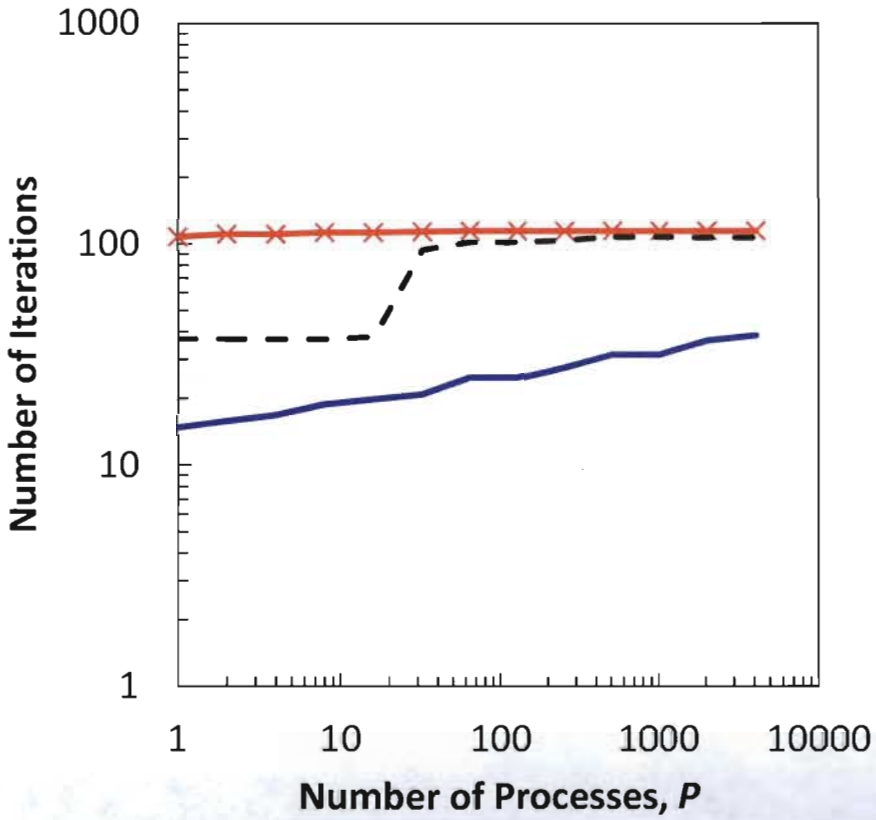

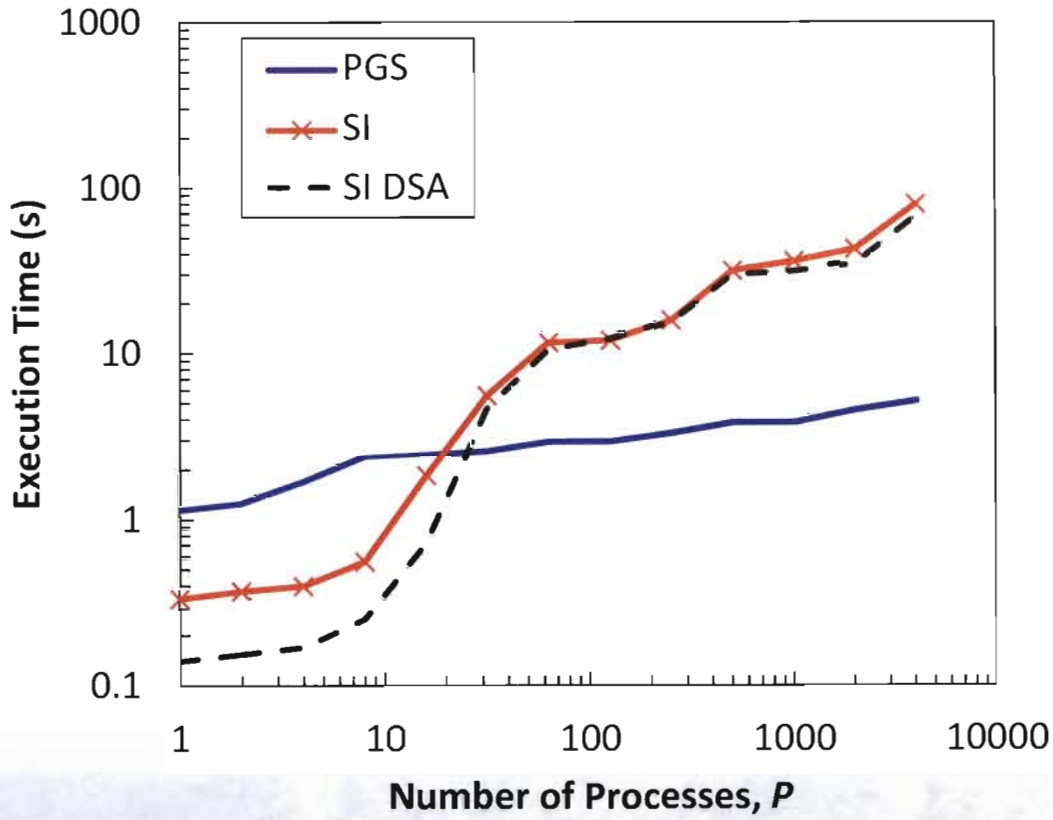

5. PGS-KBA with PHL on Jaguar

- $a = 100$ results:

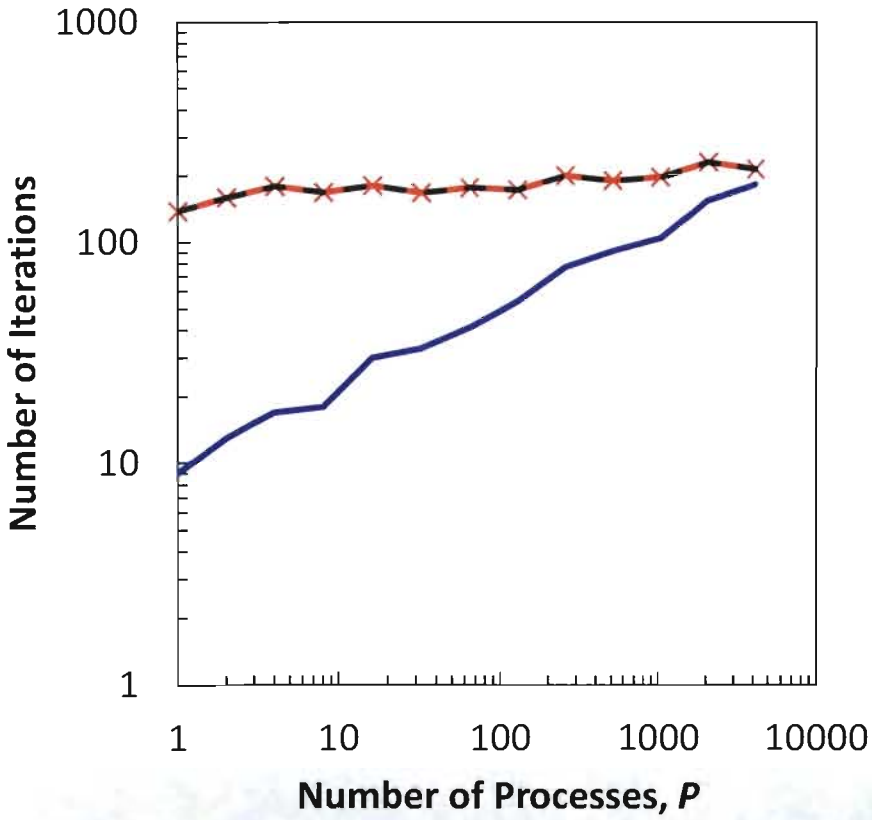

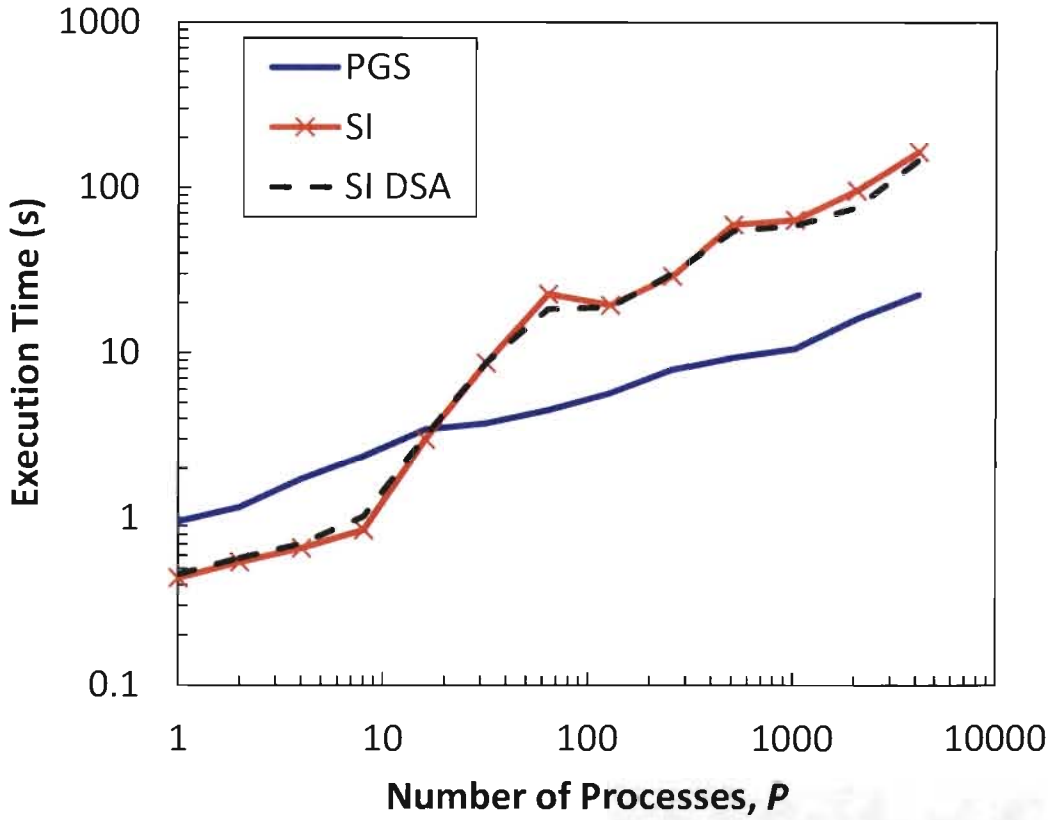

5. PGS-KBA with PHL on Jaguar

- a = 1000 results:

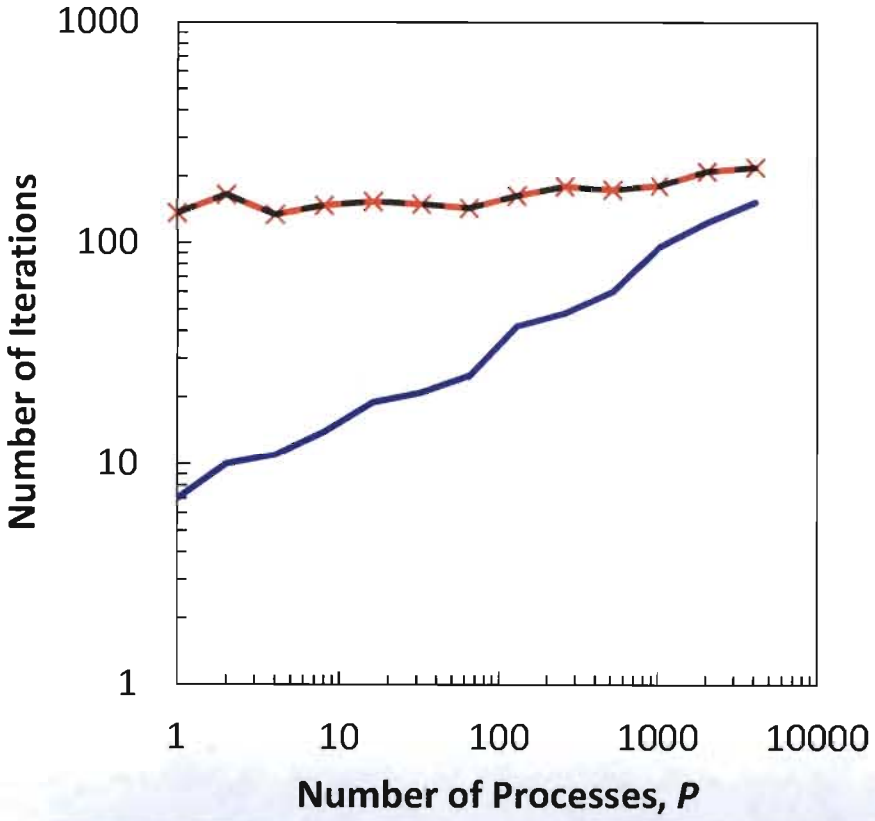

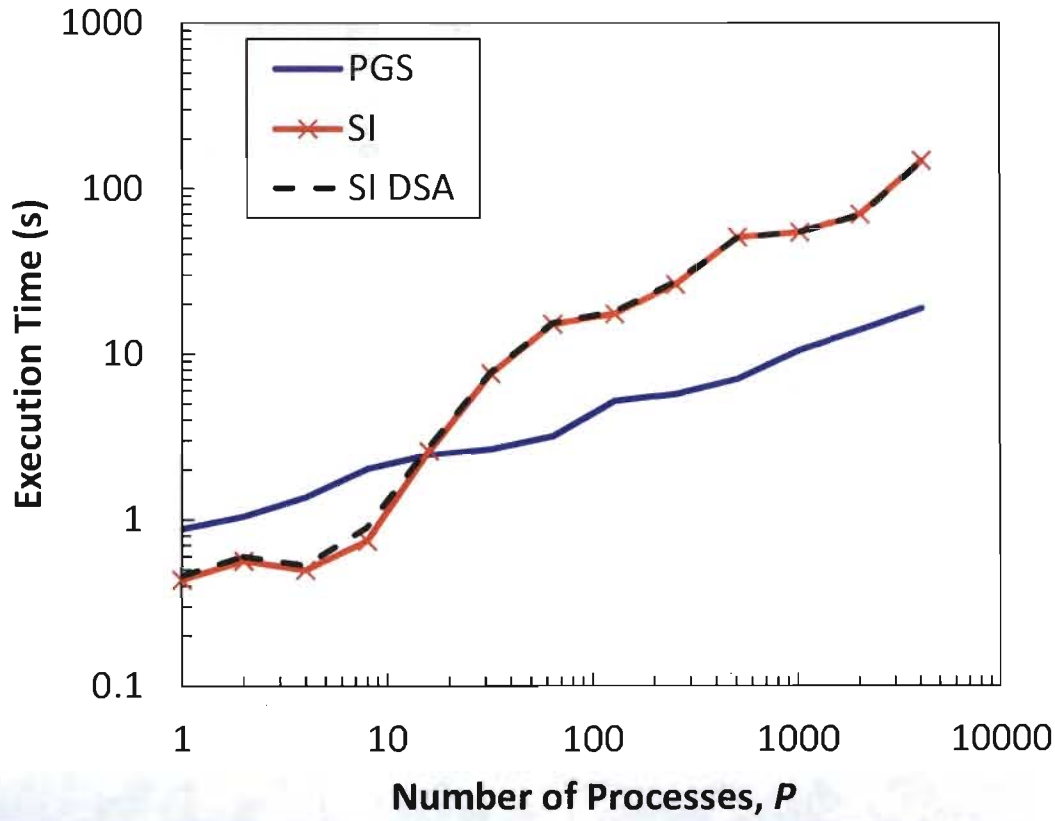

6. Summary

- ITMM operators
  - Four operators dependent on problem size
  - Constructed once per group with differential mesh sweep
- PGS compared to KBA:
  - Standard weak scaling
    - Very large differences when SI is accelerated with DSA
    - Gap closes as optical thickness and scattering ratio are increased → most difficult SI problems
    - SI and SI DSA demonstrate larger growth in execution time as *P* increases
  - PHL weak scaling
    - DSA ineffective as heterogeneity is increased
    - ITMM direct coupling resolves anisotropy faster

6. Conclusions

- Method looks very good for a special set of problems
  - Highly scattering, large optical thickness
  - Lots of material discontinuities, sharply heterogeneous
- Performance should improve with the development and inclusion of global acceleration technique
  - Preliminary efforts in spatial multigrid
  - Use of global Preconditioned GMRES solver
- Improve convergence rate and apply to higher R/B divisions
  - Super-linear relationship of operators to sub-domain size
  - Save on memory and operations

Acknowledgments and Questions

- Tom Evans of Oak Ridge National Laboratory
  - Access to JaguarPF system.
- Randal Baker and Jon Dahl of Los Alamos National Laboratory
  - Building and running PARTISN.
- This work has been sponsored by a grant from the National Nuclear Security Administration

???