CASL-U-2015-0172-000

3D Discrete Ordinates Reactor Assembly Calculations on GPUs

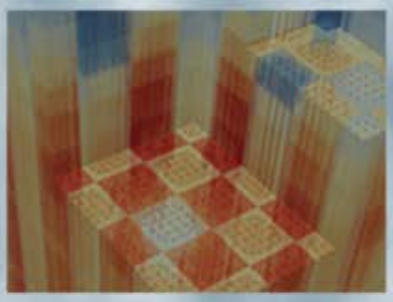

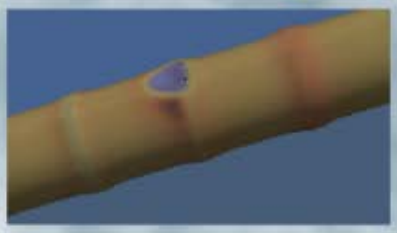

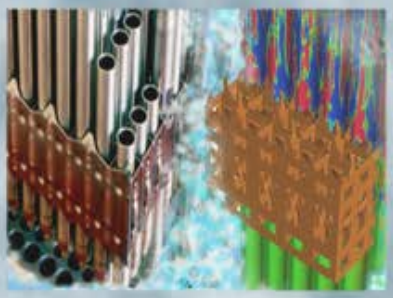

Tom Evans

Wayne Joubert, Steven Hamilton, Seth Johnson, John Turner, Greg Davidson, Tara Pandya

Oak Ridge National Laboratory

April 19, 2015

CASL-U-2015-0172-000

3D Discrete Ordinates Reactor Assembly Calculations on GPUs

Tom Evans
Wayne Joubert
Steven Hamilton
Seth Johnson
John Turner
Greg Davidson
Tara Pandya

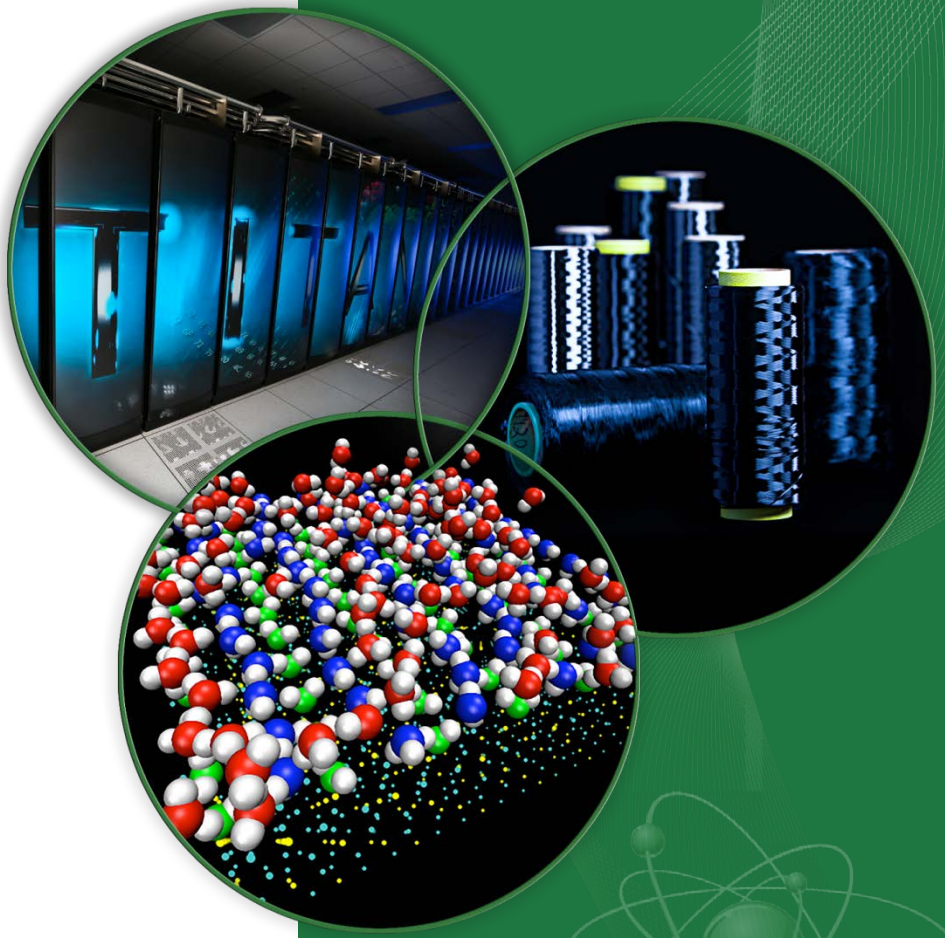

April 21, 2015

ORNL is managed by UT-Battelle
for the US Department of Energy

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Outline

- CAAR Program
- Denovo Solvers and Algorithms
- Results
  - BW 1484
  - WEC AP1000®
- Summary and Conclusions

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

2010 – 2013 CAAR Program

- Center for Accelerated Application Readiness
- Prepare applications for Titan
  - 6 representative application codes (Denovo S_N)
  - Optimize for GPU architecture
- Objectives
  - Provide an optimized GPU-kernel for k-eigenvalue calculations (tightly coupled groups)
  - GPU effort focused on S_N transport sweeps

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Denovo Deterministic Transport

- 2- and 3-D regular grids
- S_N (first-order form) and SP_N angular discretizations
  - Level symmetric, Gauss-Legendre Product, Quadruple Range quadratures
- Multigroup in energy
- Spatial discretizations
  - Tri-linear DGFEM
  - Linear DGFEM
  - Weighted-diamond, theta-weighted diamond
  - Step-characteristic (cell-based slice balance)
  - Finite volume
- Anisotropic P_N scattering

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Denovo Solvers

- Denovo has a hierarchy of solvers for S_N problems with the form

$$\mathbf{L}\psi = \mathbf{MS}\phi + \frac{1}{k}\mathbf{M}\chi\mathbf{f}^T\phi$$

- In the work here we use Arnoldi iteration for the standard eigenvalue problem

$$(\mathbf{I} - \mathbf{TMS})^{-1}\mathbf{TMF}\phi = k\phi$$

$$\mathbf{T} = \mathbf{DL}^{-1} \quad \mathbf{F} = \chi\mathbf{f}^T \quad \phi = \mathbf{D}\psi$$

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Denovo Solvers

- Implementation of Arnoldi solve ($v^l \equiv \{v^l_{g,c,n,m}\}$ is the iteration vector)

$$z^l = \mathbf{TMF}v^l$$ mat-vec and sweep

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

- GMRES(m) is used on the inner multigroup problem

$$w^k = (\mathbf{I} - \mathbf{TMS})x^k$$ mat-vec and sweep

- Each application of $\mathbf{T}$ involves a sweep over all groups

$$y = \mathbf{DL}^{-1}x$$

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Denovo Multilevel Decompostion

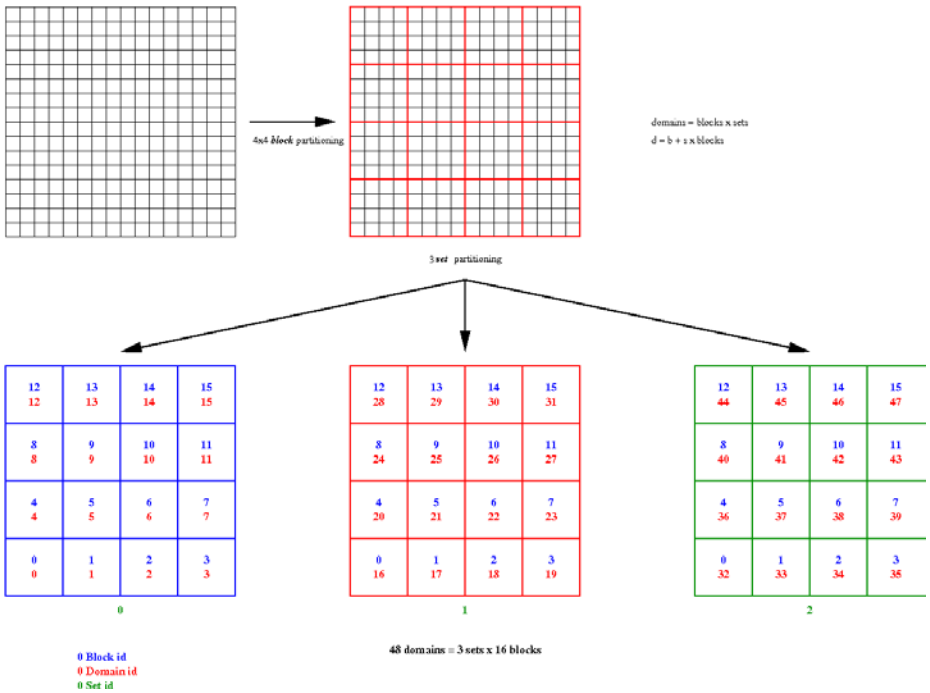

$$s_g = \mathbf{S}_{g0}x_0^k + \mathbf{S}_{g1}x_1^k + \ldots + \mathbf{S}_{gG}x_G^k$$

- Energy decomposed across sets
- Space decomposed across blocks blocks
- Domains = sets × blocks
- Requires an MPI *reduce-scatter* to initialize the scattering source ($\mathbf{S}x^k$)
- Decouples energy sweeps

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Device-Based Parallelism

- Need 4 – 8K threads for GPU to cover latencies
- Must have good memory access patterns
  - Reuse of data loaded from global memory
  - Coalesced stride-1 memory references
  - Efficient use of registers
  - Shared memory
  - Caches on GPU
- Candidate dimensions for parallelism
  - Space, energy, moment/angle, octant
  - 4 spatial unknowns (LD DGFEM)

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Device-Based Parallelism

- Energy
  - Moved to inner threaded loops (across threadblocks)
- Octants
  - KBA extended to sweep in all 8 octants concurrently (across threadblocks)
  - Different octants update same output vector – *scheduling*
- Moment/angle
  - Mapped to CUDA threads in threadblock
- Space
  - Block wavefronts assigned to threads on GPU
  - Mapped to CUDA thread in threadblock

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Summary of Mapping of Dimensions

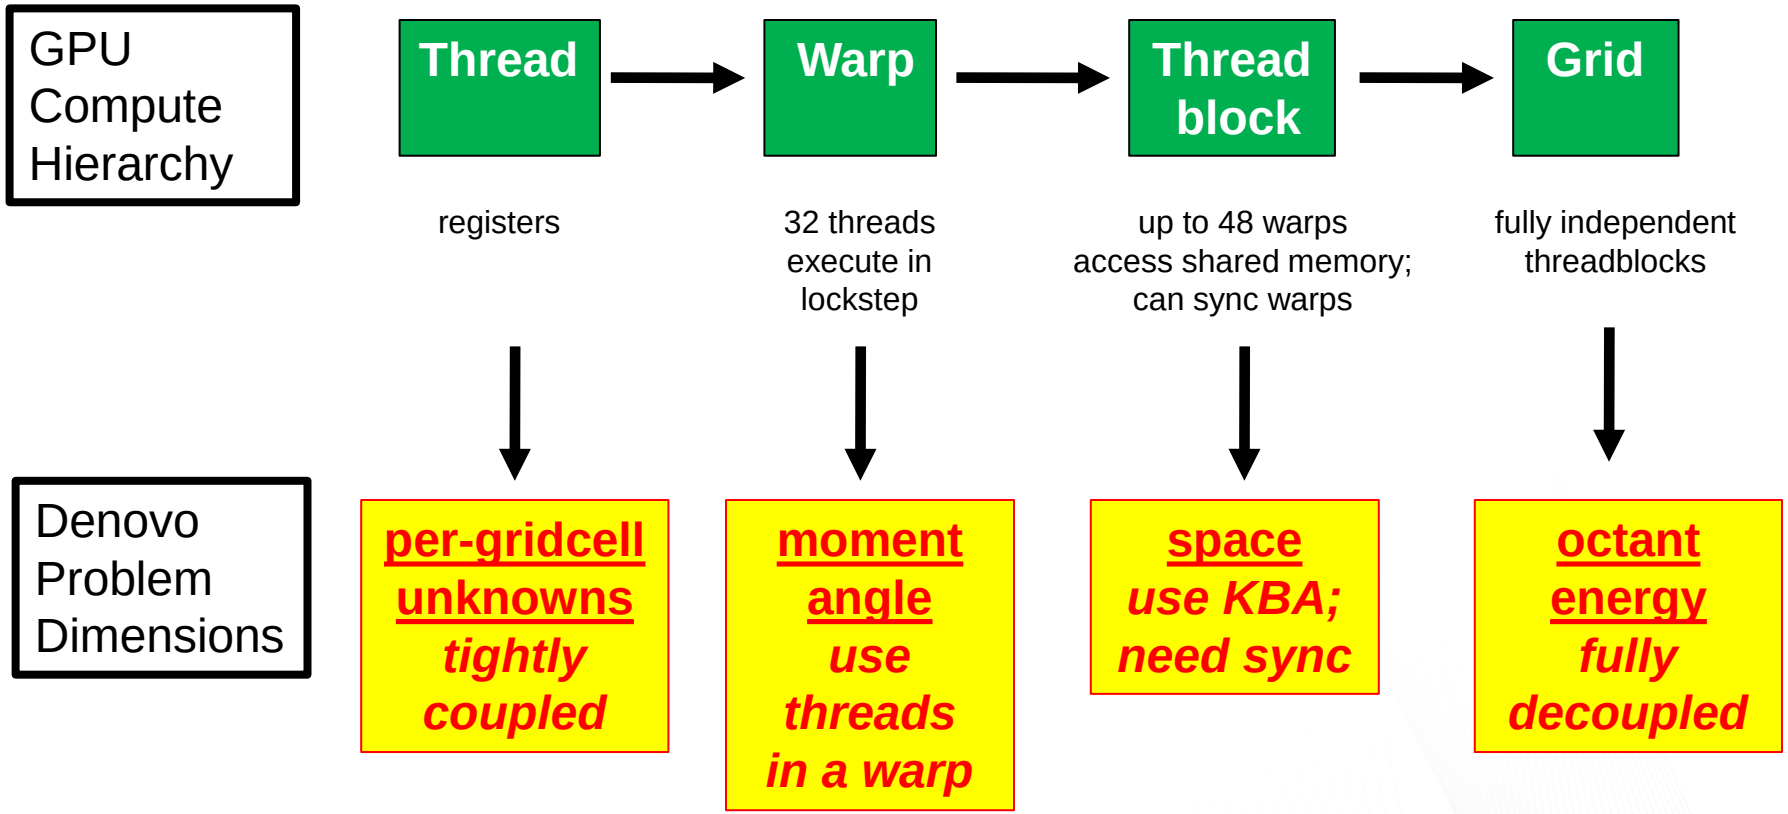

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Results

- Two problems analyzed
  - Babcock and Wilcox (B&W 1484)
  - WEC AP1000® HZP assembly
- All problems run on Titan
  - Configured for 8 cores per node
  - 1 NVIDIA Tesla K20x GPU per node
  - 4 GB per process
- Denovo executed through VERA-CS
  - Inline cross section processing using cell-based XSProc
  - Automated mesh generator (volume-weighting)

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

B&W 1484

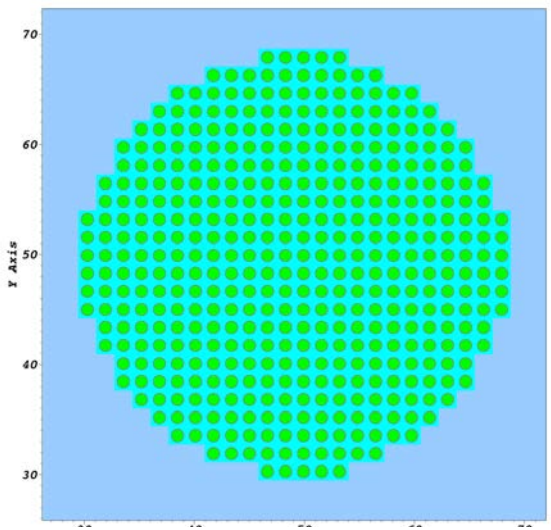

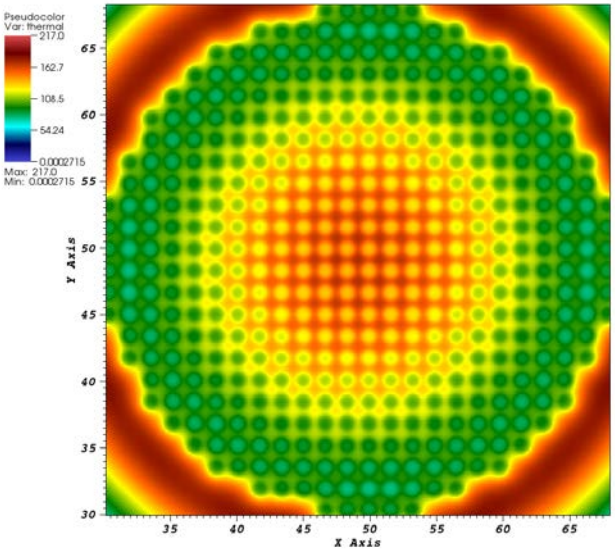

Thermal flux

Spatial differencing	Linear DGFEM
Radial mesh	20×20
Max Δz	2 cm
Mesh dimensions	552×552×84
Energy groups	56 (20 coupled)
P_N order	1
Quadrature	QR
Angles per octant	72 (6/12)
Solved unknowns	1.32×10^{13}
Spatial decomposition	92×24
Energy sets	28
Titan nodes	7,728
Total domains	61,824
Arnoldi iterations (τ=0.0001)	9
MG GMRES iterations	339

CASL-U-2015-0172-000

CPU v GPU Performance – BW 1484

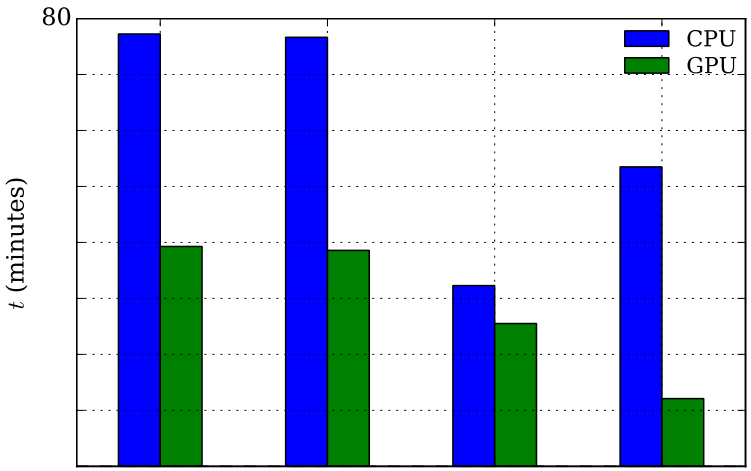

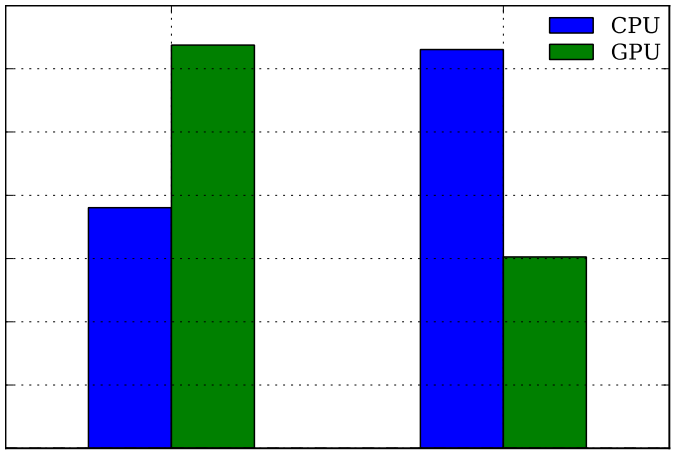

- Arnoldi iteration incurs no overhead
  - All time is spent in multigroup inners solve
- Sweep source initialization requires a *reduce-scatter* across sets
  - Sweep source initialization is 64% of runtime on GPU kernel due to reduced time spent in sweeps
- GPU sweep is 4.4× faster than CPU sweep
- Total solution time ≈ 2× faster with GPU kernel sweep

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

WEC AP1000®

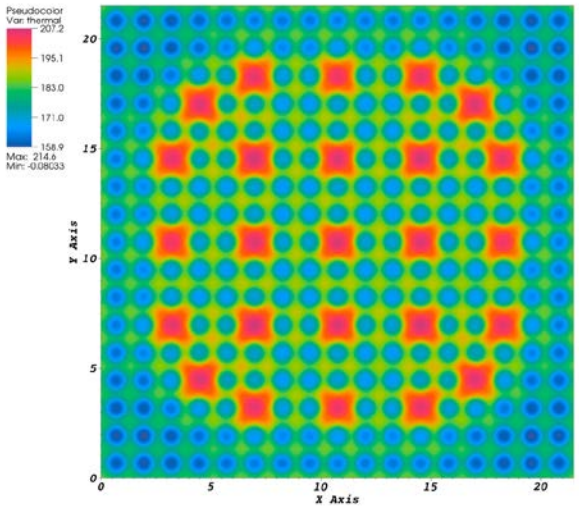

Thermal flux

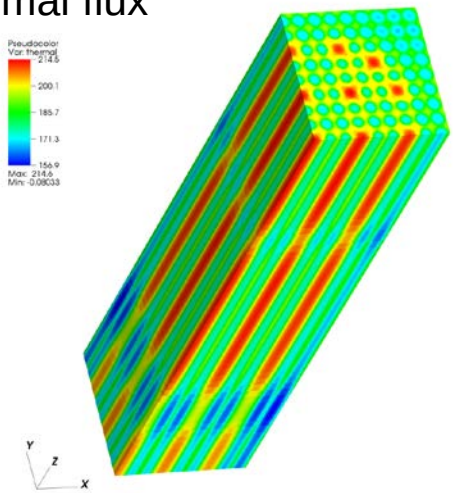

Spatial differencing	Linear DGFEM
Radial mesh	30×30
Max Δz	1 cm
Mesh dimensions	256×256×532
Energy groups	56 (20 coupled)
P_N order	1
Quadrature	QR
Angles per octant	32 (4/8)
Solved unknowns	8.0×10^{12}
Spatial decomposition	64×64
Energy sets	14
Titan nodes	7,168
Total domains	57,344
Arnoldi iterations (τ=0.0001)	17
MG GMRES iterations	526

CASL-U-2015-0172-000

CPU v GPU Performance – AP1000®

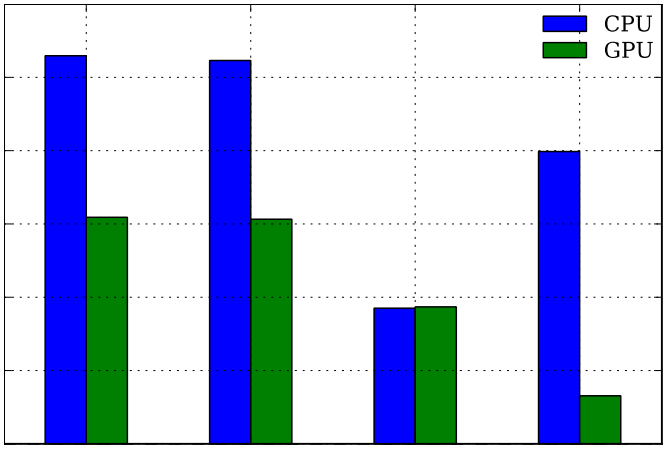

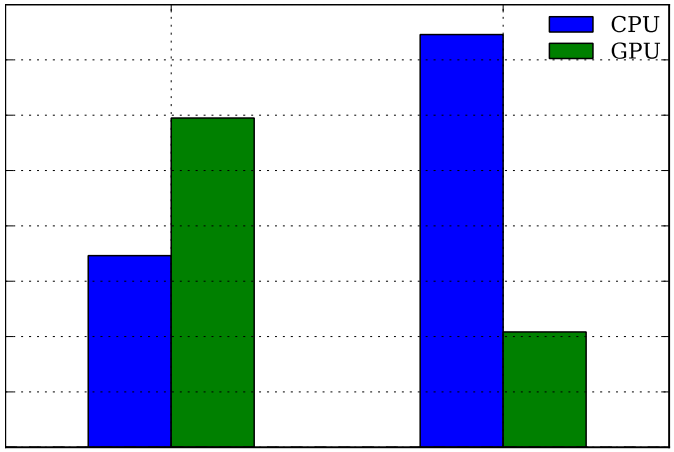

- Again Arnoldi iteration incurs no overhead
  - Reflected boundary fluxes part of inner multigroup vector
- GPU sweep is 6.1× faster than CPU sweep
  - More groups per set (4 vs 2)
  - 32 angles per octant = number of threads per warp
- Total solution time ≈ 1.7× faster with GPU kernel sweep

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Iterative Behavior

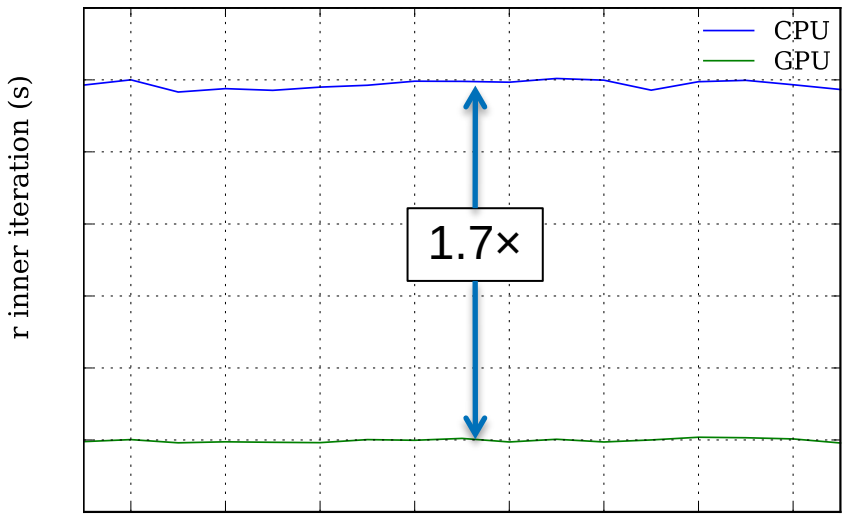

Time per inner MG iteration

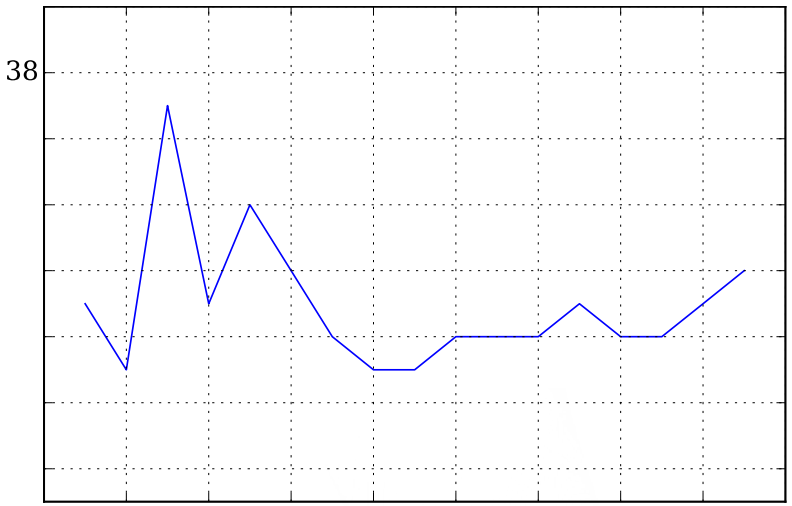

Number of inners per Arnoldi iteration

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Strong Scaling

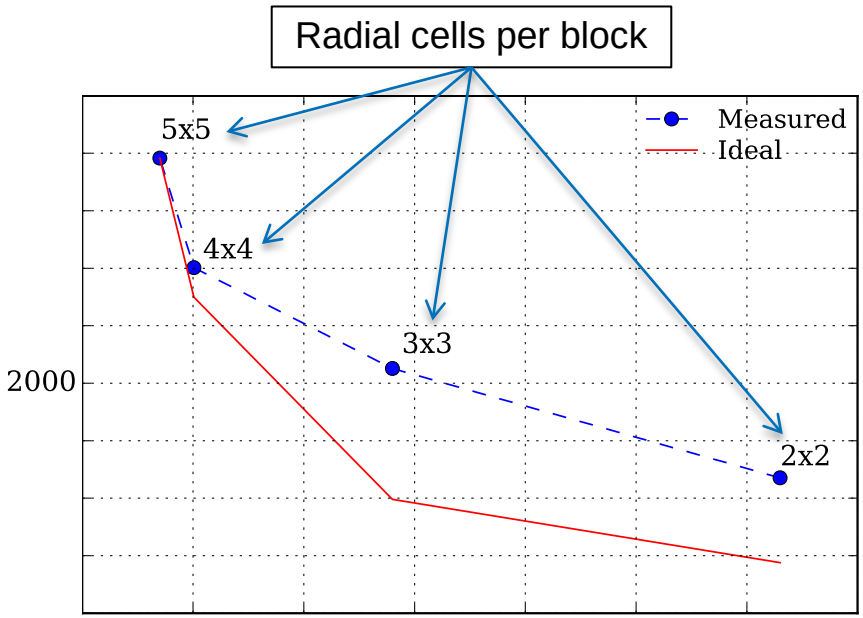

- Mesh reduced to 14×14 per pin
- 120×120×532 mesh size
- As block size reduces the work performed by GPU kernel is too small to overcome CPU – GPU transfer and MPI communication latency

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Weak Scaling

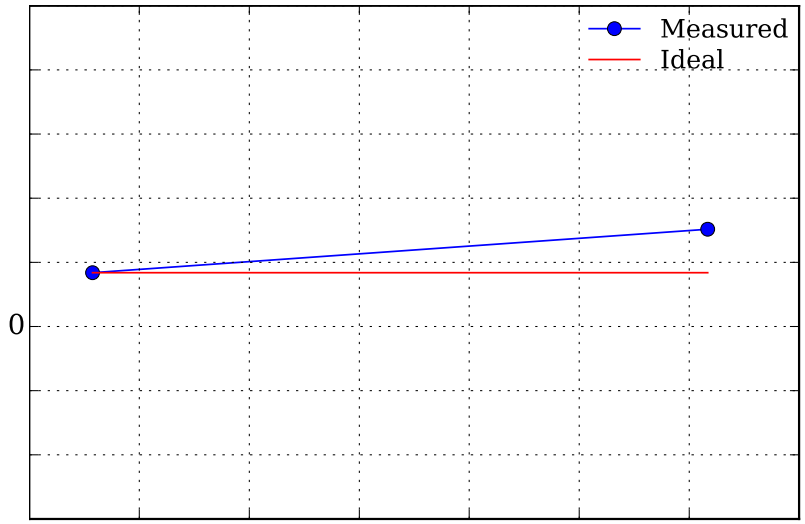

- Constant block size of 2×2×532
- 14×14 and 30×30 mesh per pin
- 85% efficiency attained from 50,400 to 229,376 cores

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Conclusions

- Overall solution time provided by GPU kernel on Titan ≈ 2×
- GPU sweep kernel improved between factor 4 – 6 over CPU-only sweeps
- Sweep source initialization (reduce-scatter) becomes dominant runtime bottleneck
- Performance sensitive to problem layout
  - Angles per warp
  - Equal size work units
  - Number of unknowns per GPU to exploit concurrency

OAK RIDGE National Laboratory
CASL-U-2015-0172-000

Acknowledgements

- This research was supported by the Consortium for Advanced Simulation of Light Water Reactors (http://www.casl.gov), an Energy Innovation Hub (http://www.energy.gov/hubs) for Modeling and Simulation of Nuclear Reactors under U.S. Department of Energy Contract No. DE-AC05-00OR22725.
- This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.
- This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

CASL-U-2015-0172-000