CATEGORY: MEDICAL IMAGING - MI04
POSTER P5210 | CONTACT NAME Tianyu Liu: liut10@rpi.edu
GPU TECHNOLOGY CONFERENCE

Development of a Medical Physics Monte Carlo Radiation Transport Code ARCHER

Tianyu Liu, Noah Wolfe, Christopher D. Carothers, X. George Xu
Nuclear Engineering Program, Rensselaer Polytechnic Institute, Troy, New York 12180, USA
Computer Science Department, Rensselaer Polytechnic Institute, Troy, New York 12180, USA
Email: xug2@rpi.edu   Website: http://RRMDG.rpi.edu

Rensselaer why not change the world?℠

Introduction

I. Accelerators such as Nvidia Graphics Processing Units (GPUs) and Intel Many Integrated Core (MIC) coprocessors are advanced computing devices with outstandingly high computing performance and energy efficiency.

II. The rapid development of accelerators has fostered abundant research activities in the field of Monte Carlo radiation transport.

III. In this NIH-funded (R01EB015478) project, we have been developing a Monte Carlo coupled photon-electron transport code ARCHER (Accelerated Radiation-transport Computation in Heterogeneous EnviRonments) for medical physics applications, and tested the code on three platforms: Intel CPU, Nvidia GPU and Intel MIC coprocessor.

Methods

I. Physics models in ARCHER

- Photon transport: Photoelectric effect, incoherent scattering, coherent scattering and pair production is accurately simulated.
- Electron transport: Class-II condensed history method is used, elastic interaction is modeled by Goudsmit and Saunderson (GS) multiple scattering theory, hard inelastic collision is sampled by Møller cross-section.

II. Concurrent execution implementation

- A generic CPU-GPU-MIC execution model is developed.
- One of the m CPU cores is oversubscribed with n MPI processes, each controlling one GPU. m-1 threads are launched within one MPI process and are bound to m-1 CPU cores. 1 MPI process is launched on MIC, containing 240 threads.

III. Software-based power analysis

Hardware	Utility
CPU (Westmere)	SPEC power_ssj2008 benchmark
CPU (Haswell)	Intel Performance Counter Monitor (PCM)
GPU	Nvidia System Management Interface (nvidia-smi)
MIC	Intel Xeon Phi coprocessor platform status panel (micsmc)

Applications and Results

I. CT dosimetry

Performance of ARCHER for a whole-body CT scan simulation

Code	Hardware	Time [min]	Speedup	Speedup
MCNPX	1 X5650 CPU	476.35		
ARCHER-CPU	1 X5650 CPU	11.22	Baseline	
ARCHER-CPU	1 E5-2697 v3 CPU (Haswell)	3.51	3.20×	Baseline
ARCHER-GPU	1 M2090 GPU	2.08	5.40×	1.69×
ARCHER-GPU	6 M2090 GPUs	0.37	30.23×	9.44×
ARCHER-GPU	1 K20 GPU	1.75	6.40×	2.00×
ARCHER-GPU	1 K40 GPU	1.03	10.89×	3.40×
ARCHER-MIC	1 5110P MIC	3.33	3.37×	1.05×

Efficiency of concurrent execution implementation. Achieved efficiency = the ratio of the actual performance (number of particles simulated per second) to the ideal performance (assuming ideal load balancing)

Hardware	Achieved efficiency
X5650 CPU + 5110P MIC + M2090 GPU	82.0%
X5650 CPU + 5110P MIC + K40 GPU	84.7%

Power profiling of ARCHER for an abdominal CT scan simulation

Code	Hardware	Average power draw [Watt]	Energy consumption [Joule]	Improvement	Improvement
ARCHER-CPU	X5650 CPU	129	9675	Baseline	
ARCHER-CPU	E5-2697 v3 CPU	144.82	3376.85	2.87×	Baseline
ARCHER-GPU	M2090 GPU	137.04	2037.81	4.75×	1.66×
ARCHER-GPU	K20 GPU	98.69	1274.72	7.59×	2.65×
ARCHER-GPU	K40 GPU	121.93	909.28	10.64×	3.71×
ARCHER-MIC	5110P MIC	149.78	3406.03	2.84×	0.99×

II. Radiotherapy dosimetry

The result of ARCHER matches well with GEANT4 (left). ARCHER is considerably faster in producing clinically desired isodose maps (right)

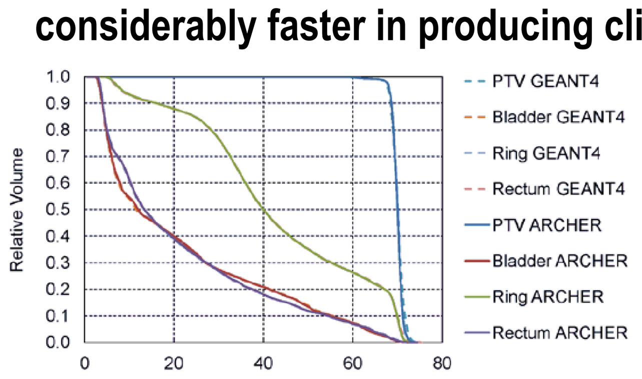

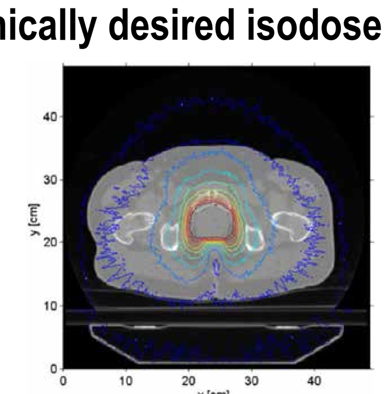

Performance of ARCHER for three cases of radiotherapy simulation

Code	Hardware	Prostate [min]	Lung [min]	Head & neck [min]	Speedup
ARCHER-CPU	1 E5-2620 CPU	729	507	876	Baseline
ARCHER-GPU	1 M2090 GPU	63.4	49.8	79.1	10.18~11.50×
ARCHER-GPU	6 M2090 GPUs	10.9	8.9	13.4	56.97~66.88×
ARCHER-GPU	1 K20 GPU	44.7	35.6	59.4	14.24~16.31×
ARCHER-GPU	1 K40 GPU	36	29.9	44.2	16.96~20.25×

Applications and Results

III. Radiation shielding design

Performance of ARCHER in solving a radiation shielding design problem

Code	Hardware	Time [min]	Speedup
MCNPX	1 E5507 CPU	36.6	
ARCHER-CPU	1 E5507 CPU	4.47	Baseline
ARCHER-GPU	1 K20 GPU	0.7	6.36×

IV. XSBench

- XSBench is a proxy neutronics application of Department Of Energy's (DOE), developed by Argonne National Laboratory (ANL). It models the macroscopic cross-section construction, the most time-consuming subroutine in Monte Carlo reactor criticality calculation.
- The optimization techniques of XSBench are applicable to photon and electron transport code.
- Scudiero (Nvidia) [1] optimized XSBench to the CUDA GPU platform.
- We optimized XSBench to CPU, GPU and MIC platforms.

Performance comparison of original and optimized XSBench on different computing platforms. H-M large problem is used in the test.

Processor	Code	Performance [Lookups/sec]	Speedup	Speedup	Speedup
Westmere	CPU (original)	1,297,450	Baseline		
	CPU (tuned)	1,602,893	1.24×	Baseline	
Haswell	CPU (original)	3,047,600	2.35×		
	CPU (tuned)	4,678,024	3.61×	2.92×	Baseline
Knights Corner	MIC (original)	3,206,490	2.47×		
	MIC (tuned)	7,405,966	5.71×	4.62×	1.58×
Kepler	GPU (original)	1,531,614	1.18×		
	GPU (tuned)	10,214,464	7.87×	6.37×	2.18×

Conclusion

I. ARCHER is an accurate and fast parallel code for Monte Carlo simulation, able to execute on the CPU, GPU and MIC.

II. In our test, Intel 14-core Haswell CPU significantly outperforms the Westmere ancestor, being comparable to Intel 60-core KNC MIC coprocessor. Nvidia Kepler GPUs on the other hand, outperform both the Haswell and KNC processors.

[1] T. Scudiero, "Monte Carlo neutron transport: simulating nuclear reactions one neutron at a time," GPU Technology Conference (GTC) 2014, Nvidia, 2014.