3-D Full Core Power Calculations for PWRs

Thomas M. Evans

Greg Davidson

Rachel Slaybaugh

SciDAC 2010, Chattanooga, TN

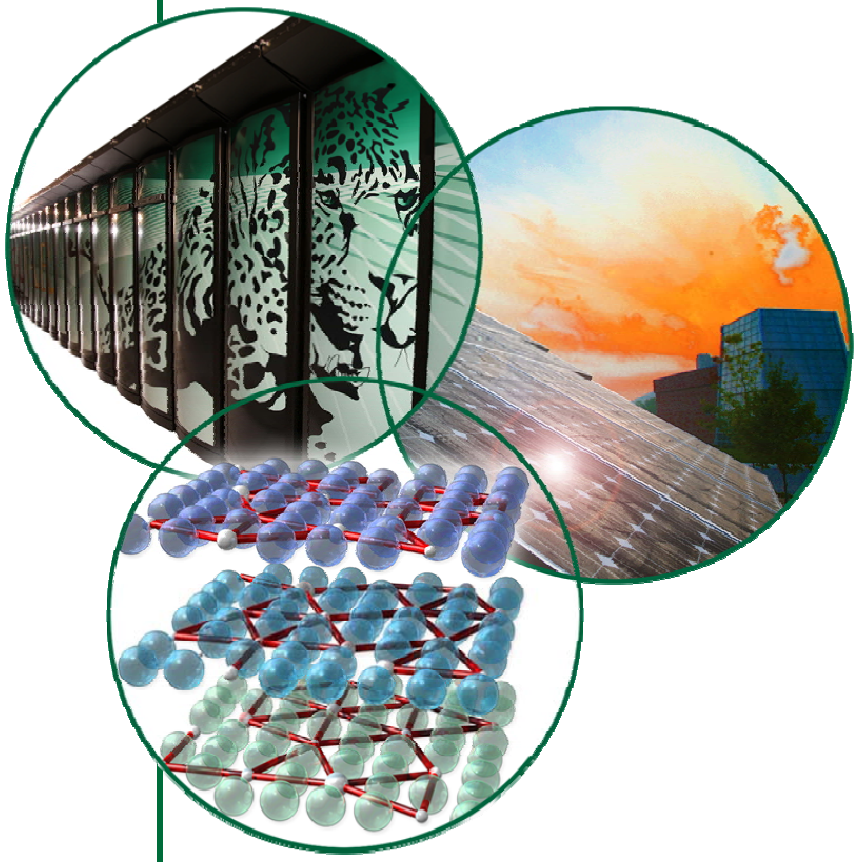

U.S. DEPARTMENT OF ENERGY

OAK RIDGE NATIONAL LABORATORY

MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

Collaborators

NSTD Staff

- Kevin Clarno
- Douglas Peplow

Support

- ORNL LDRD (2 projects)
- DTRA, DNDO, FEMA
- 2010 INCITE, 2009 Early Science
- NM&S Hub (CASL)
- OLCF-3 (ASCR)

Students

- Josh Hykes (North Carolina State)
- Alyssa Stafford (Texas A&M)
- Bruno Turcksin (Texas A&M)
- Ahmad Ibrahim (Wisconsin)

OLCF Support

- Dave Pugmire
- Sean Ahern
- Wayne Joubert
- Chris Baker

Outline

- Computational Science Drivers for Transport in NE
- Denovo Capabilities Overview
- Transport Solvers
- KBA Sweep Algorithm
- Multilevel Parallel Decompositions
- Summary

Multiscale in Space: All physics must model the massive heterogeneous structure

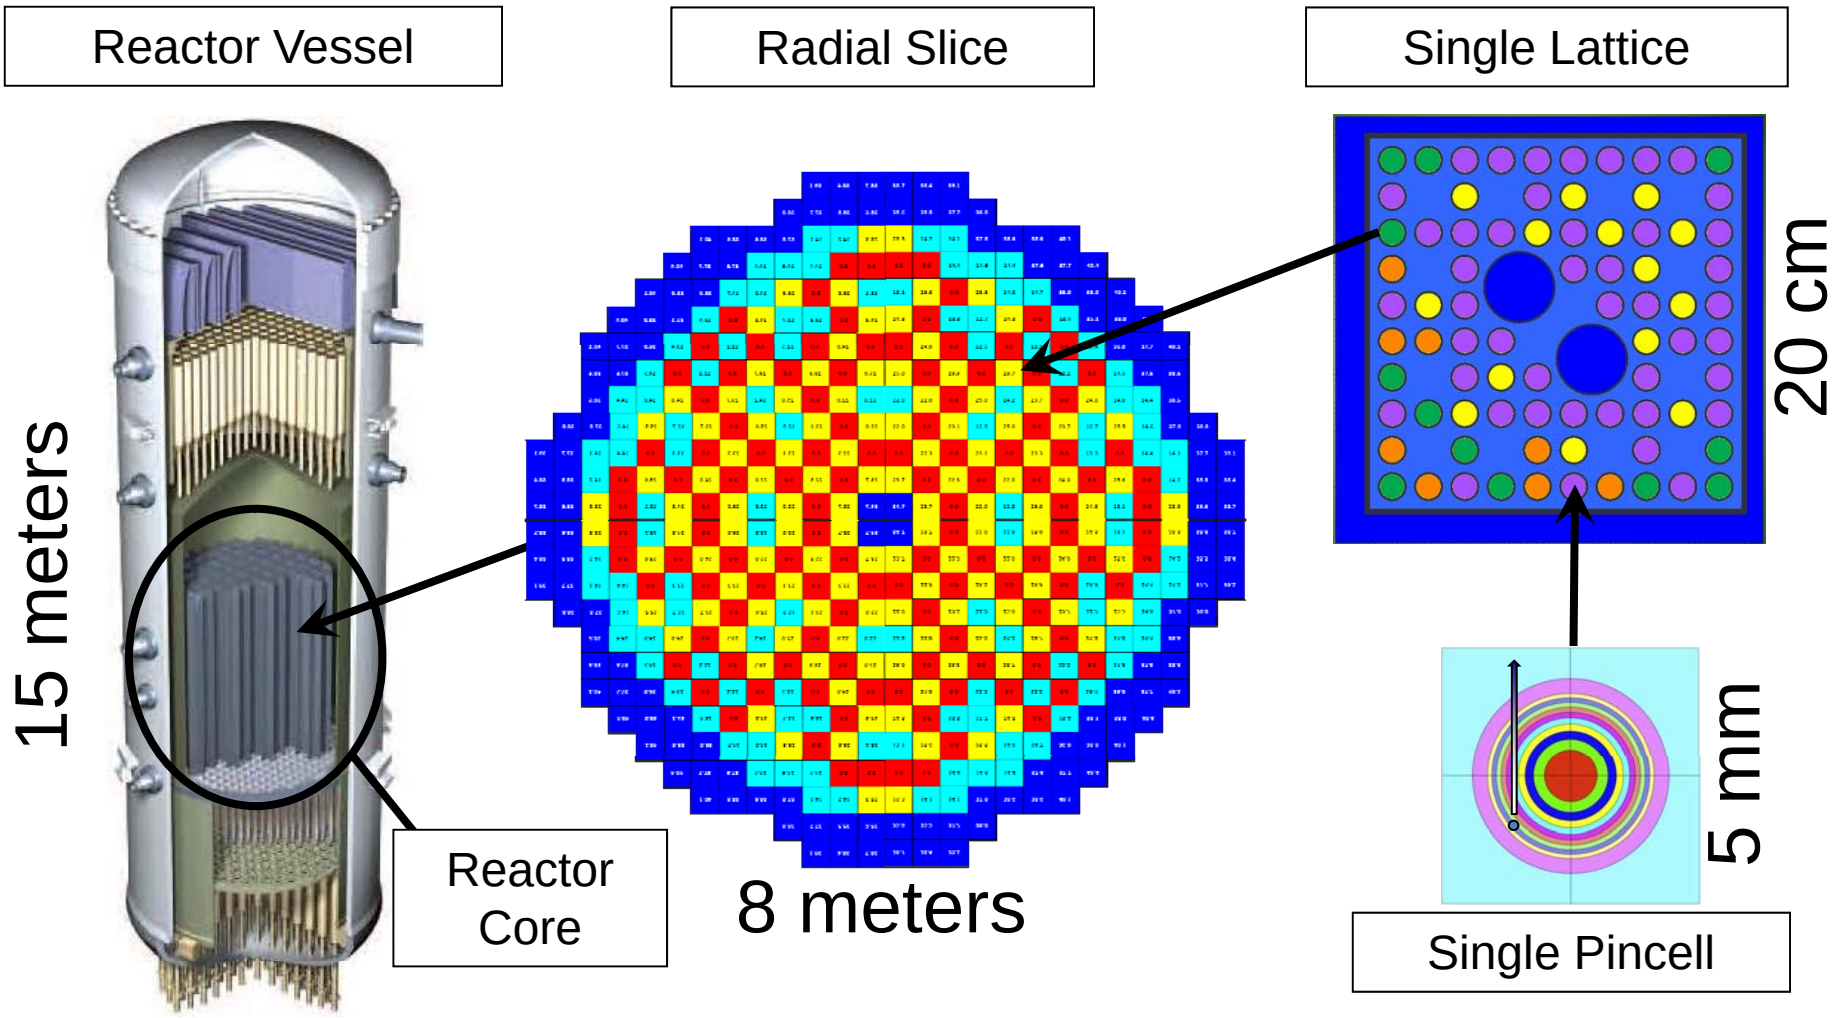

Current State-of-the-Art in Reactor Neutronics

General Electric ESBWR

pin cell

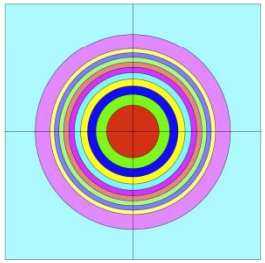

- 0/1-D transport
- High energy fidelity (10^{2-5} unknowns)
- Approximate state and BCs

lattice cell

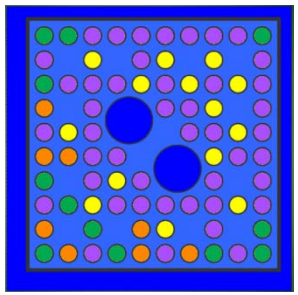

- 2-D transport
- Moderate energy fidelity (7-10^2 groups)
- Approximate state and BCs
- Depletion with spectral corrections
- Space-energy homogenization

core

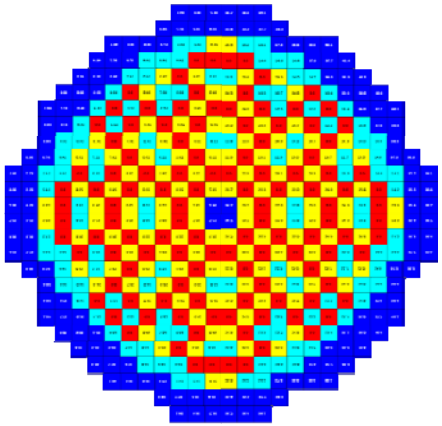

- 3-D diffusion
- Low energy fidelity (2-4 groups)
- Homogeneous lattice cells
- Heterogeneous flux reconstruction
- Coupled physics

Science Drivers for Neutronics

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

- Crud and distortion
- Control rod insertion
- *Ab initio* design

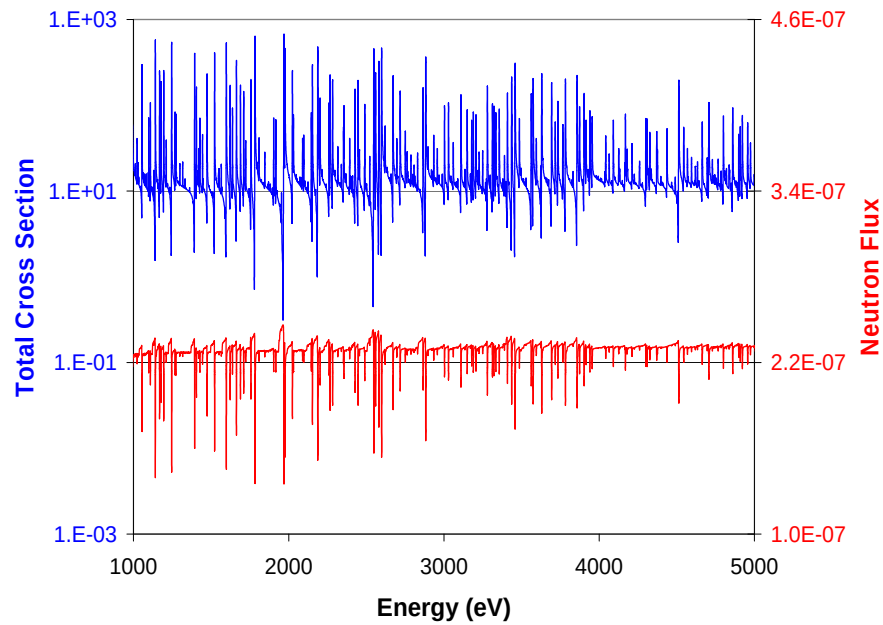

Denovo Capabilities

- State of the art transport methods
  - 3D, non-uniform, regular grid SN
  - Multigroup energy, anisotropic Pn 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
    - Galerkin

- 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)
  - Eigenvalue solvers
    - Power iteration (with rebalance)
      - CMFD in testing phase
    - Krylov (Arnoldi)
    - Shifted-inverse iteration in development

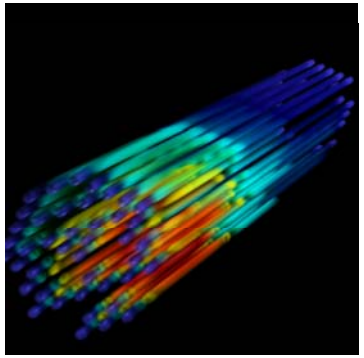

Power distribution in a BWR assembly

Denovo Capabilities

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

> 5M CPU hours on Jaguar with 2 bugs

2010 INCITE Award
Uncertainty Quantification for Three Dimensional Reactor Assembly Simulations, 8 MCPU-HOURS
2010 ASCR Joule Code
2009-2011 2 ORNL LDRDs

- Advanced visualization, run-time, and development environment
  - 3 front-ends (HPC, SCALE, Python-bindings)
  - Direct connection to SCALE geometry and data
  - Direct connection to MCNP input through ADVANTG
  - HDF5 output directly interfaced with VisIt
  - Built-in unit-testing and regression harness with DBC
  - Emacs-based code-development environment
  - Support for multiple external vendors
    - GSL, 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):

$$\hat{\mathbf{\Omega}} \cdot \nabla \psi(\mathbf{r}, \hat{\Omega}, E) + \sigma(\mathbf{r}, E)\psi(\mathbf{r}, \hat{\Omega}, E) = \int dE' \int_{4\pi} d\hat{\Omega}' \; \sigma_s(\mathbf{r}, \hat{\mathbf{\Omega}}' \cdot \mathbf{\Omega}, E' \rightarrow E)\psi(\mathbf{r}, \hat{\Omega}', E') + \frac{1}{k}\frac{\chi(E)}{4\pi} \int dE' \int_{4\pi} d\hat{\Omega}' \; \nu\sigma_f(\mathbf{r}, E')\psi(\mathbf{r}, \hat{\Omega}', E')$$

  - Fixed source form:

$$\hat{\mathbf{\Omega}} \cdot \nabla \psi(\mathbf{r}, \hat{\Omega}, E) + \sigma(\mathbf{r}, E)\psi(\mathbf{r}, \hat{\Omega}, E) = \int dE' \int_{4\pi} d\hat{\Omega}' \; \sigma_s(\mathbf{r}, \hat{\mathbf{\Omega}}' \cdot \mathbf{\Omega}, E' \rightarrow E)\psi(\mathbf{r}, \hat{\Omega}', E') + \frac{1}{4\pi} q_0(\mathbf{r}, \hat{\Omega}, E)$$

Application Areas

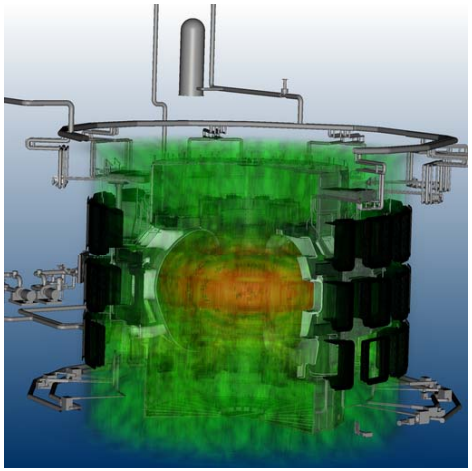

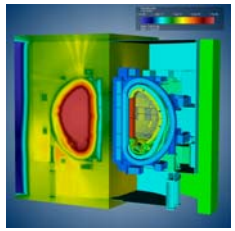

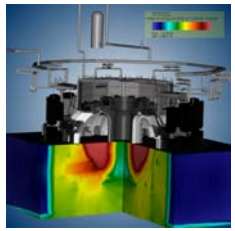

ITER component performance/shielding

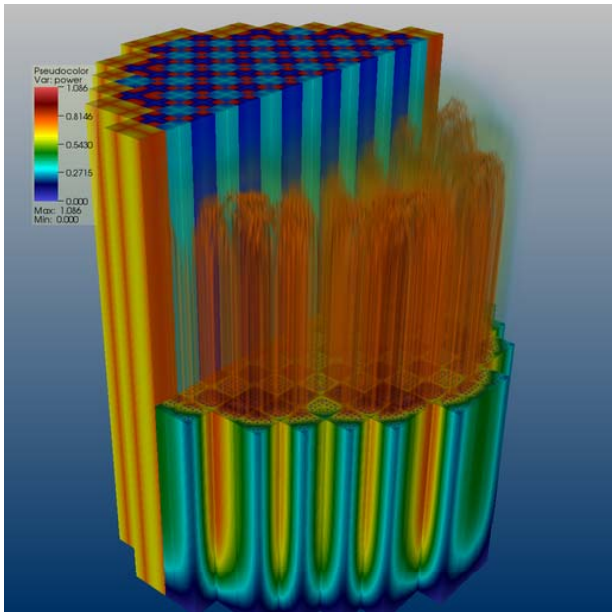

Reactor Analysis

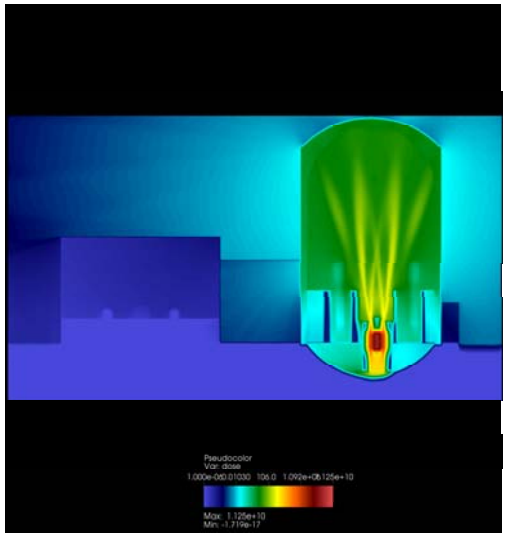

Facility shielding/dosimetry

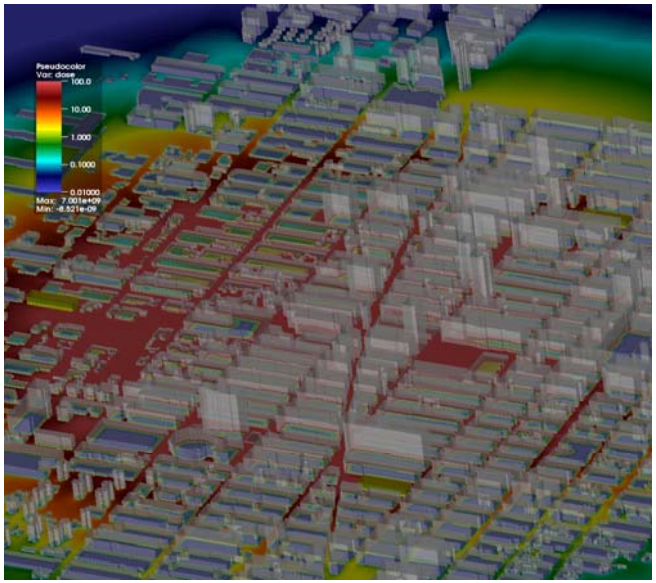

National defense/
Urban modeling

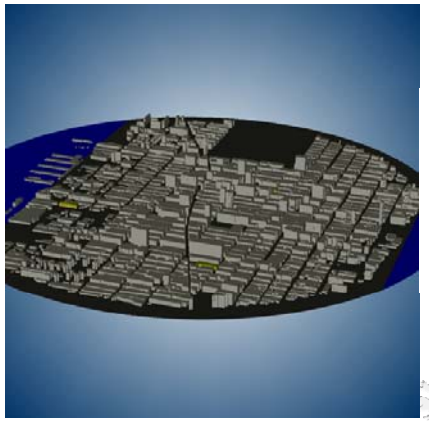

Discrete Ordinates Methods

- 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:

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

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

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

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

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

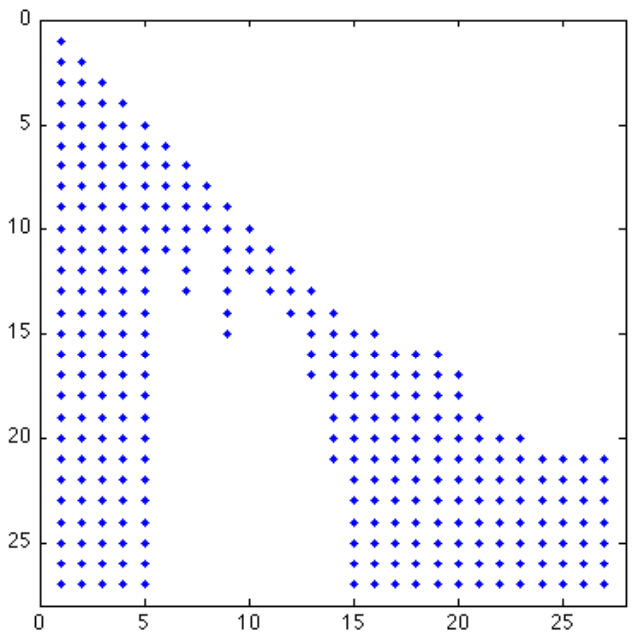

Iron-D2O-Graphite energy-space-angle S matrix

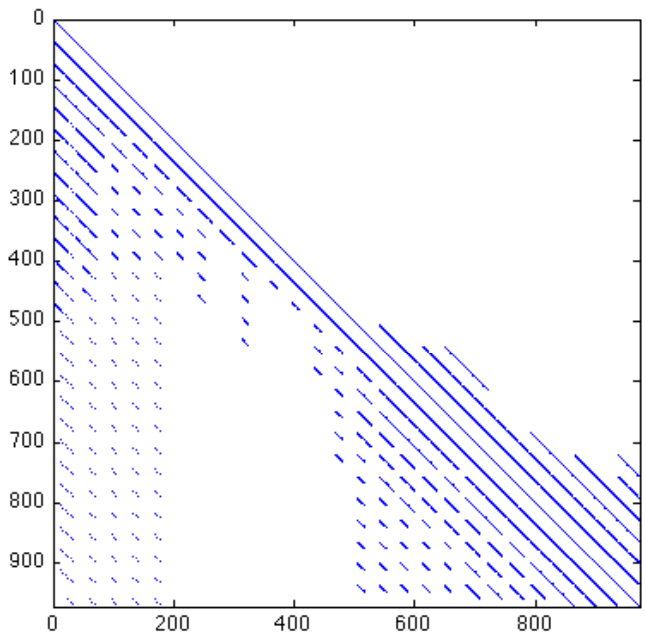

Multigroup Transport Problem

- Operate on the S_N equations by $\mathbf{T}=\mathbf{D}\mathbf{L}^{-1}$

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

- 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

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

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

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

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

$$\mathbf{A} = \mathbf{f}^T(\mathbf{I} - \mathbf{TMS})^{-1}\mathbf{TM}\chi \quad x = \mathbf{f}^T\phi$$ 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 $\quad z^k = \mathbf{TM}\chi \mathbf{f}'^T v^k$

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

- Shifted-inverse iteration (Raleigh-Quotient Iteration) is also being 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

Advanced Solvers

- Regular GMRES performs very well without acceleration for most problems

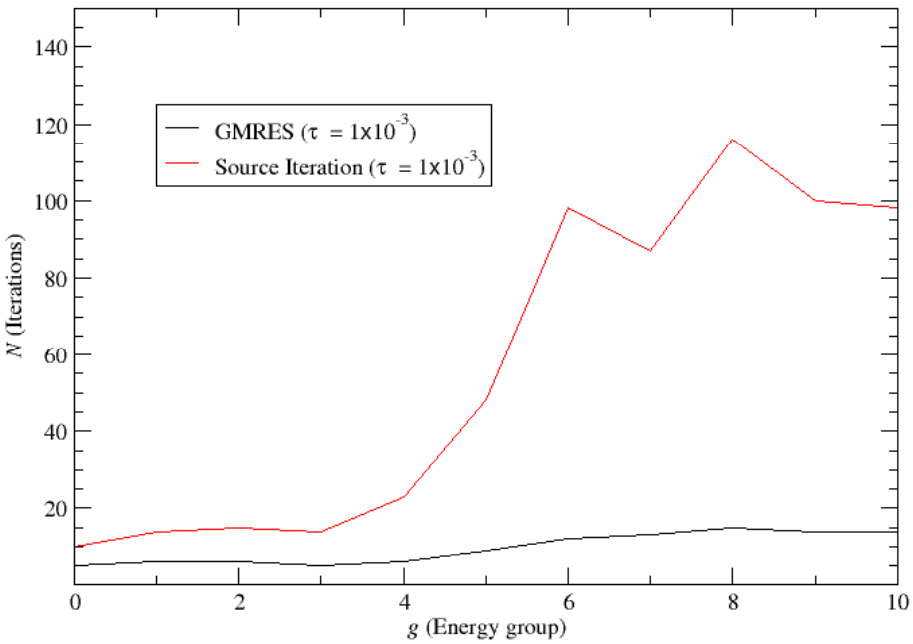

- Transport two grid (TG) acceleration of upscatter is highly efficient and stable

Method	Acc. S_N	GS Iterations	Normalized Time
Gauss-Seidel	-	175	1.000
Transport TG	8	15	0.113
Transport TG	4	14	0.097
Transport TG	2	13	0.086

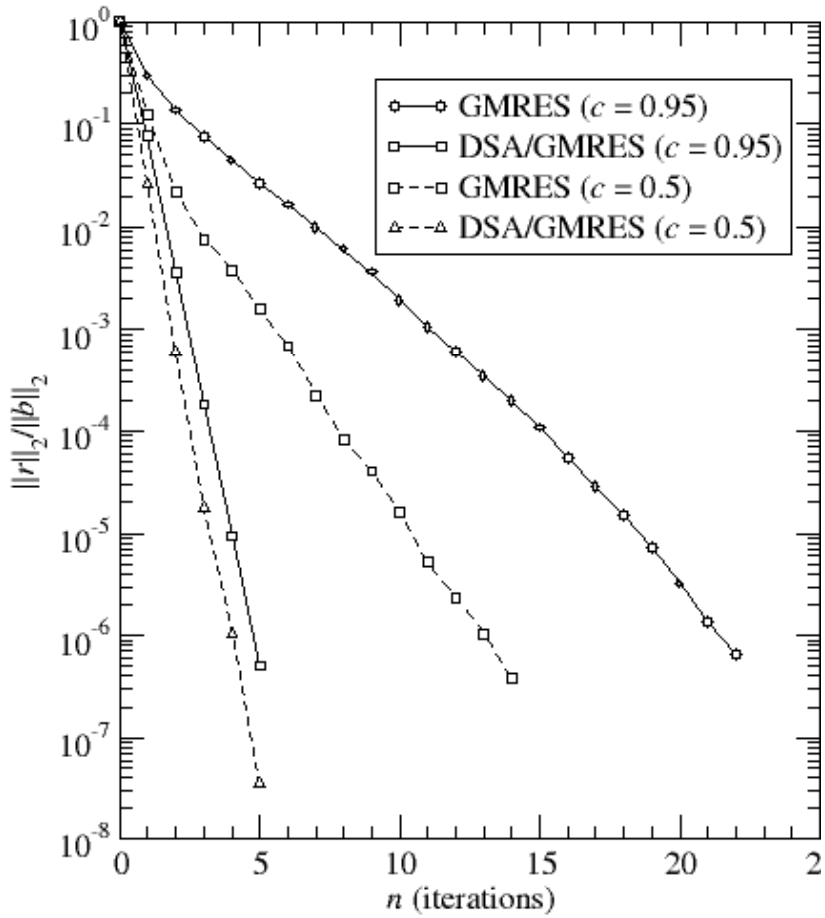

- Preconditioning Krylov with DSA is unconditionally stable
- Excellent results for high scattering and/or tight convergence

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)

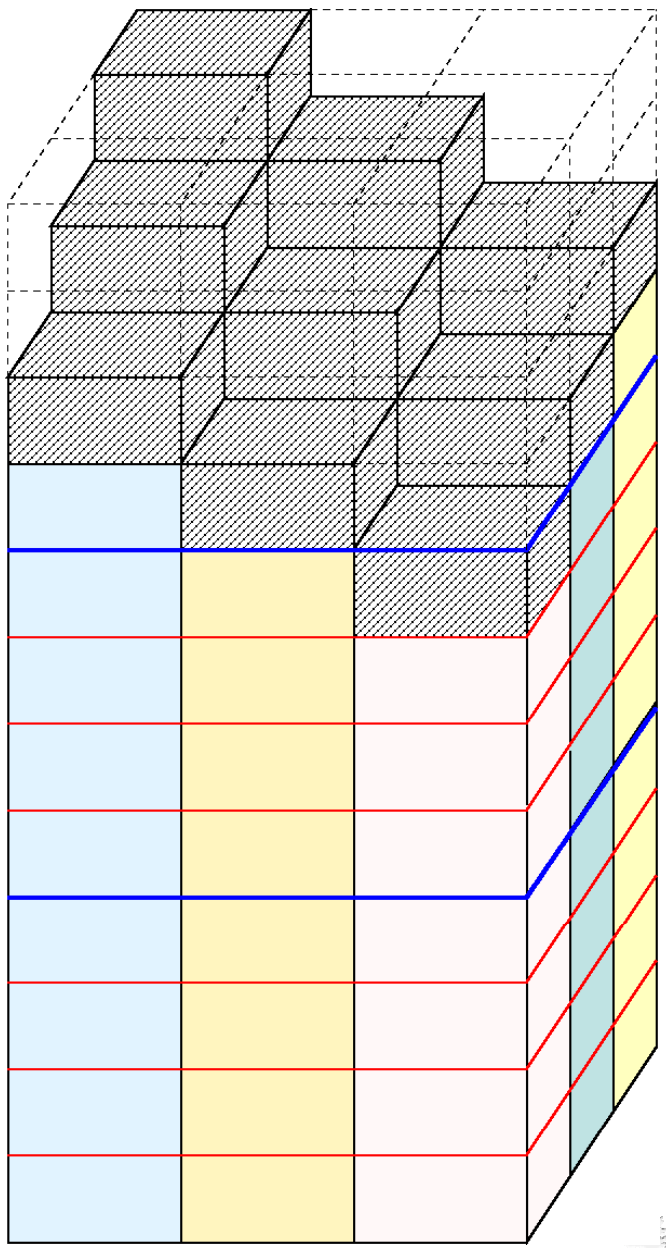

KBA Reality

KBA does not achieve close to the predicted maximum

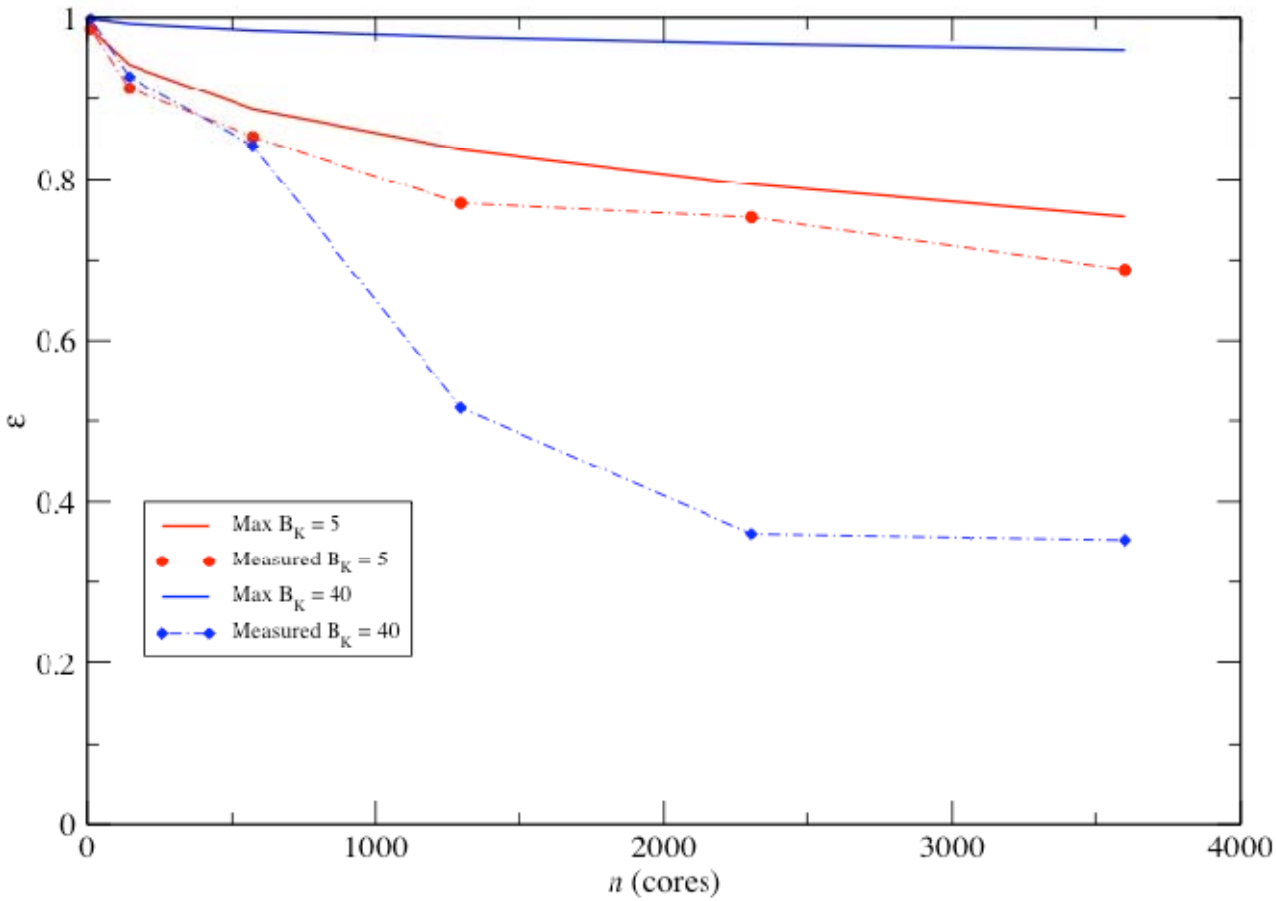

- Communication latency dominates as the block size becomes small
- Using a larger block size helps achieve the predicted efficency but,
  - Maximum achievable efficiency is lower
  - Places a fundamental limit on the number of cores that can be used for any given problem

Efficiency vs Block Size

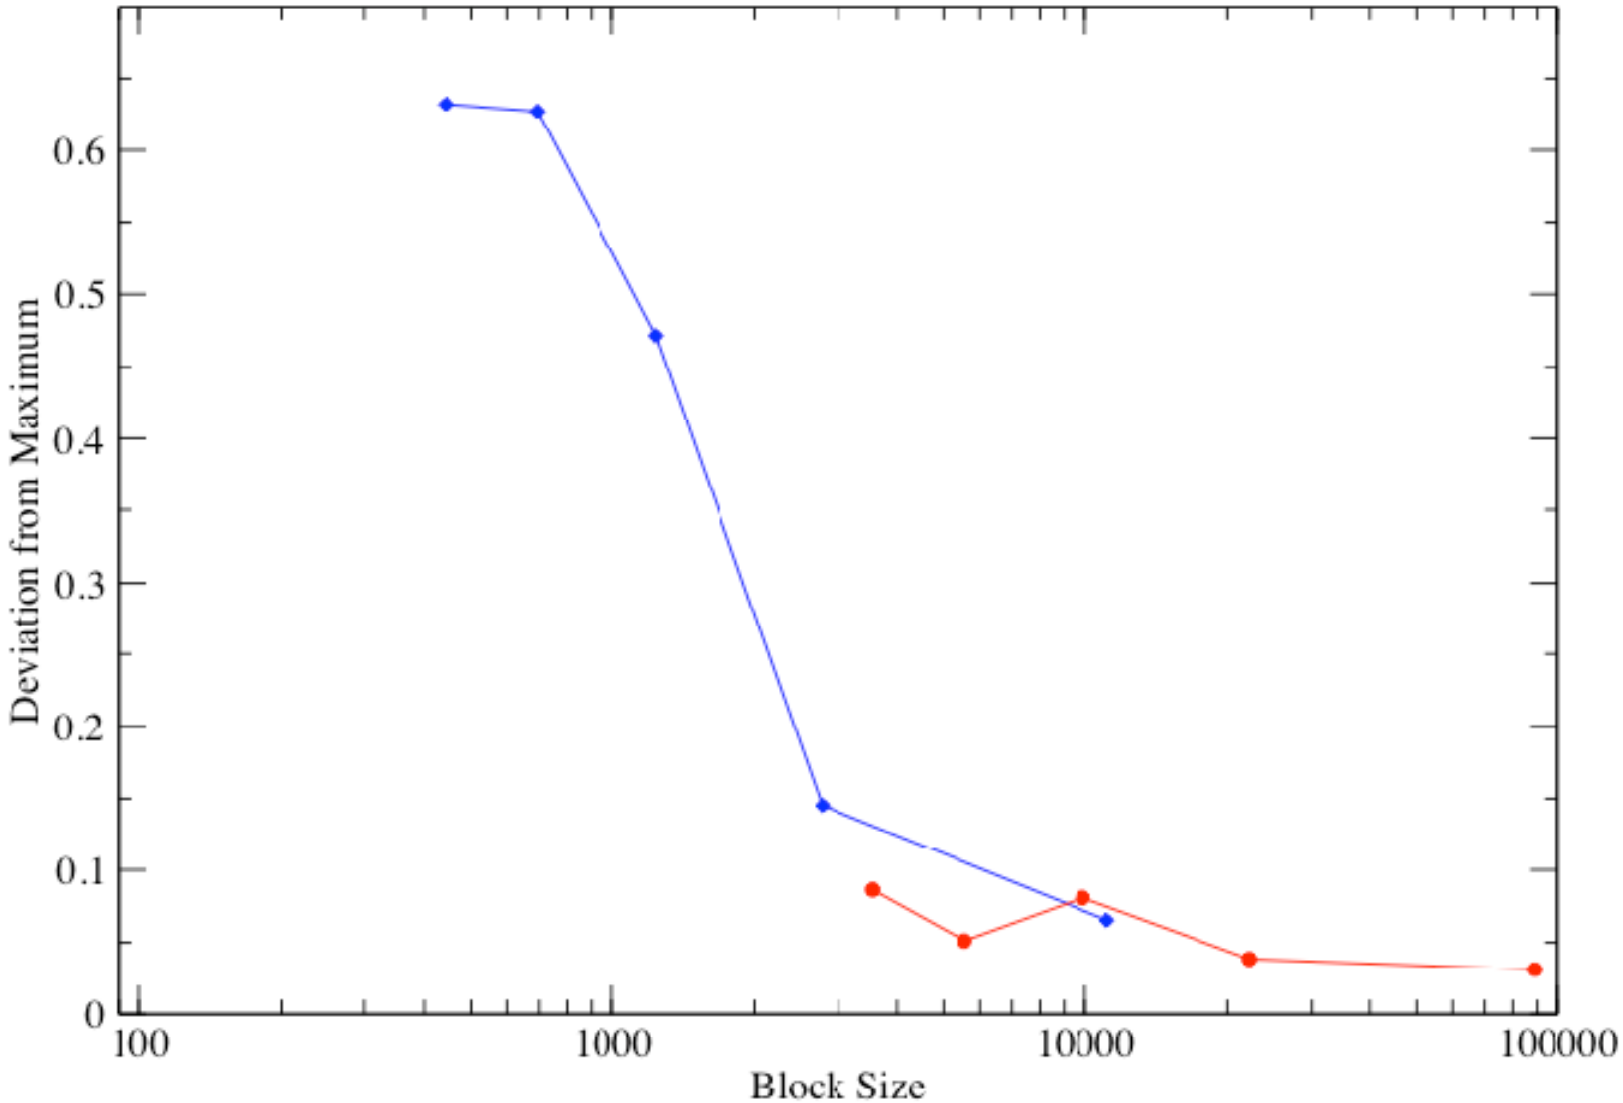

Overcoming Wavefront Challenge

- 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

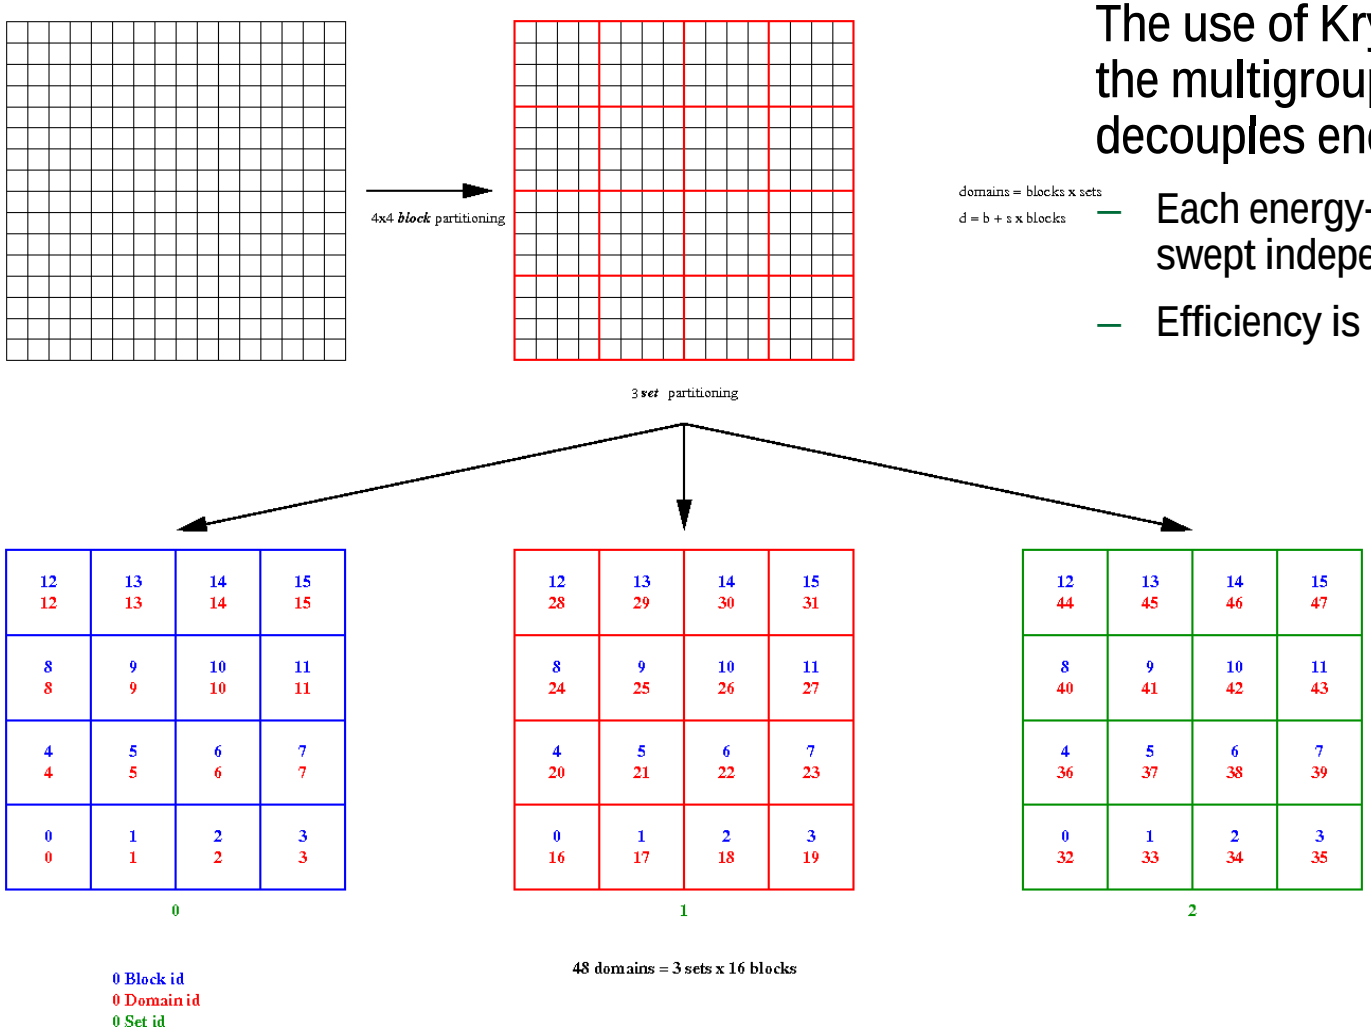

The use of Krylov methods to solve the multigroup equations effectively decouples energy

- 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
  - 168 angles (1 moment)
  - 1 spatial unknown per cell

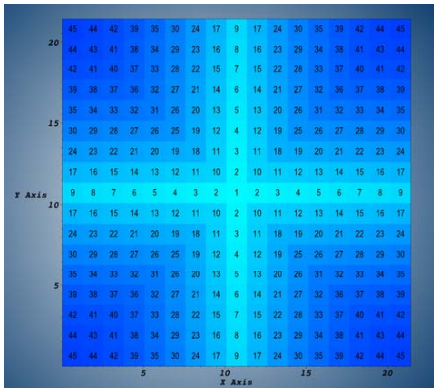

17×17 assembly

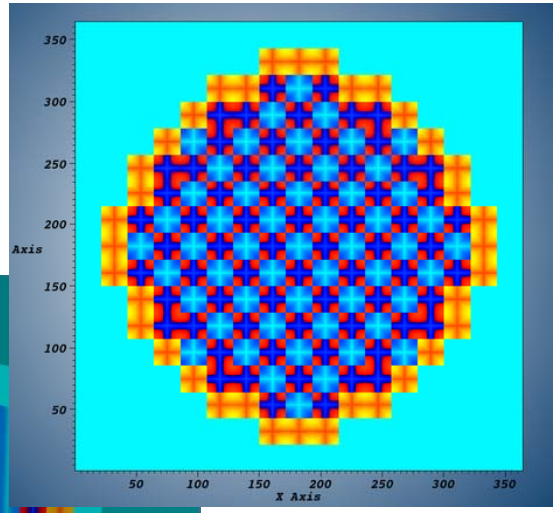

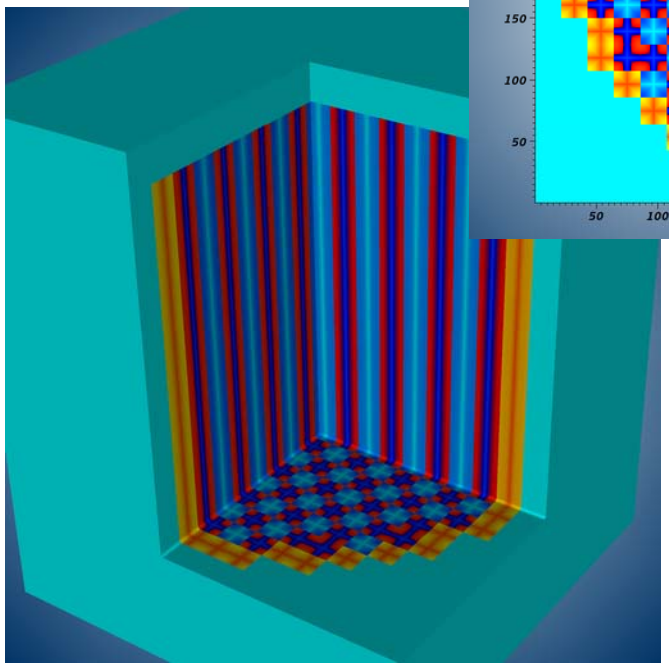

Results

Solvers	Blocks	Sets	Domains	Solver Time (min)
PI + MG GS (2-grid preconditioning)	17,424	1	17,424	11.00
PI + MG Krylov	10,200	2	20,400	3.03
Arnoldi + MG Krylov	10,200	2	20,400	2.05

Total unknowns = 78,576,556,800
Number of groups = 2
k_{eff} tolerance = 1.0e-3

- Arnoldi performs best, but is even more efficient at tighter convergence
  - 27 v 127 iterations for eigenvector convergence of 0.001
- The GS solver cannot use more computational resource for a problem of this spatial size
  - Simply using more spatial partitions will not result in more efficiency
  - Problem cannot effectively use more cores to run a higher fidelity problem in energy

Parallel Scaling

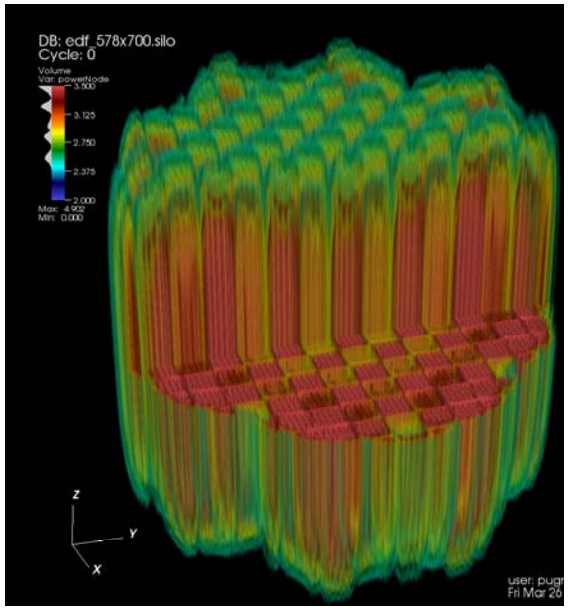

17,424 cores is effectively the maximum that can be used by KBA alone

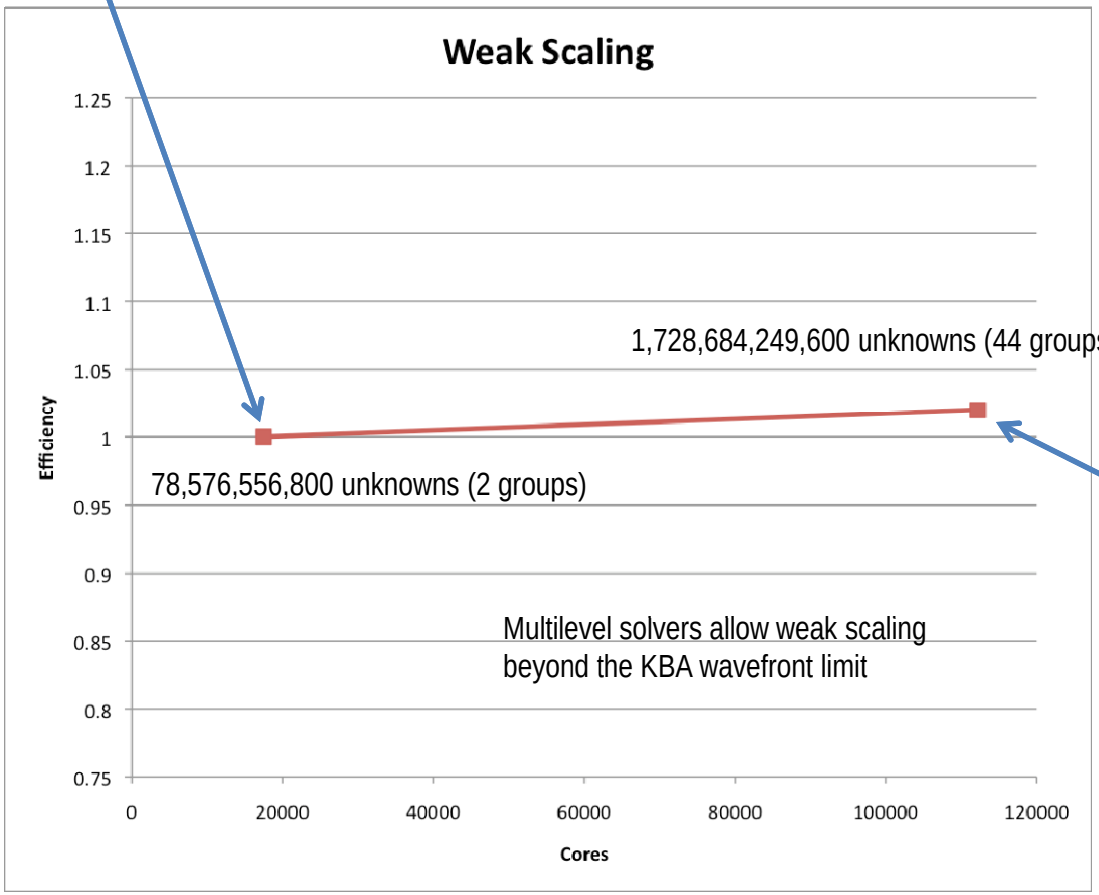

MG Krylov solver partitioned across 11 sets

Strong Scaling

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

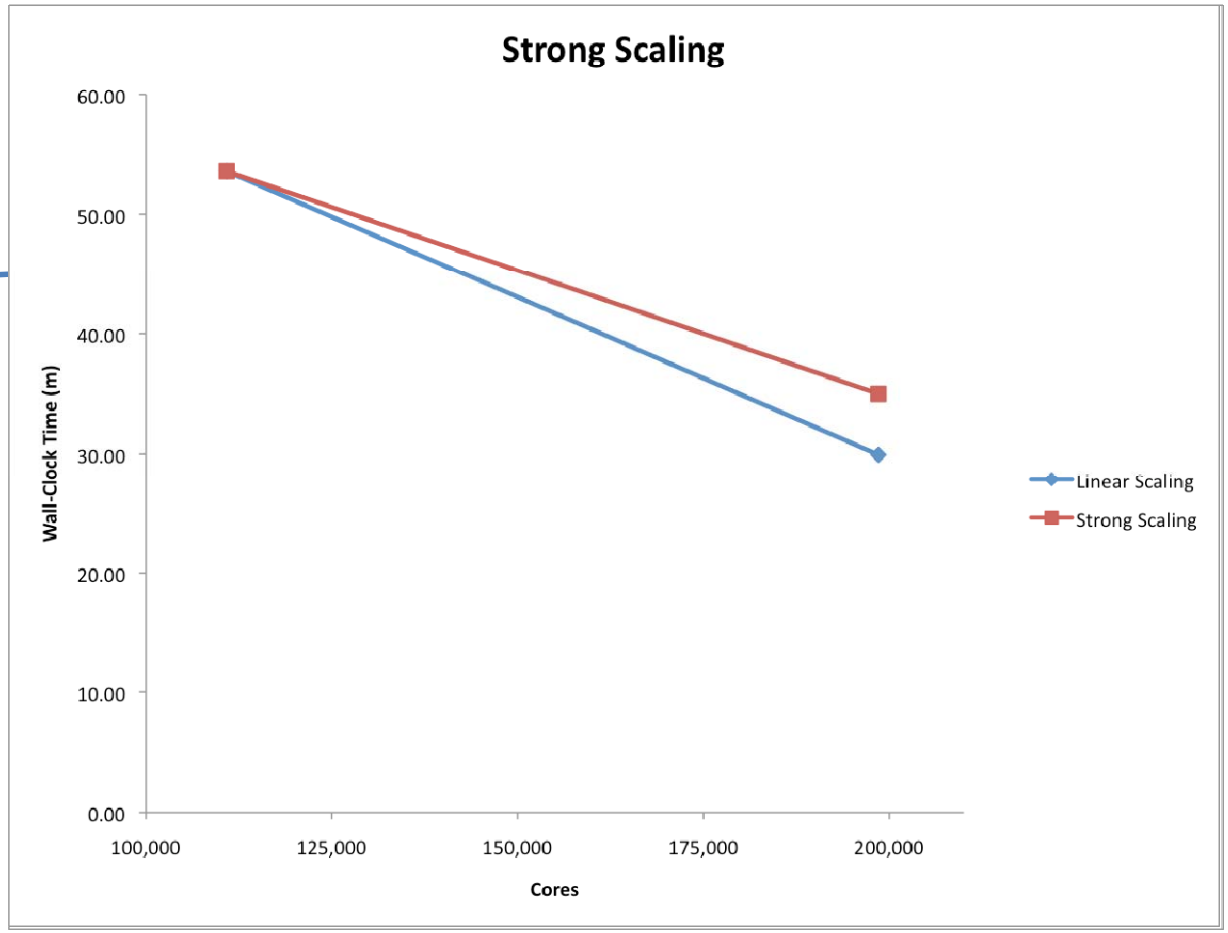

- Communication improvements were significant at 100K core level (using 11 sets).
- They do not appear to scale to 200K core. Why?
  - The problem isn't big enough to demonstrate strong scaling.
  - We are not using the optimal block decomposition.
  - Communication collision on torus across full machine.
  - Multiset communication latency across entire machine dominates.

Summary

- The next step is to combine efficient preconditioners to reduce the number of multigroup and eigenvalue iterations
  - Multigrid in energy
  - Rayleigh-Quotient iteration
- The combination of
  - Krylov solvers over energy
  - Parallel multilevel energy decomposition
  - Krylov eigenvalue solvers

  Has allowed us to break through the wavefront barrier for eigenvalue problems and multigroup fixed-source problems with upscattering.

  It also enables an extra dimension of parallelism that can be applied on GPUs in next-gen HPC hardware.

Questions?

KBA Algorithm
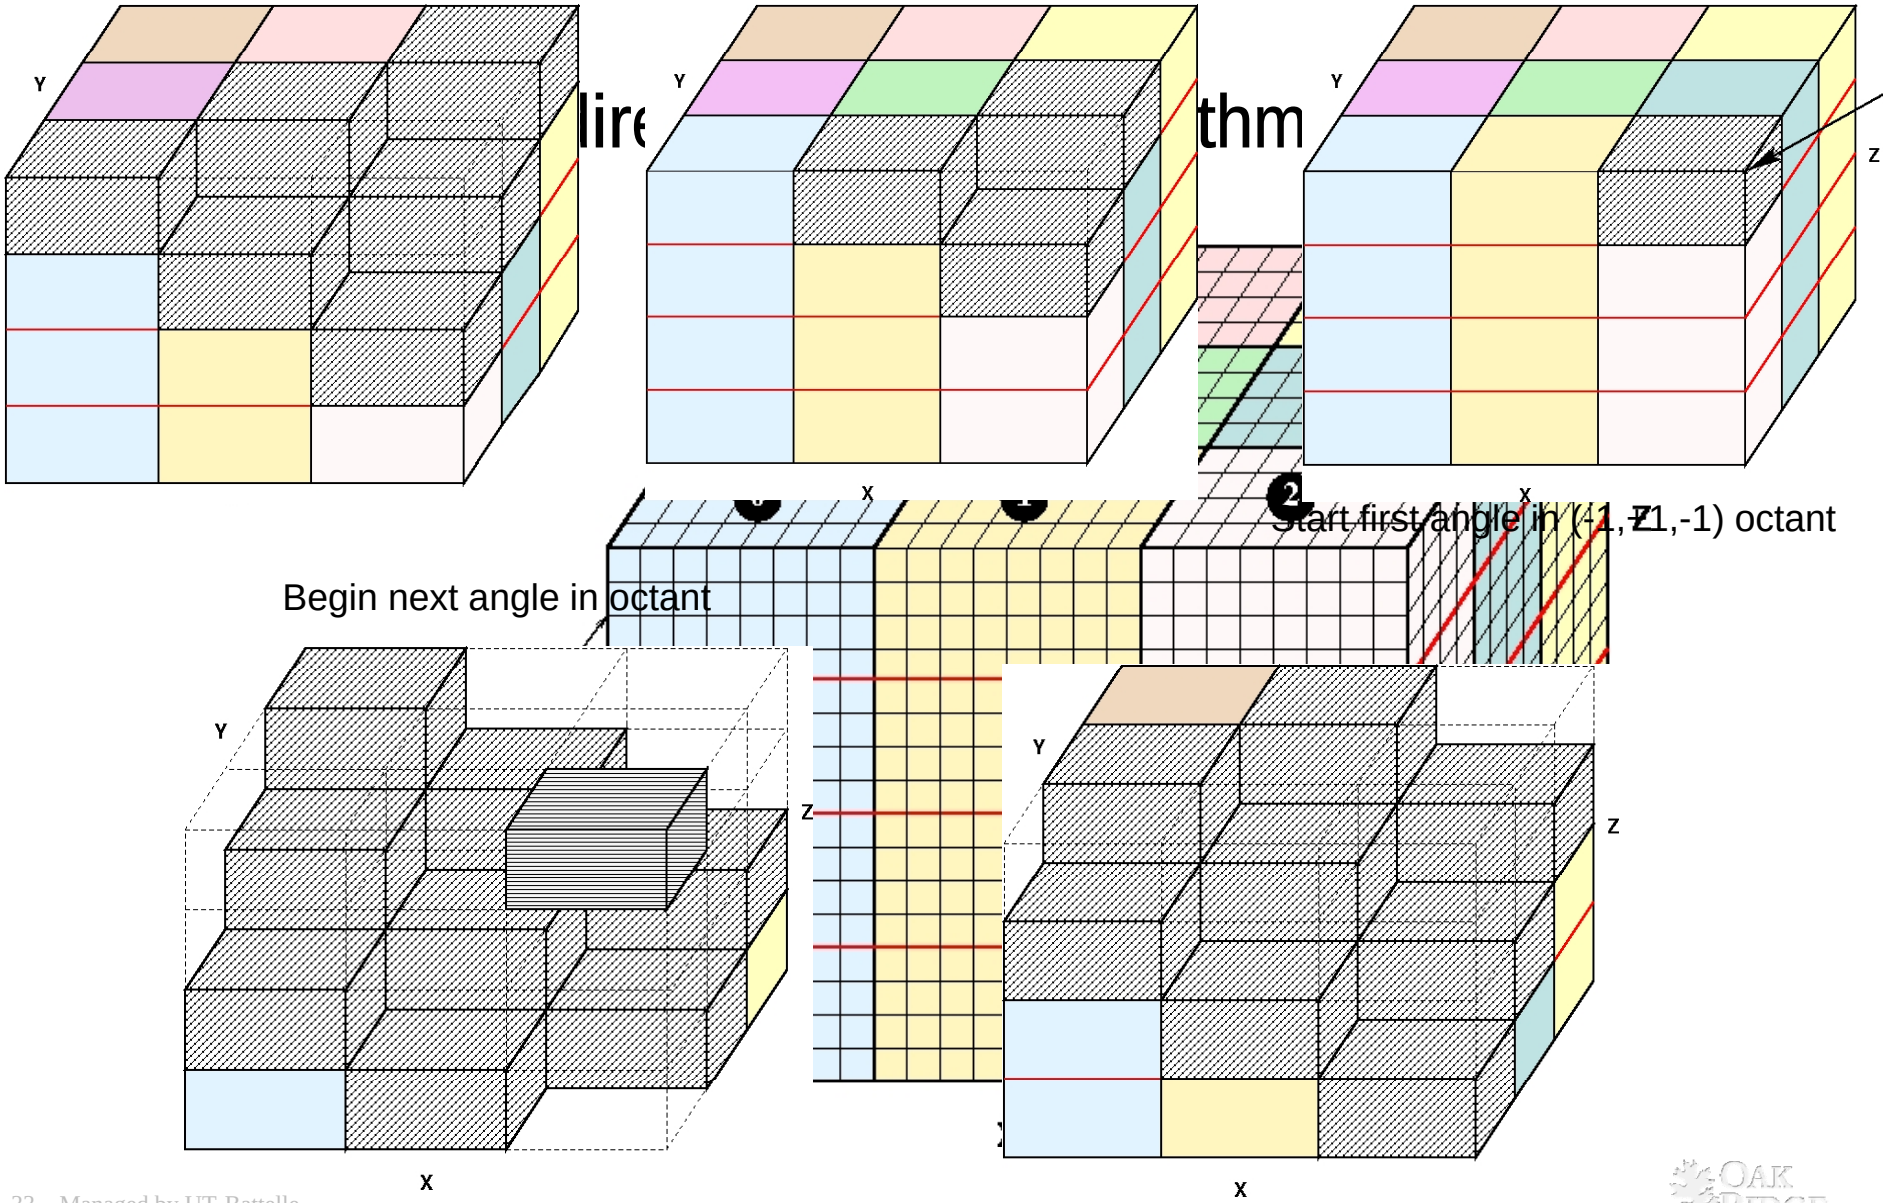
sweeping in direction of particle flow