(#47) Hardware Accelerator for Compute-Intensive Tasks in Solving Neutron Transport Problems by MOC

Dr. Thuy T. Le, Professor & Department Chair

Electrical Engineering, San José State university, San Jose, California 95192 U.S.A

https://www.sjsu.edu/ee/

The author would like to thank his graduate students Dev Desai & Priyansh Bhimani for their major contributions in this study

The Neutron Transport Equation

$$\vec{\Omega}\cdot\nabla\varphi\left(\vec{r},\vec{\Omega},E\right)+\Sigma_t\left(\vec{r},\vec{\Omega},E\right)\varphi\left(\vec{r},\vec{\Omega},E\right)=Q\left(\vec{r},\vec{\Omega},E\right)$$

$$Q\left(\vec{r},\vec{\Omega},E\right)=\int_0^{4\pi}\int_0^{\infty}\Sigma_s\left(\vec{r},\vec{\Omega}'\to\vec{\Omega},E'\to E\right)\varphi\left(\vec{r},\vec{\Omega}',E'\right)dE'd\vec{\Omega}'+\frac{\chi\left(\vec{r},E\right)}{4\pi k_{eff}}\int_0^{\infty}\nu\Sigma_f\left(\vec{r},E'\right)\int_0^{4\pi}\varphi\left(\vec{r},\vec{\Omega}',E'\right)d\vec{\Omega}'dE'$$

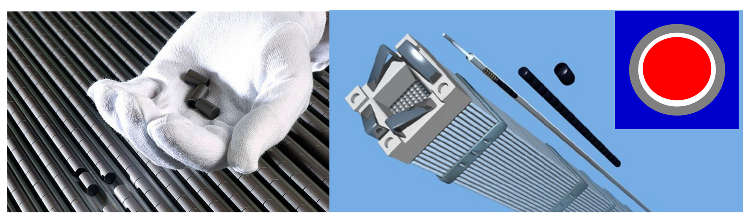

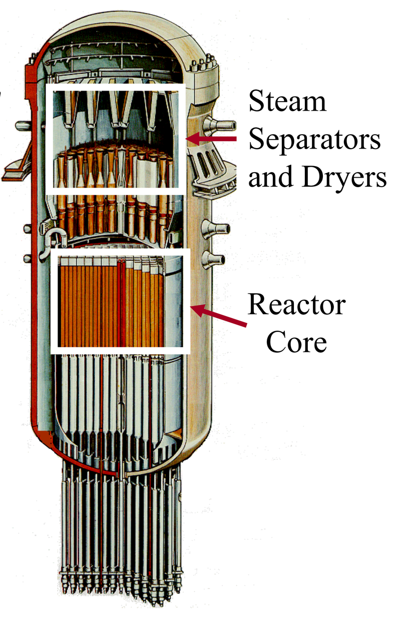

- Uranium fuel pellets are sealed inside zirconium alloy rod (fuel rod)
- About 236 fuel rods (+ other rods) in each fuel assembly
- About 200 fuel assemblies (+ other assemblies) in the reactor core

Example: A Quarter PWR Core

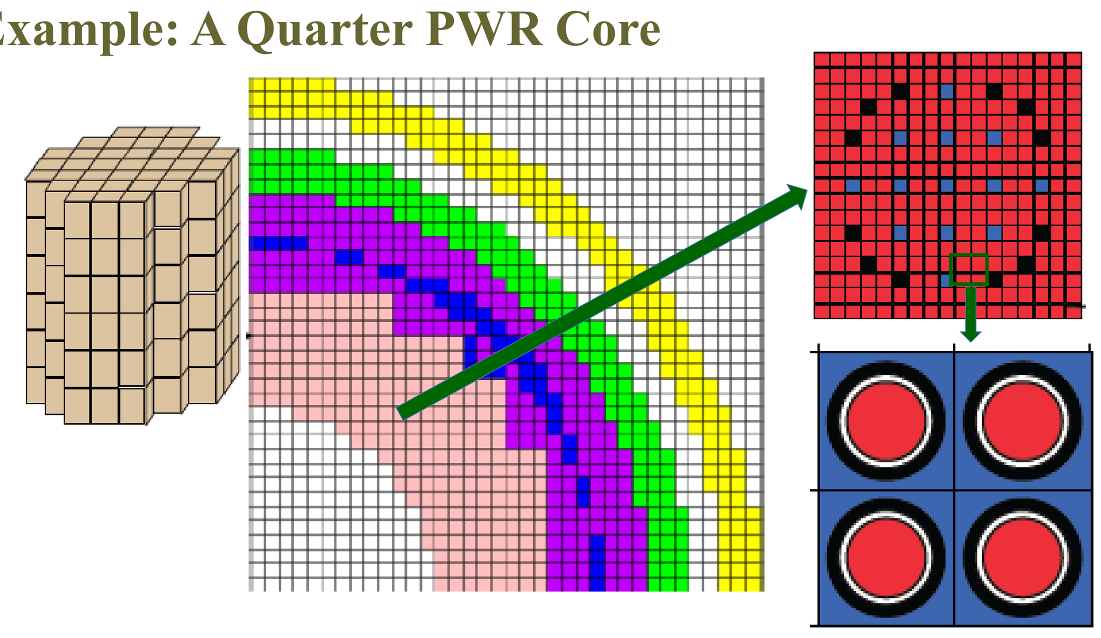

Approximations by Discretization

- ❑ Continuous space → **discrete regions** (1.2 billion regions). Material properties in each region are homogenized as constants
- ❑ Continuous neutron energy → **discrete energy** (72 energy groups typically): Material properties are function of energy groups (energy homogenization)
- ❑ Continuous neutron direction → **discrete angles** (128 angles typically): Material properties are independent of neutron travel direction

(About 10 trillion unknowns)

Method of Characteristics (MOC)

- ❑ **Deterministic ray-based algorithm** (similar to MC methods)
- ❑ For a partial differential equation (PDE)
  - MOC establishes rays (or tracks) crossing the whole spatial domain with fixed angular quadrature for each direction called **characteristics**
  - Each characteristic is **sub-divided into segments**
  - PDE becomes ODE along the characteristic lines
  - Solutions of the ODE is obtained **along the characteristics** and transformed back to PDE
- ❑ Can be structured across the domain such that high-cost calculations are **independent from the problem dimension and geometry**

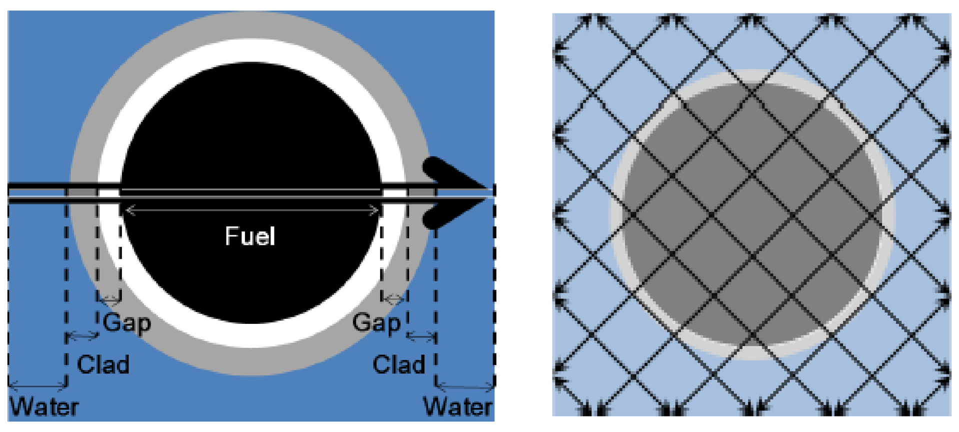

- ❑ The space and direction are transformed into "**characteristic direction**" $\vec{r_0}+s\vec{\Omega}$

$$\frac{d\varphi_g}{ds}\left(\vec{r_0}+s\vec{\Omega},\vec{\Omega}\right)+\Sigma_{t,g}\left(\vec{r_0}+s\vec{\Omega}\right)\varphi_g\left(\vec{r_0}+s\vec{\Omega},\vec{\Omega}\right)=Q_g\left(\vec{r_0}+s\vec{\Omega},\vec{\Omega}\right)$$

$$Q_g\left(\vec{r_0}+s\vec{\Omega},\vec{\Omega}\right)=\sum_{g'=1}^{G}\int_0^{4\pi}\Sigma_{s,g'\to g}\left(\vec{r_0}+s\vec{\Omega}',\vec{\Omega}'\cdot\vec{\Omega}\right)\varphi_{g'}\left(\vec{r_0}+s\vec{\Omega}',\vec{\Omega}'\right)d\vec{\Omega}'+\frac{\chi_g\left(\vec{r_0}+s\vec{\Omega}\right)}{4\pi k_{eff}}\sum_{g'=1}^{G}\int_0^{4\pi}\nu\Sigma_{f,g'}\left(\vec{r_0}+s\vec{\Omega}'\right)\varphi_{g'}\left(\vec{r_0}+s\vec{\Omega}',\vec{\Omega}'\right)d\vec{\Omega}'$$

- ❑ **Flat & isotropic** source approximations: Q_g in a region is calculated by regional average flux instead of angular flux φ_g
- ❑ MOC solvers use nested **power iteration scheme**

$$\varphi_{g,i,j}\left(\vec{r_0}+s_{i,j}\vec{\Omega}_j\right)=\varphi_{g,i,j}\left(\vec{r_0}\right)\exp\left(-\Sigma_{t,g,i}s_{i,j}\right)+\frac{Q_{g,i}}{\Sigma_{t,g,i}}\left[1-\exp\left(-\Sigma_{t,g,i}s_{i,j}\right)\right]$$

$$\phi_{g,i}=\frac{1}{\Sigma_{t,g,i}}\left[Q_{g,i}+\frac{1}{V_i}\sum_{j\in V_i}w_j\left[\varphi_{g,i,j}(0)-\varphi_{g,i,j}\left(\ell_{i,j}\right)\right]\right]$$

Inner Iteration / Transport sweep

$$Q_{g,i}=\sum_{g'=1}^{G}\Sigma_{s,g'\to g,i}\phi_{g',i}+\frac{\chi_{g,i}}{k_{eff}}\sum_{g'=1}^{G}\nu\Sigma_{f,g',i}\phi_{g',i}$$

Outer iteration

Exponential: The Most Expensive

Akio Yamamoto, Yasunori Kitamura, Yoshihiro Yamane (Annals of Nuclear Energy 31, 1027–1037)

Relative Error of Exponential Function versus Types of Table Lookup

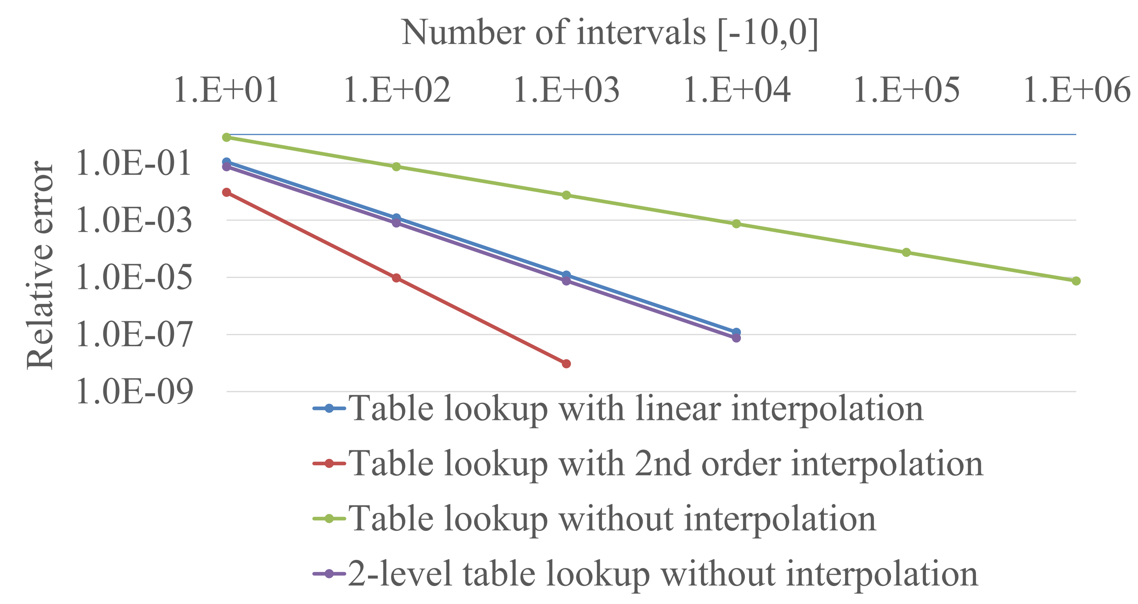

2.8 GHz Xeon processor with 2 MB L3 cache

Calculation time (ns) versus number of intervals

System exp function: 125 ns

	No. of Intervals [-10,0]			
	10	100	1000	10000
1	44	44	44	45
2	47	47	50	54
3	43	43	43	44
4	46	46	47	55

Calculation time (ns) for desired accuracies

System exp function = 125 ns

	Desired Accuracy				
	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}
1	44	44	**45**	96	155
2	47	47	**47**	49	51
3	43	43	**43**	43	43
4	46	46	**46**	46	46

1 = TL without interpolation; 2 =2-level TL without interp.
3 = TL with linear interpolation; 4 = TL with 2nd order interp.

High-Cost Computations in MOC

- ❑ High cost due to repeated calculation of the **transport sweep** (~50 iterations): **over 90% of total CPU time**
- ❑ Traditional (software) transport sweep (each iteration)

```
For all assemblies
 For all characteristics/tracks
  For all segments
   For all energy groups
    Calculate angular fluxes
    Accumulate to regional scalar fluxes
```

- ❑ Impossible to utilize subtask parallelism (pipelining) for the "energy groups" loop

Accelerator Architecture

$$fsji=\varphi_j\left(\vec{r_0}\right);\quad fsjo=\varphi_j\left(\vec{r_0}+s_j\vec{\Omega}_j\right);$$

$$expj=\exp\left(-\Sigma_t\times s_j\right);\quad src=\frac{Q}{\Sigma_t};$$

$$rinv=\frac{1}{\Sigma_t\times V};\quad wj=w_j;\quad avf=\phi;$$

$$fsjo=fsji\times expj-src\times expj+src;$$

$$avf=src+rinv\times\sum wj\left(fsji-fsjo\right)$$

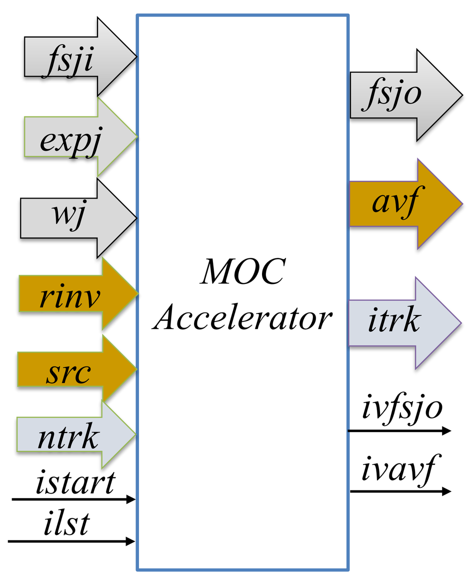

FPGA Implementations

- ❑ **Maximize levels parallelism**

```
For all assemblies
 For all energy groups
  For all number of track groups
   Number of tracks are calculated in
   parallel (hardware resources)
   Pipelining segment calculations
```

- ❑ **Bottlenecks**
  - Hardware resources
  - Pipelining segments from different tracks: input constraints
  - Managing data input, output and control signals

Two Implementations

- ❑ Implement **3-stage Pipelined Arithmetic Circuits**
  - Max. number of pipelining stages by the adder
  - Clock by the longest stage in the adder
  - Performance is limited by the FPGA input constraint
  - Pipeline depth: 18
- ❑ Implement **Using Device Arithmetic IPs**
  - No need to pipeline arithmetic circuits
  - Clock by the slowest arithmetic circuit
  - Pipeline depth: 6

FPGA Specifications

FPGA	Process	Logic Slices/ALM	DSP Slices	RAM (Kbits)
Virtex-7 (XC7V)	28 nm	91,050	1,260	28,620
Altera Arria-10	20 nm	339,620	1,518	48,460

Virtex-7: Logic slice = 4 6-input LUTs and 8 registers. Each DSP has a 25x18 multiplier and a 48-bit accumulator.

Arria-10: ALM = 8-input Adaptive Logic Module and 4 registers. Each DSP has two 18x19 multipliers and a 64-bit accumulator

Experiments

- ❑ **2D C5G7 Benchmark**
  - 4 17x17 pin cell assemblies, 7 different materials
  - 7 energy group nuclear cross-section data
  - Small model: 142,964 flat source regions
- ❑ **3D BEAVRS Benchmark**
  - Representing a Westinghouse PWR
  - 193 fuel assemblies (17x17 fuel rods per assembly)
  - Different enrichments in different assembly
  - Using 70 group cross-section library
- ❑ **Simulations**
  - Transport sweep data exported from OpenMOC runs
  - Repeat simulations for one assembly data

Experiment Results

Hardware Resources Used in Pipelined Arithmetic Implementation

FPGA	LUTs	Registers	DSP Slices
Virtex-7 (XC7V)	54,261	76,962	352
Altera Arria-10	28,448	41,952	192

Hardware Resources Used in Device IP Implementation

FPGA	LUTs	Registers	DSP Slices
Virtex-7 (XC7V)	29,261	54,464	416
Altera Arria-10	15,648	41,952	192

Iteration Rates (in million) from 2D C5G7 Benchmark

Implementation	Xilinx Virtex-7	Altera Arria-10	IBM BG/Q	Intel Xeon
Non-Pipelining	261	396	7.11	65.4
3-stage Pipelined	1,040.8	1,418.8		

Iteration Rates (in million) from 3D BEAVRS Benchmark

Implementation	Xilinx Virtex-7	Altera Arria-10
Non-Pipelining	266	409
3-stage Pipelined	1,057.1	1,434.3

Conclusions

- ❑ The design is independent from problem geometry
- ❑ The level of parallelism in the implementations defines the degree of computational speedup
  - The design mostly benefits large problems
- ❑ The level of parallelism depends on the input constraint of the hardware device and the available hardware resources
  - Limitation due to I/O constraint can be minimized by utilizing device memory for I/O transmissions
- ❑ Minor revision of the host program is required

Key References

- ❑ B. Kochunas, A hybrid parallel algorithm for the 3-D method of characteristics solution of the Boltzmann transport equation on high performance computing clusters. Ph.D. Thesis, University of Michigan, Department of Nuclear Eng. and Radiological Sci., 2013.
- ❑ B. Kelley and E. LarSsen, "2D/1D approximations to the 3D neutron transport equation," International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, Sun Valley, ID, USA, May 2013.
- ❑ W. Boyd, S. Shaner, L. Li, B. Forget, and K. Smith, "The OpenMOC Method of Characteristics Neutral Particle Transport Code," Annals of Nuclear Energy, vol. 68, pp. 43–52, 2014.
- ❑ W. Boyd, A. Siegel, S. He, B. Forget, and K. Smith, "Parallel performance results for the OpenMOC neutron transport code on multi-core platforms," International Journal of High Performance Computing Applications, vol. 30 Issue 3, pp. 360–375, 2016.

SAN JOSÉ STATE UNIVERSITY | ELECTRICAL ENGINEERING