Perspectives from the ExaSMR Project: Nuclear Data Needs and Opportunities

Paul K. Romano
Computational Scientist, Argonne National Laboratory

WANDA 2021: Advanced Computing for Nuclear Data Session
January 29, 2021

Background

- ExaSMR project
  - Joint ORNL/ANL/MIT project on coupled Monte Carlo–CFD simulations for small modular reactors funded by the Exascale Computing Project
  - Several modeling challenges: small size results in large spatial gradients, natural circulation, no operational data
  - Use exascale resources to produce “virtual experiment” datasets that can be used to validate low-order engineering simulations
- Software stack:
  - Particle transport: OpenMC (ANL/MIT) and Shift (ORNL)
  - Thermal hydraulics: Nek5000 / NekRS
  - Coupling: ENRICO

Background

- ExaSMR project
  - Joint ORNL/ANL/MIT project on coupled Monte Carlo–CFD simulations for small modular reactors funded by the Exascale Computing Project
  - Several modeling challenges: small size results in large spatial gradients, natural circulation, no operational data
  - Use exascale resources to produce “virtual experiment” datasets that can be used to validate low-order engineering simulations
- Software stack:
  - **Particle transport: OpenMC (ANL/MIT) and Shift (ORNL)**
  - Thermal hydraulics: Nek5000 / NekRS
  - Coupling: ENRICO

Nuclear data considerations are crucial for particle transport

Computational Needs

- Current resources
  - OLCF Summit: 200 PF, IBM POWER9 CPU, NVIDIA Volta V100 GPU
  - ALCF Theta: 11.7 PF, Intel Xeon Phi CPU
- Future resources
  - OLCF Frontier: > 1.5 EF, AMD EPYC CPU, AMD Radeon Instinct GPU
  - ALCF Aurora: $\geq$ 1 EF, Intel Xeon CPU, Intel Xe GPU

OpenMC Data Pipeline

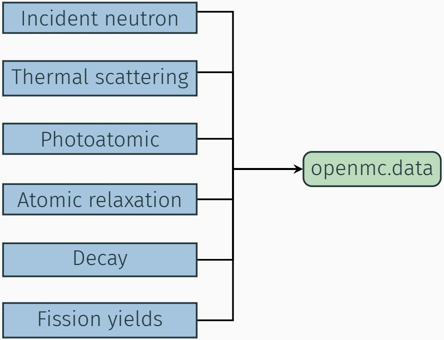

OpenMC Data Pipeline

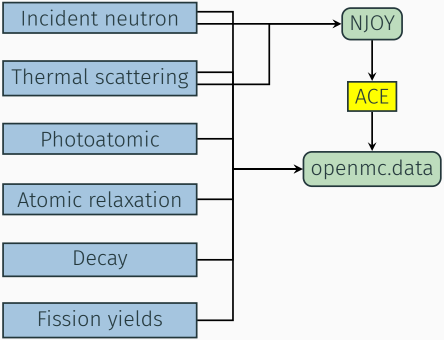

OpenMC Data Pipeline

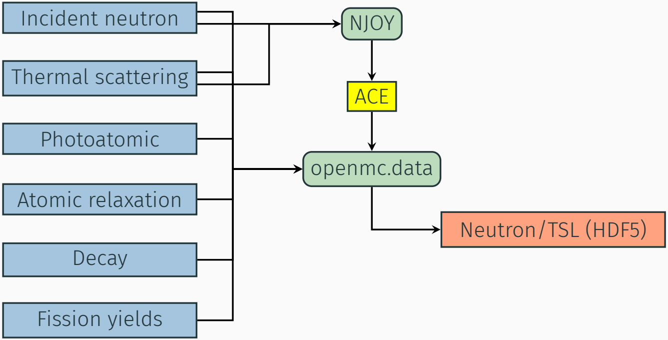

OpenMC Data Pipeline

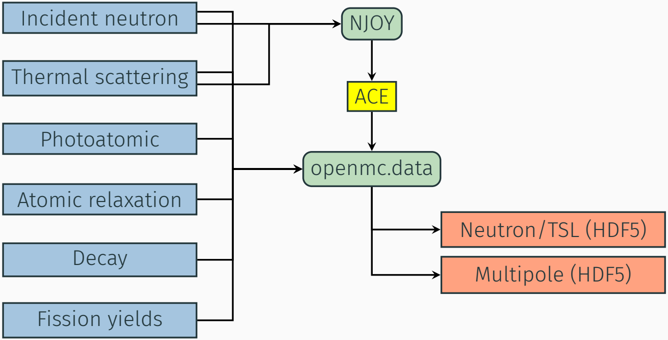

OpenMC Data Pipeline

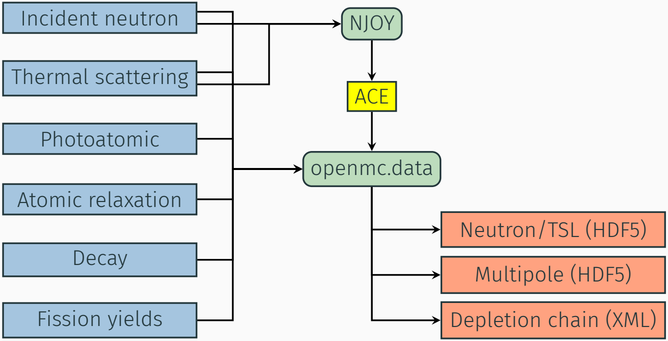

OpenMC Data Pipeline

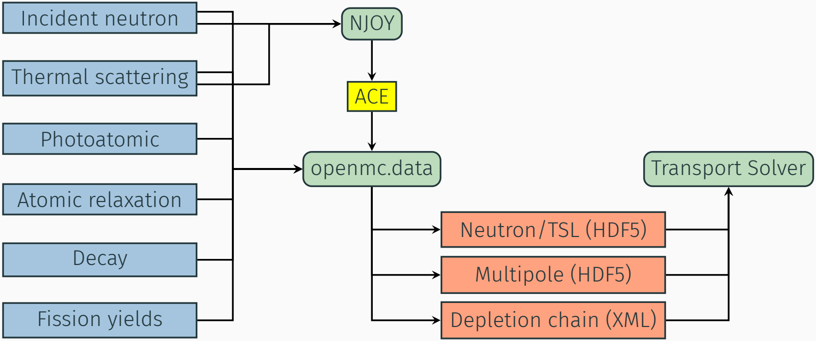

Challenges: Complex Data Hierarchies

- Both target exascale platforms are based on GPU architectures
- Nuclear data is needed at many temperatures; large size of data
- Data size can place a significant burden on I/O time
- Want to preserve fidelity of original data in ENDF; multitude of data formalisms pose programming challenges:
    - Inability to use polymorphism
    - Data transfer between host and GPU

Challenges: Temperature Dependence

- In the resolved resonance range, OpenMC/Shift rely on the *windowed multipole* method
  - Resonance parameters stored as complex poles/residues
  - Can analytically Doppler broaden cross sections
  - Significant reduce memory requirements but requires more operations to evaluate cross section
- Does not help in unresolved range or thermal energies

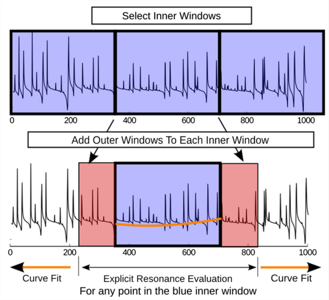

Opportunities: Machine Learning

- Temperature dependence of thermal scattering and unresolved resonance range:
  - Brute force interpolation on tables stored at many temperatures
  - Again, memory requirements quickly go up depending on temperature grid

Need innovations in methods for thermal scattering/URR $\rightarrow$ **Machine learning** may be suitable given lack of theoretical models for temperature dependence

Opportunities: Model-based Physics

- Evaluations continue to grow in size
- For HPC simulations, strong incentive to use less memory and more FLOPs
- Integrating model-based physics is very attractive for Monte Carlo transport simulations
  - Multipole format is essentially just resonance parameters
  - Fission event generators (FREYA, CGMF, GEF, etc.)
  - Thermal scattering physics with just phonon frequencies?
- Better physics *and* better performance

Thank you!