Parallel S_N: Lessons Learned

Paul Nowak

LLNL, DNT, A-Division

We are currently working toward the CY03 ASCI milestone

- High-fidelity-physics calculation of a stockpile relevant nuclear test with three-dimensional engineering features
- Focus, high-fidelity implies a multi-group [illegible] transport treatment for radiation flow
- [illegible] is a difficult challenge and your efforts [illegible] play an important role in helping us [illegible] our goal

The road is long and tedious

	Analytic	Code Comparison	AGEX Data	UGT Data
Full Physics MD				X
Full Physics 1D		X		
Multi-Physics MD		X	X	X
Multi-Physics 1D		X		
Single Physics MD	X	X		
Single Physics 1D	X	X		

We are in a transition phase

- UCBL - tet-based sweep
- Grey S_2, TSA
- B[illegible]k-Jacobi with [illegible]dary flux [illegible]teration [illegible] in space, threads [illegible]

- UCB or PWLDFEM - corner-based sweep
- Grey Diffusion, DSA
- Parallel sweep
- MPI and threads in space

UCBL has been very useful, but has several fatal flaws

- Extremely expensive - tet-based sweep
- Less than second-order accurate
- Obtains linear solution on distorted grids in the “thick” diffusion limit, but not in the “intermediate” diffusion limit

Early experience with true corner methods is encouraging

- UCB results for a radiation flow experiment compare well with 2D results and data
- Run time reduced from a week to a day
- Learned how to obtain linear solution in [illegible] the thick and intermediate diffusion [illegible]
- [illegible] seems more susceptible to cycles

TSA is not robust for real applications

- Unstable for fine meshes - exacerbated by strong material discontinuities
- CG proved to be impractical - resorted to a two-level TSA
- [illegible] forces the use of the largest β possible ([illegible]0.8) - always on the edge

We are exploring a DSA option

- Not sure that TSA can be fixed - FA of homogeneous problems alone appears to be insufficient to predict behavior
- Great success with new non-symmetric matrix solvers (HYPRE from CASC)

Block-Jacobi works well for some applications, but not all

- Boundary Flux acceleration makes BJ useful for optically thick systems
- Performs best as albedo->1
- Performs poorly when tracking wavefronts in thin systems

Speedup is linear!

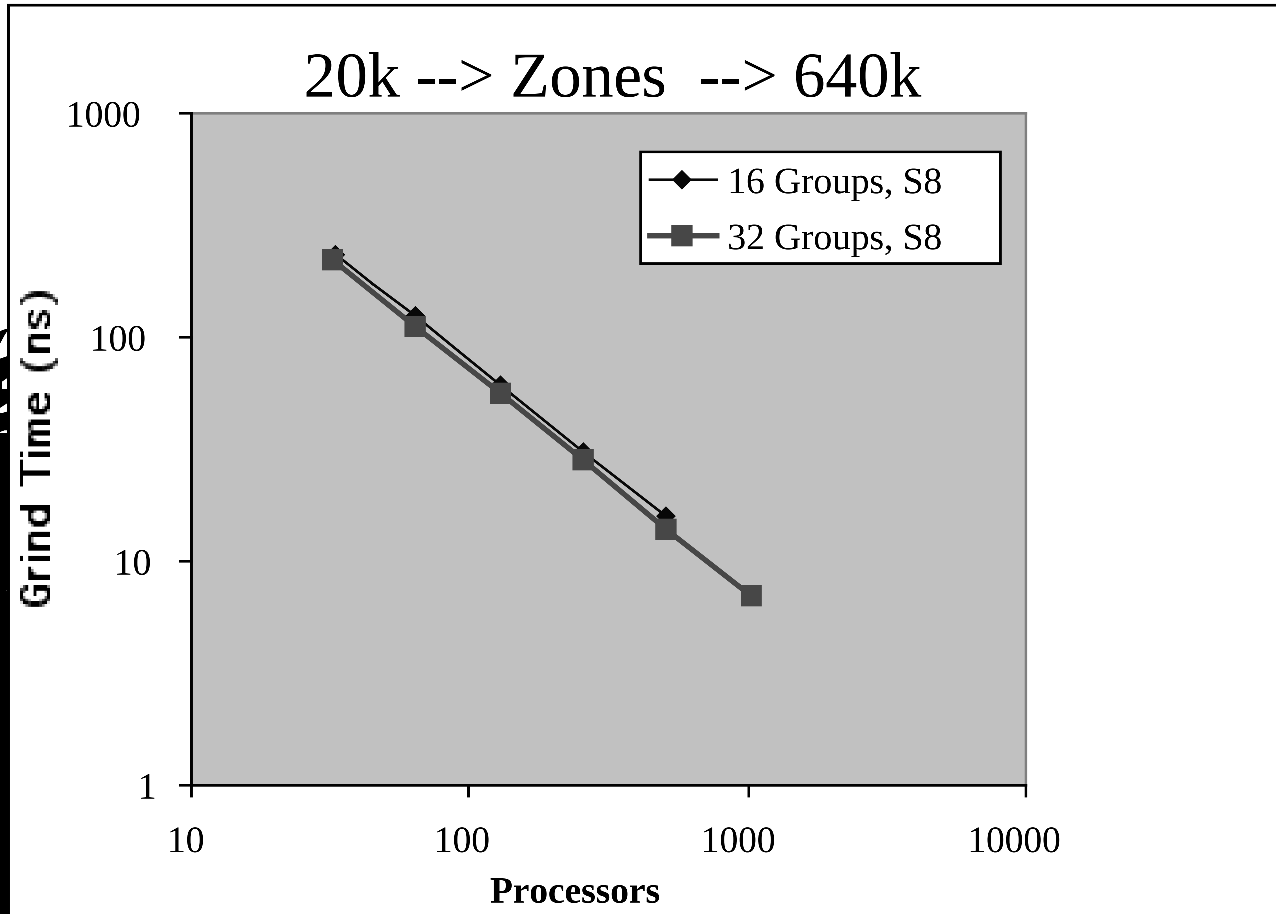

We are exploring a parallel sweep option

- Similar in style to KBA
- Will not be perfect due to re-entrant and disjoint domains
- Expect best performance gain on laser-driven aboveground experiments (AGEX)

Two level parallel strategy has worked well, but ...

- Threads do not scale well above 8 (e.g. SU is 11 out of 16 on ASCI-White)
- Limited to 8 threads if using TSA
- Reflecting boundaries can degrade performance

It’s clear that MPI and OMP belong at the same level

- Best strategy is to use the smallest mesh possible per processor - magnified by the larger cache on White
- Threads really want their own dedicated memory
- Provides the most flexibility when running on different platforms
- [illegible]es scaling to large[illegible]

We are in a transition phase

- UCBL - tet-based sweep
- Grey S_2, TSA
- Block-Jacobi with boundary flux acceleration
- [illegible] in space, threads [illegible]

- UCB or PWLDFEM - corner-based sweep
- Grey Diffusion, DSA
- Parallel sweep
- MPI and threads in space