

An Analytical Model for a GPU Architecture with Memory-level and Thread-level Parallelism Awareness

Sunpyo Hong
Electrical and Computer Engineering
Georgia Institute of Technology
shong9@gatech.edu

Hyesoon Kim
School of Computer Science
Georgia Institute of Technology
hyesoon@cc.gatech.edu

ABSTRACT

GPU architectures are increasingly important in the multi-core era due to their high number of parallel processors. Programming thousands of massively parallel threads is a big challenge for software engineers, but understanding the performance bottlenecks of those parallel programs on GPU architectures to improve application performance is even more difficult. Current approaches rely on programmers to tune their applications by exploiting the design space exhaustively without fully understanding the performance characteristics of their applications.

To provide insights into the performance bottlenecks of parallel applications on GPU architectures, we propose a simple analytical model that estimates the execution time of massively parallel programs. The key component of our model is estimating the number of parallel memory requests (we call this the memory warp parallelism) by considering the number of running threads and memory bandwidth. Based on the degree of memory warp parallelism, the model estimates the cost of memory requests, thereby estimating the overall execution time of a program. Comparisons between the outcome of the model and the actual execution time in several GPUs show that the geometric mean of absolute error of our model on micro-benchmarks is 5.4% and on GPU computing applications is 13.3%. All the applications are written in the CUDA programming language.

Categories and Subject Descriptors

C.1.4 [Processor Architectures]: Parallel Architectures
; C.4 [Performance of Systems]: Modeling techniques
; C.5.3 [Computer System Implementation]: Microcomputers

General Terms

Measurement, Performance

Keywords

Analytical model, CUDA, GPU architecture, Memory level parallelism, Warp level parallelism, Performance estimation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISCA'09, June 20–24, 2009, Austin, Texas, USA.

Copyright 2009 ACM 978-1-60558-526-0/09/06 ...\$5.00.

1. INTRODUCTION

The increasing computing power of GPUs gives them considerably higher peak computing power than CPUs. For example, NVIDIA's GTX280 GPUs [3] provide 933 Gflop/s with 240 cores, while Intel's Core2Quad processors [2] deliver only 100 Gflop/s. Intel's next generation of graphics processors will support more than 900 Gflop/s [26]. AMD/ATI's latest GPU (HD4870) provides 1.2 Tflop/s [1]. However, even though hardware is providing high performance computing, writing parallel programs to take full advantage of this high performance computing power is still a big challenge.

Recently, there have been new programming languages that aim to reduce programmers' burden in writing parallel applications for the GPUs such as Brook+ [5], CUDA [22], and OpenCL [16]. However, even with these newly developed programming languages, programmers still need to spend enormous time and effort to optimize their applications to achieve better performance [24]. Although the GPGPU community [11] provides general guidelines for optimizing applications using CUDA, *clearly* understanding various features of the underlying architecture and the associated performance bottlenecks in their applications is still remaining homework for programmers. Therefore, programmers might need to vary all the combinations to find the best performing configurations [24].

To provide insight into performance bottlenecks in massively parallel architectures, especially GPU architectures, we propose a simple analytical model. The model can be used statically without executing an application. The basic intuition of our analytical model is that estimating the cost of memory operations is the key component of estimating the performance of parallel GPU applications. The execution time of an application is dominated by the latency of memory instructions, but the latency of each memory operation can be hidden by executing multiple memory requests concurrently. By using the number of concurrently running threads and the memory bandwidth consumption, we estimate how many memory requests can be executed concurrently, which we call *memory warp*¹ *parallelism* (*MWP*). We also introduce *computation warp parallelism* (*CWP*). *CWP* represents how much computation can be done by other warps while one warp is waiting for memory values. *CWP* is similar to a metric, arithmetic intensity²[23] in the GPGPU community. Using both *MWP* and *CWP*, we estimate effective costs of memory requests, thereby estimating the overall execution time of a program.

We evaluate our analytical model based on the CUDA [20, 22]

¹A warp is a batch of threads that are internally executed together by the hardware. Section 2 describes a warp.

²Arithmetic intensity is defined as math operations per memory operation.

programming language, which is C with extensions for parallel threads. We compare the results of our analytical model with the actual execution time on several GPUs. Our results show that the geometric mean of absolute error of our model on micro-benchmarks is 5.4% and on the Merge benchmarks [17]³ is 13.3%

The contributions of our work are as follows:

1. To the best of our knowledge, we propose the first analytical model for the GPU architecture. This can be easily extended to other multithreaded architectures as well.
2. We propose two new metrics, MWP and CWP, to represent the degree of warp level parallelism that provide key insights identifying performance bottlenecks.

2. BACKGROUND AND MOTIVATION

We provide a brief background on the GPU architecture and programming model that we modeled. Our analytical model is based on the CUDA programming model and the NVIDIA Tesla architecture [3, 8, 20] used in the GeForce 8-series GPUs.

2.1 Background on the CUDA Programming Model

The CUDA programming model is similar in style to a single-program multiple-data (SPMD) software model. The GPU is treated as a coprocessor that executes data-parallel kernel functions.

CUDA provides three key abstractions, a hierarchy of thread groups, shared memories, and barrier synchronization. Threads have a three level hierarchy. A grid is a set of thread blocks that execute a kernel function. Each grid consists of blocks of threads. Each block is composed of hundreds of threads. Threads within one block can share data using shared memory and can be synchronized at a barrier. All threads within a block are executed concurrently on a multithreaded architecture.

The programmer specifies the number of threads per block, and the number of blocks per grid. A thread in the CUDA programming language is much lighter weight than a thread in traditional operating systems. A thread in CUDA typically processes one data element at a time. The CUDA programming model has two shared read-write memory spaces, the shared memory space and the global memory space. The shared memory is local to a block and the global memory space is accessible by all blocks. CUDA also provides two read-only memory spaces, the constant space and the texture space, which reside in external DRAM, and are accessed via read-only caches.

2.2 Background on the GPU Architecture

Figure 1 shows an overview of the GPU architecture. The GPU architecture consists of a scalable number of *streaming multiprocessors* (SMs), each containing eight *streaming processor* (SP) cores, two special function units (SFUs), a multithreaded instruction fetch and issue unit, a read-only constant cache, and a 16KB read/write shared memory [8].

The SM executes a batch of 32 threads together called a *warp*. Executing a warp instruction applies the instruction to 32 threads, similar to executing a SIMD instruction like an SSE instruction [14] in X86. However, unlike SIMD instructions, the concept of warp is not exposed to the programmers, rather programmers write a program for one thread, and then specify the number of parallel threads in a block, and the number of blocks in a kernel grid. The Tesla architecture forms a warp using a batch of 32 threads [13, 9] and in the rest of the paper we also use a warp as a batch of 32 threads.

³The Merge benchmarks consist of several media processing applications.

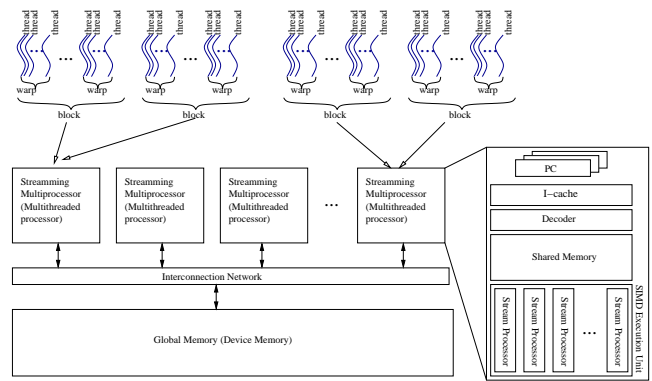


Figure 1: An overview of the GPU architecture

All the threads in one block are executed on one SM together. One SM can also have multiple concurrently running blocks. The number of blocks that are running on one SM is determined by the resource requirements of each block such as the number of registers and shared memory usage. The blocks that are running on one SM at a given time are called *active blocks* in this paper. Since one block typically has several warps (the number of warps is the same as the number of threads in a block divided by 32), the total number of active warps per SM is equal to the number of warps per block times the number of active blocks.

The shared memory is implemented within each SM multiprocessor as an SRAM and the global memory is part of the offchip DRAM. The shared memory has very low access latency (almost the same as that of register) and high bandwidth. However, since a warp of 32 threads access the shared memory together, when there is a bank conflict within a warp, accessing the shared memory takes multiple cycles.

2.3 Coalesced and Uncoalesced Memory Accesses

The SM processor executes one warp at one time, and schedules warps in a time-sharing fashion. The processor has enough functional units and register read/write ports to execute 32 threads (i.e. one warp) together. Since an SM has only 8 functional units, executing 32 threads takes 4 SM processor cycles for computation instructions.⁴

When the SM processor executes a memory instruction, it generates memory requests and switches to another warp until all the memory values in the warp are ready. Ideally, all the memory accesses within a warp can be combined into one memory transaction. Unfortunately, that depends on the memory access pattern within a warp. If the memory addresses are sequential, all of the memory requests within a warp can be coalesced into a single memory transaction. Otherwise, each memory address will generate a different transaction. Figure 2 illustrates two cases. The CUDA manual [22] provides detailed algorithms to identify types of coalesced/uncoalesced memory accesses. If memory requests in a warp are uncoalesced, the warp cannot be executed until all memory transactions from the same warp are serviced, which takes significantly longer than waiting for only one memory request (coalesced case).

⁴In this paper, a computation instruction means a non-memory instruction.




Figure 2: Memory requests from a single warp. (a) coalesced memory access (b) uncoalesced memory access

2.4 Motivating Example

To motivate the importance of a static performance analysis on the GPU architecture, we show an example of performance differences from three different versions of the same algorithm in Figure 3. The SVM benchmark is a kernel extracted from a face classification algorithm [28]. The performance of applications is measured on NVIDIA QuadroFX5600 [4]. There are three different optimized versions of the same SVM algorithm: *Naive*, *Constant*, and *Constant+Optimized*. *Naive* uses only the global memory, *Constant* uses the cached read-only constant memory⁵, and *Constant+Optimized* also optimizes memory accesses⁶ on top of using the constant memory. Figure 3 shows the execution time when the number of threads per block is varied. In this example, the number of blocks is fixed so the number of threads per block determines the total number of threads in a program. The performance improvement of *Constant+Optimized* and that of *Constant* over the *Naive* implementation are 24.36x and 1.79x respectively. Even though the performance of each version might be affected by the number of threads, once the number of threads exceeds 64, the performance does not vary significantly.

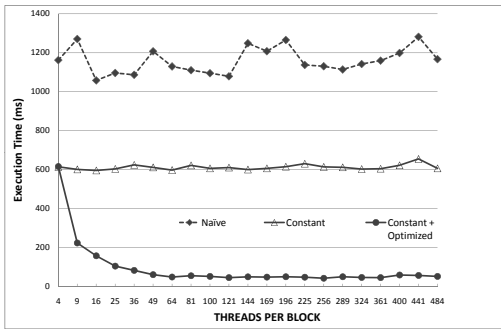


Figure 3: Optimization impacts on SVM

Figure 4 shows SM processor occupancy [22] for the three cases. The SM processor occupancy indicates the resource utilization, which has been widely used to optimize GPU computing applications. It is calculated based on the resource requirements for a given program. Typically, high occupancy (the max value is 1) is better for performance since many actively running threads would more likely hide the DRAM memory access latency. However, SM processor occupancy does not *sufficiently* estimate the performance

⁵The benefits of using the constant memory are (1) it has an on-chip cache per SM and (2) using the constant memory can reduce register usage, which might increase the number of running blocks in one SM.

⁶The programmer optimized the code to have coalesced memory accesses instead of uncoalesced memory accesses.

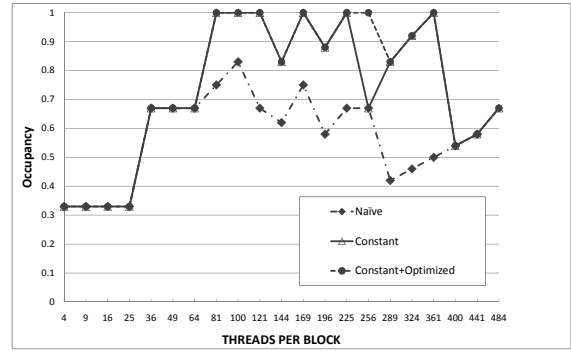


Figure 4: Occupancy values of SVM

improvement as shown in Figure 4. First, when the number of threads per block is less than 64, all three cases show the same occupancy values even though the performances of 3 cases are different. Second, even though SM processor occupancy is improved, for some cases, there is no performance improvement. For example, the performance of *Constant* is not improved at all even though the SM processor occupancy is increased from 0.35 to 1. Hence, we need other metrics to differentiate the three cases and to understand what the critical component of performance is.

3. ANALYTICAL MODEL

3.1 Introduction to MWP and CWP

The GPU architecture is a multithreaded architecture. Each SM can execute multiple warps in a time-sharing fashion while one or more warps are waiting for memory values. As a result, the execution cost of warps that are executed concurrently can be hidden. The key component of our analytical model is finding out how many memory requests can be serviced and how many warps can be executed together while one warp is waiting for memory values.

To represent the degree of warp parallelism, we introduce two metrics, *MWP* (*Memory Warp Parallelism*) and *CWP* (*Computation Warp Parallelism*). *MWP* represents the maximum number of warps per SM that can access the memory simultaneously during the time period from right after the SM processor executes a memory instruction from one warp (therefore, memory requests are just sent to the memory system) until all the memory requests from the same warp are serviced (therefore, the processor can execute the next instruction from that warp). The warp that is waiting for memory values is called a *memory warp* in this paper. The time period from right after one warp sent memory requests until all the memory requests from the same warp are serviced is called one memory warp waiting period. *CWP* represents the number of warps that the SM processor can execute during one memory warp waiting period plus *one*. A value one is added to include the warp itself that is waiting for memory values. (This means that *CWP* is always greater than or equal to 1.)

MWP is related to how much memory parallelism in the system. *MWP* is determined by the memory bandwidth, memory bank parallelism and the number of running warps per SM. *MWP* plays a very important role in our analytical model. When *MWP* is higher than 1, the cost of memory access cycles from (*MWP*-1) number of warps is all hidden, since they are all accessing the memory system together. The detailed algorithm of calculating *MWP* will be described in Section 3.3.1.

CWP is related to the program characteristics. It is similar to

an arithmetic intensity, but unlike arithmetic intensity, higher CWP means less computation per memory access. CWP also considers timing information but arithmetic intensity does not consider timing information. CWP is mainly used to decide whether the total execution time is dominated by computation cost or memory access cost. When CWP is greater than MWP, the execution cost is dominated by memory access cost. However, when MWP is greater than CWP, the execution cost is dominated by computation cost. How to calculate CWP will be described in Section 3.3.2.

3.2 The Cost of Executing Multiple Warps in the GPU architecture

To explain how executing multiple warps in each SM affects the total execution time, we will illustrate several scenarios in Figures 5, 6, 7 and 8. A computation period indicates the period when instructions from one warp are executed on the SM processor. A memory waiting period indicates the period when memory requests are being serviced. The numbers inside the computation period boxes and memory waiting period boxes in Figures 5, 6, 7 and 8 indicate a warp identification number.

3.2.1 CWP is Greater than MWP




Figure 5: Total execution time when CWP is greater than MWP: (a) 8 warps (b) 4 warps

For Case 1 in Figure 5a, we assume that all the computation and memory waiting periods are from different warps. The system can service two memory warps simultaneously. Since one computation period is roughly one third of one memory waiting warp period. (i.e., MWP is 2 and CWP is 4 for this case.) As a result, the 6 computation periods are completely overlapped with other memory waiting periods. Hence, only 2 computations and 4 memory waiting periods contribute to the total execution cycles.

For Case 2 in Figure 5b, there are four warps and each warp has two computation periods and two memory waiting periods. The second computation period can start only after the first memory waiting period of the same warp is finished. MWP and CWP are the same as Case 1. First, the processor executes four of the first computation periods from each warp one by one. By the time the processor finishes the first computation periods from all warps, two memory waiting periods are already serviced. So the processor can execute the second computation periods for these two warps. After that, there are no ready warps. The first memory waiting periods for the remaining two warps are still not finished yet. As soon as these two memory requests are serviced, the processor starts to execute the second computation periods for the other warps. Surprisingly, even though there are some idle cycles between computation periods, the total execution cycles are the same as Case 1. When CWP is higher than MWP, there are enough warps that are waiting for the memory values, so the cost of computation periods can be almost always hidden by memory access periods.

For both cases, the total execution cycles are only the sum of 2 computation periods and 4 memory waiting periods. Using MWP, the total execution cycles can be calculated using the below two equations. We divide $Comp_cycles$ by $\#Mem_insts$ to get the number of cycles in one computation period.

$$Exec_cycles = Mem_cycles \times \frac{N}{MWP} + Comp_p \times MWP \quad (1)$$

$$Comp_p = Comp_cycles / \#Mem_insts \quad (2)$$

Mem_cycles : Memory waiting cycles per each warp (see Equation (18))
 $Comp_cycles$: Computation cycles per each warp (see Equation (19))
 $Comp_p$: execution cycles of one computation period
 $\#Mem_insts$: Number of memory instructions
 N : Number of active running warps per SM

3.2.2 MWP is Greater than CWP

In general, CWP is greater than MWP. However, for some cases, MWP is greater than CWP. Let's say that the system can service 8 memory warps concurrently. Again CWP is still 4 in this scenario. In this case, as soon as the first computation period finishes, the processor can send memory requests. Hence, a memory waiting period of a warp always immediately follows the previous computation period. If all warps are independent, the processor continuously executes another warp. Case 3 in Figure 6a shows the timing information. In this case, the memory waiting periods are all overlapped with other warps except the last warp. The total execution cycles are the sum of 8 computation periods and only one memory waiting period.




Figure 6: Total execution time when MWP is greater than CWP: (a) 8 warps (b) 4 warps

Even if not all warps are independent, when CWP is higher than MWP, many of memory waiting periods are overlapped. Case 4 in Figure 6b shows an example. Each warp has two computation periods and two memory waiting periods. Since the computation time is dominant, the total execution cycles are again the sum of 8 computation periods and only one memory waiting period.

Using MWP and CWP, the total execution cycles can be calculated using the following equation:

$$Exec_cycles = Mem_p + Comp_cycles \times N \quad (3)$$

Mem_p : One memory waiting period (= Mem_L in Equation (12))

Case 5 in Figure 7 shows an extreme case. In this case, not even one computation period can be finished while one memory waiting period is completed. Hence, CWP is less than 2. Note that CWP is always greater 1. Even if MWP is 8, the application cannot take advantage of any memory warp parallelism. Hence, the total execution cycles are 8 computation periods plus one memory waiting period. Note that even this extreme case, the total execution cycles of Case 5 are the same as that of Case 4. Case 5 happens when $Comp_cycles$ are longer than Mem_cycles .




Figure 7: Total execution time when computation cycles are longer than memory waiting cycles. (8 warps)

3.2.3 Not Enough Warps Running

The previous two sections described situations when there are enough number of warps running on one SM. Unfortunately, if an application does not have enough number of warps, the system cannot take advantage of all available warp parallelism. MWP and CWP cannot be greater than the number of active warps on one SM.




Figure 8: Total execution time when MWP is equal to N: (a) 1 warp (b) 2 warps

Case 6 in Figure 8a shows when only one warp is running. All the executions are serialized. Hence, the total execution cycles are the sum of the computation and memory waiting periods. Both CWP and MWP are 1 in this case. Case 7 in Figure 8b shows there are two running warps. Let's assume that MWP is two. Even if one computation period is less than the half of one memory waiting period, because there are only two warps, CWP is still two. Because of MWP, the total execution time is roughly the half of the sum of all the computation periods and memory waiting periods.

Using MWP, the total execution cycles of the above two cases can be calculated using the following equation:

$$\begin{aligned}
 Exec_cycles &= Mem_cycles \times N/MWP + Comp_cycles \times \\
 & N/MWP + Comp_p(MWP - 1) \\
 & = Mem_cycles + Comp_cycles + Comp_p(MWP - 1)
 \end{aligned} \quad (4)$$

Note that for both cases, MWP and CWP are equal to N, the number of active warps per SM.

3.3 Calculating the Degree of Warp Parallelism

3.3.1 Memory Warp Parallelism (MWP)

MWP is slightly different from MLP [10]. MLP represents how many memory requests can be serviced together. MWP represents the maximum number of warps in each SM that can access the memory simultaneously during one memory warp waiting period. The main difference between MLP and MWP is that MWP is counting all memory requests from a warp as one unit, while MLP counts all individual memory requests separately. As we discussed in Section 2.3, one memory instruction in a warp can generate multiple memory transactions. This difference is very important because a warp cannot be executed until all values are ready.

MWP is tightly coupled with the DRAM memory system. In our analytical model, we model the DRAM system as a simple queue and each SM has its own queue. Each active SM consumes an equal amount of memory bandwidth. Figure 9 shows the memory model and a timeline of memory warps.

The latency of each memory warp is at least Mem_L cycles. $Departure_delay$ is the minimum departure distance between two consecutive memory warps. Mem_L is a round trip time to the DRAM, which includes the DRAM access time and the address and data transfer time.




Figure 9: Memory system model: (a) memory model (b) timeline of memory warps

MWP represents the number of memory warps per SM that can be handled during Mem_L cycles. MWP cannot be greater than the number of warps per SM that reach the peak memory bandwidth (MWP_peak_BW) of the system as shown in Equation (5). If fewer SMs are executing warps, each SM can consume more bandwidth than when all SMs are executing warps. Equation (6) represents MWP_peak_BW . If an application does not reach the peak bandwidth, MWP is a function of Mem_L and $departure_delay$. $MWP_Without_BW$ is calculated using Equations (10) – (17). MWP cannot be also greater than the number of active warps as shown in Equation (5). If the number of active warps is less than $MWP_Without_BW_full$, the processor does not have enough number of warps to utilize memory level parallelism.

$$MWP = MIN(MWP_Without_BW, MWP_peak_BW, N) \quad (5)$$

$$MWP_peak_BW = \frac{Mem_Bandwidth}{BW_per_warp \times \#ActiveSM} \quad (6)$$

$$BW_per_warp = \frac{Freq \times Load_bytes_per_warp}{Mem_L} \quad (7)$$




Figure 10: Illustrations of departure delays for uncoalesced and coalesced memory warps: (a) uncoalesced case (b) coalesced case

The latency of memory warps is dependent on memory access pattern (coalesced/uncoalesced) as shown in Figure 10. For uncoalesced memory warps, since one warp requests multiple number of transactions ($\#Uncoal_per_mw$), Mem_L includes departure delays for all $\#Uncoal_per_mw$ number of transactions. $Departure_delay$ also includes $\#Uncoal_per_mw$ number of $Departure_del_uncoal$ cycles. Mem_LD is a round-trip latency to the DRAM for each memory transaction. In this model, Mem_LD for uncoalesced and coalesced are considered as the same, even though a coalesced

memory request might take a few more cycles because of large data size.

In an application, some memory requests would be coalesced and some would be not. Since multiple warps are running concurrently, the analytical model simply uses the weighted average of memory latency of coalesced and uncoalesced latency for the memory latency (Mem_L). A weight is determined by the number of coalesced and uncoalesced memory requests as shown in Equations (13) and (14). MWP is calculated using Equations (10) – (17). The parameters used in these equations are summarized in Table 1. Mem_LD , $Departure_del_coal$ and $Departure_del_uncoal$ are measured with micro-benchmarks as we will show in Section 5.1.

3.3.2 Computation Warp Parallelism (CWP)

Once we calculate the memory latency for each warp, calculating CWP is straightforward. CWP_full is when there are enough number of warps. When CWP_full is greater than N (the number of active warps in one SM) CWP is N, otherwise, CWP_full becomes CWP .

$$CWP_full = \frac{Mem_cycles + Comp_cycles}{Comp_cycles} \quad (8)$$

$$CWP = MIN(CWP_full, N) \quad (9)$$

3.4 Putting It All Together in CUDA

So far, we have explained our analytical model without strongly being coupled with the CUDA programming model to simplify the model. In this section, we extend the analytical model to consider the CUDA programming model.

3.4.1 Number of Warps per SM

The GPU SM multithreading architecture executes 100s of threads concurrently. Nonetheless, not all threads in an application can be executed at the same time. The processor fetches a few blocks at one time. The processor fetches additional blocks as soon as one block retires. $\#Rep$ represents how many times a single SM executes multiple active number of blocks. For example, when there are 40 blocks in an application and 4 SMs. If each SM can execute 2 blocks concurrently, then $\#Rep$ is 5. Hence, the total number of warps per SM is $\#Active_warps_per_SM$ (N) times $\#Rep$. N is determined by machine resources.

3.4.2 Total Execution Cycles

Depending on MWP and CWP values, total execution cycles for an entire application ($Exec_cycles_app$) are calculated using Equations (22),(23), and (24). Mem_L is calculated in Equation (12). Execution cycles that consider synchronization effects will be described in Section 3.4.6.

3.4.3 Dynamic Number of Instructions

Total execution cycles are calculated using the number of dynamic instructions. The compiler generates intermediate assembler-level instruction, the NVIDIA PTX instruction set [22]. PTX instructions translate nearly one to one with native binary microinstructions later.⁷ We use the number of PTX instructions for the dynamic number of instructions.

The total number of instructions is proportional to the number of data elements. Programmers must decide the number of threads and blocks for each input data. The number of total instructions per thread is related to how many data elements are computed in one thread, programmers must know this information. If we know

⁷Since some PTX instructions expand to multiple binary instructions, using PTX instruction count could be one of the error sources in the analytical model.

the number of elements per thread, counting the number of total instructions per thread is simply counting the number of computation instructions and the number of memory instructions per element. The detailed algorithm to count the number of instructions from PTX code is provided in an extended version of this paper [12].

3.4.4 Cycles Per Instruction (CPI)

Cycles per Instruction (CPI) is commonly used to represent the cost of each instruction. Using total execution cycles, we can calculate Cycles Per Instruction using Equation (25). Note that, CPI is the cost when an instruction is executed by all threads in one warp.

$$CPI = \frac{Exec_cycles_app}{\#Total_insts \times \frac{\#Threads_per_block}{\#Threads_per_warp} \times \frac{\#Blocks}{\#Active_SMs}} \quad (25)$$

3.4.5 Coalesced/Uncoalesced Memory Accesses

As Equations (15) and (12) suggest, the latency of memory instruction is heavily dependent on memory access type. Whether memory requests inside a warp can be coalesced or not is dependent on the microarchitecture of the memory system and memory access pattern in a warp. The GPUs that we evaluated have two coalesced/uncoalesced policies, specified by the Compute capability version. The CUDA manual [22] describes when memory requests in a warp can be coalesced or not in more detail. Earlier compute capability versions have two differences compared with the later version(1.3): (1) stricter rules are applied to be coalesced, (2) when memory requests are uncoalesced, one warp generates 32 memory transactions. In the latest version (1.3), the rules are more relaxed and all memory requests are coalesced into as few memory transactions as possible.⁸

The detailed algorithms to detect coalesced/uncoalesced memory accesses and to count the number of memory transactions per each warp at static time are provided in an extended version of this paper [12].

3.4.6 Synchronization Effects




Figure 11: Additional delay effects of thread synchronization: (a) no synchronization (b) thread synchronization after each memory access period

The CUDA programming model supports thread synchronization through the `__syncthreads()` function. Typically, all the threads are executed asynchronously whenever all the source operands in a warp are ready. However, if there is a barrier, the processor cannot execute the instructions after the barrier until all the threads

⁸In the CUDA manual, compute capability 1.3 says all requests are coalesced because all memory requests within each warp are always combined into as few transactions as possible. However, in our analytical model, we use the coalesced memory access model only if all memory requests are combined into one memory transaction.

$$Mem_L_Uncoal = Mem_LD + (\#Uncoal_per_mw - 1) \times Departure_del_uncoal \quad (10)$$

$$Mem_L_Coal = Mem_LD \quad (11)$$

$$Mem_L = Mem_L_Uncoal \times Weight_uncoal + Mem_L_Coal \times Weight_coal \quad (12)$$

$$Weight_uncoal = \frac{\#Uncoal_Mem_insts}{(\#Uncoal_Mem_insts + \#Coal_Mem_insts)} \quad (13)$$

$$Weight_coal = \frac{\#Coal_Mem_insts}{(\#Coal_Mem_insts + \#Uncoal_Mem_insts)} \quad (14)$$

$$Departure_delay = (Departure_del_uncoal \times \#Uncoal_per_mw) \times Weight_uncoal + Departure_del_coal \times Weight_coal \quad (15)$$

$$MWP_Without_BW_full = Mem_L / Departure_delay \quad (16)$$

$$MWP_Without_BW = MIN(MWP_Without_BW_full, \#Active_warps_per_SM) \quad (17)$$

$$Mem_cycles = Mem_L_Uncoal \times \#Uncoal_Mem_insts + Mem_L_Coal \times \#Coal_Mem_insts \quad (18)$$

$$Comp_cycles = \#Issue_cycles \times (\#total_insts) \quad (19)$$

$$N = \#Active_warps_per_SM \quad (20)$$

$$\#Rep = \frac{\#Blocks}{\#Active_blocks_per_SM \times \#Active_SMs} \quad (21)$$

If (MWP is N warps per SM) and (CWP is N warps per SM)

$$Exec_cycles_app = (Mem_cycles + Comp_cycles + \frac{Comp_cycles}{\#Mem_insts} \times (MWP - 1)) \times \#Rep \quad (22)$$

If (CWP >= MWP) or (Comp_cycles > Mem_cycles)

$$Exec_cycles_app = (Mem_cycles \times \frac{N}{MWP} + \frac{Comp_cycles}{\#Mem_insts} \times (MWP - 1)) \times \#Rep \quad (23)$$

If (MWP > CWP)

$$Exec_cycles_app = (Mem_L + Comp_cycles \times N) \times \#Rep \quad (24)$$

*All the parameters are summarized in Table 1.

reach the barrier. Hence, there will be additional delays due to a thread synchronization. Figure 11 illustrates the additional delay effect. Surprisingly, the additional delay is less than one waiting period. Actually, the additional delay per synchronization instruction in one block is the multiple of *Departure_delay* and (MWP-1). Since the synchronization occurs as a block granularity, we need to account for the number of blocks in each SM. The final execution cycles of an application with synchronization delay effect can be calculated by Equation (27).

$$Synch_cost = Departure_delay \times (MWP - 1) \times \#synch_insts \times \#Active_blocks_per_SM \times \#Rep \quad (26)$$

$$Exec_cycles_with_synch = Exec_cycles_app + Synch_cost \quad (27)$$

3.5 Limitations of the Analytical Model

Our analytical model does not consider the cost of cache misses such as I-cache, texture cache, or constant cache. The cost of cache misses is negligible due to almost 100% cache hit ratio.

The current G80 architecture does not have a hardware cache for the global memory. Typical stream applications running on the GPUs do not have strong temporal locality. However, if an application has temporal locality and a future architecture provides a hardware cache, the model should include a model of cache. In future work, we will include cache models.

The cost of executing branch instructions is not modeled in detail. Double counting the number of instructions in both paths will probably provide an upper bound of execution cycles.

3.6 Code Example

To provide a concrete example, we apply the analytical model for a tiled matrix multiplication example in Figure 12 to a system that has 80GB/s memory bandwidth, 1GHz frequency and 16 SM processors. Let's assume that the programmer specified 128 threads

```

1: MatrixMulKernel<<<80, 128>>> (M, N, P);
2: ...
3: MatrixMulKernel(Matrix M, Matrix N, Matrix P)
4: {
5:     // init code ...
6:
7:     for (int a=starta, b=startb, iter=0; a<=enda;
8:          a+=stepa, b+=stepb, iter++)
9:     {
10:         __shared__ float Msub[BLOCKSIZE][BLOCKSIZE];
11:         __shared__ float Nsub[BLOCKSIZE][BLOCKSIZE];
12:
13:         Msub[ty][tx] = M.elements[a + wM * ty + tx];
14:         Nsub[ty][tx] = N.elements[b + wN * ty + tx];
15:
16:         __syncthreads();
17:
18:         for (int k=0; k < BLOCKSIZE; ++k)
19:             subsum += Msub[ty][k] * Nsub[k][tx];
20:
21:         __syncthreads();
22:     }
23:
24:     int index = wN * BLOCKSIZE * by + BLOCKSIZE
25:     P.elements[index + wN * ty + tx] = subsum;
26: }

```

Figure 12: CUDA code of tiled matrix multiplication

per block (4 warps per block), and 80 blocks for execution. And 5 blocks are actively assigned to each SM (*Active_blocks_per_SM*) instead of 8 maximum blocks⁹ due to high resource usage.

We assume that the inner loop is iterated only once and the outer loop is iterated 3 times to simplify the example. Hence, *#Comp_insts* is 27, which is 9 computation (Figure 13 lines 5, 7, 8, 9, 10, 11, 13,

⁹Each SM can have up to 8 blocks at a given time.

Table 1: Summary of Model Parameters

| | Model Parameter | Definition | Obtained |
|----|--------------------------|--|--|
| 1 | #Threads_per_warp | Number of threads per warp | 32 [22] |
| 2 | Issue_cycles | Number of cycles to execute one instruction | 4 cycles [13] |
| 3 | Freq | Clock frequency of the SM processor | Table 3 |
| 4 | Mem_Bandwidth | Bandwidth between the DRAM and GPU cores | Table 3 |
| 5 | Mem_LD | DRAM access latency (machine configuration) | Table 6 |
| 6 | Departure_del_uncoal | Delay between two uncoalesced memory transactions | Table 6 |
| 7 | Departure_del_coal | Delay between two coalesced memory transactions | Table 6 |
| 8 | #Threads_per_block | Number of threads per block | Programmer specifies inside a program |
| 9 | #Blocks | Total number of blocks in a program | Programmer specifies inside a program |
| 10 | #Active_SMs | Number of active SMs | Calculated based on machine resources |
| 11 | #Active_blocks_per_SM | Number of concurrently running blocks on one SM | Calculated based on machine resources [22] |
| 12 | #Active_warps_per_SM (N) | Number of concurrently running warps on one SM | Active_blocks_per_SM x Number of warps per block |
| 13 | #Total_insts | (#Comp_insts + #Mem_insts) | |
| 14 | #Comp_insts | Total dynamic number of computation instructions in one thread | Source code analysis |
| 15 | #Mem_insts | Total dynamic number of memory instructions in one thread | Source code analysis |
| 16 | #Uncoal_Mem_insts | Number of uncoalesced memory type instructions in one thread | Source code analysis |
| 17 | #Coal_Mem_insts | Number of coalesced memory type instructions in one thread | Source code analysis |
| 18 | #Synch_insts | Total dynamic number of synchronization instructions in one thread | Source code analysis |
| 19 | #Coal_per_mw | Number of memory transactions per warp (coalesced access) | 1 |
| 20 | #Uncoal_per_mw | Number of memory transactions per warp (uncoalesced access) | Source code analysis[12](Table 3) |
| 21 | Load_bytes_per_warp | Number of bytes for each warp | Data size (typically 4B) x #Threads_per_warp |

```

1: ... // Init Code
2:
3: $OUTERLOOP:
4: ld.global.f32 %f2, [%rd23+0]; //
5: st.shared.f32 [%rd14+0], %f2; //
6: ld.global.f32 %f3, [%rd19+0]; //
7: st.shared.f32 [%rd15+0], %f3; //
8: bar.sync 0; // Synchronization
9: ld.shared.f32 %f4, [%rd8+0]; // Innerloop unrolling
10: ld.shared.f32 %f5, [%rd6+0]; //
11: mad.f32 %f1, %f4, %f5, %f1; //
12: // the code of unrolled loop is omitted
13: bar.sync 0; // synchronization
14: setp.le.s32 %p2, %r21, %r24; //
15: @%p2 bra $OUTERLOOP; // Branch
16: ... // Index calculation
17: st.global.f32 [%rd27+0], %f1; // Store in P.elements

```

Figure 13: PTX code of tiled matrix multiplication

14, and 15) instructions times 3. Note that `ld.shared` instructions in Figure 13 lines 9 and 10 are also counted as a computation instruction since the latency of accessing the shared memory is almost as fast as that of the register file. Lines 13 and 14 in Figure 12 show global memory accesses in the CUDA code. Memory indexes $(a+wM*ty+tx)$ and $(b+wN*ty+tx)$ determine memory access coalescing within a warp. Since a and b are more likely not a multiple of 32, we treat that all the global loads are uncoalesced [12]. So $\#Uncoal_Mem_insts$ is 6, and $\#Coal_Mem_insts$ is 0.

Table 2 shows the necessary model parameters and intermediate calculation processes to calculate the total execution cycles of the program. Since CWP is greater than MWP, we use Equation (23) to calculate $Exec_cycles_app$. Note that in this example, the execution cost of synchronization instructions is a significant part of the total execution cost. This is because we simplified the example. In most real applications, the number of dynamic synchronization instructions is much less than other instructions, so the synchronization cost is not that significant.

4. EXPERIMENTAL METHODOLOGY

4.1 The GPU Characteristics

Table 3 shows the list of GPUs used in this study. GTX280 supports 64-bit floating point operations and also has a later computing version (1.3) that improves uncoalesced memory accesses. To measure the GPU kernel execution time, `cudaEventRecord` API that uses GPU Shader clock cycles is used. All the measured execution time is the average of 10 runs.

4.2 Micro-benchmarks

All the benchmarks are compiled with NVCC [22]. To test the analytical model and also to find memory model parameters, we design a set of micro-benchmarks that simply repeat a loop for 1000 times. We vary the number of load instructions and computation instructions per loop. Each micro-benchmark has two memory access patterns: coalesced and uncoalesced memory accesses.

4.3 Merge Benchmarks

To test how our analytical model can predict typical GPGPU applications, we use 6 different benchmarks that are mostly used in the Merge work [17]. Table 5 explains the description of each benchmark and summarizes the characteristics of each benchmark. The number of registers used per thread and shared memory usage per block are statically obtained by compiling the code with `-cubin` flag. The number of dynamic PTX instructions is calculated using program's input values [12]. The rest of the characteristics are statically determined and can be found in PTX code. Note that, since we estimate the number dynamic instructions just based on static information and an input size, the number counted is an approximated value. To simplify the evaluation, depending on the majority load type, we treat all memory access as either coalesced or uncoalesced for each benchmark. For the Mat. (tiled) benchmark, the number of memory instructions and computation instructions change with respect to the number of warps per block, which the programmers specify. This is because the number of inner loop iterations for each thread depends on blocksize (i.e., the tile size).

Table 5: Characteristics of the Merge Benchmarks (Arith. intensity means arithmetic intensity.)

| Benchmark | Description | Input size | Comp insts | Mem insts | Arith. intensity | Registers | Shared Mem |
|-------------------|---|---------------|--------------|--------------------------|------------------|-----------|------------|
| Sepia [17] | Filter for artificially aging images | 7000 x 7000 | 71 | 6 (uncoalesced) | 11.8 | 7 | 52B |
| Linear [17] | Image filter for computing the avg. of 9-pixels | 10000 x 10000 | 111 | 30 (uncoalesced) | 3.7 | 15 | 60B |
| SVM [17] | Kernel from a SVM-based algorithm | 736 x 992 | 10871 | 819 (coalesced) | 13.3 | 9 | 44B |
| Mat. (naive) | Naive version of matrix multiplication | 2000 x 2000 | 12043 | 4001 (uncoalesced) | 3 | 10 | 88B |
| Mat. (tiled) [22] | Tiled version of matrix multiplication | 2000 x 2000 | 9780 - 24580 | 201 - 1001 (uncoalesced) | 48.7 | 18 | 3960B |
| Blackscholes [22] | European option pricing | 9000000 | 137 | 7 (uncoalesced) | 19 | 11 | 36B |

Table 2: Applying the Model to Figure 12

| Model Parameter | Obtained | Value |
|-----------------------|-----------------------------------|---|
| Mem_LD | Machine conf. | 420 |
| Departure_del_uncoal | Machine conf. | 10 |
| #Threads_per_block | Figure 12 Line 1 | 128 |
| #Blocks | Figure 12 Line 1 | 80 |
| #Active_blocks_per_SM | Occupancy [22] | 5 |
| #Active_SMs | Occupancy [22] | 16 |
| #Active_warps_per_SM | $128/32(\text{Table 1}) \times 5$ | 20 |
| #Comp_insts | Figure 13 | 27 |
| #Uncoal_Mem_insts | Figure 12 Lines 13, 14 | 6 |
| #Coal_Mem_insts | Figure 12 Lines 13, 14 | 0 |
| #Synch_insts | Figure 12 Lines 16, 21 | $6 = 2 \times 3$ |
| #Coal_per_mw | see Sec. 3.4.5 | 1 |
| #Uncoal_per_mw | see Sec. 3.4.5 | 32 |
| Load_bytes_per_warp | Figure 13 Lines 4, 6 | $128B = 4B \times 32$ |
| Departure_delay | Equation (15) | $320 = 32 \times 10$ |
| Mem_L | Equations (10), (12) | $730 = 420 + (32 - 1) \times 10$ |
| MWP_without_BW_full | Equation (16) | $2.28 = 730/320$ |
| BW_per_warp | Equation (7) | $0.175 \text{GB/s} = \frac{1G \times 128B}{730}$ |
| MWP_peak_BW | Equation (6) | $28.57 = \frac{80 \text{GB/s}}{0.175 \text{GB} \times 16}$ |
| MWP | Equation (5) | $2.28 = \text{MIN}(2.28, 28.57, 20)$ |
| Comp_cycles | Equation (19) | $132 \text{ cycles} = 4 \times (27 + 6)$ |
| Mem_cycles | Equation (18) | $4380 = (730 \times 6)$ |
| CWP_full | Equation (8) | $34.18 = (4380 + 132)/132$ |
| CWP | Equation (9) | $20 = \text{MIN}(34.18, 20)$ |
| #Rep | Equation (21) | $1 = 80/(16 \times 5)$ |
| Exec_cycles_app | Equation (23) | $38450 = 4380 \times \frac{20}{2.28} + \frac{132}{6} \times (2.28 - 1)$ |
| Synch_cost | Equation (26) | $12288 = 320 \times (2.28 - 1) \times 6 \times 5$ |
| Final Time | Equation (27) | $50738 = 38450 + 12288$ |

5. RESULTS

5.1 Micro-benchmarks

The micro-benchmarks are used to measure the constant variables that are required to model the memory system. We vary three parameters (*Mem_LD*, *Departure_del_uncoal*, and *Departure_del_coal*) for each GPU to find the best fitting values. FX5600, 8800GTX and 8800GT use the same model parameters. Table 6 summarizes the results. *Departure_del_coal* is related to the memory access time to a single memory block. *Departure_del_uncoal* is longer than *Departure_del_coal*, due to the overhead of 32 small memory access requests. *Departure_del_uncoal* for GTX280 is much longer than that of FX5600. GTX280 coalesces 32 thread memory requests per warp into the minimum number of memory access requests, and the overhead per access request is higher, with fewer accesses.

Using the parameters in Table 6, we calculate CPI for the micro-benchmarks. Figure 14 shows the average CPI of the micro-benchmarks for both measured value and estimated value using the analytical model. The results show that the average geometric mean of the error is 5.4%. As we can predict, as the benchmark has more number

Table 3: The specifications of GPUs used in this study

| Model | 8800GTX | Quadro FX5600 | 8800GT | GTX280 |
|----------------------|-----------|---------------|-----------|------------|
| #SM | 16 | 16 | 14 | 30 |
| (SP) Processor Cores | 128 | 128 | 112 | 240 |
| Graphics Clock | 575 MHz | 600 MHz | 600 MHz | 602 MHz |
| Processor Clock | 1.35 GHz | 1.35GHz | 1.5 GHz | 1.3 GHz |
| Memory Size | 768 MB | 1.5 GB | 512 MB | 1 GB |
| Memory Bandwidth | 86.4 GB/s | 76.8 GB/s | 57.6 GB/s | 141.7 GB/s |
| Peak Gflop/s | 345.6 | 384 | 336 | 933 |
| Computing Version | 1.0 | 1.0 | 1.1 | 1.3 |
| #Uncoal_per_mw | 32 | 32 | 32 | [12] |
| #Coal_per_mw | 1 | 1 | 1 | 1 |

Table 4: The characteristics of micro-benchmarks

| # inst. per loop | Mb1 | Mb2 | Mb3 | Mb4 | Mb5 | Mb6 | Mb7 |
|------------------|---------|--------|---------|--------|--------|--------|--------|
| Memory | 0 | 1 | 1 | 2 | 2 | 4 | 6 |
| Comp. (FP) | 23 (20) | 17 (8) | 29 (20) | 27(12) | 35(20) | 47(20) | 59(20) |

of load instructions, the CPI increases. For the coalesced load cases (Mb1_C – Mb7_C), the cost of load instructions is almost hidden because of high MWP but for uncoalesced load cases (Mb1_UC – Mb7_UC), the cost of load instructions linearly increases as the number of load instructions increases.

5.2 Merge Benchmarks

Figure 15 and Figure 16 show the measured and estimated execution time of the Merge benchmarks on FX5600 and GTX280. The number of threads per block is varied from 4 to 512, (512 is the maximum value that one block can have in the evaluated CUDA programs.) Even though the number of threads is varied, the programs calculate the same amount data elements. In other words, if we increase the number of threads in a block, the total number of blocks is also reduced to process the same amount of data in one application. That is why the execution times are mostly the same. For the Mat.(tiled) benchmark, as we increase the number of threads the execution time reduces, because the number of active warps per SM increases.

Figure 17 shows the average of the measured and estimated CPIs across four GPUs in Figures 15 and 16 configurations. The average value of CWP and MWP per SM are also shown in Figures 18, and 19 respectively. 8800GT has the least amount of bandwidth

Table 6: Results of the Memory Model Parameters

| Model | FX5600 | GTX280 |
|----------------------|--------|--------|
| Mem_LD | 420 | 450 |
| Departure_del_uncoal | 10 | 40 |
| Departure_del_coal | 4 | 4 |

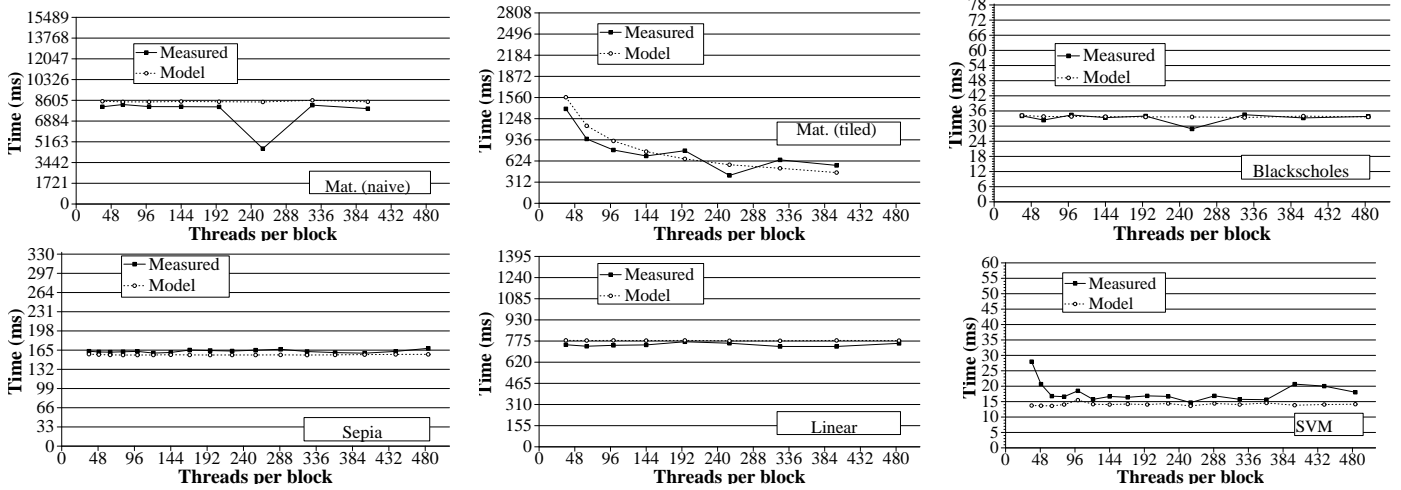


Figure 15: The total execution time of the Merge benchmarks on FX5600

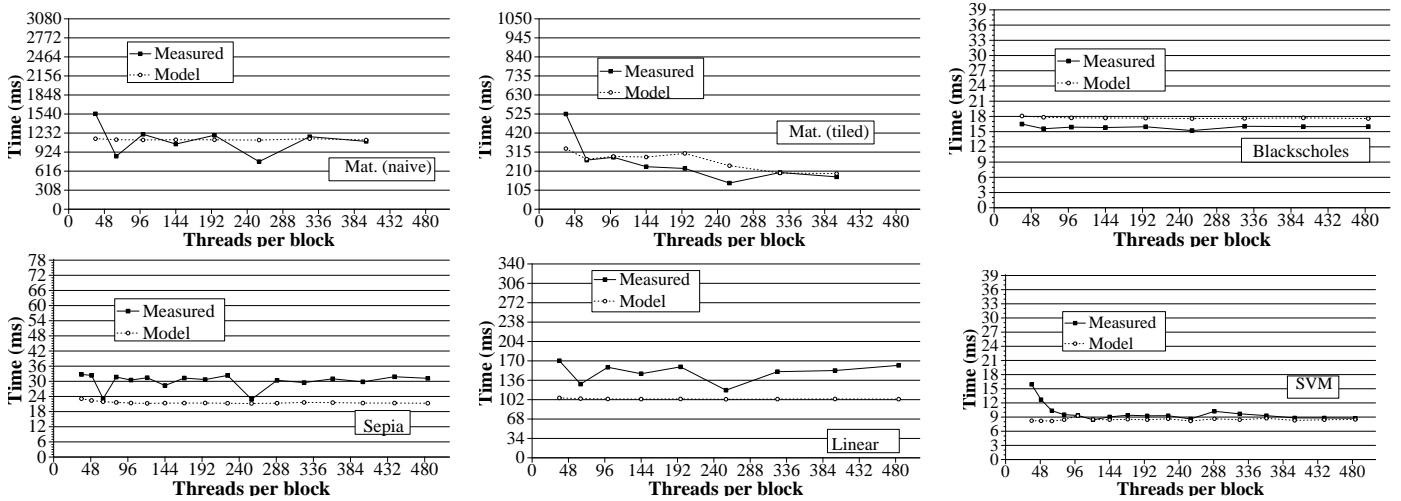


Figure 16: The total execution time of the Merge benchmarks on GTX280

compared to other GPUs, resulting in the highest CPI in contrast to GTX280. Generally, higher arithmetic intensity means lower CPI (lower CPI is higher performance). However, even though the Mat.(tiled) benchmark has the highest arithmetic intensity, SVM has the lowest CPI value. SVM has higher MWP and CWP than those of Mat.(tiled) as shown in Figures 18 and 19. SVM has the highest MWP and the lowest CPI because only SVM has fully coalesced memory accesses. MWP in GTX280 is higher than the rest of GPUs because even though most memory requests are not fully coalesced, they are still combined into as few requests as possible, which results in higher MWP. All other benchmarks are limited by *departure_delay*, which makes all other applications never reach the peak memory bandwidth.

Figure 20 shows the average occupancy of the Merge benchmarks. Except Mat.(tiled) and Linear, all other benchmarks have higher occupancy than 70%. The results show that occupancy is less correlated to the performance of applications.

The final geometric mean of the estimated CPI error on the Merge benchmarks in Figure 17 over all four different types of GPUs is 13.3%. Generally the error is higher for GTX 280 than others, be-

cause we have to estimate the number of memory requests that are generated by partially coalesced loads per warp in GTX280, unlike other GPUs which have the fixed value 32. On average, the model estimates the execution cycles of FX5600 better than others. This is because we set the machine parameters using FX5600.

There are several error sources in our model: (1) We used a very simple memory model and we assume that the characteristics of the memory behavior are similar across all the benchmarks. We found out that the outcome of the model is very sensitive to MWP values. (2) We assume that the DRAM memory scheduler schedules memory requests equally for all warps. (3) We do not consider the bank conflict latency in the shared memory. (4) All computation instructions have the same latency even though some special functional unit instructions have longer latency than others. (5) For some applications, the number of threads per block is not always a multiple of 32. (6) The SM retires warps as a block granularity. Even though there are free cycles, the SM cannot start to fetch new blocks, but the model assumes on average active warps.

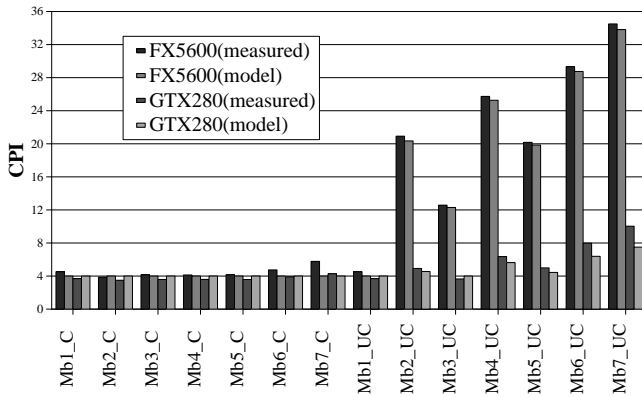


Figure 14: CPI on the micro-benchmarks

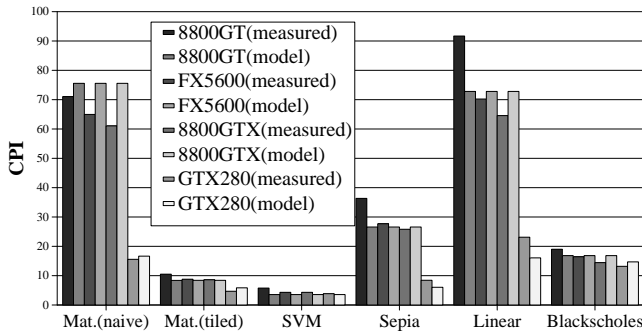


Figure 17: CPI on the Merge benchmarks

6. RELATED WORK

We discuss research related to our analytical model in the areas of performance analytical modeling, and GPU performance estimation. No previous work we are aware of proposed a way of accurately predicting GPU performance or multithreaded program performance at compile-time using only static time available information. Our cost estimation metrics provide a new way of estimating the performance impacts.

6.1 Analytical Modeling

There have been many existing analytical models proposed for superscalar processors [21, 19, 18]. Most work did not consider memory level parallelism or even cache misses. Karkhanis and Smith [15] proposed a first-order superscalar processor model to

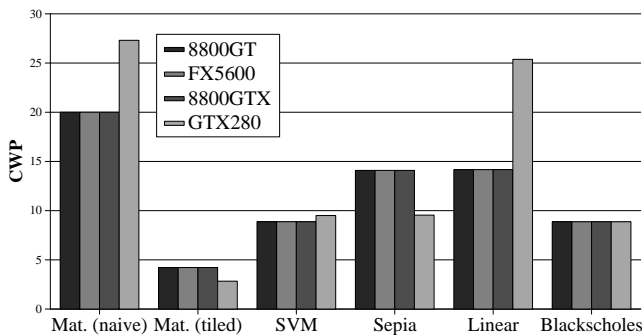


Figure 18: CWP per SM on the Merge benchmarks




Figure 19: MWP per SM on the Merge benchmarks

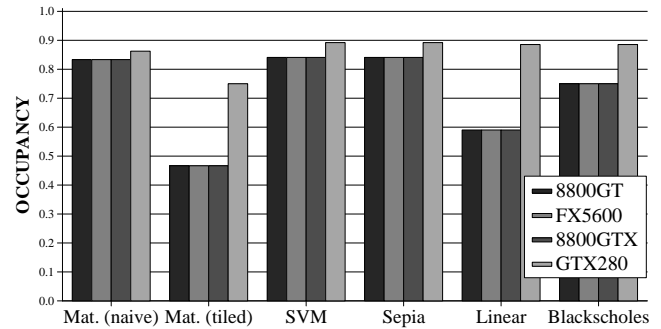


Figure 20: Occupancy on the Merge benchmarks

analyze the performance of processors. They modeled long latency cache misses and other major performance bottleneck events using a first-order model. They used different penalties for dependent loads. Recently, Chen and Aamodt [7] improved the first-order superscalar processor model by considering the cost of pending hits, data prefetching and MSHRs (Miss Status/Information Holding Registers). They showed that not modeling prefetching and MSHRs can increase errors significantly in the first-order processor model. However, they only showed memory instructions' CPI results comparing with the results of a cycle accurate simulator.

There is a rich body of work that predicts parallel program performance prediction using stochastic modeling or task graph analysis, which is beyond the scope of our work. Saavedra-Barrera and Culler [25] proposed a simple analytical model for multithreaded machines using stochastic modeling. Their model uses memory latency, switching overhead, the number of threads that can be interleaved and the interval between thread switches. Their work provided insights into the performance estimation on multithreaded architectures. However, they have not considered synchronization effects. Furthermore, the application characteristics are represented with statistical modeling, which cannot provide detailed performance estimation for each application. Their model also provided insights into a saturation point and an efficiency metric that could be useful for reducing the optimization spaces even though they did not discuss that benefit in their work.

Sorin et al. [27] developed an analytical model to calculate throughput of processors in the shared memory system. They developed a model to estimate processor stall times due to cache misses or resource constraints. They also discussed coalesced memory effects inside the MSHR. The majority of their analytical model is also based on statistical modeling.

6.2 GPU Performance Modeling

Our work is strongly related with other GPU optimization techniques. The GPGPU community provides insights into how to optimize GPGPU code to increase memory level parallelism and thread level parallelism [11]. However, all the heuristics are qualitatively discussed without using any analytical models. The most relevant metric is an occupancy metric that provides only general guidelines as we showed in our Section 2.4. Recently, Ryoo et al. [24] proposed two metrics to reduce optimization spaces for programmers by calculating utilization and efficiency of applications. However, their work focused on non-memory intensive workloads. We thoroughly analyzed both memory intensive and non-intensive workloads to estimate the performance of applications. Furthermore, their work just provided optimization spaces to reduce program tuning time. In contrast, we predict the actual program execution time. Bakhoda et al. [6] recently implemented a GPU simulator and analyzed the performance of CUDA applications using the simulation output.

7. CONCLUSIONS

This paper proposed and evaluated a memory parallelism aware analytical model to estimate execution cycles for the GPU architecture. The key idea of the analytical model is to find the maximum number of memory warps that can execute in parallel, a metric which we called MWP, to estimate the effective memory instruction cost. The model calculates the estimated CPI (cycles per instruction), which could provide a simple performance estimation metric for programmers and compilers to decide whether they should perform certain optimizations or not. Our evaluation shows that the geometric mean of absolute error of our analytical model on micro-benchmarks is 5.4% and on GPU computing applications is 13.3%. We believe that this analytical model can provide insights into how programmers should improve their applications, which will reduce the burden of parallel programmers.

Acknowledgments

Special thanks to John Nickolls for insightful and detailed comments in preparation of the final version of the paper. We thank the anonymous reviewers for their comments. We also thank Chi-keung Luk, Philip Wright, Guru Venkataramani, Gregory Diamos, and Eric Sprangle for their feedback on improving the paper. We gratefully acknowledge the support of Intel Corporation, Microsoft Research, and the equipment donations from NVIDIA.

8. REFERENCES

- [1] ATI Mobility RadeonTM HD4850/4870 Graphics-Overview. <http://ati.amd.com/products/radeonhd4800>.
- [2] Intel Core2 Quad Processors. <http://www.intel.com/products/processor/core2quad>.
- [3] NVIDIA GeForce series GTX280, 8800GTX, 8800GT. <http://www.nvidia.com/geforce>.
- [4] NVIDIA Quadro FX5600. <http://www.nvidia.com/quadro>.
- [5] Advanced Micro Devices, Inc. AMD Brook+. <http://ati.amd.com/technology/streamcomputing/AMD-Brookplus.pdf>.
- [6] A. Bakhoda, G. Yuan, W. W. L. Fung, H. Wong, and T. M. Aamodt. Analyzing cuda workloads using a detailed GPU simulator. In *IEEE ISPASS*, April 2009.
- [7] X. E. Chen and T. M. Aamodt. A first-order fine-grained multithreaded throughput model. In *HPCA*, 2009.
- [8] E. Lindholm, J. Nickolls, S. Oberman and J. Montrym. NVIDIA Tesla: A Unified Graphics and Computing Architecture. *IEEE Micro*, 28(2):39–55, March-April 2008.
- [9] M. Fatica, P. LeGresley, I. Buck, J. Stone, J. Phillips, S. Morton, and P. Micikevicius. High Performance Computing with CUDA, SC08, 2008.
- [10] A. Glew. MLP yes! ILP no! In *ASPLOS Wild and Crazy Idea Session '98*, Oct. 1998.
- [11] GPGPU. General-Purpose Computation Using Graphics Hardware. <http://www.gpgpu.org/>.
- [12] S. Hong and H. Kim. An analytical model for a GPU architecture with memory-level and thread-level parallelism awareness. Technical Report TR-2009-003, Atlanta, GA, USA, 2009.
- [13] W. Hwu and D. Kirk. Ece 498a1: Programming massively parallel processors, fall 2007. <http://courses.ece.uiuc.edu/ece498/a1/>.
- [14] Intel SSE / MMX2 / KNI documentation. <http://www.intel80386.com/simd/mmx2-doc.html>.
- [15] T. S. Karkhanis and J. E. Smith. A first-order superscalar processor model. In *ISCA*, 2004.
- [16] Khronos. Opencl - the open standard for parallel programming of heterogeneous systems. <http://www.khronos.org/opencl/>.
- [17] M. D. Linderman, J. D. Collins, H. Wang, and T. H. Meng. Merge: a programming model for heterogeneous multi-core systems. In *ASPLOS XIII*, 2008.
- [18] P. Michaud and A. Seznec. Data-flow prescheduling for large instruction windows in out-of-order processors. In *HPCA*, 2001.
- [19] P. Michaud, A. Seznec, and S. Jourdan. Exploring instruction-fetch bandwidth requirement in wide-issue superscalar processors. In *PACT*, 1999.
- [20] J. Nickolls, I. Buck, M. Garland, and K. Skadron. Scalable Parallel Programming with CUDA. *ACM Queue*, 6(2):40–53, March-April 2008.
- [21] D. B. Noonburg and J. P. Shen. Theoretical modeling of superscalar processor performance. In *MICRO-27*, 1994.
- [22] NVIDIA Corporation. *CUDA Programming Guide, Version 2.1*.
- [23] M. Pharr and R. Fernando. *GPU Gems 2*. Addison-Wesley Professional, 2005.
- [24] S. Ryoo, C. Rodrigues, S. Stone, S. Bagsorkhi, S. Ueng, J. Stratton, and W. Hwu. Program optimization space pruning for a multithreaded gpu. In *CGO*, 2008.
- [25] R. H. Saavedra-Barrera and D. E. Culler. An analytical solution for a markov chain modeling multithreaded. Technical report, Berkeley, CA, USA, 1991.
- [26] L. Seiler, D. Carmean, E. Sprangle, T. Forsyth, M. Abrash, P. Dubey, S. Junkins, A. Lake, J. Sugerma, R. Cavin, R. Espasa, E. Grochowski, T. Juan, and P. Hanrahan. Larrabee: a many-core x86 architecture for visual computing. *ACM Trans. Graph.*, 2008.
- [27] D. J. Sorin, V. S. Pai, S. V. Adve, M. K. Vernon, and D. A. Wood. Analytic evaluation of shared-memory systems with ILP processors. In *ISCA*, 1998.
- [28] C. A. Waring and X. Liu. Face detection using spectral histograms and SVMs. *Systems, Man, and Cybernetics, Part B, IEEE Transactions on*, 35(3):467–476, June 2005.