

Global Memory Access Modelling for Efficient Implementation of the Lattice Boltzmann Method on Graphics Processing Units*

Christian Obrecht¹, Frédéric Kuznik¹, Bernard Tourancheau², and Jean-Jacques Roux¹

¹ Centre de Thermique de Lyon
(UMR 5008 CNRS, INSA-Lyon, Université de Lyon)
Bât. Sadi Carnot, 9 rue de la Physique, 69621 Villeurbanne cedex
`{christian.obrecht, frederic.kuznik, jean-jacques.roux}@insa-lyon.fr`

² Laboratoire de l'Informatique du Parallélisme
(UMR 5668 CNRS, ENS de Lyon, INRIA, UCB Lyon 1)
École Normale Supérieure de Lyon, 46 allée d'Italie, 69364 Lyon cedex 07
`bernard.tourancheau@ens-lyon.fr`

Abstract. In this work, we investigate the global memory access mechanism on recent GPUs. For the purpose of this study, we created specific benchmark programs, which allowed us to explore the scheduling of global memory transactions. Thus, we formulate a model capable of estimating the execution time for a large class of applications. Our main goal is to facilitate optimisation of regular data-parallel applications on GPUs. As an example, we finally describe our CUDA implementations of LBM flow solvers on which our model was able to estimate performance with less than 5% relative error.

Keywords: GPU computing, CUDA, lattice Boltzmann method, CFD

Addressed topics: parallel computing, performance analysis, multi-physics problems

Introduction

State-of-the-art graphics processing units (GPU) have proven to be extremely efficient on regular data-parallel algorithms [1]. For many of these applications, like lattice Boltzmann method (LBM) fluid flow solvers, the computational cost is entirely hidden by global memory access. The present study intends to give some insight on the global memory access mechanism of the nVidia's GT200 GPU. The obtained results led us to optimisation elements which we used for our implementations of the LBM.

* Candidate to the Best Student Paper Award.

The structure of this paper is as follows. First, we briefly review nVidia’s compute unified device architecture (CUDA) technology and the algorithmic aspects of the LBM. Then, we describe our measurement methodology and results. To conclude, we present our CUDA implementations of the LBM.

1 Compute Unified Device Architecture

CUDA capable GPUs, i.e. the G8x, G9x, and GT200 processors consist in a variable amount of texture processor clusters (TPC) containing two (G8x, G9x) or three (GT200) streaming multiprocessors (SM), texture units and caches [2]. Each SM contains eight scalar processors (SP), two special functions units (SFU), a register file, and shared memory. Registers and shared memory are fast but in rather limited amount, e.g. 64 KB and 16 KB per SM for the GT200. On the other hand, the off-chip global memory is large but suffers from high latency and low throughput compared to registers or shared memory.

The CUDA programming language is an extension to C/C++. Functions intended for GPU execution are named *kernels*, which are invoked on an execution grid specified at runtime. The execution grid is formed of blocks of threads. The blocks may have up to three dimensions, the grid two. During execution, blocks are dispatched to the SMs and split into warps of 32 threads.

CUDA implementations of data intensive applications are usually bound by global memory throughput. Hence, to achieve optimal efficiency, the number of global memory transactions should be minimal. Global memory transactions within a half-warp are coalesced into a single memory access whenever all the requested addresses lie in the same aligned segment of size 32, 64, or 128 bytes. Thus, improving the data access pattern of a CUDA application may dramatically increase performance.

2 Lattice Boltzmann Method

The Lattice Boltzmann Method is a rather innovative approach in computational fluid dynamics [3,4,5]. It is proven to be a valid alternative to the numerical integration of the Navier-Stokes equations. With the LBM, space is usually represented by a regular lattice. The physical behaviour of the simulated fluid is determined by a finite set of *mass fractions* associated to each node. From an algorithmic standpoint, the LBM may be summarised as:

```

for each time step do
  for each lattice node do
    if boundary node then
      apply boundary conditions
    end if
    compute new mass fractions
    propagate to neighbouring nodes
  end for

```

end for

The propagation phase follows some specific stencil. Figure 1 illustrates D3Q19, the most commonly used three-dimensional stencil, in which each node is linked to 18 of its 27 immediate neighbours.¹




Fig. 1. The D3Q19 stencil

CUDA implementations of the LBM may take advantage of its inherent data parallelism by assigning a thread to each node, the data being stored in global memory. Since there is no efficient global synchronisation barrier, a kernel has to be invoked for each time step [6]. CPU implementations of the LBM usually adopt an array of structures (AoS) data layout, which improves locality of mass fractions belonging to a same node [7]. On the other hand, CUDA implementations benefit from structure of arrays (SoA) data layouts, which allows coalesced global memory accesses [8]. However, this approach is not sufficient to ensure optimal memory transactions, since propagation corresponds to one unit shifts of global memory addresses for the minor spatial dimension. In other words, for most mass fractions, the propagation phase yields misalignments. A way to solve this issue consists in performing propagation partially in shared memory [9]. Yet, as shown in [10], this approach is less efficient than using carefully chosen propagation schemes in global memory.

3 Methodology

To study transactions between global memory and registers, we used kernels performing the following operations :

1. Store time t_0 in a register.

¹ Taking the stationary mass fraction into account, the number of mass fractions per node amounts to 19, hence D3Q19.

2. Read N words from global memory, with possibly L misalignments.
3. Store time t_1 in a register.
4. Write N words to global memory, with possibly M misalignments.
5. Store time t_2 in a register.
6. Write t_2 to global memory.

Time is accurately determined using the CUDA `clock()` function which gives access to counters that are incremented at each clock cycle. Our observations enabled us to confirm that these counters are per TPC, as described in [11], and not per SM as stated in [2]. Step 6 may influence the timings, but we shall see that it can be neglected under certain circumstances.

The parameters of our measurements are N , L , M , and k , the number of warps concurrently affected to each SM. Number k is proportional to the occupancy rate α , which is the ratio of active warps to the maximum number of warps supported on one SM. With the GT200, this maximum number being 32, we have: $k = 32\alpha$.

We used a one-dimensional grid and one-dimensional blocks containing one single warp. Since the maximum number of blocks supported on one SM is 8, the occupancy rate is limited to 25%. Nonetheless, this rate is equivalent to the one obtained with actual CUDA applications.

We chose to create a script generating the kernels rather than using runtime parameters and loops, since the layout of the obtained code is closer to the one of actual computation kernels. We processed the CUDA binaries using `decuda` [12] to check whether the compiler had reliably translated our code. We carried out our measurements on a GeForce GTX 295 graphics board, featuring two GT200 processors.²

4 Modelling

At kernel launch, blocks are dispatched to the TPCs one by one up to k blocks per SM [13]. Since the GT200 contains ten TPCs, blocks affected to the same TPC have identical `blockIdx.x` unit digit. This enables to extract information about the scheduling of global memory access at TPC level. In order to compare the measurements, as the clock registers are peculiar to each TPC [11], we shifted the origin of the time scale to the minimal t_0 . We noticed that the obtained timings are coherent on each of the TPCs.

For a number of words read and written $N \leq 20$, we observed that:

- Reads and writes are performed in one stage, hence storing of t_2 has no noticeable influence.
- Warps 0 to 8 are launched at once (in a determined but apparently incoherent order).

² In the CUDA environment, the GPUs of the GTX 295 are considered as two distinct devices. It should be noted that our benchmark programs involve only one of those devices.

- Subsequent warps are launched one after the other every ~ 63 clock cycles.

For $N > 20$, reads and writes are performed in two stages. One can infer the following behaviour: if the first n warps in a SM read at least 4,096 words, where $n \in \{4, 5, 6\}$, then the processing of the subsequent warps is postponed. The number of words read by the first n warps being $n \times 32N$, this occurs whenever $n \times N \geq 128$. Hence, $n = 4$ yields $N \geq 32$, $n = 5$ yields $N \geq 26$, and $n = 6$ yields $N \geq 21$.

Time t_0 for the first $3n$ warps of a TPC follow the same pattern as in the first case. We also noticed a slight overlapping of the two stages, all the more as storing t_2 should here be taken into account. Nonetheless, the read time for the first warp in the second stage is noticeably larger than for the next ones. Therefore, we may consider, as a first approximation, that the two stages are performed sequentially.

In the targeted applications, the global amount of threads is very large. Moreover, when a set of blocks is affected to the SMs, the scheduler waits until all blocks are completed before providing new ones. Hence, knowing the average processing time T of k warps per SM allows to estimate the global execution time.

For $N \leq 20$, we have $T = \ell + T_R + T_W$, where ℓ is time t_0 for the last launched warp, T_R is read time, and T_W is write time. Time ℓ only depends on k . For $N > 20$, we have $T = T_0 + \ell' + T'_R + T'_W$, where T_0 is the processing time of the first stage, $\ell'(i) = \ell(i - 3n + 9)$ with $i = 3k - 1$, T'_R and T'_W are read and write times for the second stage.




Fig. 2. Launch delay in respect of warp rank

To estimate ℓ , we averaged t_0 over a large number of warps. Figure 2 shows, in increasing order, the obtained times in cycles. Numerically, we have $\ell(i) \approx 0$ for $i \leq 9$ and $\ell(i) \approx 63(i - 10) + 13$ otherwise.

5 Throughput

5.1 $N \leq 20$

Figures 3 and 4 show the distribution of read and write times for 96,000 warps with $N = 19$. The bimodal shape of the read time distribution is due to translation look-aside buffer (TLB) misses [14]. This aspect is reduced when adding misalignments, since the number of transactions increases while the number of misses remains constant. Using the average read time to approximate T is acceptable provided no special care is taken to avoid TLB misses.




Fig. 3. Read time for $N = 19$




Fig. 4. Write time for $N = 19$

We observed that average read and write times depend linearly of N . Numerically, with $k = 8$, we obtained:

$$T_R \approx 317(N - 4) + 440 \quad T_W \approx 562(N - 4) + 1,178$$

$$T_{R'} \approx 575(N - 4) + 291 \quad T_{W'} \approx 983(N - 4) + 2,030$$

where $T_{R'}$ and $T_{W'}$ are read and write times with $L = N$ and $M = N$ misalignments. Hence, we see that writes are more expensive than reads. Likewise, misalignments in writes are more expensive than misalignments in reads.

5.2 $21 \leq N \leq 39$

As shown in figures 5 and 6, T_0 , T'_R , and T'_W depend linearly of N in the three intervals $\{21, \dots, 25\}$, $\{26, \dots, 32\}$, and $\{33, \dots, 39\}$. As an example, for the third interval, we obtain:

$$T_0 \approx 565(N - 32) + 15,164$$

$$T'_R \approx 112(N - 32) + 2,540 \quad T'_W \approx 126(N - 32) + 3,988$$




Fig. 5. First stage duration




Fig. 6. Timings in second stage

5.3 Complementary studies

We also investigated the impact of misalignments and occupancy rate on average read and write times. Figures 7 and 8 show obtained results for $N = 19$.




Fig. 7. Misaligned reads




Fig. 8. Occupancy impact

For misaligned reads, we observe that the average write time remains approximately constant. Read time increases linearly with the number of misalignments until some threshold is reached. From then on, the average read time is maximal. Similar conclusion can be drawn for misaligned writes.

Average read and write times seem to depend quadratically on k . Since the amount of data transferred depends only linearly on k , this leads to think that the scheduling cost of each warp is itself proportional to k .

6 Implementations

We implemented several LBM fluid flow solvers: a D3Q19 LBGK [4], a D3Q19 MRT [5], and a double population thermal model requiring 39 words per node [15]. Our global memory access study lead us to multiple optimisations. For each implementation, we used a SoA like data layout, and a two-dimensional grid of one-dimensional blocks. Since misaligned writes are more expensive than misaligned reads, we experimented several propagation schemes in which misalignments are deferred to the read phase of the next time step. The most efficient appears to be the reversed scheme where propagation is entirely performed at reading, as outlined in figure 9. For the sake of simplicity, the diagram shows a two-dimensional version.




Fig. 9. Reversed propagation scheme

Performance of a LBM based application is usually given in million lattice node updates per second (MLUPS). Our global memory access model enables us to give an estimate of the time T (in clock cycles) required to process k warps per SM. On the GT200, where the number of SMs is 30 and the warp size is 32, k warps per SM amounts to $K = 30 \times k \times 32 = 960k$ threads. Since one thread takes care of one single node, T is therefore the number of clock cycles needed to perform K lattice node updates. Hence, using the global memory frequency F in MHz, the expected performance in MLUPS is: $P = (K/T) \times F$.

With our D3Q19 implementations, for instance, we have $N = 19$ reads and writes, $L = 10$ misaligned reads, no misaligned writes, and 25% occupancy (thus $k = 8$). Using the estimation provided by our measurements, we obtain: $T = \ell + T_R + T_W = 15,594$. Since $K = 7,680$ and $F = 999$ MHz, we have $P = 492$ MLUPS.

To summarize, table 1 gives both the actual and estimated performances for our implementations on a 128^3 lattice. Our estimations appear to be rather accurate, thus validating our model.

Model	Occupancy	Actual	Estimated	Relative error
D3Q19 LBGK	25%	481	492	2.3%
D3Q19 MRT	25%	516	492	4.6%
Thermal LBM	12.5%	195	196	1.0%

Table 1. Performance of LBM implementations (in MLUPS)

Summary and discussion

In this work, we present an extensive study of the global memory access mechanism between global memory and GPU for the GT200. A description of the

scheduling of global memory accesses at hardware level is given. We express a model which allows to estimate the global execution time of a regular data-parallel application on GPU. The cost of individual memory transactions and the impact of misalignments is investigated as well.

We believe our model is applicable to other GPU applications provided certain conditions are met:

- The application should be data-parallel and use a regular data layout in order to ensure steady data throughput.
- The computational cost should be negligible as compared with the cost of global memory reads and writes.
- The kernel should make moderate use of branching in order to avoid branch divergence, which can dramatically impact performance. This would probably not be the case with an application dealing, for instance, with complex boundaries.

On the other hand, our model does not take possible TLB optimisation into account. Hence, some finely tuned applications may slightly outvalue our performance estimation.

The insight provided by our study, turned out to be useful in our attempts to optimize CUDA implementations of the LBM. It may contribute to efficient implementations of other applications on GPU.

Acknowledgement

We thank the reviewers for their helpful comments improving the clarity of our paper.

References

1. Dongarra, J., Moore, S., Peterson, G., Tomov, S., Allred, J., Natoli, V., Richie, D.: Exploring new architectures in accelerating CFD for Air Force applications. In: Proceedings of HPCMP Users Group Conference, Citeseer (2008) 14–17
2. nVidia: Compute Unified Device Architecture Programming Guide version 2.3.1. (August 2009)
3. McNamara, G.R., Zanetti, G.: Use of the Boltzmann Equation to Simulate Lattice-Gas Automata. *Phys. Rev. Lett.* **61** (1988) 2332–2335
4. Qian, Y.H., d’Humières, D., Lallemand, P.: Lattice BGK models for Navier-Stokes equation. *Europhys. Lett* **17**(6) (1992) 479–484
5. d’Humières, D., Ginzburg, I., Krafczyk, M., Lallemand, P., Luo, L.: Multiple-relaxation-time lattice Boltzmann models in three dimensions. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences* (2002) 437–451
6. Ryoo, S., Rodrigues, C.I., Baghsorkhi, S.S., Stone, S.S., Kirk, D.B., Hwu, W.W.: Optimization principles and application performance evaluation of a multithreaded GPU using CUDA. In: Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming, ACM New York, NY, USA (2008) 73–82

7. Pohl, T., Kowarschik, M., Wilke, J., Iglberger, K., Råde, U.: Optimization and Profiling of the Cache Performance of Parallel Lattice Boltzmann Codes. *Parallel Processing Letters* **13**(4) (2003) 549–560
8. Kuznik, F., Obrecht, C., Rusaouën, G., Roux, J.J.: LBM Based Flow Simulation Using GPU Computing Processor. *Computers and Mathematics with Applications* (27) (June 2009)
9. Tölke, J., Krafczyk, M.: TeraFLOP computing on a desktop PC with GPUs for 3D CFD. *International Journal of Computational Fluid Dynamics* **22**(7) (2008) 443–456
10. Obrecht, C., Kuznik, F., Tourancheau, B., Roux, J.J.: A new approach to the lattice Boltzmann method for graphics processing units. *Computers and Mathematics with Applications* (in press) (2010)
11. Papadopoulou, M., Sadooghi-Alvandi, M., Wong, H.: Micro-benchmarking the GT200 GPU
12. van der Laan, W.J.: Decuda G80 disassembler version 0.4. (2007)
13. Collange, S., Defour, D., Tisserand, A.: Power Consumption of GPUs from a Software Perspective. In: *Proceedings of the 9th International Conference on Computational Science: Part I*, Springer (2009) 923
14. Volkov, V., Demmel, J.: Benchmarking GPUs to tune dense linear algebra. In: *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, IEEE Press Piscataway, NJ, USA (2008)
15. Peng, Y., Shu, C., Chew, Y.T.: A 3D incompressible thermal lattice Boltzmann model and its application to simulate natural convection in a cubic cavity. *Journal of Computational Physics* **193**(1) (2004) 260–274