Research on GPU-accelerated algorithm in 3D finite difference neutron diffusion calculation method Postprint

Authors: XU Qi, Gang-Lin Yu, Kan Wang, SUN Jia-Long

Date: 2023-06-18T00:00:00+00:00

Abstract

In this paper, the adaptability of the neutron diffusion numerical algorithm on GPUs was studied, and a GPU-accelerated multi-group 3D neutron diffusion code based on finite difference method was developed. The IAEA 3D PWR benchmark problem was calculated in the numerical test. The results demonstrate both high efficiency and adequate accuracy of the GPU implementation for neutron diffusion equation.

Full Text

Preamble

NUCLEAR SCIENCE AND TECHNIQUES 25, 010501 (2014)

Research on GPU-accelerated algorithm in 3D finite difference neutron diffusion calculation method

XU Qi (徐琪),[1],† YU Gang-Lin (余纲林),[1] WANG Kan (王侃),[1] and SUN Jia-Long (孙嘉龙)[1]

[1]Department of Engineering Physics, Tsinghua University, Beijing 100084, China

(Received March 14, 2013; accepted in revised form September 20, 2013; published online February 20, 2014)

Abstract: In this paper, the adaptability of the neutron diffusion numerical algorithm on GPUs was studied, and a GPU-accelerated multi-group 3D neutron diffusion code based on finite difference method was developed. The IAEA 3D PWR benchmark problem was calculated in the numerical test. The results demonstrate both high efficiency and adequate accuracy of the GPU implementation for neutron diffusion equation.

Keywords: Neutron diffusion, Finite difference, Graphics Processing Unit (GPU), CUDA, Acceleration

DOI: 10.13538/j.1001-8042/nst.25.010501

Introduction

In the field of reactor physics, numerical solutions of 3-dimensional neutron diffusion equation are always required. Compared with the coarse mesh nodal techniques, the finite difference method is considered simpler and more precise; however, it costs unendurable computer time when analyzing a full-size reactor core.

Since 2006, NVIDIA' s GPUs (Graphics Processing Units) have provided us with tremendous computational horsepower because of the release of CUDA [1]. In the field of nuclear reactor physics, the importance of the GPU+CPU heterogeneous platform has been growing gradually. Prayudhatama et al. [2] implemented a 1-D finite difference diffusion code on GPUs in 2010, and obtained up to 70× speedup compared to a corresponding CPU code. In 2011, Kodama et al. [3] ported the code SCOPE2 to GPUs, achieving about 3 times speedup. In the same year, Gong et al. [4] exploited the parallelism of GPUs for the Sn code Sweep3D, which was speeded up by about 2 to 8 times.

In this work, a GPU-accelerated multi-group 3D neutron diffusion code based on finite difference method was implemented and optimized. The IAEA 3D PWR benchmark problem [5] was utilized to prove the high computational efficiency and accuracy of the GPU version code. The result in this work shows a bright future of GPU applications in nuclear reactor analysis.

Neutron Diffusion Equation

According to the neutron diffusion theory, we have the multi-group neutron diffusion equation [6] as below:

$$-\nabla \cdot D_g \nabla \phi_g + \Sigma_g^t(\mathbf{r})\phi_g = \sum_{g'=1}^{G} \Sigma_{g'g}^s \phi_{g'} + \frac{1}{k_{\text{eff}}} \sum_{g'=1}^{G} \nu \Sigma_{g'}^f \phi_{g'}$$

where g is the energy group number, ranging from 1 to G, k_{eff} is the effective multiplication factor, and ϕ_g is the gth neutron flux. We solve this equation by the source iteration methodology [6], which includes inner and outer iterations. The inner iteration computes a group of linear algebra equations in the form of $AX = B$ given the neutron scattering source and fission source. In the outer iteration, we use the neutron flux to update the neutron source and prepare for the next inner iteration.

In this work, we focus on accelerating the inner iteration. Suppose the neutron source on the right hand of Eq. (1) is known, then after discretization on an

XYZ grid, we can obtain Eq. (2), which is a linear equation with a 7-diagonal positive definite matrix as the coefficient:

$$a_{i,j,k}\phi_{i-1,j,k}+b_{i,j,k}\phi_{i,j-1,k}+c_{i,j,k}\phi_{i,j,k-1}+d_{i,j,k}\phi_{i,j,k}+e_{i,j,k}\phi_{i+1,j,k}+f_{i,j,k}\phi_{i,j+1,k}+h_{i,j,k}\phi_{i,j,k+1}=s_{i,j,k},\quad 1\leq i\leq$$

Eq. (2) represents a large-scale sparse matrix problem for full-size reactor analysis. From the perspective of numerical mathematics, Jacobi iteration is inefficient; thus, arithmetic techniques such as CG, SOR, LSOR, and ADI (Alternating Direction Implicit method) are needed for efficient calculation.

GPU Implementation Details

In order to test the computational capacity of GPUs, we do not resort to any mathematical skill; instead, the Jacobi iteration method is adopted for the inner iteration. The inner iteration is ported to a GTX TITAN GPU. In the inner iteration, the neutron flux is estimated according to the neutron source calculated with the neutron flux from the last source iteration. The outer iteration remains on CPUs to calculate the neutron effective multiplication factor according to the neutron fission source from GPUs. After an outer iteration, the effective multiplication factor is transferred from CPU memory to GPU memory to obtain the neutron source for the next source iteration. Fig. 1 [Figure 1: see original paper] shows the task distribution and data transfer between GPUs and CPUs during one source iteration.

A. Solving Neutron Flux

The neutron flux is solved via Jacobi inner iteration on GPUs. Because of the natural parallelism of the Jacobi iteration, there is high potential to implement this algorithm on GPUs with exciting speedups.

According to CUDA, GPUs have two levels of parallelism: the first level is called grids of thread blocks, while the second level is blocks of threads. One thread block is designed to be mapped to an SM (Streaming Multiprocessor) on GPU chips, and one thread to be mapped to an SP (Streaming Processor) in an SM. In Jacobi iteration, the main computing tasks are production and addition operations at each flux point, as shown by Eq. (2). To speed up such iterations, the operations at each flux point (i, j, k) should be allocated to a specific GPU thread so that the computation tasks can be spread among the SPs on GPUs. Fig. 2 [Figure 2: see original paper] demonstrates the mapping relationships between flux points and threads. As can be seen in Fig. 2, one flux point is mapped to one GPU thread, and each thread is responsible for updating the flux at that point.

For a large-scale 3D reactor model, there will be millions or even tens of millions of flux points that need to be updated using the surrounding old flux; however, the hardware resources of a GPU chip are limited and cannot create as many

threads as flux points. To solve this problem, we update the neutron flux layer by layer as illustrated in Fig. 2. This approach provides sufficient computing resources for a GPU to accelerate the inner iteration procedure for each layer of flux points.

B. Generating Sources and Data Movements

When the neutron flux is solved after inner iterations, the fission source and the neutron source can be determined by the following equations:

$$S_{\text{fission}} = \sum_{g'=1}^{G} \nu \Sigma_{g'}^{f} \phi_{g'}$$

$$S_{\text{neutron},g} = \frac{S_{\text{fission}}}{k_{\text{eff}}} + \sum_{g'=1}^{G} \Sigma_{g'g}^{s} \phi_{g'}$$

where S_{fission} stands for the fission source and $S_{\text{neutron},g}$ stands for the neutron source of energy group g, both of which are calculated from the newly updated neutron flux. In order to reduce data exchange between CPU and GPU memories, these two sources are obtained on GPUs in parallel.

As shown in Fig. 1, there are three data movements during one source iteration. The first data transfer occurs after the fission source is created, moving the fission source from device memory to host memory to calculate the effective multiplication factor k_{eff} by accumulating the fission source of each flux point. The second data movement is for comparison between old and new neutron flux, during which the new neutron flux is transferred from device to host. The third transfer moves k_{eff}, a double-type variable, back to device memory to obtain the neutron source for the next iteration.

C. Data Storage

For the fine grid finite difference method, a large number of flux points leads to substantial memory space requirements. Suppose there are N_g energy groups, and N_x, N_y, N_z flux points in the X, Y, Z directions respectively, then the memory space to store the eight coefficients (including the neutron source $S_{i,j,k}$) would be $4 \times N_g \times N_x \times N_y \times N_z \times 8$ bytes, and the memory for the neutron flux would be $4 \times N_g \times N_x \times N_y \times N_z$ bytes. When analyzing 3D full-size reactors using GPUs, all the above data should be allocated to GPU memory, which has limited volume. For GTX TITAN, the device memory is up to 6 GB under 64-bit operating systems.

Under CUDA, a programmer is allowed to manage seven different kinds of memory space, among which only global memory and texture memory can be utilized to store coefficient data and flux data. Because texture memory has a texture

cache and higher bandwidth than global memory, it is advantageous for frequently accessed data. The only limitation of texture memory is that it is read-only. Thus, the coefficient data can be placed in texture memory, and the flux data allocated to global memory.

Performance Test

In this section, we demonstrate the accuracy and efficiency of the GPU-accelerated code. Additionally, we discuss a method of performance improvement through overclocking GPU processors.

A. Experiment Platform and Benchmark Problem

The accuracy of the GPU version diffusion code is tested by comparing the neutron flux computed by CITATION [7]. To prove the efficiency of the GPU code, we measure the performance of three diffusion codes listed in Table 1 . 3DFD-CPU is a serial CPU version code which uses the Jacobi iteration method for inner iterations. 3DFD-GPU is obtained by accelerating the inner iteration part of 3DFD-CPU utilizing GPUs. HYPRE-8CORE [8] is a parallel diffusion code running on an 8-core CPU. The computing hardware for these codes is also shown in Table 1.

The IAEA PWR benchmark problem, shown in Fig. 3 [Figure 3: see original paper], is used for the numerical experiment. This is an important benchmark problem widely used to test the performance of neutron deterministic codes. The core is composed of 177 fuel assemblies, 9 of which are fully rodded and 4 of which are partially rodded. There are 64 reflector assemblies surrounding the core. The size of the assemblies is 20 cm $\times$ 20 cm $\times$ 340 cm, while the size of 1/4 core is 170 cm $\times$ 170 cm $\times$ 380 cm.

B. Accuracy of the GPU Code

To prove the accuracy of GPU computation, we compare the power distribution of 3DFD-CPU with that of CITATION. CITATION, developed by ORNL, is an industrial-class code for solving the neutron diffusion equation. The comparison results are shown in Fig. 4 [Figure 4: see original paper]. The convergence criterion is set so that the simulation ends when the effective multiplication factor relative error is less than 1.0×10^{-6} and the maximum point flux relative error is less than 1.0×10^{-5}. The computing grid size used in Fig. 4 is 2 cm, resulting in 1,372,750 spatial flux points.

In Fig. 4, the relative error represents the difference between the code result and the benchmark result. The power distribution of the GPU version code is close to that of CITATION. The accuracy comparison demonstrates that there is no need to worry about the accuracy and reliability of GPUs.

C. Efficiency of the GPU Code

We use the codes listed in Table 1 to testify to the computing power of GPUs. First, 3DFD-GPU is compared with the 8-core CPU parallelized code HYPRE-8CORE, and then a comparison of computing time between 3DFD-GPU and 3DFD-CPU is made.

According to Ref. [8], the author utilized the MPI-based parallelized linear algebra library HYPRE [9] to accelerate the diffusion code. Here we refer to the corresponding code in Ref. [8] as HYPRE-8CORE. HYPRE is a library developed by LLNL for solving large sparse linear systems of equations on massively parallel computers. On an 8-core tower server, the inner iteration part of the diffusion code is accelerated by the parallelized Conjugate Gradient algorithm. During simulation, the computing grid size is set to 2.5 cm, and the convergence standard is that k_{eff} relative error converges to 1.0×10^{-5} and the maximum point flux relative error to 1.0×10^{-4}. The computation speed comparison between 3DFD-GPU and HYPRE-8CORE is shown in Table 2 .

In Table 2, although HYPRE-8CORE is accelerated by an 8-core server, 3DFD-GPU performs better.

The performance comparison of 3DFD-GPU and 3DFD-CPU is shown in Fig. 5 [Figure 5: see original paper]. We use six grid sizes, from 5 cm × 5 cm × 5 cm to 1 cm × 1 cm × 1 cm, to demonstrate the acceleration characteristic of GPUs for Jacobi iteration. Table 3 lists the grid sizes and corresponding grid numbers. The convergence criterion is that k_{eff} relative error converges to 1.0×10^{-6} and the maximum point flux relative error to 1.0×10^{-5}.

Figure 5 shows the remarkable accelerating power of GPUs compared with CPUs, especially when the grid size is set to 2 cm, where a speedup factor of 86 was obtained. This phenomenon is caused by latency hiding: when the problem scales up and the amount of data increases, all cores on the GPU work at full capacity, allowing data transfer from GPU memory by part of thread blocks to operate while other blocks execute computational tasks. However, as seen from Fig. 5, it should be noted that oversized data amounts may decrease speedups because the communication overhead between host and device increases and the sequential part of the code may play an increasingly important role in the whole process.

D. Performance Improvement by Overclocking

In order to achieve the same performance with lower energy consumption, NVIDIA decreased the base clock of GPUs in the Kepler series while increasing the number of streaming processors in streaming multiprocessors (SMX). The base core clock of GTX TITAN is 837 MHz, which is lower than that of GTX 580 (Fermi architecture, 1544 MHz). We use the overclocking utility NVIDIA Inspector to set the core clock to 1166 MHz and the memory clock to 3334 MHz. Fig. 6 [Figure 6: see original paper] shows the performance improvement

after overclocking, where the runtime and speedup factor of 3DFD-GPU before and after overclocking are compared, with the speedup factor relative to the runtime of 3DFD-CPU.

Through overclocking, the GPU acceleration effect is improved. The performance improvement depends on the scale of the analyzed problem; that is, more obvious performance enhancement can be obtained when the grid number increases.

Conclusion

In this work, a GPU-accelerated multi-group 3D neutron diffusion code based on finite difference method was developed to speed up the finite difference methodology and examine GPU performance. The IAEA 3D PWR benchmark problem was used as the problem model in the numerical experiment. By comparing the power distribution obtained from 3DFD-GPU and CITATION, we proved the accuracy of GPU computing. The performance advantage of GPUs was also demonstrated by comparing the runtime of 3DFD-GPU, 3DFD-CPU, and HYPRE-8CORE.

Regarding future work, mathematical accelerating techniques such as the Conjugate Gradient method and the Chebyshev extrapolation method will be adopted to reduce the runtime of the GPU-based finite difference method to the same order of magnitude as the coarse mesh nodal methodology.

References

[1] NVIDIA Corporation. CUDA C Programming Guide. 2012, 3-4 and 71-75.

[2] Prayudhatama D, Waris A, Kurniasih N, et al. Proceedings of AIP Conference Proceedings, 2010, 1244: 121-126.

[3] Kodama Y, Tatsumi M, Ohoka Y. Study on GPU Computing for SCOPE2 with CUDA. Proceedings of International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C2011), Brazil, 2001.

[4] Gong C, Liu J, Chi L, et al. GPU Accelerated Simulations of 3D Deterministic Particle Transport Using Discrete Ordinates Method. Journal of Computational Physics, 2011, 230(15): 6116-6132.

[5] Argonne National Laboratory. Benchmark Problem Book. ANL-7416, Suppl.2, 1977, 277-280.

[6] Duderstadt J J and Hamilton L J, Nuclear Reactor Analysis. New York (USA): John Wiley & Sons, Inc.,1976, 285-314.

[7] Fowler T B, Vondy D R, Cunningham G W. Nuclear Reactor Core Analysis Code CITATION, ORNL-TM-2496, Supplement 2. ORNL, 1972, 104-140.

[8] Wu W B, Li Q, Wang K. Parallel Solution of 3D Neutron Diffusion Equation Based on HYPRE. Science and Technology on Reactor System Design Technology Laboratory Annual Report. Chengdu, China, 2010, 35-40 (in Chinese).

[9] Lawrence Livermore National Laboratory. HYPRE User' s Manual (Version 2.7.0b). 2011, 1-6.

Note: Figure translations are in progress. See original paper for figures.