

Monte Carlo Neutron Transport on GPUs

The Neutron Transport Problem

- ▶ Neutron transport equation solutions give us reactor power distributions and nuclide inventories

$$\frac{1}{v(E)} \frac{\partial \phi}{\partial t} + \hat{\Omega} \cdot \nabla \phi + \Sigma(E, r, t) \phi = q(\hat{\Omega}, E, r, t) + \int_0^\infty dE' \int_{4\pi} d\hat{\Omega}' \Sigma_s(\hat{\Omega}' \rightarrow \hat{\Omega}, E' \rightarrow E, r, t) \phi|_{\hat{\Omega}', E'} \quad (1)$$

- ▶ **Quadrillion unknowns** for direct PWR solution
- ▶ We obtain **unbiased** approximate solutions using the **Monte Carlo** method relatively easily

Monte Carlo Methods

- ▶ Track neutrons from birth to death, provably equivalent to transport equation in the limit




Figure: Monte Carlo neutron tracks through a fuel assembly

- ▶ **History** algorithm uses one neutron per CPU core
- ▶ **Event** algorithm allocates a large array of particles, and the above operations execute in unison
- ▶ Event method found to work best in continuous energy problems on GPUs, vector computers

GPU Programming 101

- ▶ Execute 1000s of compute threads
- ▶ Prefer having threads execute same task
- ▶ Challenge in neutron physics with cornucopia of reactions, geometries
- ▶ Random memory access 10x slower than aligned




Figure: An Nvidia P100




Figure: "Thread divergence" can catastrophically impair GPU computations

CUDA OpenMC

- ▶ World's first GPU MC code sharing most of its code with the CPU codebase of OpenMC [1]
- ▶ Can use particle queue sorting and linear XS search as proposed by resp. [2] and [3]




Figure: Parallel compaction, unique to GPU OpenMC, reduces need for atomic operations




Figure: Parallel compaction speedup

Results

- ▶ Code tested on fresh fuel pin cell (37 nuclides) and depleted fuel assembly (150 nuclides) case with Nvidia V100 GPU




Figure: Running more particles at once saturates GPU resources




Figure: Collision processing matters less as more particles are ran; XS lookup always dominates




Figure: Speedups are problem-dependent; particle advance best aided by GPU for pin cell, and XS lookup best aided for the assembly.

References

- [1] P. Romano and B. Forget. "The OpenMC Monte Carlo Particle Transport Code." Annals of Nuclear Energy 51 (January 2013).
- [2] J. Tramm and P. Romano. "A21 MC Kernel Implementation." Milestone ECP-SE-08-103. DOE Technical Report. September 30, 2019.
- [3] S. Hamilton and T. Evans. "Continuous-Energy Monte Carlo Neutron Transport on GPUs in the Shift Code." Annals of Nuclear Energy 128 (June 2019).