# CASL: Consortium for Advanced Simulation of Light Water Reactors

# Neutronics and 3D $S_N$ Transport


Thomas M. Evans

ORNL


Accelerating Computational Science Symposium

March 29

Washington, D.C.

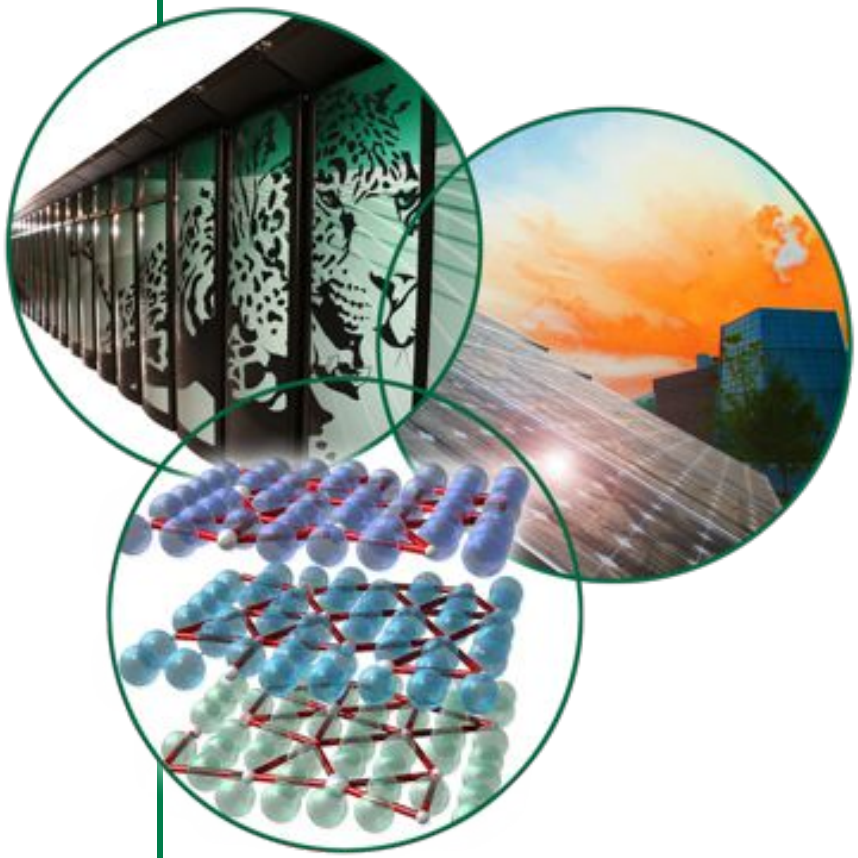

U.S. DEPARTMENT OF ENERGY

OAK RIDGE NATIONAL LABORATORY

MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

# Questions

# Contributors

## *ORNL Staff*

- Greg Davidson
- Josh Jarrell
- Steve Hamilton
- Chris Baker
- Andrew Godfrey
- Kevin Clarno
- Douglas Peplow
- Scott Mosher

## *CASL*

- Roger Pawloski
- Brian Adams

## *Students and PostDocs*

- Rachel Slaybaugh (Wisconsin)
- Stuart Slattery (Wisconsin)
- Josh Hykes (North Carolina State)
- Todd Evans (North Carolina State)
- Cyrus Proctor (North Carolina State)

## *OLCF (NCCS) Support*

- Dave Pugmire
- Sean Ahern
- Wayne Joubert

## *Other Funding Support*

- INCITE/ASCR/NRC/NNSA

# Outline

- Neutronics
- Deterministic Transport
- Parallel Algorithms and Solvers
- Performance
- Projections

# VERA (Virtual Environment for Reactor Applications)

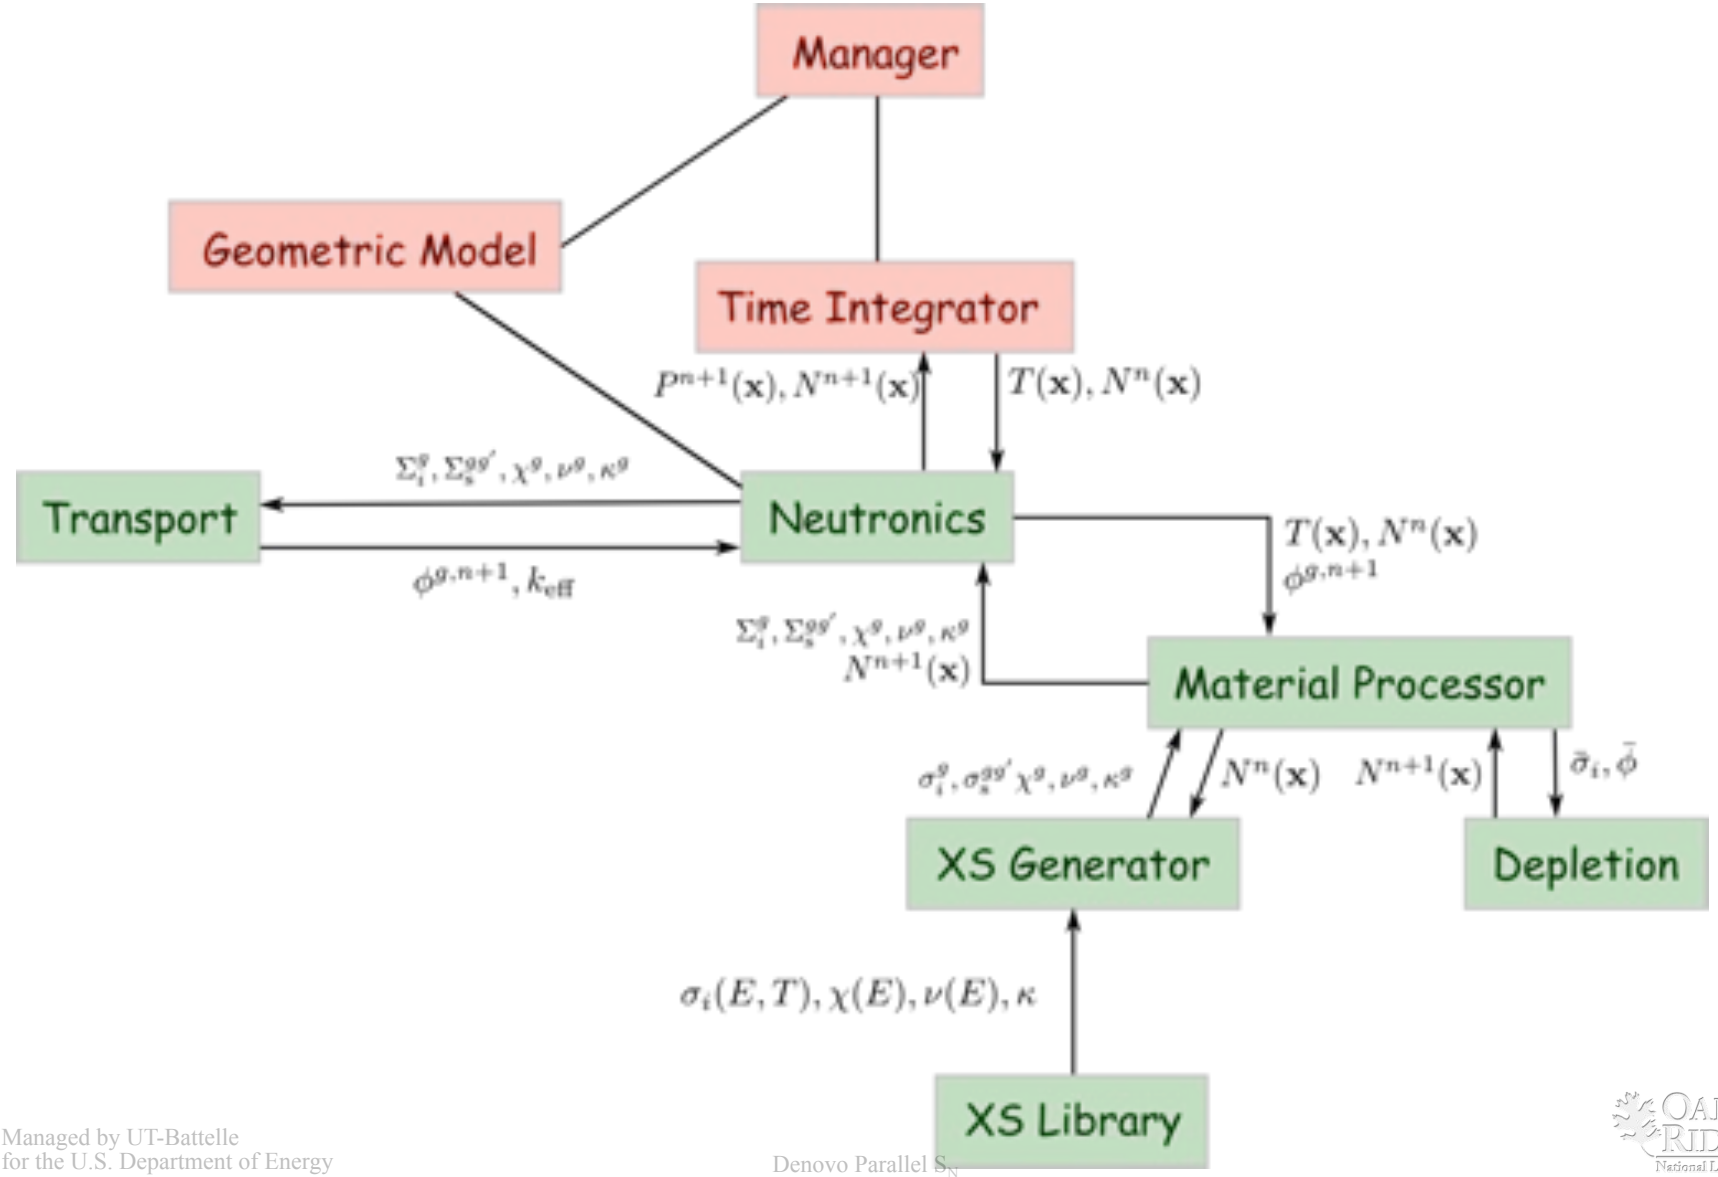

# Science Drivers for Neutronics

- **Spatial resolution**
  - **To resolve the geometry**
    - **$10^{9\text{-}12}$ unknowns**
    - **$mm^3$ cells in a $m^3$ vessel**
  - **Depletion makes it harder**
- **Energy resolution**
  - **To resolve resonances**
    - **$10^{4\text{-}6}$ unknowns**
    - **Done in 0D or 1D today**
- **Angular resolution**
  - **To resolve streaming**
    - **$10^{2\text{-}4}$ unknowns**
  - **Space-energy resolution make it harder**

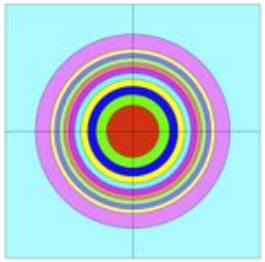


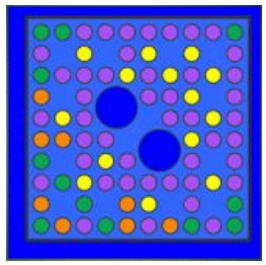


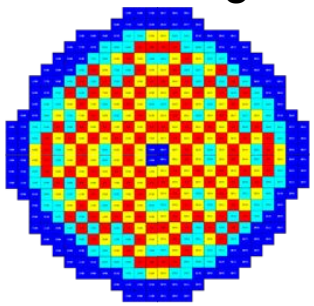


BWR and PWR cores have similar dimension, but much different compositions and features

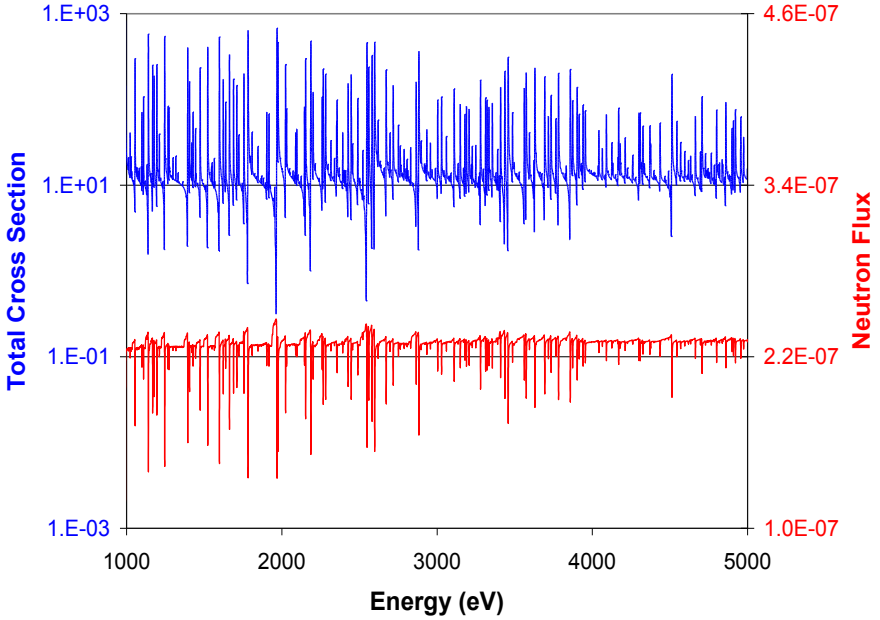

# CASL Test Problems

- CASL AMA Focus Area has defined 10 test problems that drive requirements for Core-Simulation
  - Required in order to do Challenge Problems
- VERA-CS for FY12 is targeting the first five
  - ✓ 2D Host Zero Power (HZP) Pin Cell
  - ✓ 2D HZP Lattice
  - ✓ 3D HZP Assembly
  - ❑ HZP 3x3 Assembly Control Rod Worth
  - ❑ Physical Reactor Zero Power Physics Test (ZPPT)
- TH-feedback starts in problem 6
- Depletion starts in problem 8

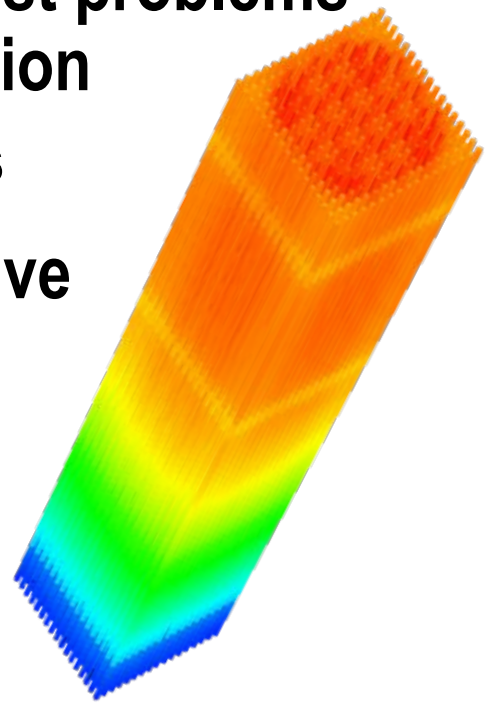

# Denovo Capabilities

- State of the art transport methods
  - 3D/2D, non-uniform, regular grid $S_N$
  - 2D MoC solver option
  - Multigroup energy, anisotropic $P_N$ scattering
  - Forward/Adjoint
  - Fixed-source/*k*-eigenvalue
  - 6 spatial discretization algorithms
    - Linear and Trilinear discontinuous FE, step-characteristics, theta-weighted diamond, weighted diamond + flux-fixup
  - Parallel first-collision
    - Analytic ray-tracing (DR)
    - Monte Carlo (DR and DD)
  - Multiple quadratures
    - Level-symmetric
    - Generalized Legendre Product
    - Quadruple Range

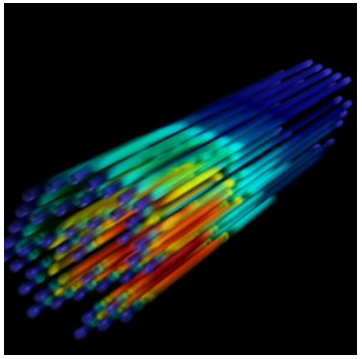

Denovo Parallel $S_N$

- Modern, Innovative, High-Performance Solvers
  - Within-group solvers
    - Krylov (GMRES, BiCGStab) and source iteration
    - DSA preconditioning (SuperLU/ML-preconditioned CG/PCG)
  - Multigroup solvers
    - Transport Two-Grid upscatter acceleration of Gauss-Seidel
    - Krylov (GMRES, BiCGtab)
      - Multigrid preconditioning
  - Eigenvalue solvers
    - Power iteration (with rebalance)
      - CMFD acceleration (for MoC)
    - Krylov (Arnoldi)
    - RQI with multigrid preconditioning

Power distribution in a BWR assembly

# Denovo Capabilities

- Parallel Algorithms
  - Koch-Baker-Alcouffe (KBA) wavefront decomposition
  - Domain-replicated (DR) and domain-decomposed first-collision solvers
  - Multilevel energy decomposition
  - Parallel I/O built on SILO/HDF5


**Core Neutronics Package in VERA Toolset**
**2012-13 INCITE Award**
*The Solution of 3D PWR Neutronics Benchmark Problems for CASL*, 19 MCPU-HOURS
**2010-11 INCITE Award**
*Uncertainty Quantification for Three Dimensional Reactor Assembly Simulations*, 26 MCPU-HOURS
**2010 ASCR Joule Code**
**2009-2011 2 ORNL LDRDs**

- Advanced visualization, run-time, and development environment
  - multiple front-ends (HPC, SCALE, Python-bindings, core-neutronics)
  - Automated mesh generation from reactor metadata and combinatorial geometry
  - Direct connection to SCALE geometry and data (MG cross section processing)
  - Direct connection to MCNP input through ADVANTG
  - HDF5 output directly interfaced with VisIt
  - Built-in unit-testing and regression harness with DBC (353 separate tests)
  - Emacs-based code-development environment
  - Support for multiple external vendors
    - BLAS/LAPACK, TRILINOS (required)
    - BRLCAD, SUPERLU/METIS, SILO/HDF5 (optional)
    - MPI (toggle for parallel/serial builds)
    - SPRNG (required for MC module)
    - PAPI (optional instrumentation)

# Discrete Ordinates Methods

- **We solve the first-order form of the transport equation:**
  - **Eigenvalue form for multiplying media (fission):**

$$\hat{\mathbf{\Omega}} \cdot \nabla \psi(\mathbf{r}, \mathbf{\Omega}, E) + \Sigma(\mathbf{r}, E, T)\psi(\mathbf{r}, \mathbf{\Omega}, E) =$$
$$\int dE' \int_{4\pi} d\mathbf{\Omega}' \, \Sigma_{\mathrm{s}}(\mathbf{r}, \hat{\mathbf{\Omega}}' \cdot \hat{\mathbf{\Omega}}, E' \rightarrow E, T)\psi(\mathbf{r}, \mathbf{\Omega}', E') +$$
$$\frac{1}{k}\frac{\chi(E)}{4\pi} \int dE' \int_{4\pi} d\mathbf{\Omega}' \, \nu\Sigma_{\mathrm{f}}(\mathbf{r}, E', T)\psi(\mathbf{r}, \mathbf{\Omega}', E')$$

  - **TH coupling comes through the temperature-dependent material cross sections**

# Discrete Ordinates Methods

- The $S_N$ method is a collocation method in angle.
  - Energy is discretized in groups.
  - Scattering is expanded in Spherical Harmonics.
  - Multiple spatial discretizations are used (DGFEM, Characteristics, Cell-Balance).

$$\mathbf{L}\psi = \mathbf{MS}\phi + Q$$

$$\phi = \mathbf{D}\psi$$

- Dimensionality of operators:

$$t = N_g \times N_c \times N_u \times N_m$$

$$n = N_g \times N_c \times N_u \times N_a$$

$$(n \times n)(n \times 1) = (n \times t)(t \times t)(t \times 1) + (n \times 1)$$

# Degrees of Freedom

- **Total number of unknowns in solve:**

$$\text{unknowns} = N_g \times N_c \times N_u \times N_a \times N_m$$

- **An ideal (conservative) estimate.**

$$N_g = 238$$
$$N_c = 1 \times 10^9$$
$$N_u = 4$$
$$N_m = 16$$
$$N_a = 288$$

$$\boxed{\text{unknowns} \geq 4 \times 10^{15}}$$

# Traditional $S_N$ Solution Methods

- Traditional $S_N$ solutions are divided into outer iterations over energy and inner iterations over space-angle.
- Generally, accelerated Gauss-Seidel or SOR is used for outer iterations.
- Eigenvalue forms of the equation are solved using *Power Iteration*
- In Denovo we are motivated to look at more advanced solvers
  - Improved robustness
  - Improved efficiency
  - Improved parallelism

# Reformulating the Problem

$$\phi^{n+1} = \mathbf{DL}^{-1}(\mathbf{MS}\phi^{n} + q)$$

operate by $\mathbf{DL}^{-1}$ to get Source Iteration

$$x^{n+1} = (\mathbf{I} - \mathbf{A})x^{n} + b$$

which is really fixed-point (Richardson) iteration

$$\mathbf{I} - \mathbf{A} = \mathbf{DL}^{-1}\mathbf{MS}$$

iteration matrix for Source Iteration

$$(\mathbf{I} - \mathbf{DL}^{-1}\mathbf{MS})\phi = \mathbf{DL}^{-1}q$$

put in form $\mathbf{A}x = b$, we can use non-stationary iterative methods (**Krylov subspace**) to solve this linear problem

The inversion of $\mathbf{L}$ is done using a wavefront solver that is implemented by solving for $\phi$ in the direction of particle flow → Transport Sweep.

# Krylov Methods

- Krylov methods are more robust than stationary solvers
  - Uniformly stable (preconditioned and unpreconditioned)
- Can be implemented *matrix-free*
- More efficient
  - Source iteration spectral radius $\rho(0)\epsilon = \frac{\sigma_s}{\sigma}\epsilon$
  - Gauss-Seidel spectral radius $\rho(0)\epsilon = (\mathbf{T} - \mathbf{S}_D)^{-1}\mathbf{S}_U\epsilon$
- There is no coupling in Krylov methods
  - Gauss-Seidel imposes coupling between rows in the matrix
  - Krylov has no coupling; opportunities for enhanced parallelism

# Physics Dictates Convergence

- The Gauss-Seidel spectral radius for uniform graphite is 0.9812 = slow convergence
- Systems that are block-dense in energy are sparse in energy-space-angle
- Ideal candidates for Krylov methods

Iron-D2O-Graphite block energy S matrix

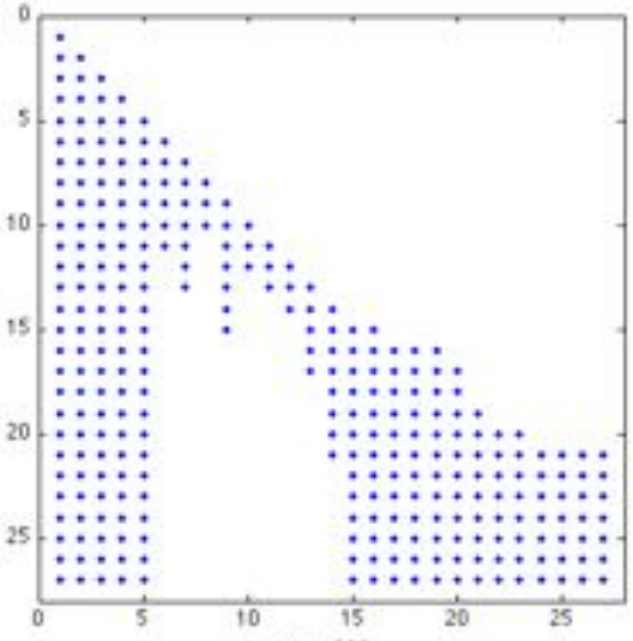


Iron-D2O-Graphite energy-space-angle S matrix

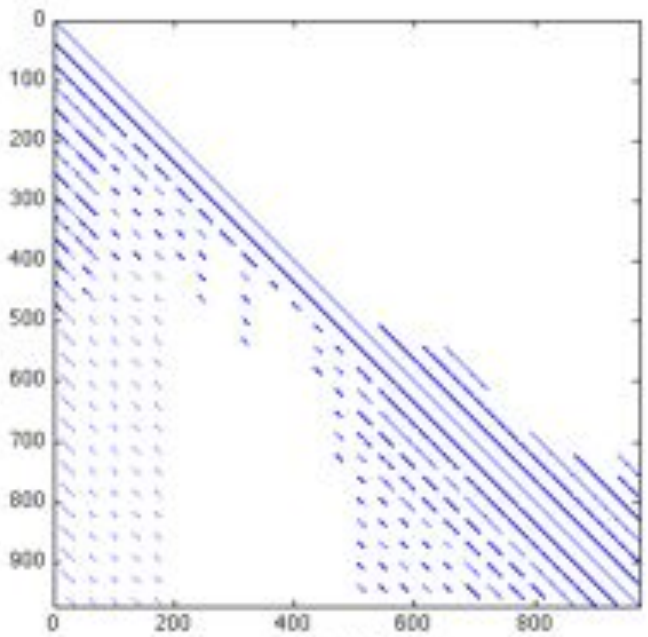

# Multigroup Transport Problem

- Using Gauss-Seidel requires the solution of *G* within-group equations (using Krylov iteration) in each GS iteration
- Alternatively, the full energy system can be solved by Krylov iteration ($\mathbf{T}=\mathbf{DL}^{-1}$)

$$(\mathbf{I}-\mathbf{TMS})\phi=\mathbf{T}q$$

$$\left(\mathbf{I}-\begin{bmatrix}\mathbf{T}_0 & \cdots & 0\\ 0 & \ddots & 0\\ 0 & 0 & \mathbf{T}_G\end{bmatrix}\begin{bmatrix}\mathbf{M} & \cdots & 0\\ 0 & \ddots & 0\\ 0 & 0 & \mathbf{M}\end{bmatrix}\begin{bmatrix}\mathbf{S}_{00} & \cdots & \mathbf{S}_{0G}\\ \vdots & \ddots & \vdots\\ \mathbf{S}_{G0} & \cdots & \mathbf{S}_{GG}\end{bmatrix}\right)\begin{bmatrix}\phi_0\\ \vdots\\ \phi_G\end{bmatrix}=\begin{bmatrix}\mathbf{T}_0 q_0\\ \vdots\\ \mathbf{T}_G q_G\end{bmatrix}$$

# Eigenvalue Problem

- **The eigenvalue problem has the following form**

$$(\mathbf{I} - \mathbf{TMS})\phi = \frac{1}{k}\mathbf{TM}\chi\mathbf{f}^T\phi$$

- **Expressed in standard form**

$$\mathbf{A}x = kx$$

$$\mathbf{A} = (\mathbf{I} - \mathbf{TMS})^{-1}\mathbf{TM}\chi\mathbf{f}^T \quad x = \phi \qquad \text{Energy-dependent}$$

$$\mathbf{A} = \mathbf{f}^T(\mathbf{I} - \mathbf{TMS})^{-1}\mathbf{TM}\chi \quad x = \mathbf{f}^T\phi \qquad \text{Energy-indepedent}$$

- **The traditional way to solve this problem is with *Power Iteration***

# Advanced Eigenvalue Solvers

- **We can use Krylov (Arnoldi) iteration to solve the eigenvalue problem more efficiently**

$$y^k = \mathbf{A}v^k$$

Matrix-vector multiply and sweep $z^k = \mathbf{TM}\chi\mathbf{f}^T v^k$

Multigroup fixed-source solve $(\mathbf{I} - \mathbf{TMS})y^k = z^k$

- **Shifted-inverse iteration (Raleigh-Quotient Iteration) has been developed (using Krylov to solve the shifted multigroup problem in each eigenvalue iteration)**

$$(\mathbf{I} - \mathbf{TM}\underbrace{(\mathbf{S} + \mu\mathbf{F})}_{\text{block-dense}})\phi = (\lambda - \mu)\mathbf{TMF}\phi$$

# Solver Taxonomy

**The innermost part of each solver are transport sweeps**

$$y = \mathbf{T}z = \mathbf{D}\underbrace{\mathbf{L}^{-1}z}$$

$$\mathbf{L}\psi = z$$

"It's turtles all the way down…"

**Eigenvalue Solvers**

Power iteration
Arnoldi
Shifted-inverse

**Multigroup Solvers**

Gauss-Seidel
Residual Krylov
Gauss-Seidel + Krylov

**Within-group Solvers**

Krylov
Residual Krylov
Source iteration

# KBA Algorithm

sweeping in direction of particle flow

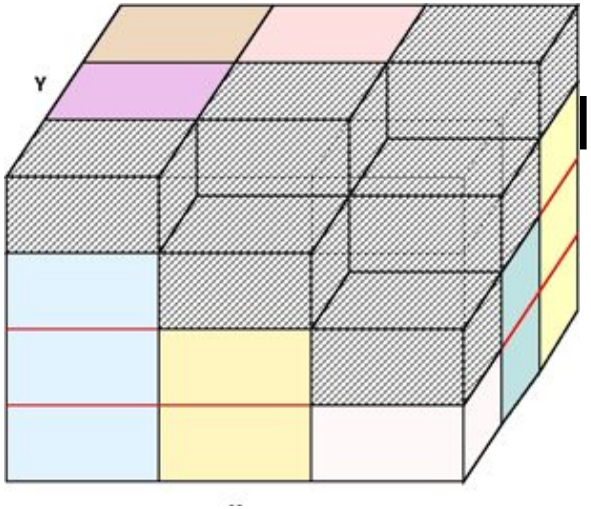


lire

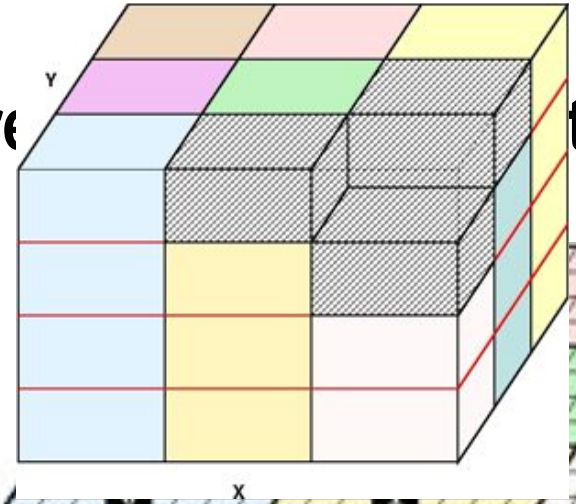


thm

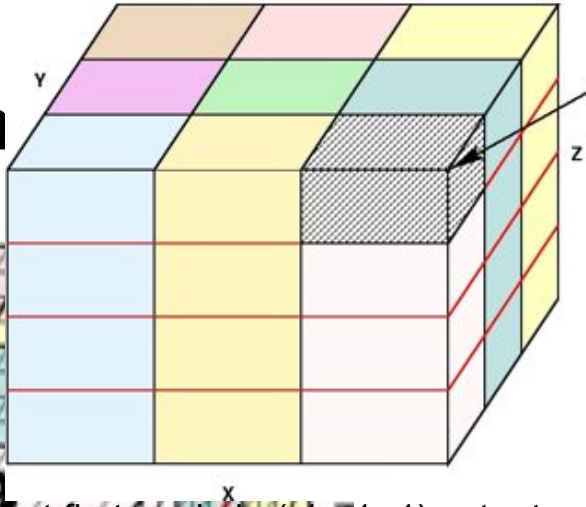


Start first angle in (-1,-1,-1) octant

Begin next angle in octant

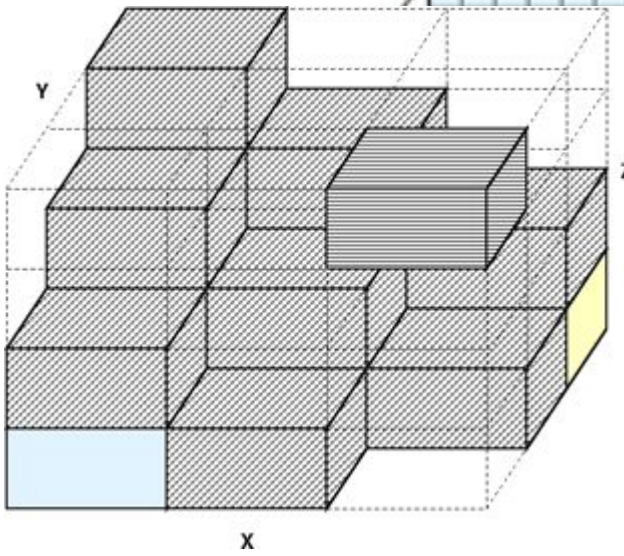


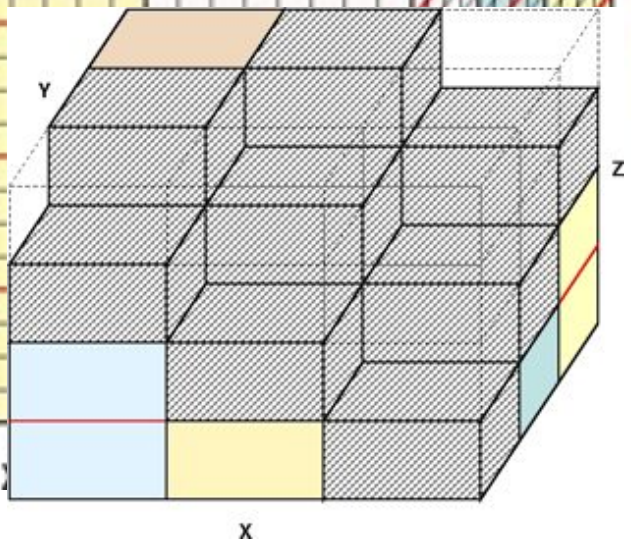

# Parallel Performance

## Angular Pipelining

- Angles in ± z directions are pipelined
- Results in 2×M pipelined angles per octant
- Quadrants are ordered to reduce latency

$$\epsilon_{\max} = \frac{2MB_K}{2MB_K + P_I + P_J - 2}$$

6 angle pipeline ($S_4$; M = 3)

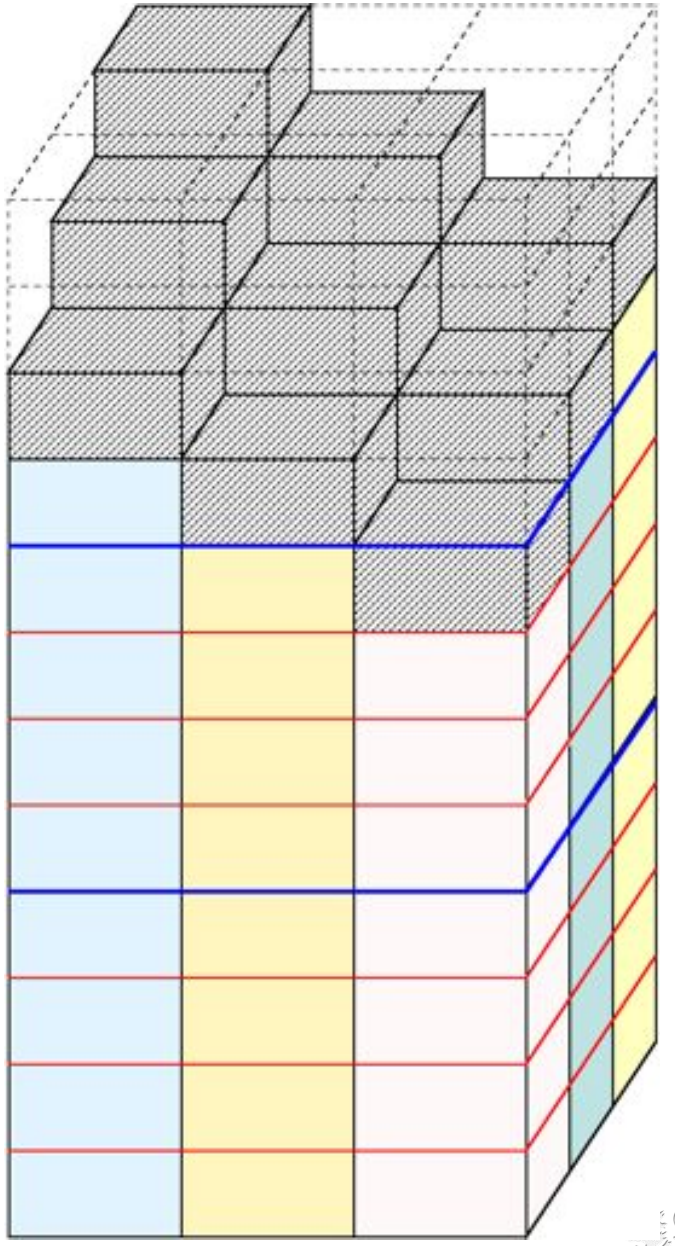

# KBA Reality

**KBA does not achieve close to the predicted maximum**

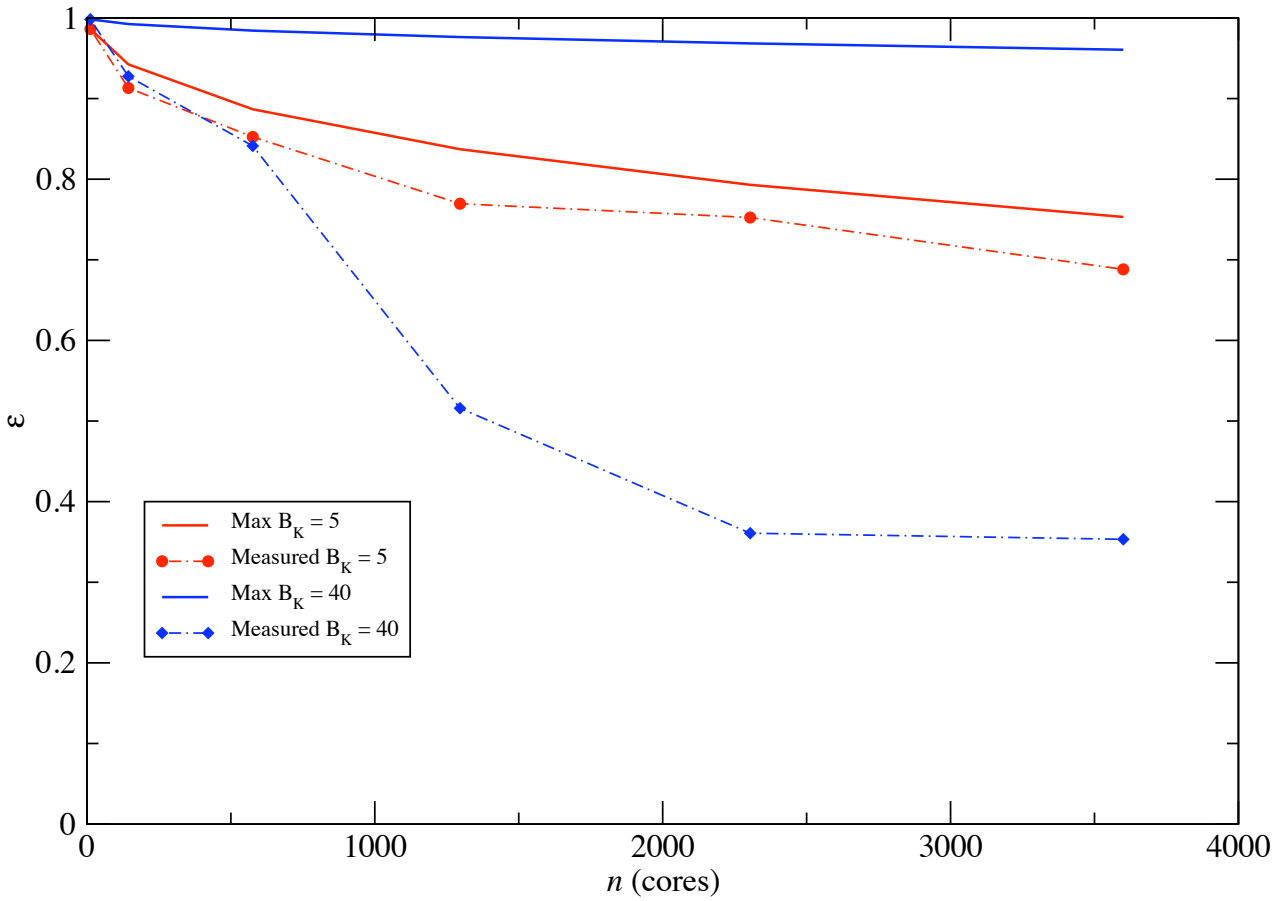


- **Communication latency dominates as the block size becomes small**
- **Using a larger block size helps achieve the predicted efficiency but,**
  - **Maximum achievable efficiency is lower**
  - **Places a fundamental limit on the number of cores that can be used for any given problem**

# Efficiency vs Block Size

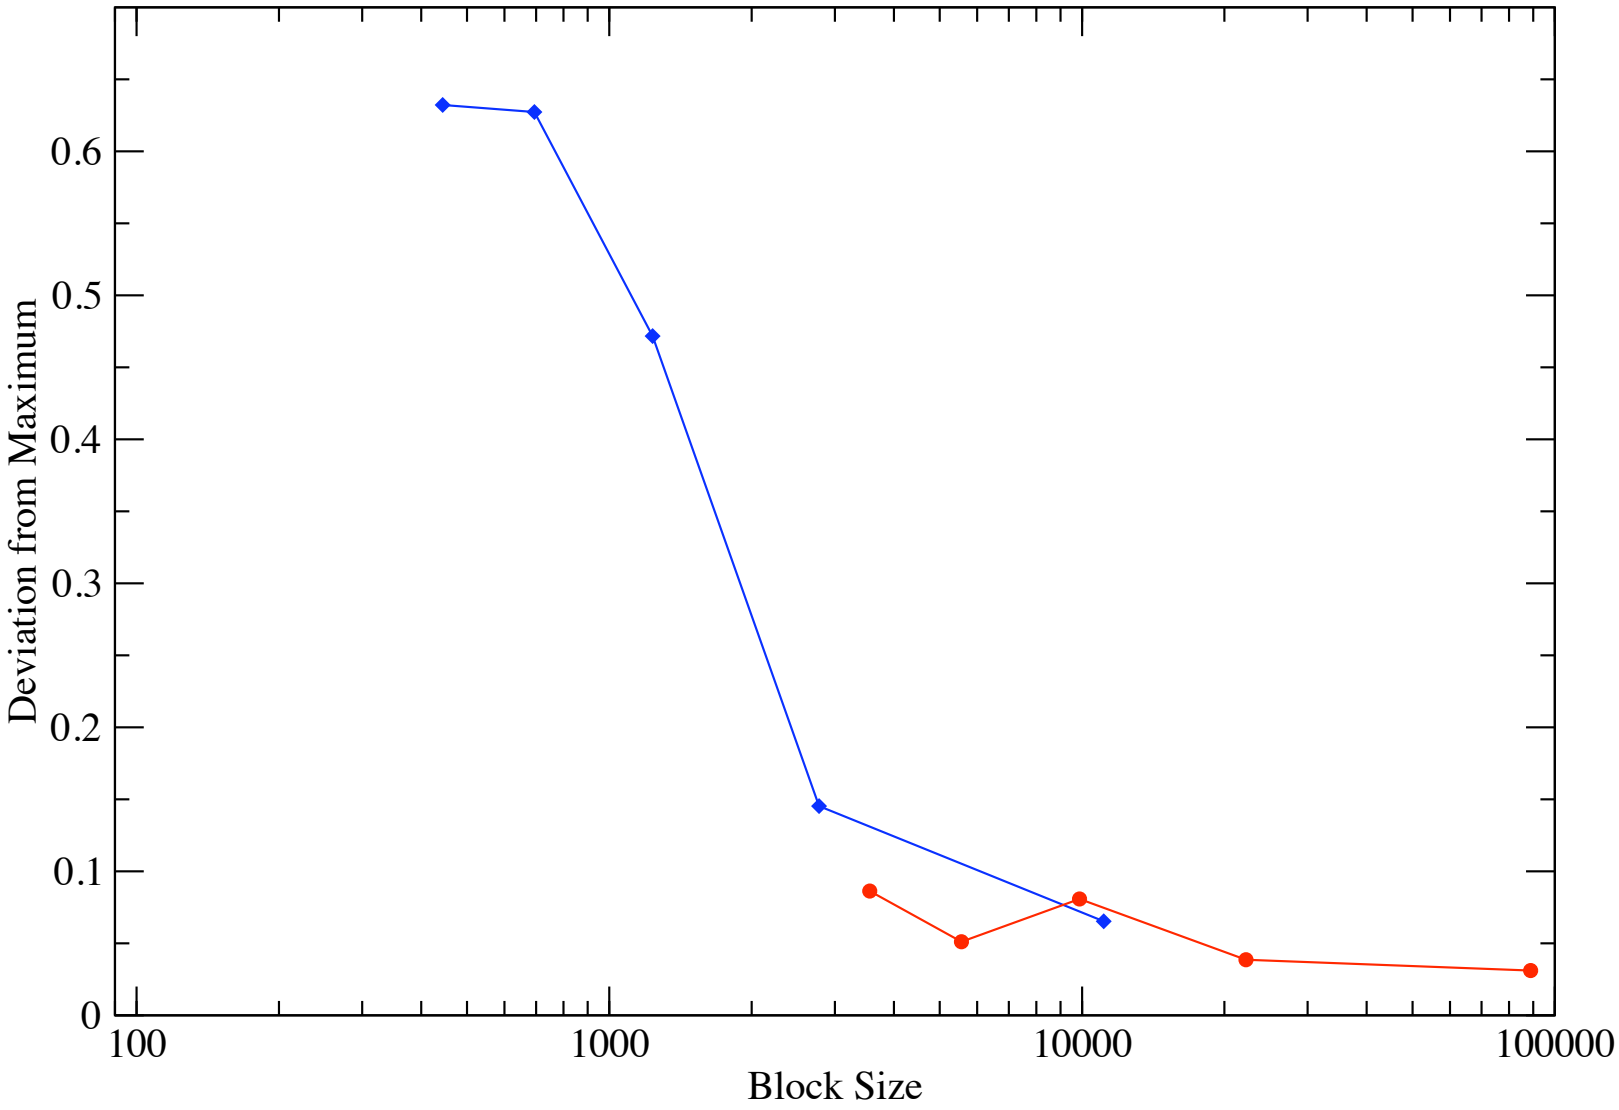

# Overcoming Wavefront Challenge

- This behavior is systemic in any wavefront-type problem
  - Hyberbolic aspect of transport operator
- We need to exploit parallelism beyond space-angle
  - Energy
  - Time
- Amortize the inefficiency in KBA while still retaining direct inversion of the transport operator

# Multilevel Energy Decomposition

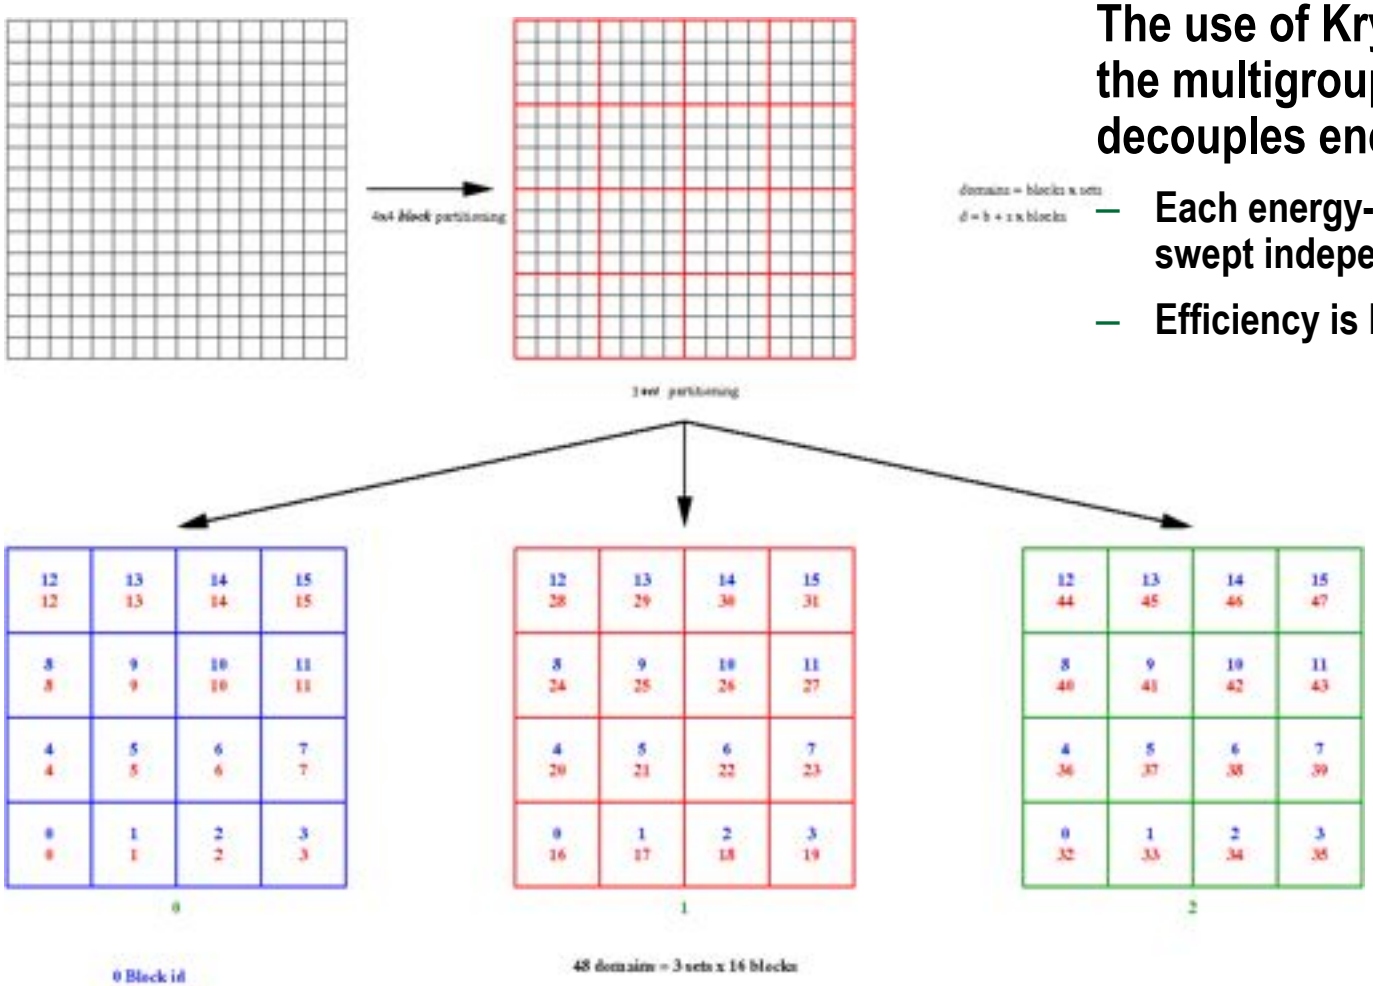


**The use of Krylov methods to solve the multigroup equations effectively decouples energy**

- Each energy-group $S_N$ equation can be swept independently
- Efficiency is better than Gauss-Seidel

# Multilevel Summary

- Energy decomposed into sets.
- Each set contains blocks constituting the entire spatial mesh.
- The total number of domains is

$$\text{domains} = \text{sets} \times \text{blocks}$$

- KBA is performed for each group in a set across all of the blocks.
  - Not required to scale beyond *O(1000)* cores.
- Scaling in energy across sets should be linear.
- Allows scaling to O(100K) cores and enhanced parallelism on accelerators.

# Whole Core Reactor Problem

**PWR-900 Whole Core Problem**

- 2 and 44-group, homogenized fuel pins
- 2×2 spatial discretization per fuel pin
- 17×17 fuel pins per assembly
- 289 assemblies (157 fuel, 132 reflector) – high, med, low enrichments
- Space-angle unknowns:
  - 233,858,800 cells
  - 128 angles (1 moment)
  - 1 spatial unknown per cell

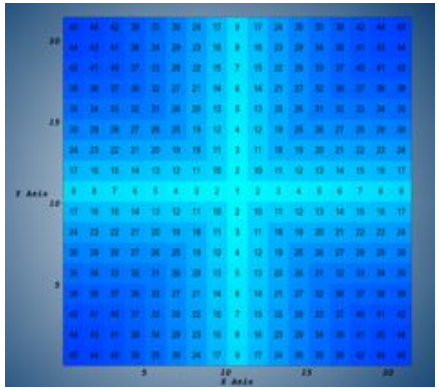

17×17 assembly

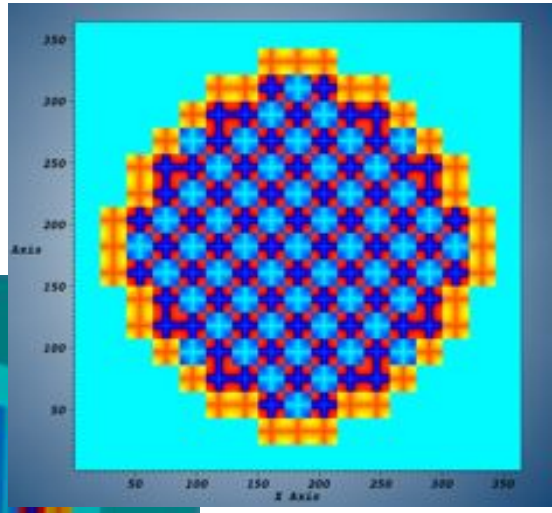

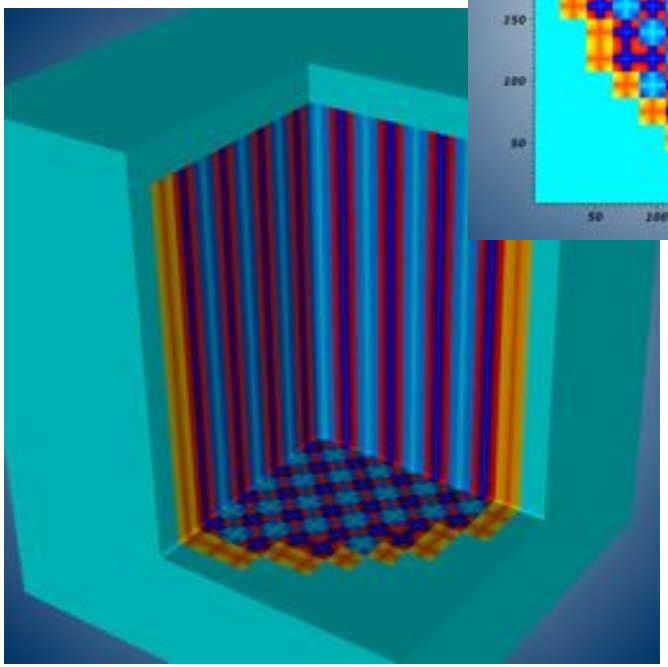

# Results

| Solvers | Blocks | Sets | Domains | Solver Time (min) |
|---|---|---|---|---|
| PI + MG GS (2-grid preconditioning) | 17,424 | 1 | 17,424 | 150.15 |
| PI + MG Krylov | 17,424 | 1 | 17,424 | 52.99 |
| Arnoldi + MG Krylov | 17,424 | 1 | 17,424 | 23.62 |
| Arnoldi + MG Krylov | 17,424 | 2 | 34,848 | 12.81 |

Total unknowns = 59,867,852,800
Number of groups = 2
$k_{eff}$ tolerance = 1.0e-5

- **The GS solver cannot use more computational resource for a problem of this spatial size**
  - **Simply using more spatial partitions will not reduce time to solution**
  - **Problem cannot effectively use more cores to run a higher fidelity problem in energy**
- **PI + MG Krylov will scale with sets similarly to Arnoldi, they just use different outer iteration strategies**

# Strong Scaling on XT5

Optimized communication gave performance boost to 100K core job, number of sets = 11

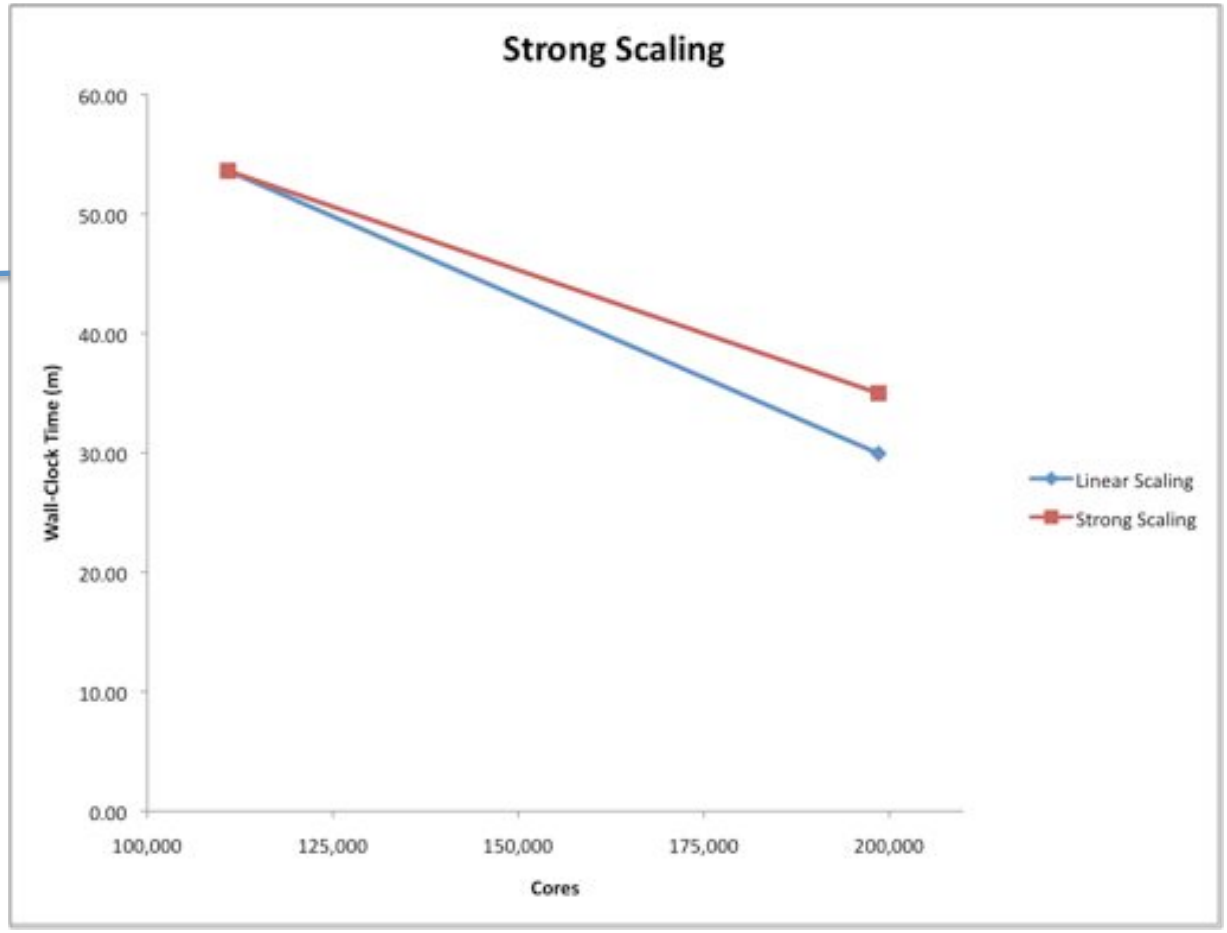


- Communication improvements were significant at 100K core level (using 11 sets).
- They do not appear to scale to 200K core. Why?
  - Multiset reduction each iteration imposes a constant cost!

# Scaling Limitations

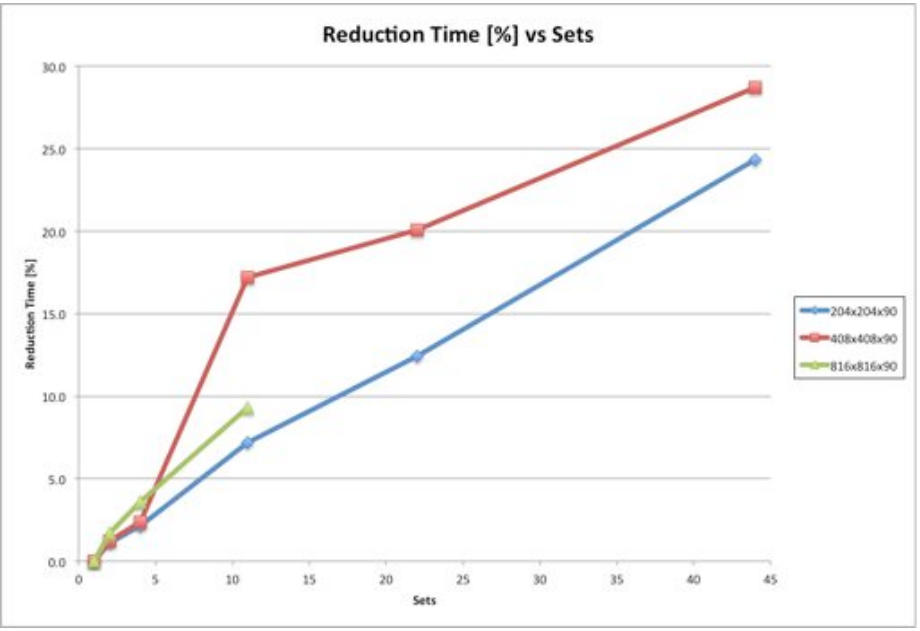


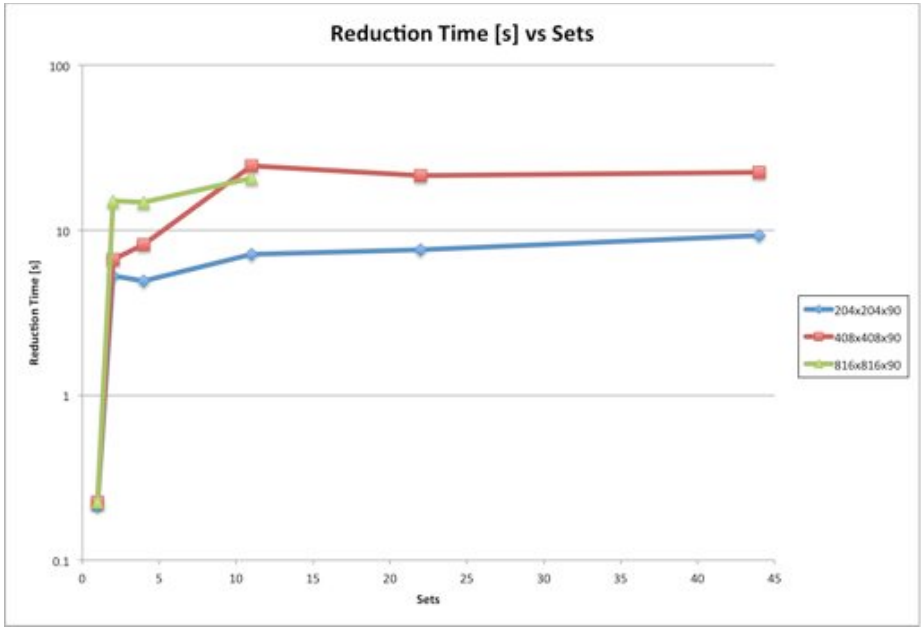


- Reduction across groups each iteration imposes a “flat” cost
- Only way to reduce this cost is to increase the work per set each iteration (more angles)
  - Generally the work in space will not increase because we attempt to keep the number of blocks per domain constant
- **However, we were able to replace a global-reduction with a reduced-scatter that considerably reduced the reduction cost per outer iteration**

# Improved Scaling on XK6

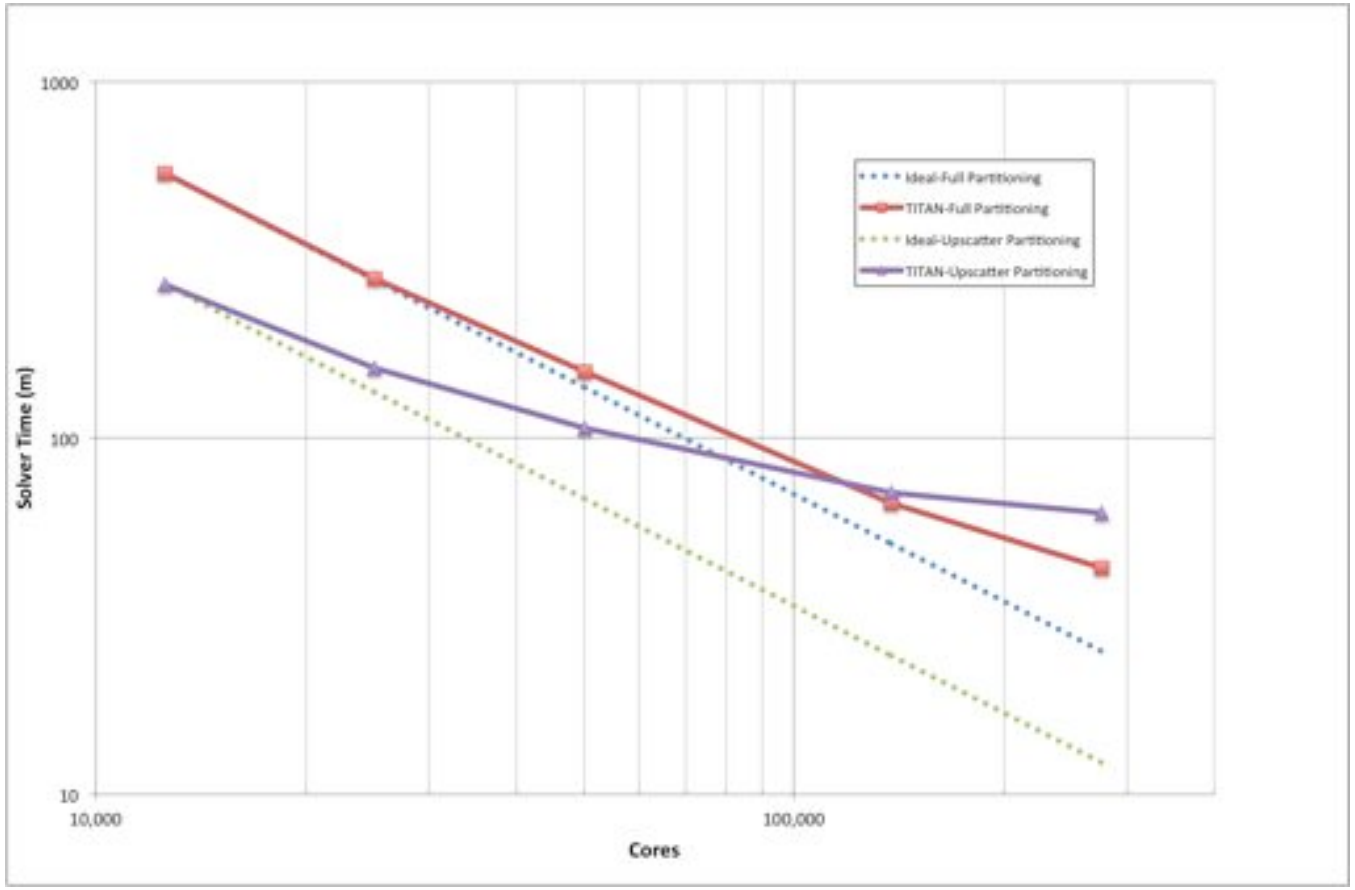


- Constant number of blocks = 12,544
- 44 total groups/22 coupled groups

Full partitioning scales well to 275K cores

Improved interconnects + reduce-scatter have dramatically reduced global reduction cost

Upscatter partitioning more efficient at lower set counts

Roll-over occurs between 4 and 11 sets (5 and 2 groups per set) where serial work in GS solver dominates

# Peak Performance on XK6

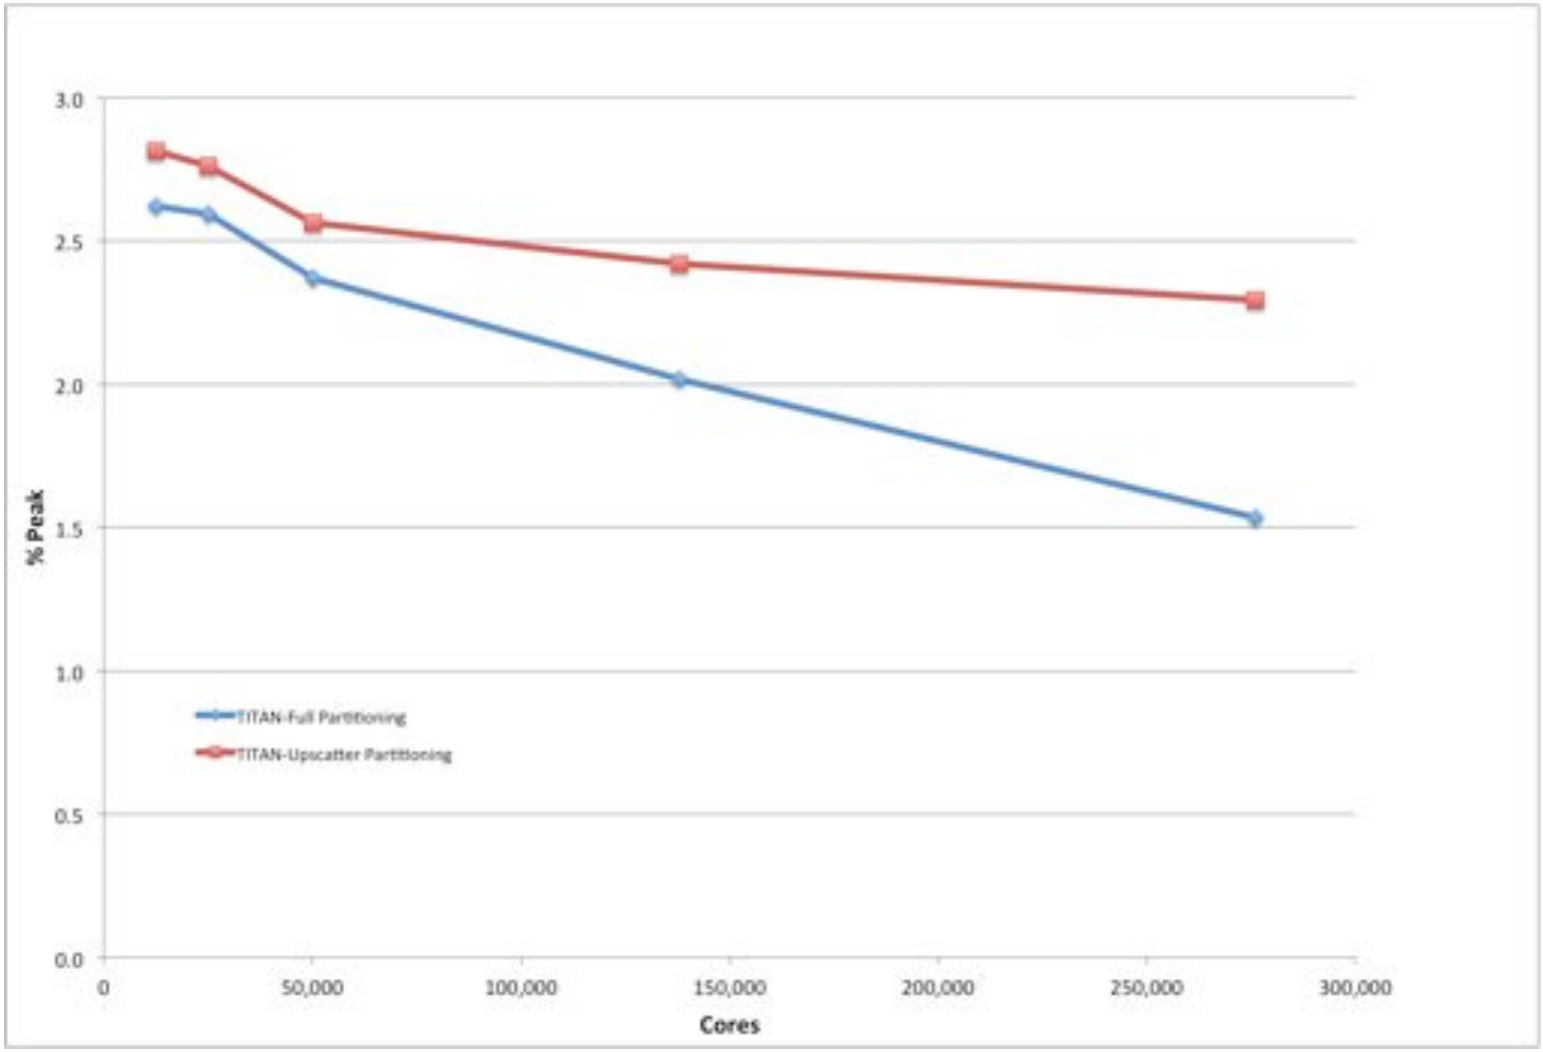

# GPU Sweep Kernel

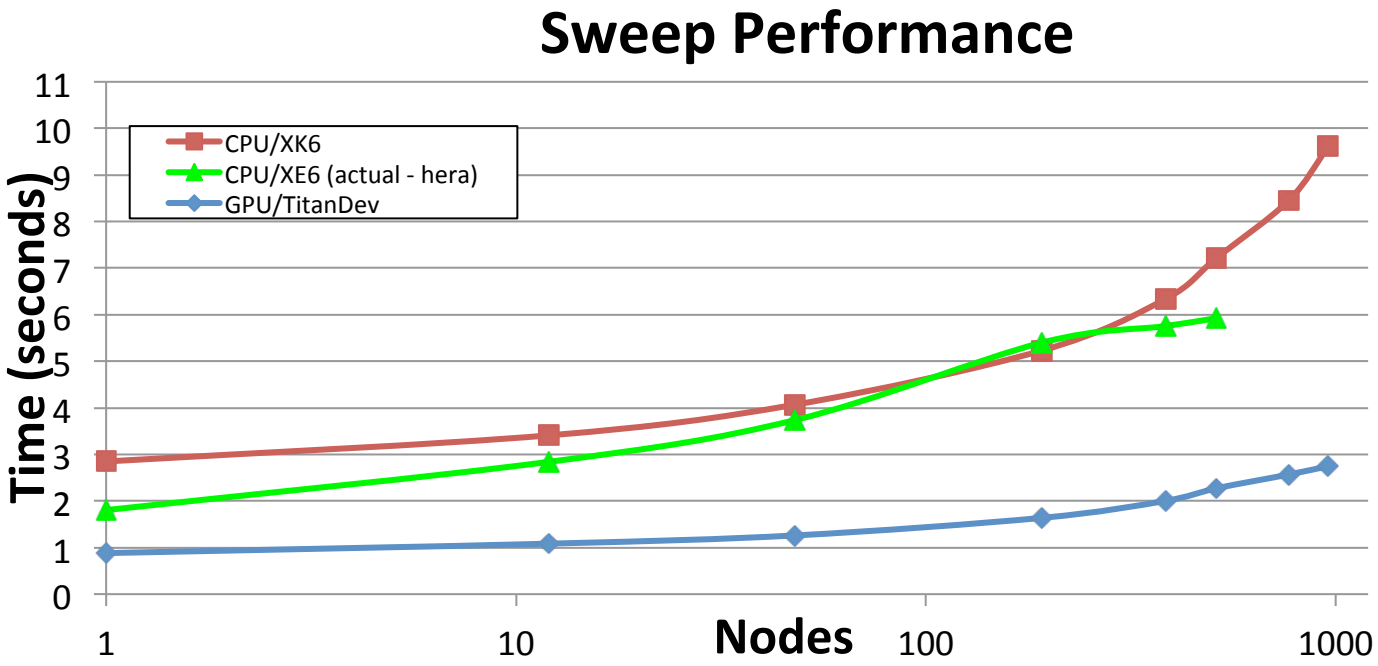


| Performance Improvement factors | | GPU |
|---|---|---|
| | | *XK6 Fermi* |
| CPU | *XK6 / Interlagos* | **3.5** |
| | *XE6 / dual Interlagos* | **3.3** |

- Krylov multigroup solvers allow the space-angle sweeps to be performed over all groups concurrently.
- Ideal for exploiting thread-based concurrency on GPUs
- We do a space-angle sweep for all groups on the GPU.

# CASL Quarter Core Simulations

- We have run (XT5) a set of 3-D ¼ core simulations on real reactor models
- Varying numerics parameters we found that it is feasible to run fully consistent transport on 3-D cores
  - The problems run
  - We get good solutions
- However…


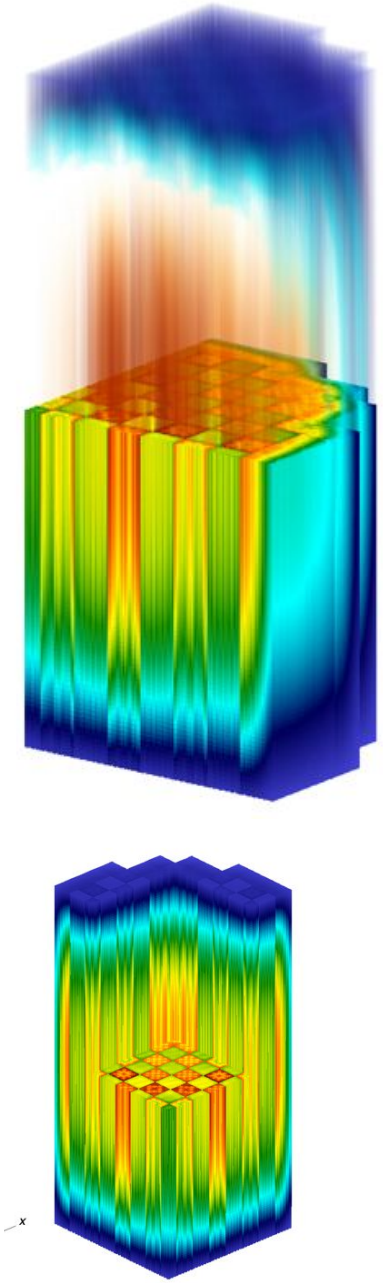

# Projections

**Where we want to be**

- Reproduce fidelity of 2D calculations using consistent 3D methods
- Produce all state-points for a depletion cycle in *O(8 hours)*
- *O(72)* state points per cycle
- Steady-state, coupled Neutronics simulation with TH feedback = $O(10^{19})$ unknowns

**Where we are**

- Assuming 2% peak, we can solve $1.7 \times 10^{13}$ unknowns/hour (XT5)
- This means we can solve a much-reduced 3D problem ($O(10^{15})$ unknowns) in 175 hours
- This assumes status quo on a 1PF XT5 machine

# Projections

## What this means

- To reach 2D fidelity at 3D we need to solve ~$10^4$ × more unknowns
- So to run all state points in a full day at this fidelity using existing code and methods requires
  - ~ 141 Eflops

## However, there is hope

- Consulting with industry, a fully consistent 3D calculation in 1 week would be acceptable (factor of 7)
- We can still gain very valuable insight into challenge problems without reproducing full 2D fidelity (factor of 150-200)
- Yet, we still need >100 PF to run a full depletion cycle

# Conclusion

## The final piece

- Utilize GPUs to get more efficiency out of the hardware
- If early projections hold, we can potentially get a factor of 3-4 improvement by exploiting sweep kernels on the GPU
- Further solver research (multigrid-in-energy) shows promise for reducing iteration counts as well
- Getting a factor of 3-4 from GPUs means that a 30-40 PF machine could allow fully consistent, 3-D neutronics simulations that could be used to address CASL challenge problems

# Questions

# One Group Iterations

**Gauss-Seidel** Iteration in energy

$$\mathbf{L}_g \psi_g^{k+1} = \mathbf{MS}_{gg} \phi_g^{k+1} + \sum_{g=0}^{g-1} \mathbf{MS}_{gg'} \phi_{g'}^{k+1} + \sum_{g=g+1}^{G} \mathbf{MS}_{gg'} \phi_{g'}^{k} + Q_g$$

reduces to a series of one-group solves
(within-group **inner** iterations)

$$\mathbf{L}_g \psi_g = \mathbf{MS}_{gg} \phi_g + \bar{Q}_g$$

↓

up and down-scatter rolled into source

inners have the general form

$$\mathbf{L}\psi = \mathbf{MS}\phi + q$$

# MG Krylov Preconditioning

- Each MG Krylov iteration involves two-steps
  - preconditioning: $\mathbf{G}z^k = v^k$
  - matrix-vector multiply: $v^{k+1} = (\mathbf{I} - \mathbf{TM\bar{S}})z^k$
- At end of iteration we must apply the preconditioner one last time to recover $w^{k+1}$
- We use a simple 1-D multigrid preconditioner in energy:
  - 1-pass V-cycle

$$z \leftarrow \bar{\mathbf{G}}(z^{2h}, v^{2h})$$

# V-Cycle Relaxation

- We are investigating both weighted-Jacobi

$$(\mathbf{I} - \mathbf{TMS}_D)z^{n+1} = \mathbf{TM}(w\mathbf{S} - \mathbf{S}_D)z^n + wv + (1-w)z^n$$

- And weighted-Richardson relaxation schemes

$$z^{n+1} = w\mathbf{TMS}z^n + wv + (1-w)z^n$$

- Energy-parallelism is largely preserved

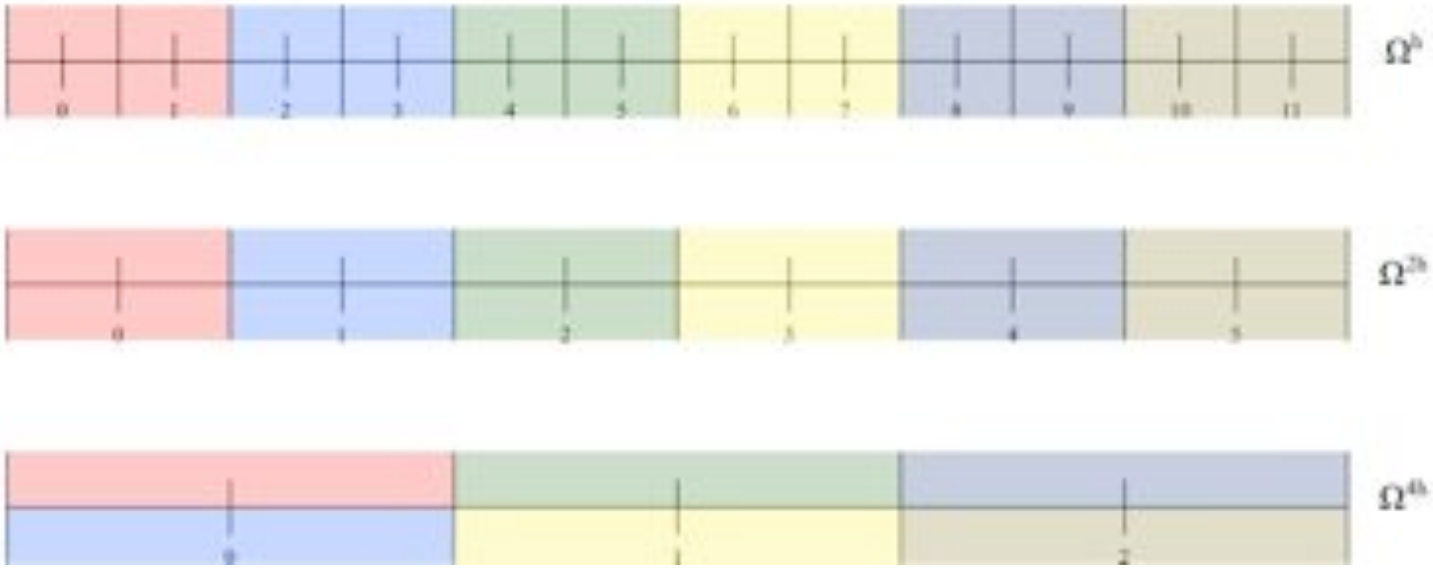

# Virtual Reactor Simulation

- **Neutronics is one part of a complete reactor simulation**