2018 Building Performance Modeling Conference and
SimBuild co-organized by ASHRAE and IBPSA-USA
Chicago, IL
September 26-28, 2018

ENERGYPLUS INTERIOR RADIANT HEAT EXCHANGE RUNTIME PERFORMANCE IMPROVEMENTS

ABSTRACT

EnergyPlus is the flagship whole-building energy simulation program developed by the US Department of Energy. This paper describes the various approaches to improve the performance of the interior radiant heat exchange algorithm. Vectorization, optimized Basic Linear Algebra Subprograms (BLAS) library, multithreading, and Graphical Processing Unit (GPU) computation are all investigated. The approach with the best performance, while maintaining drop-in compatibility with existing EnergyPlus code, is the optimized BLAS library. However, GPU computing has the potential to reduce simulation time by orders of magnitude with major refactoring to EnergyPlus.

INTRODUCTION

EnergyPlus began in 1995 to replace DOE-2 and is currently the US Department of Energy's (DOE's) flagship whole-building energy simulation program (Drury B. Crawley 2001). Since that time, DOE has invested over $65 million in adding new building technologies and modern simulation capabilities. The primary users of EnergyPlus are architects and engineers assessing energy impacts in design, major software vendors offering simulation-based services for buildings, and agencies interested in code and policy impacts on building energy use. EnergyPlus is released under a commercial-friendly, open-source license on GitHub. Building Energy Modeling (BEM) has multiple use cases, both established and emerging:

- Design: architecture, HVAC system selection & sizing
- Operations: HVAC fault diagnosis, dynamic control, model predictive control, & demand response
- Market: code development & compliance, ratings, incentives, M&V, policy, etc.

The EnergyPlus development team is active in the American Society of Heating and Air-Conditioning Engineers (ASHRAE), International Building Performance Simulation Association (IBPSA), and the American Institute of Architects (AIA). The project team also provides public outreach, training, feature requests, troubleshooting, and other support via UnmetHours, email, phone, user group, helpdesk, and other user support. EnergyPlus is capable of modeling most building materials, constructions, and equipment with output from annual to 1 minute resolution for energy use, temperature, relative humidity, and other fields of interest. Many algorithms of varying fidelity exist for modeling certain phenomena within the simulation engine, allowing the user to occasionally define the trade-off between more accurate simulations and longer runtime. EnergyPlus consists of 600,000 lines of Fortran code and was cross-compiled to 700,000 lines of C++ for version 8.0 and is at version 8.8 as of the time of this writing.

To accurately model the build physics, EnergyPlus calculates the heat transfer in the building every timestep. One heat transfer calculation, interior radiant heat exchange, is computationally intensive in its current form, especially for large buildings with many surfaces in each thermal zone. EnergyPlus calculates the net long wave radiation (NLWR) from all surfaces in a thermal zone that have a direct line of sight with another surface. EnergyPlus uses Hottels ScriptF method to approximate the gray body heat exchange of each surface (Hottel 1954). Due to the high computational intensity for calculating ScriptF factors and NLWR for each surface, these were investigated for different methods to speed up the computation. Multiple Central Processing Unit (CPU)-based and Graphical Processing Unit (GPU)-based improvements were studied. To help facilitate easier testing of the various approaches, as well as data structure changes, the original EnergyPlus code for calculating interior radiant heat exchange was ported to a simplified C codebase specific only to this calculation.

CPU-BASED IMPROVEMENTS

This first section looks at various CPU-based performance enhancements to the test code. Since the existing code is serial, CPU-only code, this is a logical starting point for drop-in ready improvements.

As shown in Table 1, auto-vectorization with data structure improvements and BLAS yielded similarly good performance. The auto-vectorized code takes better advan-

Table 1: Comparison of performance improvement techniques using timestep of 15 minutes (35,040 iterations), 256 thermal zones, and 128 surfaces per zone

	Time (minutes)	**Improvement (%)**
Naive	4.48	-
Altered Data Structure	2.04	54.5
BLAS (single threaded)	2.25	49.8
Hand Vectorize	4.09	8.7

tage of the machine used for testing, which is a Haswell-based AVX2 with a 256-bit-wide vector unit and is typical of current Intel-based processors. These techniques have the potential for a 4 speed-up using double precision. Altering the data structure allows: (1) reduced number of calculations for improved arithmetic intensity, and (2) reduced number of memory lookups for improved cache coherence of the loop. The hand-vectorized results show the difficulty in writing high performance code that takes full advantage of the hardware and pipeline. A primary difference between hand- vs. auto-vectorized code was that auto-vectorized code performed aggressive loop unrolling while the hand-vectorized code had none. Writing high performance code under all code uses and machines is challenging; therefore, using a highly optimized BLAS library is often the best approach to improve the performance of an application. The single-threaded BLAS implementation was quicker to implement and less intrusive than the data structure changes while completing in nearly the same time.

Table 2: Comparison of performance improvement techniques using timestep of 15 minutes (35,040 iterations), 1 thermal zone, and 1,024 surfaces per zone with 8 threads (4 cores)

	Time (seconds)	**Improvement (%)**
Naive	67.2	-
BLAS (multithreaded)	6.67	90
OpenMP	13.586	79.8
Altered Data Structure	19.93	70.3

Table 2 quantifies multithreading improvement for a code base that is already vectorized and pipelined. This test stressed the $O(n^2)$ algorithm when calculating the NLWR, where n is the number of surfaces. Comparing the vectorized code improvements in Table 1 and Table 2, the runtime improvement increases from 54.5% to 70.3% as the number of surfaces per zone increases. This is due to the improvements in loop unrolling, cache coherency, and pipelining. The OpenMP version of the vectorized code saw an improvement, but it was only 1.47 times faster when it should have been 8 times faster if there was linear scaling. It is anticipated that this algorithm is memory bound and acknowledge that the OpenMP parallelization may not have been at an optimal location. The multi-threaded BLAS did not have linear scaling, but had the best performance in this test case.

Running Lawrence Berkeley National Laboratorys Empirical Roofline Tool on the test machine gives the roofline model shown in Figure 1 (Lo 2015). The interior radiant heat exchange function has an arithmetic intensity of 1 for the auto-vectorized code. This means the code is memory bound. The optimized BLAS library is cache aware and cache sensitive, so it still improves the performance of this memory-bound problem.

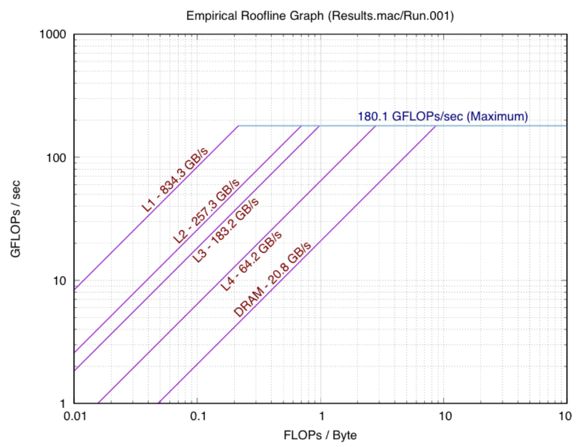

Figure 1: CPU roofline plot for 2013 i7 Haswell MacBook Pro.

GPU-BASED IMPROVEMENTS

A GPU accelerator has the potential to significantly improve the performance of some programs. GPUs typically have higher computational power, relative to traditional CPUs, but relatively slow data transfer mechanisms. GPUs require computational kernels with low-data, high-compute algorithms to attain their full potential. NVIDIAs Compute Unified Device Architecture (CUDA) was used on a GeForce GT 750M graphics card with 384 streaming multiprocessors and 2 GB GDDR5 RAM with roofline model, as shown in Figure 2 (Nickolls 2008). Given the same arithmetic intensity of 1 for the GPU accelerated code, the code is compute bound on this GPU device.

In CUDA programming, a GPU kernel is a grid of blocks and within each block is a grid of threads. This threading hierarchy is how GPUs achieve their massively parallel computation when computational work is divided across all blocks and threads. The direct port of the CPU code to

CUDA has one kernel called every timestep (35,040 iterations). This kernel uses one block and assigns the number of threads to the number of zones. With this setup, buildings with larger numbers of thermal zones take better advantage of the 384 GPU cores.. However, low numbers of thermal zones can lead to worse performance, or modest improvements for this test compared to the serial CPU code. The overhead of repeated GPU kernel calls can overwhelmed parallelization and performance improvements from GPU computation. In order to highlight GPU potential, a test was created using a single GPU kernel for all iterations in all zones. This showed a 96.2% reduction in simulation time, as shown in Table 3. However, this improvement is not realizable within the current structure of EnergyPlus.

Table 3: Comparison of GPU performance improvement techniques using timestep of 15 minutes (35,040 iterations), 16 thermal zone, and 16 surfaces per zone

	Time (seconds)	**Improvement (%)**
CPU-vectorized code	135	-
One kernel total	5.1	96.2
One kernel per iteration	11.1	17.8

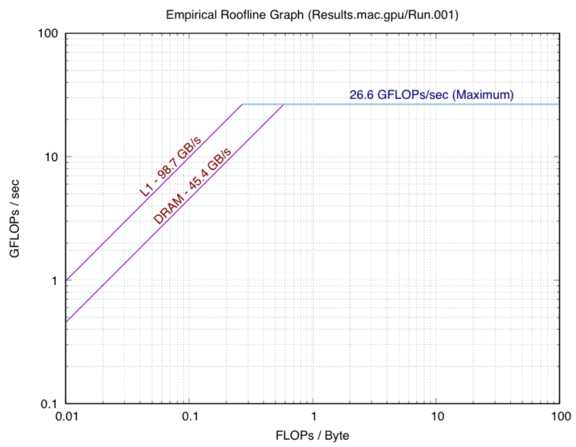

Figure 2: GPU roofline plot for 2013 i7 Haswell MacBook Pro with NVIDIA GT 750M.

FUTURE WORK

The potential of GPU computation within EnergyPlus should be further investigated. Interior radiation heat balance may not be the best place for performance enhancement, other locations in the code are potentially well-suited to GPU computation. Specifically, solar shading calculations seem to lend themselves to significant, GPU-based acceleration.

Replacing hand-derived linear algebra and matrix inversions within EnergyPluss codebase with calls to an optimized BLAS library would likely yield significant performance increases with relatively minor code changes.

CONCLUSION

The highly optimized BLAS library shows the best performance improvements while maintaining the best drop-in compatibility with existing EnergyPlus code and avoids issues that may arise from hand- vectorizing code or relying on the compiler to auto-vectorize code. GPU computation has the potential to improve performance by orders of magnitude; however, this would require a larger refactoring effort to achieve this performance.

ACKNOWLEDGMENT

REFERENCES

Drury B. Crawley, et al. 2001. “EnergyPlus: creating a new-generation building energy simulation program.” *Energy and Buildings* 33 (4): 319 – 331.

Hottel, H. C. 1954. *Radiant heat transmission*. 3.

Lo, Yu Jung, et al. 2015. “Roofline Model Toolkit: A Practical Tool for Architectural and Program Analysis.” *High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation.* Cham: Springer International Publishing, 129–148.

Nickolls, John, et al. 2008. “Scalable Parallel Programming with CUDA.” *Queue* 6 (2): 40–53 (March).

NOMENCLATURE

BEM	Building Energy Modeling
BLAS	Basic Linear Algebra Subprograms
CPU	Central Processing Unit
GPU	Graphical Processing Unit
DDR5	Double Data Rate Type 5 Synchronous Graphics Random-Access Memory
RAM	Random Access Memory
CUDA	Compute Unified Device Architecture
NLWR	Net Long Wave Radiation