

A Framework for Accelerating Bottlenecks in GPU Execution with Assist Warps

Nandita Vijaykumar Gennady Pekhimenko Adwait Jog[†] Saugata Ghose Abhishek Bhowmick
Rachata Ausavarungnirun Chita Das[†] Mahmut Kandemir[†] Todd C. Mowry Onur Mutlu

Carnegie Mellon University

[†] Pennsylvania State University

{nandita,ghose,abhowmick,rachata,onur}@cmu.edu

{gpekhime,tcm}@cs.cmu.edu

{adwait,das,kandemir}@cse.psu.edu

Abstract

Modern Graphics Processing Units (GPUs) are well provisioned to support the concurrent execution of thousands of threads. Unfortunately, different bottlenecks during execution and heterogeneous application requirements create imbalances in utilization of resources in the cores. For example, when a GPU is bottlenecked by the available off-chip memory bandwidth, its computational resources are often overwhelmingly idle, waiting for data from memory to arrive.

This work describes the Core-Assisted Bottleneck Acceleration (CABA) framework that employs idle on-chip resources to alleviate different bottlenecks in GPU execution. CABA provides flexible mechanisms to automatically generate “assist warps” that execute on GPU cores to perform specific tasks that can improve GPU performance and efficiency.

CABA enables the use of idle computational units and pipelines to alleviate the memory bandwidth bottleneck, e.g., by using assist warps to perform data compression to transfer less data from memory. Conversely, the same framework can be employed to handle cases where the GPU is bottlenecked by the available computational units, in which case the memory pipelines are idle and can be used by CABA to speed up computation, e.g., by performing memoization using assist warps.

We provide a comprehensive design and evaluation of CABA to perform effective and flexible data compression in the GPU memory hierarchy to alleviate the memory bandwidth bottleneck. Our extensive evaluations show that CABA, when used to implement data compression, provides an average performance improvement of 41.7% (as high as 2.6X) across a variety of memory-bandwidth-sensitive GPGPU applications.

We believe that CABA is a flexible framework that enables the use of idle resources to improve application performance with different optimizations and perform other useful tasks. We discuss how CABA can be used, for example, for memoization, prefetching, handling interrupts, profiling, redundant multi-threading, and speculative precomputation.

1. Introduction

Modern Graphics Processing Units (GPUs) play an important role in delivering high performance and energy efficiency for many classes of applications and different computational platforms. GPUs employ fine-grained multi-threading to hide the high memory access latencies with thousands of concurrently running threads [58]. GPUs are well provisioned with different resources (e.g., SIMD-like computational units, large register files) to support the execution of a large number of

these hardware contexts. Ideally, if the demand for all types of resources is properly balanced, all these resources should be fully utilized by the application. Unfortunately, this balance is very difficult to achieve in practice.

As a result, bottlenecks in program execution, e.g., limitations in memory or computational bandwidth, lead to long stalls and idle periods in the shader pipelines of modern GPUs [51, 52, 79, 103]. Alleviating these bottlenecks with optimizations implemented in dedicated hardware requires significant engineering cost and effort. Fortunately, the resulting under-utilization of on-chip computational and memory resources from these imbalances in application requirements, offers some new opportunities. For example, we can use these resources for efficient integration of *hardware-generated threads* that perform useful work to accelerate the execution of the primary threads. Similar *helper threading* ideas have been proposed in the context of general-purpose processors [22, 23, 27, 31, 32, 97, 122] to either extend the pipeline with more contexts or use spare hardware contexts to precompute useful information that aids main code execution (e.g., to aid branch prediction, prefetching, etc.).

We believe that the general idea of helper threading can lead to even more powerful optimizations and new opportunities in the context of modern GPUs than in CPUs because (1) the abundance of on-chip resources in a GPU obviates the need for idle hardware contexts [27, 28] or the addition of more storage (registers, rename tables, etc.) and compute units [22, 75] required to handle more contexts and (2) the relative simplicity of the GPU pipeline avoids the complexities of handling register renaming, speculative execution, precise interrupts, etc. [23]. However, GPUs that execute and manage thousands of thread contexts at the same time pose new challenges for employing helper threading, which must be addressed carefully. First, the numerous regular program threads executing in parallel could require an equal or larger number of helper threads that need to be managed at low cost. Second, the compute and memory resources are dynamically partitioned between threads in GPUs, and resource allocation for helper threads should be cognizant of resource interference and overheads. Third, lock-step execution and complex scheduling—which are characteristic of GPU architectures—exacerbate the complexity of fine-grained management of helper threads.

In this work, we describe a new, flexible framework for bottleneck acceleration in GPUs via helper threading (called *Core-Assisted Bottleneck Acceleration* or CABA), which exploits the aforementioned new opportunities while effectively han-

dling the new challenges. CABA performs acceleration by generating special warps—*assist warps*—that can execute code to speed up application execution and system tasks. To simplify the support of the numerous assist threads with CABA, we manage their execution at the granularity of a *warp* and use a centralized mechanism to track the progress of each *assist warp* throughout its execution. To reduce the overhead of providing and managing new contexts for each generated thread, as well as to simplify scheduling and data communication, an assist warp *shares the same context* as the regular warp it assists. Hence, the regular warps are overprovisioned with *available registers* to enable each of them to host its own assist warp.

Use of CABA for compression. We illustrate an important use case for the CABA framework: alleviating the memory bandwidth bottleneck by enabling *flexible data compression* in the memory hierarchy. The basic idea is to have assist warps that (1) compress cache blocks before they are written to memory, and (2) decompress cache blocks before they are placed into the cache.

CABA-based compression/decompression provides several benefits over a purely hardware-based implementation of data compression for memory. First, CABA primarily employs hardware that is already available on-chip but is otherwise underutilized. In contrast, hardware-only compression implementations require *dedicated logic* for specific algorithms. Each new algorithm (or a modification of an existing one) requires engineering effort and incurs hardware cost. Second, different applications tend to have distinct data patterns [87] that are more efficiently compressed with different compression algorithms. CABA offers versatility in algorithm choice as we find that many existing hardware-based compression algorithms (e.g., Base-Delta-Immediate (BDI) compression [87], Frequent Pattern Compression (FPC) [4], and C-Pack [25]) can be implemented using different assist warps with the CABA framework. Third, not all applications benefit from data compression. Some applications are constrained by other bottlenecks (e.g., oversubscription of computational resources), or may operate on data that is not easily compressible. As a result, the benefits of compression may not outweigh the cost in terms of additional latency and energy spent on compressing and decompressing data. In these cases, compression can be easily disabled by CABA, and the CABA framework can be used in other ways to alleviate the current bottleneck.

Other uses of CABA. The generality of CABA enables its use in alleviating other bottlenecks with different optimizations. We discuss two examples: (1) using assist warps to perform *memoization* to eliminate redundant computations that have the same or similar inputs [13, 29, 106], by storing the results of frequently-performed computations in the main memory hierarchy (i.e., by converting the computational problem into a storage problem) and, (2) using the idle memory pipeline to perform opportunistic *prefetching* to better overlap computation with memory access. Assist warps offer a hardware/software interface to implement hybrid prefetching algorithms [35] with varying degrees of complexity. We also briefly discuss other uses of CABA for (1) redundant

multithreading, (2) speculative precomputation, (3) handling interrupts, and (4) profiling and instrumentation.

Contributions. This work makes the following contributions:

- It introduces the *Core-Assisted Bottleneck Acceleration (CABA) Framework*, which can mitigate different bottlenecks in modern GPUs by using underutilized system resources for *assist warp* execution.
- It provides a detailed description of how our framework can be used to enable effective and flexible data compression in GPU memory hierarchies.
- It comprehensively evaluates the use of CABA for data compression to alleviate the memory bandwidth bottleneck. Our evaluations across a wide variety applications from Mars [44], CUDA [83], Lonestar [20], and Rodinia [24] benchmark suites show that CABA-based compression on average (1) reduces memory bandwidth by 2.1X, (2) improves performance by 41.7%, and (3) reduces overall system energy by 22.2%.
- It discusses at least six other use cases of CABA that can improve application performance and system management, showing that CABA is a primary general framework for taking advantage of underutilized resources in modern GPU engines.

2. Background

A GPU consists of multiple simple cores, also called streaming multiprocessors (SMs) in NVIDIA terminology or compute units (CUs) in AMD terminology. Our example architecture (shown in Figure 1) consists of 15 cores each with a SIMT width of 32, and 6 memory controllers. Each core is associated with a private L1 data cache and read-only texture and constant caches along with a low latency, programmer-managed shared memory. The cores and memory controllers are connected via a crossbar and every memory controller is associated with a slice of the shared L2 cache. This architecture is similar to many modern GPU architectures, including NVIDIA Fermi [84] and AMD Radeon [10].




Figure 1: Baseline GPU architecture. Figure reproduced from [115].

A typical GPU application consists of many kernels. Each kernel is divided into groups of threads, called thread-blocks (or cooperative thread arrays (CTAs)). After a kernel is launched and its necessary data is copied to the GPU memory, the thread-block scheduler schedules available CTAs onto all the available cores [17]. Once the CTAs are launched onto the cores, the warps associated with the CTAs are scheduled onto the cores’ SIMT pipelines. Each core is capable of concurrently

executing many warps, where a warp is typically defined as a group of threads that are executed in lockstep. In modern GPUs, a warp can contain 32 threads [84], for example. The maximum number of warps that can be launched on a core depends on the available core resources (e.g., available shared memory, register file size etc.). For example, in a modern GPU as many as 64 warps (i.e., 2048 threads) can be present on a single GPU core.

For more details on the internals of modern GPU architectures, we refer the reader to [53, 61].

3. Motivation

We observe that different bottlenecks and imbalances during program execution leave resources unutilized within the GPU cores. We motivate our proposal, CABA, by examining these inefficiencies. CABA leverages these inefficiencies as an opportunity to perform useful work.

Unutilized Compute Resources. A GPU core employs fine-grained multithreading [105, 112] of *warps*, i.e., groups of threads executing the same instruction, to hide long memory and ALU operation latencies. If the number of available warps is insufficient to cover these long latencies, the core stalls or becomes idle. To understand the key sources of inefficiency in GPU cores, we conduct an experiment where we show the breakdown of the applications’ execution time spent on either useful work (*Active Cycles*) or stalling due to one of four reasons: *Compute*, *Memory*, *Data Dependence Stalls* and *Idle Cycles*. We also vary the amount of available off-chip memory bandwidth: (i) half (1/2xBW), (ii) equal to (1xBW), and (iii) double (2xBW) the peak memory bandwidth of our baseline GPU architecture. Section 6 details our baseline architecture and methodology.

Figure 2 shows the percentage of total issue cycles, divided into five components (as described above). The first two components—*Memory and Compute Stalls*—are attributed to the main memory and ALU-pipeline structural stalls. These stalls are because of backed-up pipelines due to oversubscribed resources that prevent warps from being issued to the respective pipelines. The third component (*Data Dependence Stalls*) is due to data dependence stalls. These stalls prevent warps from issuing new instruction(s) when the previous instruction(s) from the same warp are stalled on long-latency operations (usually memory load operations). In some applications (e.g., *dmr*), special-function-unit (SFU) ALU operations that may take tens of cycles to finish are also the source of data dependence stalls. The fourth component, *Idle Cycles*, refers to idle cycles when either all the available warps are issued to the pipelines and not ready to execute their next instruction or the instruction buffers are flushed due to a mispredicted branch. All these components are sources of inefficiency that cause the cores to be underutilized. The last component, *Active Cycles*, indicates the fraction of cycles during which at least one warp was successfully issued to the pipelines.

We make two observations from Figure 2. First, *Compute*, *Memory*, and *Data Dependence Stalls* are the major sources of underutilization in many GPU applications. We distinguish

applications based on their primary bottleneck as either *Memory* or *Compute Bound*. We observe that a majority of the applications in our workload pool (17 out of 27 studied) are *Memory Bound*, and bottlenecked by the off-chip memory bandwidth.

Second, for the *Memory Bound* applications, we observe that the *Memory* and *Data Dependence* stalls constitute a significant fraction (61%) of the total issue cycles on our baseline GPU architecture (1xBW). This fraction goes down to 51% when the peak memory bandwidth is doubled (2xBW), and increases significantly when the peak bandwidth is halved (1/2xBW), indicating that limited off-chip memory bandwidth is a critical performance bottleneck for *Memory Bound* applications. Some applications, e.g., *BFS*, are limited by the interconnect bandwidth. In contrast, the *Compute Bound* applications are primarily bottlenecked by stalls in the ALU pipelines. An increase or decrease in the off-chip bandwidth has little effect on the performance of these applications.

Unutilized On-chip Memory. The *occupancy* of any GPU Streaming Multiprocessor (SM), i.e., the number of threads running concurrently, is limited by a number of factors: (1) the available registers and shared memory, (2) the hard limit on the number of threads and thread blocks per core, (3) the number of thread blocks in the application kernel. The limiting resource from the above, leaves the other resources underutilized. This is because it is challenging, in practice, to achieve a perfect balance in utilization of all of the above factors for different workloads with varying characteristics. Very often, the factor determining the occupancy is the thread or thread block limit imposed by the architecture. In this case, there are many registers that are left unallocated to any thread block. Also, the number of available registers may not be a multiple of those required by each thread block. The remaining registers are not enough to schedule an entire extra thread block, which leaves a significant fraction of the register file and shared memory unallocated and unutilized by the thread blocks. Figure 3 shows the fraction of statically unallocated registers in a 128KB register file (per SM) with a 1536 thread, 8 thread block occupancy limit, for different applications. We observe that on average 24% of the register file remains unallocated. This phenomenon has previously been observed and analyzed in detail in [3, 39, 40, 41, 63]. We observe a similar trend with the usage of shared memory (not graphed).

Our Goal. We aim to exploit the underutilization of compute resources, registers and on-chip shared memory as an opportunity to enable different optimizations to accelerate various bottlenecks in GPU program execution. To achieve this goal, we would like to enable efficient helper threading for GPUs to dynamically generate threads in hardware that use the available on-chip resources for various purposes. In the next section, we present the detailed design of our CABA framework that enables the generation and management of these threads.

thread.

On the other hand, a pure hardware solution would make register allocation for the assist warps and the data communication between the helper threads and primary threads more difficult. Registers are allocated to each thread block by the compiler and are then mapped to the sections of the hardware register file at runtime. Mapping registers for helper threads and enabling data communication between those registers and the primary thread registers would be non-trivial. Furthermore, a fully hardware approach would make offering the programmer a flexible interface more challenging.

Hardware support enables simpler fine-grained management of helper threads, aware of micro-architectural events and runtime program behavior. Compiler/runtime support enables simpler context management for helper threads and more flexible programmer interfaces. Thus, to get the best of both worlds, we propose a *hardware/software cooperative approach*, where the hardware manages the scheduling and execution of helper thread subroutines, while the compiler performs the allocation of shared resources (e.g., register file and shared memory) for the helper threads and the programmer or the microarchitect provides the helper threads themselves.

4.2.1. Hardware-based management of threads. To use the available on-chip resources the same way that thread blocks do during program execution, we dynamically insert sequences of instructions into the execution stream. We track and manage these instructions at the granularity of a warp, and refer to them as *Assist Warps*. An assist warp is a set of instructions issued into the core pipelines. Each instruction is executed in lock-step across all the SIMT lanes, just like any regular instruction, with an active mask to disable lanes as necessary. The assist warp does *not* own a separate context (e.g., registers, local memory), and instead shares both a context and a warp ID with the regular warp that invoked it. In other words, each assist warp is coupled with a *parent warp*. In this sense, it is different from a regular warp and does not reduce the number of threads that can be scheduled on a single SM. Data sharing between the two warps becomes simpler, since the assist warps share the register file with the parent warp. Ideally, an assist warp consumes resources and issue cycles that would otherwise be idle. We describe the structures required to support hardware-based management of assist warps in Section 4.3.

4.2.2. Register file/shared memory allocation. Each helper thread subroutine requires a different number of registers depending on the actions it performs. These registers have a short lifetime, with no values being preserved between different invocations of an assist warp. To limit the register requirements for assist warps, we impose the restriction that only one instance of each helper thread routine can be active for each thread. All instances of the same helper thread for each parent thread use the same registers, and the registers are allocated to the helper threads statically by the compiler. One of the factors that determines the runtime SM occupancy is the number of registers required by a thread block (i.e., per-block register requirement). For each helper thread subroutine that is enabled, we add its register requirement to the per-block

register requirement, to ensure the availability of registers for both the parent threads as well as every assist warp. The registers that remain unallocated after allocation among the parent thread blocks should suffice to support the assist warps. If not, register-heavy assist warps may limit the parent thread block occupancy in SMs or increase the number of register spills in the parent warps. Shared memory resources are partitioned in a similar manner and allocated to each assist warp as and if needed.

4.2.3. Programmer/developer interface. The assist warp subroutine can be written in two ways. First, it can be supplied and annotated by the programmer/developer using CUDA extensions with PTX instructions and then compiled with regular program code. Second, the assist warp subroutines can be written by the microarchitect in the internal GPU instruction format. These helper thread subroutines can then be enabled or disabled by the application programmer. This approach is similar to that proposed in prior work (e.g., [22]). It offers the advantage of potentially being highly optimized for energy and performance while having flexibility in implementing optimizations that are not trivial to map using existing GPU PTX instructions. The instructions for the helper thread subroutine are stored in an on-chip buffer (described in Section 4.3).

Along with the helper thread subroutines, the programmer also provides: (1) the *priority* of the assist warps to enable the warp scheduler to make informed decisions, (2) the trigger conditions for each assist warp, and (3) the live-in and live-out variables for data communication with the parent warps.

Assist warps can be scheduled with different priority levels in relation to parent warps by the warp scheduler. Some assist warps may perform a function that is required for correct execution of the program and are *blocking*. At this end of the spectrum, the *high priority* assist warps are treated by the scheduler as always taking higher precedence over the parent warp execution. Assist warps should be given a high priority only when they are required for correctness. *Low priority* assist warps, on the other hand, are scheduled for execution only when computational resources are available, i.e., during idle cycles. There is no guarantee that these assist warps will execute or complete.

The programmer also provides the conditions or events that need to be satisfied for the deployment of the assist warp. This includes a specific point within the original program and/or a set of other microarchitectural events that could serve as a *trigger* for starting the execution of an assist warp.

4.3. Main Hardware Additions

Figure 4 shows a high-level block diagram of the GPU pipeline [43]. To support assist warp execution, we add three new components: (1) an Assist Warp Store to hold the assist warp code, (2) an Assist Warp Controller to perform the deployment, tracking, and management of assist warps, and (3) an Assist Warp Buffer to stage instructions from triggered assist warps for execution.

Assist Warp Store (AWS). Different assist warp subroutines are possible based on the purpose of the optimization.

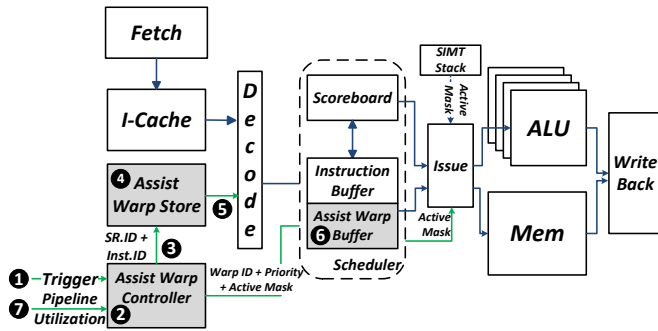


Figure 4: CABA framework flow within a typical GPU pipeline [43]. The shaded blocks are the components introduced for the framework. Figure reproduced from [115].

These code sequences for different types of assist warps need to be stored on-chip. An on-chip storage structure called the Assist Warp Store (AWS) is preloaded with these instructions before application execution. It is indexed using the subroutine index (SR.ID) along with the instruction ID (Inst.ID).

Assist Warp Controller (AWC). The AWC (2) is responsible for the triggering, tracking, and management of assist warp execution. It stores a mapping between trigger events and a subroutine index in the AWS, as specified by the programmer. The AWC monitors for such events, and when they take place, triggers the fetch, decode and execution of instructions from the AWS for the respective assist warp.

Deploying all the instructions within an assist warp, back-to-back, at the trigger point may require increased fetch/decode bandwidth and buffer space after decoding [23]. To avoid this, at each cycle, only a few instructions from an assist warp, at most equal to the available decode/issue bandwidth, are decoded and staged for execution. Within the AWC, we simply track the next instruction that needs to be executed for each assist warp and this is stored in the Assist Warp Table (AWT), as depicted in Figure 5. The AWT also tracks additional metadata required for assist warp management, which is described in more detail in Section 4.4.

Assist Warp Buffer (AWB). Fetched and decoded instructions (2) belonging to the assist warps that have been triggered need to be buffered until the assist warp can be selected for issue by the scheduler. These instructions are then staged in the Assist Warp Buffer (6) along with their warp IDs. The AWB is contained within the *instruction buffer (IB)*, which holds decoded instructions for the parent warps. The AWB makes use of the existing IB structures. The IB is typically partitioned among different warps executing in the SM. Since each assist warp is associated with a parent warp, the assist warp instructions are directly inserted into the *same partition* within the IB as that of the parent warp. This simplifies warp scheduling, as the assist warp instructions can now be issued as if they were parent warp instructions with the same warp ID. In addition, using the existing partitions avoids the cost of separate dedicated instruction buffering for assist warps. We do, however, provision a small additional partition with two entries within the IB, to hold non-blocking *low priority* assist warps that are scheduled only during idle cycles. This additional partition allows the scheduler to distinguish *low pri-*



Figure 5: Fetch Logic: Assist Warp Table (contained in the AWC) and the Assist Warp Store (AWS). Figure reproduced from [115].

ority assist warp instructions from the parent warp and *high priority* assist warp instructions, which are given precedence during scheduling, allowing them to make progress.

4.4. The Mechanism

Trigger and Deployment. An assist warp is triggered (1) by the AWC (2) based on a specific set of architectural events and/or a triggering instruction (e.g., a load instruction). When an assist warp is triggered, its specific instance is placed into the Assist Warp Table (AWT) within the AWC (Figure 5). Every cycle, the AWC selects an assist warp to deploy in a round-robin fashion. The AWS is indexed (3) based on the subroutine ID (SR.ID)—which selects the instruction sequence to be executed by the assist warp, and the instruction ID (Inst.ID)—which is a pointer to the next instruction to be executed within the subroutine (Figure 5). The selected instruction is entered (5) into the AWB (6) and, at this point, the instruction enters the active pool with other active warps for scheduling. The Inst.ID for the assist warp is updated in the AWT to point to the next instruction in the subroutine. When the end of the subroutine is reached, the entry within the AWT is freed.

Execution. Assist warp instructions, when selected for issue by the scheduler, are executed in much the same way as any other instructions. The scoreboard tracks the dependencies between instructions within an assist warp in the same way as any warp, and instructions from different assist warps are interleaved in execution in order to hide latencies. We also provide an active mask (stored as a part of the AWT), which allows for statically disabling/enabling different lanes within a warp. This is useful to provide flexibility in lock-step instruction execution when we do not need all threads within a warp to execute a specific assist warp subroutine.

Dynamic Feedback and Throttling. Assist warps, if not properly controlled, may stall application execution. This can happen due to several reasons. First, assist warps take up issue cycles, and only a limited number of instructions may be issued per clock cycle. Second, assist warps require structural resources: the ALU units and resources in the load-store pipelines (if the assist warps consist of computational and memory instructions, respectively). We may, hence, need to throttle assist warps to ensure that their performance benefits outweigh the overhead. This requires mechanisms to appropriately balance and manage the aggressiveness of assist warps at runtime.

The overheads associated with assist warps can be controlled in different ways. First, the programmer can statically specify the priority of the assist warp. Depending on the criticality of the assist warps in making forward progress, the

assist warps can be issued either in idle cycles or with varying levels of priority in relation to the parent warps. For example, warps performing *decompression* are given a high priority whereas warps performing *compression* are given a low priority. Low priority assist warps are inserted into the dedicated partition in the IB, and are scheduled only during idle cycles. This priority is statically defined by the programmer. Second, the AWC can control the number of times the assist warps are deployed into the AWB. The AWC monitors the utilization of the functional units (FU) and idleness of the cores to decide when to throttle assist warp deployment.

Communication and Control. An assist warp may need to communicate data and status with its parent warp. For example, memory addresses from the parent warp need to be communicated to assist warps performing decompression or prefetching. The IDs of the registers containing the live-in data for each assist warp are saved in the AWT when an assist warp is triggered. Similarly, if an assist warp needs to report results to its parent warp (e.g., in the case of memoization), the register IDs are also stored in the AWT. When the assist warps execute, *MOVE* instructions are first executed to copy the live-in data from the parent warp registers to the assist warp registers. Live-out data is communicated to the parent warp in a similar fashion, at the end of assist warp execution.

Assist warps may need to be *killed* when they are not required (e.g., if the data does not require decompression) or when they are no longer beneficial. In this case, the entries in the AWT and AWB are simply flushed for the assist warp.

4.5. Applications of the CABA Framework

We envision multiple applications for the CABA framework, e.g., data compression [4, 25, 87, 118], memoization [13, 29, 106], data prefetching [15, 38, 54, 86, 108]. In Section 5, we provide a detailed case study of enabling data compression with the framework, discussing various tradeoffs. We believe CABA can be useful for many other optimizations, and we discuss some of them briefly in Section 8.

5. A Case for CABA: Data Compression

Data compression is a technique that exploits the redundancy in the applications’ data to reduce capacity and bandwidth requirements for many modern systems by saving and transmitting data in a more compact form. Hardware-based data compression has been explored in the context of on-chip caches [4, 11, 25, 33, 49, 87, 89, 99, 118], interconnect [30], and main memory [2, 37, 88, 90, 91, 104, 114] as a means to save storage capacity as well as memory bandwidth. In modern GPUs, memory bandwidth is a key limiter to system performance in many workloads (Section 3). As such, data compression is a promising technique to help alleviate this bottleneck. Compressing data enables less data to be transferred from/to DRAM and the interconnect.

In bandwidth-constrained workloads, idle compute pipelines offer an opportunity to employ CABA to enable data compression in GPUs. We can use assist warps to (1) decompress data, before loading it into the caches and

registers, and (2) compress data, before writing it back to memory. Since assist warps execute instructions, CABA offers some flexibility in the compression algorithms that can be employed. Compression algorithms that can be mapped to the general GPU execution model can be flexibly implemented with the CABA framework.

5.1. Mapping Compression Algorithms into Assist Warps

In order to employ CABA to enable data compression, we need to map compression algorithms into instructions that can be executed within the GPU cores. For a compression algorithm to be amenable for implementation with CABA, it ideally needs to be (1) reasonably parallelizable and (2) simple (for low latency). Decompressing data involves reading the encoding associated with each cache line that defines how to decompress it, and then triggering the corresponding decompression subroutine in CABA. Compressing data, on the other hand, involves testing different encodings and saving data in the compressed format.

We perform compression at the granularity of a cache line. The data needs to be decompressed before it is used by any program thread. In order to utilize the full SIMD width of the GPU pipeline, we would like to decompress/compress all the words in the cache line in parallel. With CABA, helper thread routines are managed at the warp granularity, enabling fine-grained triggering of assist warps to perform compression/decompression when required. However, the SIMT execution model in a GPU imposes some challenges: (1) threads within a warp operate in lock-step, and (2) threads operate as independent entities, i.e., they do not easily communicate with each other.

In this section, we discuss the architectural changes and algorithm adaptations required to address these challenges and provide a detailed implementation and evaluation of *Data Compression* within the CABA framework using the *Base-Delta-Immediate compression* algorithm [87]. Section 5.1.3 discusses implementing other compression algorithms.

5.1.1. Algorithm Overview. Base-Delta-Immediate compression (BDI) is a simple compression algorithm that was originally proposed in the context of caches [87]. It is based on the observation that many cache lines contain data with low dynamic range. BDI exploits this observation to represent a cache line with low dynamic range using a common *base* (or multiple bases) and an array of *deltas* (where a delta is the difference of each value within the cache line and the common base). Since the *deltas* require fewer bytes than the values themselves, the combined size after compression can be much smaller. Figure 6 shows the compression of an example 64-byte cache line from the *PageViewCount (PVC)* application using BDI. As Figure 6 indicates, in this case, the cache line can be represented using two bases (an 8-byte base value, 0x8001D000, and an implicit zero value base) and an array of eight 1-byte differences from these bases. As a result, the entire cache line data can be represented using 17 bytes instead of 64 bytes (1-byte metadata, 8-byte base, and eight 1-byte

deltas), saving 47 bytes of the originally used space.




Figure 6: Cache line from PVC compressed with BDI. Figure reproduced from [115].

Our example implementation of the BDI compression algorithm [87] views a cache line as a set of fixed-size values i.e., 8 8-byte, 16 4-byte, or 32 2-byte values for a 64-byte cache line. For the size of the deltas, it considers three options: 1, 2 and 4 bytes. The key characteristic of BDI, which makes it a desirable compression algorithm to use with the CABA framework, is its fast parallel decompression that can be efficiently mapped into instructions that can be executed on GPU hardware. Decompression is simply a masked vector addition of the deltas to the appropriate bases [87].

5.1.2. Mapping BDI to CABA. In order to implement BDI with the CABA framework, we need to map the BDI compression/decompression algorithms into GPU instruction subroutines (stored in the AWS and deployed as assist warps).

Decompression. To decompress the data compressed with BDI, we need a simple addition of deltas to the appropriate bases. The CABA decompression subroutine first loads the words within the compressed cache line into assist warp registers, and then performs the base-delta additions in parallel, employing the wide ALU pipeline.¹ The subroutine then writes back the uncompressed cache line to the cache. It skips the addition for the lanes with an implicit base of zero by updating the active lane mask based on the cache line encoding. We store a separate subroutine for each possible BDI encoding that loads the appropriate bytes in the cache line as the base and the deltas. The high-level algorithm for decompression is presented in Algorithm 1.

Algorithm 1 BDI: Decompression

- 1: load *base*, *deltas*
 - 2: $uncompressed_data = base + deltas$
 - 3: store *uncompressed_data*
-

Compression. To compress data, the CABA compression subroutine tests several possible encodings (each representing a different size of base and deltas) in order to achieve a high compression ratio. The first few bytes (2–8 depending on the encoding tested) of the cache line are always used as the base. Each possible encoding is tested to check whether the cache line can be successfully encoded with it. In order to perform compression at a warp granularity, we need to check whether all of the words at every SIMD lane were successfully compressed. In other words, if any one word cannot be compressed, that encoding cannot be used across the warp.

¹Multiple instructions are required if the number of deltas exceeds the width of the ALU pipeline. We use a 32-wide pipeline.

We can perform this check by adding a global predicate register, which stores the logical AND of the per-lane predicate registers. We observe that applications with homogeneous data structures can typically use the same encoding for most of their cache lines [87]. We use this observation to reduce the number of encodings we test to just one in many cases. All necessary operations are done in parallel using the full width of the GPU SIMD pipeline. The high-level algorithm for compression is presented in Algorithm 2.

Algorithm 2 BDI: Compression

- 1: **for** each *base_size* **do**
 - 2: load *base*, *values*
 - 3: **for** each *delta_size* **do**
 - 4: $deltas = abs(values - base)$
 - 5: **if** $size(deltas) \leq delta_size$ **then**
 - 6: store *base*, *deltas*
 - 7: **exit**
 - 8: **end if**
 - 9: **end for**
 - 10: **end for**
-

5.1.3. Implementing Other Algorithms. The BDI compression algorithm is naturally amenable towards implementation using assist warps because of its data-parallel nature and simplicity. The CABA framework can also be used to realize other algorithms. The challenge in implementing algorithms like FPC [5] and C-Pack [25], which have variable-length compressed words, is primarily in the placement of compressed words within the compressed cache lines. In BDI, the compressed words are in *fixed* locations within the cache line and, for each encoding, all the compressed words are of the same size and can, therefore, be processed in parallel. In contrast, C-Pack may employ multiple dictionary values as opposed to just one base in BDI. In order to realize algorithms with *variable length words* and *dictionary values* with assist warps, we leverage the coalescing/address generation logic [81, 85] already available in the GPU cores. We make two minor modifications to these algorithms [5, 25] to adapt them for use with CABA. First, similar to prior works [5, 25, 37], we observe that few encodings are sufficient to capture almost all the data redundancy. In addition, the impact of any loss in compressibility due to fewer encodings is minimal as the benefits of bandwidth compression are only at multiples of a single DRAM burst (e.g., 32B for GDDR5 [47]). We exploit this to reduce the number of supported encodings. Second, we place all the metadata containing the compression encoding at the *head* of the cache line to be able to determine how to decompress the entire line *upfront*. In the case of C-Pack, we place the dictionary entries after the metadata.

We note that it can be challenging to implement complex algorithms efficiently with the simple computational logic available in GPU cores. Fortunately, there are already Special Function Units (SFUs) [21, 66] present in the GPU SMs, used to perform efficient computations of elementary mathematical functions. SFUs could potentially be extended to implement primitives that enable the fast iterative compar-

isons performed frequently in some compression algorithms. This would enable more efficient execution of the described algorithms, as well as implementation of more complex compression algorithms, using CABA. We leave the exploration of an SFU-based approach to future work.

We now present a detailed overview of mapping the FPC and C-PACK algorithms into assist warps.

5.1.4. Implementing the FPC (Frequent Pattern Compression) Algorithm. For FPC, the cache line is treated as set of fixed-size words and each word within the cache line is compressed into a simple prefix or encoding and a compressed word if it matches a set of frequent patterns, e.g. narrow values, zeros or repeated bytes. The word is left uncompressed if it does not fit any pattern. We refer the reader to the original work [5] for a more detailed description of the original algorithm.

The challenge in mapping assist warps to the FPC decompression algorithm is in the serial sequence in which each word within a cache line is decompressed. This is because in the original proposed version, each compressed word can have a different size. To determine the location of a specific compressed word, it is necessary to have decompressed the previous word. We make some modifications to the algorithm in order to parallelize the decompression across different lanes in the GPU cores. First, we move the word prefixes (meta-data) for each word to the front of the cache line, so we know *upfront* how to decompress the rest of the cache line. Unlike with BDI, each word within the cache line has a different encoding and hence a different compressed word length and encoding pattern. This is problematic as statically storing the sequence of decompression instructions for every combination of patterns for all the words in a cache line would require very large instruction storage. In order to mitigate this, we break each cache line into a number of segments. Each segment is compressed independently and all the words within each segment are compressed using the *same encoding* whereas different segments may have different encodings. This creates a trade-off between simplicity/parallelizability versus compressibility. Consistent with previous works [5], we find that this doesn't significantly impact compressibility.

Decompression. The high-level algorithm we use for decompression is presented in Algorithm 3. Each segment within the compressed cache line is loaded in series. Each of the segments is decompressed in parallel—this is possible because all the compressed words within the segment have the same encoding. The decompressed segment is then stored before moving onto the next segment. The location of the next compressed segment is computed based on the size of the previous segment.

Compression. Similar to the BDI implementation, we loop through and test different encodings for each segment. We also compute the address offset for each segment at each iteration to store the compressed words in the appropriate location in the compressed cache line. Algorithm 4 presents the high-level FPC compression algorithm we use.

Implementing the C-Pack Algorithm. C-Pack [25] is a dictionary based compression algorithm where frequent

Algorithm 3 FPC: Decompression

```

1: for each segment do
2:   load compressed words
3:   pattern specific decompression (sign extension/zero value)
4:   store decompressed words
5:   segment-base-address = segment-base-address + segment-size
6: end for

```

Algorithm 4 FPC: Compression

```

1: load words
2: for each segment do
3:   for each encoding do
4:     test encoding
5:     if compressible then
6:       segment-base-address = segment-base-address +
       segment-size
7:     store compressed words
8:     break
9:   end if
10:  end for
11: end for

```

"dictionary" values are saved at the beginning of the cache line. The rest of the cache line contains encodings for each word which may indicate zero values, narrow values, full or partial matches into the dictionary or any word that is uncompressible.

In our implementation, we reduce the number of possible encodings to partial matches (only last byte mismatch), full word match, zero value and zero extend (only last byte) and we limit the number of dictionary values to 4. This enables fixed compressed word size within the cache line. A fixed compressed word size enables compression and decompression of different words within the cache line in parallel. If the number of required dictionary values or uncompressed words exceeds 4, the line is left decompressed. This is, as in BDI and FPC, a trade-off between simplicity and compressibility. In our experiments, we find that it does not significantly impact the compression ratio—primarily due the 32B minimum data size and granularity of compression.

Decompression. As described, to enable parallel decompression, we place the encodings and dictionary values at the head of the line. We also limit the number of encodings to enable quick decompression. We implement C-Pack decompression as a series of instructions (one per encoding used) to load all the registers with the appropriate dictionary values. We define the active lane mask based on the encoding (similar to the mechanism used in BDI) for each load instruction to ensure the correct word is loaded into each lane's register. Algorithm 5 provides the high-level algorithm for C-Pack decompression.

Compression. Compressing data with C-Pack involves determining the dictionary values that will be used to compress the rest of the line. In our implementation, we serially add each word from the beginning of the cache line to be a dictionary value if it was not already covered by a previous dictionary value. For each dictionary value, we test whether the rest of the words within the cache line is compressible. The

granularity of a full cache line, but stores can be at granularities smaller than the size of the cache line. This poses some additional difficulty if the destination cache line for a store is already compressed in main memory. Partial writes into a compressed cache line would require the cache line to be decompressed first, then updated with the new data, and written back to main memory. The common case—where the cache line that is being written to is uncompressed initially—can be easily handled. However, in the worst case, the cache line being partially written to is already in the compressed form in memory. We now describe the mechanism to handle both these cases.

Initially, to reduce the store latency, we assume that the cache line is uncompressed, and issue a store to the lower levels of the memory hierarchy, while buffering a copy in L1. If the cache line is found in L2/memory in the uncompressed form (❶), the assumption was correct. The store then proceeds normally and the buffered stores are evicted from L1. If the assumption is incorrect, the cache line is retrieved (❷) and decompressed before the store is retransmitted to the lower levels of the memory hierarchy.

5.3. Realizing Data Compression

Supporting data compression requires additional support from the main memory controller and the runtime system, as we describe below.

5.3.1. Initial Setup and Profiling. Data compression with CABA requires a one-time data setup before the data is transferred to the GPU. We assume initial software-based data preparation where the input data is stored in CPU memory in the compressed form with an appropriate compression algorithm before transferring the data to GPU memory. Transferring data in the compressed form can also reduce PCIe bandwidth usage.²

Memory-bandwidth-limited GPU applications are the best candidates for employing data compression using CABA. The compiler (or the runtime profiler) is required to identify those applications that are most likely to benefit from this framework. For applications where memory bandwidth is not a bottleneck, data compression is simply disabled (this can be done statically or dynamically within the framework).

5.3.2. Memory Controller Changes. Data compression reduces off-chip bandwidth requirements by transferring the same data in fewer DRAM bursts. The memory controller (MC) needs to know whether the cache line data is compressed and how many bursts (1–4 bursts in GDDR5 [47]) are needed to transfer the data from DRAM to the MC. Similar to prior work [88, 100], we require metadata information for every cache line that keeps track of how many bursts are needed to transfer the data. Similar to prior work [100], we simply reserve 8MB of GPU DRAM space for the metadata (~0.2% of all available memory). Unfortunately, this simple design would require an additional access for the metadata for every access to DRAM effectively doubling the required bandwidth. To avoid this, a simple *metadata (MD) cache* that

²This requires changes to the DMA engine to recognize compressed lines.

keeps frequently-accessed metadata on chip (near the MC) is required. Note that this metadata cache is similar to other metadata storage and caches proposed for various purposes in the memory controller, e.g., [42, 67, 74, 88, 94, 101]. Our experiments show that a small 8 KB 4-way associative MD cache is sufficient to provide a hit rate of 85% on average (more than 99% for many applications) across all applications in our workload pool.³ Hence, in the common case, a second access to DRAM to fetch compression-related metadata can be avoided.

6. Methodology

We model the CABA framework in GPGPU-Sim 3.2.1 [16]. Table 1 provides the major parameters of the simulated system. We use GPUWattch [65] to model GPU power and CACTI [113] to evaluate the power/energy overhead associated with the MD cache (Section 5.3.2) and the additional components (AWS and AWC) of the CABA framework. We implement BDI [87] using the Synopsys Design Compiler with 65nm library (to evaluate the energy overhead of compression/decompression for the dedicated hardware design for comparison to CABA), and then use ITRS projections [50] to scale our results to the 32nm technology node.

System Overview	15 SMs, 32 threads/warp, 6 memory channels
Shader Core Config	1.4GHz, GTO scheduler [96], 2 schedulers/SM
Resources / SM	48 warps/SM, 32768 registers, 32KB Shared Memory
L1 Cache	16KB, 4-way associative, LRU replacement policy
L2 Cache	768KB, 16-way associative, LRU replacement policy
Interconnect	1 crossbar/direction (15 SMs, 6 MCs), 1.4GHz
Memory Model	177.4GB/s BW, 6 GDDR5 Memory Controllers (MCs), FR-FCFS scheduling, 16 banks/MC
GDDR5 Timing [47]	$t_{CL} = 12, t_{RP} = 12, t_{RC} = 40, t_{RAS} = 28,$ $t_{RCD} = 12, t_{RRD} = 6, t_{CLDR} = 5, t_{WR} = 12$

Table 1: Major parameters of the simulated systems.

Evaluated Applications. We use a number of CUDA applications derived from CUDA SDK [83] (*BFS, CONS, JPEG, LPS, MUM, RAY, SLA, TRA*), Rodinia [24] (*hs, nw*), Mars [44] (*KM, MM, PVC, PVR, SS*) and lonestar [20] (*bfs, bh, mst, sp, sssp*) suites. We run all applications to completion or for 1 billion instructions (whichever comes first). CABA-based data compression is beneficial mainly for memory-bandwidth-limited applications. In computation-resource limited applications, data compression is not only unrewarding, but it can also cause significant performance degradation due to the computational overheads associated with assist warps. We rely on static profiling to identify memory-bandwidth-limited applications and disable CABA-based compression for the others. In our evaluation (Section 7), we demonstrate detailed results for applications that exhibit some compressibility in memory bandwidth (at least 10%). Applications without compressible data (e.g., *sc, SCP*) do not gain any performance from the

³For applications where MD cache miss rate is low, we observe that MD cache misses are usually also TLB misses. Hence, most of the overhead of MD cache misses in these applications is outweighed by the cost of page table lookups.

CABA framework, and we verified that these applications do not incur any performance degradation (because the assist warps are *not* triggered for them).

Evaluated Metrics. We present Instruction per Cycle (*IPC*) as the primary performance metric. We also use *average bandwidth utilization*, defined as the fraction of total DRAM cycles that the DRAM data bus is busy, and *compression ratio*, defined as the ratio of the number of DRAM bursts required to transfer data in the compressed vs. uncompressed form. As reported in prior work [87], we use decompression/compression latencies of 1/5 cycles for the hardware implementation of BDI.

7. Results

To evaluate the effectiveness of using CABA to employ data compression, we compare five different designs: (i) *Base* - the baseline system with no compression, (ii) *HW-BDI-Mem* - hardware-based *memory bandwidth compression* with dedicated logic (data is stored compressed in main memory but uncompressed in the last-level cache, similar to prior works [88, 100]), (iii) *HW-BDI* - hardware-based *interconnect and memory bandwidth compression* (data is stored uncompressed only in the L1 cache) (iv) *CABA-BDI* - Core-Assisted Bottleneck Acceleration (CABA) framework (Section 4) with all associated overheads of performing compression (for both interconnect and memory bandwidth), (v) *Ideal-BDI* - compression (for both interconnect and memory) with no latency/power overheads for compression or decompression. This section provides our major results and analyses.

7.1. Effect on Performance and Bandwidth Utilization

Figures 8 and 9 show, respectively, the normalized performance (vs. *Base*) and the memory bandwidth utilization of the five designs. We make three major observations.




Figure 8: Normalized performance. Figure reproduced from [115].




Figure 9: Memory bandwidth utilization. Figure reproduced from [115].

First, all compressed designs are effective in providing high performance improvement over the baseline. Our approach (CABA-BDI) provides a 41.7% average improvement, which

is only 2.8% less than the ideal case (Ideal-BDI) with none of the overheads associated with CABA. CABA-BDI’s performance is 9.9% better than the previous [100] hardware-based memory bandwidth compression design (HW-BDI-Mem), and only 1.6% worse than the purely hardware-based design (HW-BDI) that performs both interconnect and memory bandwidth compression. We conclude that our framework is effective at enabling the benefits of compression without requiring specialized hardware compression and decompression logic.

Second, performance benefits, in many workloads, correlate with the reduction in memory bandwidth utilization. For a fixed amount of data, compression reduces the bandwidth utilization, and, thus, increases the effective available bandwidth. Figure 9 shows that CABA-based compression 1) reduces the average memory bandwidth utilization from 53.6% to 35.6% and 2) is effective at alleviating the memory bandwidth bottleneck in most workloads. In some applications (e.g., *bfs* and *mst*), designs that compress *both* the on-chip interconnect and the memory bandwidth, i.e. CABA-BDI and HW-BDI, perform better than the design that compresses only the memory bandwidth (HW-BDI-Mem). Hence, CABA seamlessly enables the mitigation of the interconnect bandwidth bottleneck as well, since data compression/decompression is flexibly performed at the cores.

Third, for some applications, CABA-BDI performs slightly (within 3%) better than Ideal-BDI and HW-BDI. The reason for this counter-intuitive result is the effect of warp oversubscription [14, 56, 57, 96]. In these cases, too many warps execute in parallel, polluting the last level cache. CABA-BDI sometimes reduces pollution as a side effect of performing more computation in assist warps, which slows down the progress of the parent warps.

We conclude that the CABA framework can effectively enable data compression to reduce both on-chip interconnect and off-chip memory bandwidth utilization, thereby improving the performance of modern GPGPU applications.

7.2. Effect on Energy

Compression decreases energy consumption in two ways: 1) by reducing bus energy consumption, 2) by reducing execution time. Figure 10 shows the normalized energy consumption of the five systems. We model the static and dynamic energy of the cores, caches, DRAM, and all buses (both on-chip and off-chip), as well as the energy overheads related to compression: metadata (MD) cache and compression/decompression logic. We make two major observations.




Figure 10: Normalized energy consumption. Figure reproduced from [115].

First, CABA-BDI reduces energy consumption by as much as 22.2% over the baseline. This is especially noticeable for memory-bandwidth-limited applications, e.g., *PVC*, *mst*. This is a result of two factors: (i) the reduction in the amount of data transferred between the LLC and DRAM (as a result of which we observe a 29.5% average reduction in DRAM power) and (ii) the reduction in total execution time. This observation agrees with several prior works on bandwidth compression [88, 104]. We conclude that the CABA framework is capable of reducing the overall system energy, primarily by decreasing the off-chip memory traffic.

Second, CABA-BDI’s energy consumption is only 3.6% more than that of the HW-BDI design, which uses dedicated logic for memory bandwidth compression. It is also only 4.0% more than that of the Ideal-BDI design, which has no compression-related overheads. CABA-BDI consumes more energy because it schedules and executes assist warps, utilizing on-chip register files, memory and computation units, which is less energy-efficient than using dedicated logic for compression. However, as results indicate, this additional energy cost is small compared to the performance gains of CABA (recall, 41.7% over Base), and may be amortized by using CABA for other purposes as well (see Section 8).

Power Consumption. CABA-BDI increases the system power consumption by 2.9% over the baseline (not graphed), mainly due to the additional hardware and higher utilization of the compute pipelines. However, the power overhead enables energy savings by reducing bandwidth use and can be amortized across other uses of CABA (Section 8).

Energy-Delay product. Figure 11 shows the product of the normalized energy consumption and normalized execution time for the evaluated GPU workloads. This metric simultaneously captures two metrics of interest—energy dissipation and execution delay (inverse of performance). An optimal feature would simultaneously incur low energy overhead while also reducing the execution delay. This metric is useful in capturing the efficiencies of different architectural designs and features which may expend differing amounts of energy while producing the same performance speedup or vice-versa. Hence, a lower Energy-Delay product is more desirable. We observe that CABA-BDI has a 45% lower Energy-Delay product than the baseline. This reduction comes from energy savings from reduced data transfers as well as lower execution time. On average, CABA-BDI is within only 4% of Ideal-BDI which incurs none of the energy and performance overheads of the CABA framework.

7.3. Effect of Enabling Different Compression Algorithms

The CABA framework is *not limited to a single compression algorithm*, and can be effectively used to employ other hardware-based compression algorithms (e.g., FPC [4] and C-Pack [25]). The effectiveness of other algorithms depends on two key factors: (i) how efficiently the algorithm maps to GPU instructions, (ii) how compressible the data is with the algorithm. We map the FPC and C-Pack algorithms to the CABA framework




Figure 11: Energy-Delay product. and evaluate the framework’s efficacy.

Figure 12 shows the normalized speedup with four versions of our design: *CABA-FPC*, *CABA-BDI*, *CABA-C-Pack*, and *CABA-BestOfAll* with the FPC, BDI, C-Pack compression algorithms. *CABA-BestOfAll* is an idealized design that selects and uses the best of all three algorithms in terms of compression ratio for *each cache line*, assuming no selection overhead. We make three major observations.




Figure 12: Speedup with different compression algorithms. Figure reproduced from [115].

First, CABA significantly improves performance with any compression algorithm (20.7% with FPC, 35.2% with C-Pack). Similar to CABA-BDI, the applications that benefit the most are those that are both memory-bandwidth-sensitive (Figure 9) and compressible (Figure 13). We conclude that our proposed framework, CABA, is general and flexible enough to successfully enable different compression algorithms.




Figure 13: Compression ratio of algorithms with CABA. Figure reproduced from [115].

Second, applications benefit differently from each algorithm. For example, *LPS*, *JPEG*, *MUM*, *nw* have higher compression ratios with FPC or C-Pack, whereas *MM*, *PVC*, *PVR* compress better with BDI. This motivates the necessity of having *flexible data compression* with different algorithms within the same system. Implementing multiple compression algorithms completely in hardware is expensive as it adds significant area overhead, whereas CABA can flexibly enable the use of different algorithms via its general assist warp framework.

Third, the design with the best of three compression algorithms, CABA-BestOfAll, can sometimes improve performance more than each individual design with just one compression algorithm (e.g., for *MUM* and *KM*). This happens because even within an application, different cache lines compress better with different algorithms. At the same time, different compression related overheads of different algorithms can cause one to have higher performance than another even though the latter may have a higher compression ratio. For example, CABA-BDI provides higher performance on *LPS* than CABA-FPC, even though BDI has a lower compression ratio than FPC for *LPS*, because BDI’s compression/decompression latencies are much lower than FPC’s. Hence, a mechanism that selects the best compression algorithm based on *both* compression ratio and the relative cost of compression/decompression is desirable to get the best of multiple compression algorithms. The CABA framework can flexibly enable the implementation of such a mechanism, whose design we leave for future work.

7.4. Sensitivity to Peak Main Memory Bandwidth

As described in Section 3, main memory (off-chip) bandwidth is a major bottleneck in GPU applications. In order to confirm that CABA works for different designs with varying amounts of available memory bandwidth, we conduct an experiment where CABA-BDI is used in three systems with 0.5X, 1X and 2X amount of bandwidth of the baseline.

Figure 14 shows the results of this experiment. We observe that, as expected, each CABA design (*-CABA) significantly outperforms the corresponding baseline designs with the same amount of bandwidth. The performance improvement of CABA is often equivalent to the doubling the off-chip memory bandwidth. We conclude that CABA-based bandwidth compression, on average, offers almost all the performance benefit of doubling the available off-chip bandwidth with only modest complexity to support assist warps.




Figure 14: Sensitivity of CABA to memory bandwidth. Figure reproduced from [115].

7.5. Selective Cache Compression with CABA

In addition to reducing bandwidth consumption, data compression can also increase the *effective capacity* of on-chip caches. While compressed caches can be beneficial—as higher effective cache capacity leads to lower miss rates—supporting cache compression requires several changes in the cache design [4, 25, 87, 89, 99].

Figure 15 shows the effect of four cache compression designs using CABA-BDI (applied to both L1 and L2 caches with 2x or 4x the number of tags of the baseline⁴) on performance.

⁴The number of tags limits the effective compressed cache size [4, 87].

We make two major observations. First, several applications from our workload pool are not only bandwidth sensitive, but also cache capacity sensitive. For example, *bfs* and *sssp* significantly benefit from L1 cache compression, while *TRA* and *KM* benefit from L2 compression. Second, L1 cache compression can severely degrade the performance of some applications, e.g., *hw* and *LPS*. The reason for this is the overhead of decompression, which can be especially high for L1 caches as they are accessed very frequently. This overhead can be easily avoided by disabling compression at any level of the memory hierarchy.




Figure 15: Speedup of cache compression with CABA.

7.6. Other Optimizations

We also consider several other optimizations of the CABA framework for data compression: (i) avoiding the overhead of decompression in L2 by storing data in the uncompressed form and (ii) optimized load of *only useful* data.

Uncompressed L2. The CABA framework allows us to store compressed data selectively at different levels of the memory hierarchy. We consider an optimization where we avoid the overhead of decompressing data in L2 by storing data in uncompressed form. This provides another tradeoff between the savings in on-chip traffic (when data in L2 is compressed – default option), and savings in decompression latency (when data in L2 is uncompressed). Figure 16 depicts the performance benefits from this optimization. Several applications in our workload pool (e.g., *RAY*) benefit from storing data uncompressed as these applications have high hit rates in the L2 cache. We conclude that offering the choice of enabling or disabling compression at different levels of the memory hierarchy can provide higher levels of the software stack (e.g., applications, compilers, runtime system, system software) with an additional performance knob.

Uncoalesced requests. Accesses by scalar threads from the same warp are coalesced into fewer memory transactions [82]. If the requests from different threads within a warp span two or more cache lines, multiple lines have to be retrieved and decompressed before the warp can proceed its execution. Uncoalesced requests can significantly increase the number of assist warps that need to be executed. An alternative to decompressing each cache line (when only a few bytes from each line may be required), is to enhance the coalescing unit to supply only the correct *deltas* from within each compressed cache line. The logic that maps bytes within a cache line to the appropriate registers will need to be enhanced to take into account the encoding of the compressed

line to determine the size of the *base* and the *deltas*. As a result, we do not decompress the entire cache lines and only extract the data that is needed. In this case, the cache line is not inserted into the L1D cache in the uncompressed form, and hence every line needs to be decompressed even if it is found in the L1D cache.⁵ *Direct-Load* in Figure 16 depicts the performance impact from this optimization. The overall performance improvement is 2.5% on average across all applications (as high as 4.6% for *MM*).




Figure 16: Effect of different optimizations (Uncompressed data in L2 and Direct Load) on applications’ performance.

8. Other Uses of the CABA Framework

The CABA framework can be employed in various ways to alleviate system bottlenecks and increase system performance and energy efficiency. In this section, we discuss two other potential applications of CABA, focusing on two: *Memoization* and *Prefetching*. We leave the detailed evaluations and analysis of these use cases of CABA to future work.

8.1. Memoization

Hardware memoization is a technique used to avoid redundant computations by reusing the results of previous computations that have the same or similar inputs. Prior work [8, 13, 98] observed redundancy in inputs to data in GPU workloads. In applications limited by available compute resources, memoization offers an opportunity to trade off computation for storage, thereby enabling potentially higher energy efficiency and performance. In order to realize memoization in hardware, a look-up table (LUT) is required to dynamically cache the results of computations as well as the corresponding inputs. The granularity of computational reuse can be at the level of fragments [13], basic blocks or functions [7, 9, 29, 45, 106], or long-latency instructions [26]. The CABA framework provides a natural way to implement such an optimization. The availability of on-chip memory lends itself for use as the LUT. In order to cache previous results in on-chip memory, look-up tags (similar to those proposed in [41]) are required to index correct results. With applications tolerant of approximate results (e.g., image processing, machine learning, fragment rendering kernels), the computational inputs can be hashed to reduce the size of the LUT. Register values, texture/constant memory or global memory sections that are not subject to change are potential inputs. An assist warp can be employed

⁵This optimization also benefits cache lines that might *not* have many uncoalesced accesses, but have poor data reuse in the L1D.

to perform memoization in the following way: (1) compute the hashed value for look-up at predefined trigger points, (2) use the load/store pipeline to save these inputs in available shared memory, and (3) eliminate redundant computations by loading the previously computed results in the case of a hit in the LUT.

8.2. Prefetching

Prefetching has been explored in the context of GPUs [12, 51, 52, 63, 64, 73, 102] with the goal of reducing effective memory latency. With memory-latency-bound applications, the load/store pipelines can be employed by the CABA framework to perform opportunistic prefetching into GPU caches. The CABA framework can potentially enable the effective use of prefetching in GPUs due to several reasons: (1) Even simple prefetchers such as the stream [54, 86, 109] or stride [15, 38] prefetchers are non-trivial to implement in GPUs since access patterns need to be tracked and trained at the fine granularity of warps [64, 102]. CABA could enable fine-grained book-keeping by using spare registers and assist warps to save metadata for each warp. The computational units could then be used to continuously compute strides in access patterns both within and across warps. (2) It has been demonstrated that software prefetching and helper threads [1, 19, 28, 48, 48, 63, 69, 77, 78, 111] are very effective in performing prefetching for irregular access patterns. Assist warps offer the hardware/software interface to implement application-specific prefetching algorithms with varying degrees of complexity without the additional overheads of various hardware implementations. (3) In bandwidth-constrained GPU systems, uncontrolled prefetching could potentially flood the off-chip buses, delaying demand requests. CABA can enable flexible prefetch throttling (e.g., [34, 36, 109]) by scheduling assist warps that perform prefetching, *only* when the memory pipelines are idle or underutilized. (4) Prefetching with CABA entails using load or prefetch instructions, which not only enables prefetching to the hardware-managed caches, but also could simplify the usage of unutilized shared memory or register file as prefetch buffers.

8.3. Other Uses

Redundant Multithreading. Reliability of GPUs is a key concern, especially today when they are popularly employed in many supercomputing systems. Ensuring hardware protection with dedicated resources can be expensive [71]. Redundant multithreading [76, 93, 116, 117] is an approach where redundant threads are used to replicate program execution. The results are compared at different points in execution to detect and potentially correct errors. The CABA framework can be extended to redundantly execute portions of the original program via the use of such approaches to increase the reliability of GPU architectures.

Speculative Precomputation. In CPUs, speculative multithreading ([72, 92, 107]) has been proposed to speculatively parallelize serial code and verify the correctness later. Assist

warps can be employed in GPU architectures to speculatively pre-execute sections of code during idle cycles to further improve parallelism in the program execution. Applications tolerant to approximate results could particularly be amenable towards this optimization [119].

Handling Interrupts and Exceptions. Current GPUs do not implement support for interrupt handling except for some support for timer interrupts used for application time-slicing [84]. CABA offers a natural mechanism for associating architectural events with subroutines to be executed in throughput-oriented architectures where thousands of threads could be active at any given time. Interrupts and exceptions can be handled by special assist warps, without requiring complex context switching or heavy-weight kernel support.

Profiling and Instrumentation. Profiling and binary instrumentation tools like Pin [70] and Valgrind [80] proved to be very useful for development, performance analysis and debugging on modern CPU systems. At the same time, there is a lack⁶ of tools with same/similar capabilities for modern GPUs. This significantly limits software development and debugging for modern GPU systems. The CABA framework can potentially enable easy and efficient development of such tools, as it is flexible enough to invoke user-defined code on specific architectural events (e.g., cache misses, control divergence).

9. Related Work

To our knowledge, this paper is the first to (1) propose a flexible and general framework for employing idle GPU resources for useful computation that can aid regular program execution, and (2) use the general concept of *helper threading* to perform memory and interconnect bandwidth compression. We demonstrate the benefits of our new framework by using it to implement multiple compression algorithms on a throughput-oriented GPU architecture to alleviate the memory bandwidth bottleneck. In this section, we discuss related works in helper threading, and memory bandwidth optimizations, and memory compression.

Helper Threading. Previous works [1, 19, 22, 23, 27, 28, 31, 32, 48, 55, 59, 68, 69, 77, 78, 111, 120, 121, 122] demonstrated the use of *helper threads* in the context of Simultaneous Multithreading (SMT) and multi-core and single-core processors, primarily to speed up single-thread execution by using idle SMT contexts, idle cores in CPUs, or idle cycles during which the main program is stalled on a single thread context. These works typically use helper threads (generated by the software, the hardware, or cooperatively) to pre-compute useful information that aids the execution of the primary thread (e.g., by prefetching, branch outcome pre-computation, and cache management). No previous work discussed the use of helper threads for memory/interconnect bandwidth compression or cache compression.

These works primarily use helper threads to capture data flow and pre-compute useful information to aid in the execution of the primary thread. In these prior works, helper threads

are either generated with the help of the compiler [59, 68, 120] or completely in hardware [27, 77, 78]. These threads are used to perform prefetching in [1, 19, 22, 28, 48, 48, 60, 69, 77, 78, 111] where the helper threads predict future load addresses by doing some computation and then prefetch the corresponding data. Simultaneous Subordinate Multithreading [22] employs hardware generated helper threads to improve branch prediction accuracy and cache management. Speculative multithreading [62, 72, 92, 107] involves executing different sections of the serial program in parallel and then later verifying the run-time correctness. Assisted execution [31, 32] is an execution paradigm where tightly-coupled *nanothreads* are generated using *nanotrap handlers* and execute routines to enable optimizations like prefetching. In Slipstream Processors [48], one thread runs ahead of the program and executes a reduced version of the program. In runahead execution [77, 78], the main thread is executed speculatively solely for prefetching purposes when the program is stalled due to a cache miss.

While our work was inspired by these prior studies of helper threading in latency-oriented architectures (CPUs), developing a framework for helper threading (or *assist warps*) in throughput-oriented architectures (GPUs) enables new opportunities and poses new challenges, both due to the massive parallelism and resources present in a throughput-oriented architecture (as discussed in Section 1). Our CABA framework exploits these new opportunities and addresses these new challenges, including (1) low-cost management of a large number of assist warps that could be running concurrently with regular program warps, (2) means of state/context management and scheduling for assist warps to maximize effectiveness and minimize interference, and (3) different possible applications of the concept of assist warps in a throughput-oriented architecture.

In the GPU domain, CudaDMA [18] is a recent proposal that aims to ease programmability by decoupling execution and memory transfers with specialized DMA warps. This work does *not* provide a general and flexible hardware-based framework for using GPU cores to run warps that aid the main program.

Memory Latency and Bandwidth Optimizations in GPUs. A lot of prior works focus on optimizing for memory bandwidth and memory latency in GPUs. Jog et al. [51] aim to improve memory latency tolerance by coordinating prefetching and warp scheduling policies. Lakshminarayana et al. [63] reduce effective latency in graph applications by using spare registers to store prefetched data. In OWL [52] and [56], intelligent scheduling is used to improve DRAM bank-level parallelism and bandwidth utilization and Rhu et al. [95] propose a locality-aware memory to improve memory throughput. Kayiran et al. [56] propose GPU throttling techniques to reduce memory contention in heterogeneous systems. Ausavarangniran et al. [14] leverage heterogeneity in warp behavior to design more intelligent policies at the cache and memory controller. These works do not consider data compression and are orthogonal to our proposed framework.

Compression. Several prior works [6, 11, 88, 89, 90, 91, 100, 104, 114] study memory and cache compression with

⁶With the exception of one recent work [110].

several different compression algorithms [4, 11, 25, 49, 87, 118], in the context of CPUs or GPUs.

Alameldeen et al. [6] investigated the possibility of bandwidth compression with FPC [5]. The authors show that significant decrease in pin bandwidth demand can be achieved with FPC-based bandwidth compression design. Sathish et al. [100] examine the GPU-oriented memory link compression using C-Pack [25] compression algorithm. The authors make the observation that GPU memory (GDDR3 [46]) indeed allows transfer of data in small bursts and propose to store data in compressed form in memory, but without capacity benefits. Thuresson et al. [114] consider a CPU-oriented design where a compressor/decompressor logic is located on the both ends of the main memory link. Pekhimenko et al. [88] propose Linearly Compressed Pages (LCP) with the primary goal of compressing main memory to increase capacity.

Our work is the first to demonstrate how one can adapt some of these algorithms for use in a general helper threading framework for GPUs. As such, compression/decompression using our new framework is more flexible since it does not require a specialized hardware implementation for any algorithm and instead utilizes the existing GPU core resources to perform compression and decompression. Finally, as discussed in Section 8, our CABA framework is applicable beyond compression and can be used for other purposes.

10. Conclusion

This paper makes a case for the Core-Assisted Bottleneck Acceleration (CABA) framework, which employs *assist warps* to alleviate different bottlenecks in GPU execution. CABA is based on the key observation that various imbalances and bottlenecks in GPU execution leave on-chip resources, i.e., computational units, register files and on-chip memory, underutilized. CABA takes advantage of these idle resources and employs them to perform useful work that can aid the execution of the main program and the system.

We provide a detailed design and analysis of how CABA can be used to perform flexible data compression in GPUs to mitigate the memory bandwidth bottleneck. Our extensive evaluations across a variety of workloads and system configurations show that the use of CABA for memory compression significantly improves system performance (by 41.7% on average on a set of bandwidth-sensitive GPU applications) by reducing the memory bandwidth requirements of both the on-chip and off-chip buses.

We conclude that CABA is a general substrate that can alleviate the memory bandwidth bottleneck in modern GPU systems by enabling flexible implementations of data compression algorithms. We believe CABA is a general framework that can have a wide set of use cases to mitigate many different system bottlenecks in throughput-oriented architectures, and we hope that future work explores both new uses of CABA and more efficient implementations of it.

Acknowledgments

We thank the reviewers for their valuable suggestions. We thank the members of the SAFARI group for their feedback and the stimulating research environment they provide. Special thanks to Evgeny Bolotin and Kevin Hsieh for their feedback during various stages of this project. We acknowledge the support of our industrial partners: Facebook, Google, IBM, Intel, Microsoft, Nvidia, Qualcomm, VMware, and Samsung. This research was partially supported by NSF (grants 0953246, 1065112, 1205618, 1212962, 1213052, 1302225, 1302557, 1317560, 1320478, 1320531, 1409095, 1409723, 1423172, 1439021, 1439057), the Intel Science and Technology Center for Cloud Computing, and the Semiconductor Research Corporation. Gennady Pekhimenko is supported in part by a Microsoft Research Fellowship and an Nvidia Graduate Fellowship. Rachata Ausavarungnirun is supported in part by the Royal Thai Government scholarship. This article is a revised and extended version of our previous ISCA 2015 paper [115].

References

- [1] T. M. Aamodt et al. Hardware support for prescient instruction prefetch. In *HPCA*, 2004.
- [2] B. Abali et al. Memory Expansion Technology (MXT): Software Support and Performance. *IBM J.R.D.*, 2001.
- [3] M. Abdel-Majeed et al. Warped register file: A power efficient register file for gpgpus. In *HPCA*, 2013.
- [4] A. Alameldeen et al. Adaptive Cache Compression for High-Performance Processors. In *ISCA*, 2004.
- [5] A. Alameldeen et al. Frequent Pattern Compression: A Significance-Based Compression Scheme for L2 Caches. Technical report, U. Wisconsin, 2004.
- [6] A. Alameldeen et al. Interactions between compression and prefetching in chip multiprocessors. In *HPCA*, 2007.
- [7] C. Alvarez et al. On the potential of tolerant region reuse for multimedia applications. In *ICS*, 2001.
- [8] C. Alvarez et al. Fuzzy memoization for floating-point multimedia applications. *IEEE Trans. Comput.*, 2005.
- [9] C. Alvarez et al. Dynamic tolerance region computing for multimedia. *IEEE Trans. Comput.*, 2012.
- [10] AMD. Radeon GPUs. <http://www.amd.com/us/products/desktop/graphics/amd-radeon-hd-6000/Pages/amd-radeon-hd-6000.aspx>.
- [11] A. Arelakis et al. SC2: A Statistical Compression Cache Scheme. In *ISCA*, 2014.
- [12] J. Arnau et al. Boosting mobile GPU performance with a decoupled access/execute fragment processor. In *ISCA*, 2012.
- [13] J. Arnau et al. Eliminating Redundant Fragment Shader Executions on a Mobile GPU via Hardware Memoization. In *ISCA*, 2014.
- [14] R. Ausavarangnirun et al. Exploiting Inter-Warp Heterogeneity to Improve GPGPU Performance. In *PACT*, 2014.
- [15] J. Baer et al. Effective hardware-based data prefetching for high-performance processors. *IEEE Trans. Comput.*, 1995.
- [16] A. Bakhoda et al. Analyzing CUDA Workloads Using a Detailed GPU Simulator. In *ISPASS*, 2009.
- [17] A. Bakhoda, G.L. Yuan, W.W.L. Fung, H. Wong, and T.M. Aamodt. Analyzing CUDA Workloads Using a Detailed GPU Simulator. In *ISPASS*, 2009.
- [18] M. Bauer et al. CudaDMA: Optimizing GPU memory bandwidth via warp specialization. In *SC*, 2011.
- [19] J. A. Brown et al. Speculative precomputation on chip multiprocessors. In *MTEAC*, 2001.
- [20] M. Burtscher et al. A quantitative study of irregular programs on gpus. In *IISWC*, 2012.
- [21] D. De Caro et al. High-performance special function unit for programmable 3-d graphics processors. *Trans. Cir. Sys. Part I*, 2009.
- [22] R. S. Chappell et al. Simultaneous subordinate microthreading (SSMT). In *ISCA*, 1999.
- [23] R. S. Chappell et al. Microarchitectural support for precomputation microthreads. In *MICRO*, 2002.
- [24] S. Che et al. Rodinia: A Benchmark Suite for Heterogeneous Computing. In *IISWC*, 2009.
- [25] X. Chen et al. C-pack: A high-performance microprocessor cache compression algorithm. In *IEEE Trans. on VLSI Systems*, 2010.
- [26] D. Citron et al. Accelerating multi-media processing by implementing memoing in multiplication and division units. In *ASPLOS*, 1998.
- [27] J. D. Collins et al. Dynamic speculative precomputation. In *MICRO*, 2001.
- [28] J. D. Collins et al. Speculative Precomputation: Long-range Prefetching of Delinquent Loads. *ISCA*, 2001.
- [29] D. A. Connors et al. Compiler-directed dynamic computation reuse: rationale and initial results. In *MICRO*, 1999.
- [30] R. Das et al. Performance and Power Optimization through Data Compression in Network-on-Chip Architectures. In *HPCA*, 2008.
- [31] M. Dubois. Fighting the memory wall with assisted execution. In *CF*, 2004.
- [32] M. Dubois et al. Assisted execution. Technical report, USC, 1998.
- [33] J. Dusser et al. Zero-content augmented caches. In *ICS*, 2009.
- [34] E. Ebrahimi et al. Coordinated Control of Multiple Prefetchers in Multi-core Systems. In *MICRO*, 2009.
- [35] E. Ebrahimi et al. Techniques for Bandwidth-efficient Prefetching of Linked Data Structures in Hybrid Prefetching Systems. In *HPCA*, 2009.
- [36] E. Ebrahimi et al. Prefetch-aware shared resource management for multi-core systems. *ISCA*, 2011.
- [37] M. Ekman et al. A Robust Main-Memory Compression Scheme. In *ISCA-32*, 2005.
- [38] J. W. C. Fu et al. Stride directed prefetching in scalar processors. In *MICRO*, 1992.
- [39] M. Gebhart et al. A compile-time managed multi-level register file hierarchy. In *MICRO*, 2011.
- [40] M. Gebhart et al. Energy-efficient Mechanisms for Managing Thread Context in Throughput Processors. In *ISCA*, 2011.
- [41] M. Gebhart et al. Unifying primary cache, scratch, and register file memories in a throughput processor. In *MICRO*, 2012.
- [42] M. Ghosh et al. Smart refresh: An enhanced memory controller design for reducing energy in conventional and 3d die-stacked drams. *MICRO*, 2007.
- [43] GPGPU-Sim v3.2.1. GPGPU-Sim Manual.
- [44] B. He et al. Mars: A MapReduce Framework on Graphics Processors. In *PACT*, 2008.
- [45] J. Huang et al. Exploiting basic block value locality with block reuse. In *HPCA*, 1999.
- [46] Hynix. 512M (16mx32) GDDR3 SDRAM hy5rs123235fp.
- [47] Hynix. Hynix GDDR5 SGRAM Part H5GQ1H24AFR Revision 1.0.
- [48] K. Z. Ibrahim et al. Slipstream execution mode for cmp-based multiprocessors. In *HPCA*, 2003.
- [49] M. Islam et al. Zero-Value Caches: Cancelling Loads that Return Zero. In *PACT*, 2009.
- [50] ITRS. International technology roadmap for semiconductors. 2011.
- [51] A. Jog et al. Orchestrated Scheduling and Prefetching for GPGPUs. In *ISCA*, 2013.
- [52] A. Jog et al. OWL: Cooperative Thread Array Aware Scheduling Techniques for Improving GPGPU Performance. In *ASPLOS*, 2013.
- [53] John L. Hennessy and David A. Patterson. *Computer Architecture, A Quantitative Approach*. Morgan Kaufmann, 2010.
- [54] N. Jouppi. Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers. In *ISCA*, 1990.
- [55] M. Kamruzzaman et al. Inter-core Prefetching for Multicore Processors Using Migrating Helper Threads. In *ASPLOS*, 2011.
- [56] O. Kayiran et al. Neither More Nor Less: Optimizing Thread-Parallelism for GPGPUs. In *PACT*, 2013.
- [57] O. Kayiran et al. Managing GPU Concurrency in Heterogeneous Architectures. In *MICRO*, 2014.
- [58] S. W. Keckler et al. GPUs and the future of parallel computing. *IEEE Micro*, 2011.
- [59] D. Kim et al. Design and Evaluation of Compiler Algorithms for Pre-execution. In *ASPLOS*, 2002.
- [60] Dongkeun Kim et al. Physical experimentation with prefetching helper threads on intel's hyper-threaded processors. In *CGO*, 2004.
- [61] David B. Kirk and W. Hwu. *Programming massively parallel processors: a hands-on approach*. Morgan Kaufmann, 2010.
- [62] Venkata Krishnan and Josep Torrellas. A chip-multiprocessor architecture with speculative multithreading. *IEEE Trans. Comput.*, 1999.
- [63] N. Lakshminarayana et al. Spare register aware prefetching for graph algorithms on GPUs. In *HPCA*, 2014.
- [64] J. Lee et al. Many-Thread Aware Prefetching Mechanisms for GPGPU Applications. In *MICRO*, 2010.
- [65] J. Leng et al. GPUWattch: Enabling Energy Optimizations in GPGPUs. In *ISCA*, 2013.
- [66] E. Lindholm et al. Nvidia tesla: A unified graphics and computing architecture. *IEEE Micro*, 2008.
- [67] J. Liu et al. Raidr: Retention-aware intelligent dram refresh. *ISCA*, 2012.
- [68] J. Lu et al. Dynamic Helper Threaded Prefetching on the Sun Ultra-SPARC CMP Processor. In *MICRO*, 2005.
- [69] C. Luk. Tolerating memory latency through software-controlled pre-execution in simultaneous multithreading processors. In *ISCA*, 2001.
- [70] C. Luk et al. Pin: Building Customized Program Analysis Tools with Dynamic Instrumentation. In *PLDI*, 2005.
- [71] Y. Luo et al. Characterizing Application Memory Error Vulnerability to Optimize Data Center Cost via Heterogeneous-Reliability Memory. *DSN*, 2014.
- [72] Pedro Marcuello et al. Speculative multithreaded processors. In *ICS*, 1998.
- [73] J. Meng et al. Dynamic warp subdivision for integrated branch and memory divergence tolerance. In *ISCA*, 2010.
- [74] J. Meza et al. Enabling efficient and scalable hybrid memories. *IEEE CAL*, 2012.
- [75] A. Moshovos et al. Slice-processors: An implementation of operation-based prediction. In *ICS*, 2001.
- [76] S. Mukherjee et al. Detailed design and evaluation of redundant multithreading alternatives. *ISCA*.
- [77] O. Mutlu et al. Runahead execution: An alternative to very large instruction windows for out-of-order processors. In *HPCA*, 2003.
- [78] O. Mutlu et al. Techniques for efficient processing in runahead execution engines. *ISCA*, 2005.
- [79] V. Narasiman et al. Improving GPU performance via large warps and two-level warp scheduling. In *MICRO*, 2011.
- [80] N. Nethercote et al. Valgrind: A Framework for Heavyweight Dynamic Binary Instrumentation. In *PLDI*, 2007.
- [81] B. S. Nordquist et al. Apparatus, system, and method for coalescing parallel memory requests, 2009. US Patent 7,492,368.
- [82] NVIDIA. Programming Guide.
- [83] NVIDIA. CUDA C/C++ SDK Code Samples, 2011.
- [84] NVIDIA. Fermi: NVIDIA's Next Generation CUDA Compute Architecture, 2011.
- [85] L. Nyland et al. Systems and methods for coalescing memory accesses of parallel threads, 2011. US Patent 8,086,806.

- [86] S. Palacharla et al. Evaluating stream buffers as a secondary cache replacement. In *ISCA*, 1994.
- [87] G. Pekhimenko et al. Base-Delta-Immediate Compression: Practical Data Compression for On-Chip Caches. In *PACT*, 2012.
- [88] G. Pekhimenko et al. Linearly Compressed Pages: A Low Complexity, Low Latency Main Memory Compression Framework. In *MICRO*, 2013.
- [89] G. Pekhimenko et al. Exploiting Compressed Block Size as an Indicator of Future Reuse. In *HPCA*, 2015.
- [90] G. Pekhimenko et al. Toggle-Aware Compression for GPUs. In *IEEE CAL*, 2015.
- [91] G. Pekhimenko et al. A Case for Toggle-Aware Compression in GPUs. In *HPCA*, 2016.
- [92] Carlos García Quiñones et al. Mitosis Compiler: An Infrastructure for Speculative Threading Based on Pre-computation Slices. In *PLDI*, 2005.
- [93] M. Qureshi et al. Microarchitecture-based introspection: A technique for transient-fault tolerance in microprocessors. *DSN*, 2005.
- [94] M. Qureshi et al. Fundamental latency trade-off in architecting dram caches: Outperforming impractical sram-tags with a simple and practical design. *MICRO*, 2012.
- [95] Minsoo Rhu, Michael Sullivan, Jingwen Leng, and Mattan Erez. A Locality-Aware Memory Hierarchy for Energy-Efficient GPU Architectures. In *MICRO*, 2013.
- [96] T. G. Rogers et al. Cache-Conscious Wavefront Scheduling. In *MICRO*, 2012.
- [97] A. Roth et al. Speculative data-driven multithreading. In *HPCA*, 2001.
- [98] M. Samadi et al. Sage: Self-tuning approximation for graphics engines. In *MICRO*, 2013.
- [99] S. Sardashti et al. Decoupled Compressed Cache: Exploiting Spatial Locality for Energy-optimized Compressed Caching. In *MICRO*, 2013.
- [100] V. Sathish et al. Lossless and Lossy Memory I/O Link Compression for Improving Performance of GPGPU Workloads. In *PACT*, 2012.
- [101] V. Seshadri et al. Page overlays: An enhanced virtual memory framework to enable fine-grained memory management. *ISCA*, 2015.
- [102] A. Sethia et al. Apogee: adaptive prefetching on gpus for energy efficiency. In *PACT*, 2013.
- [103] A. Sethia et al. Equalizer: Dynamic tuning of gpu resources for efficient execution. In *MICRO*, 2014.
- [104] A. Shafiee et al. MemZip: Exploring Unconventional Benefits from Memory Compression. In *HPCA*, 2014.
- [105] B. Smith. A pipelined, shared resource MIMD computer. *Advance Computer Architecture*, 1986.
- [106] A. Sodani et al. Dynamic Instruction Reuse. In *ISCA*, 1997.
- [107] Sohi et al. Multiscalar Processors. In *ISCA*, 1995.
- [108] S. Srinath et al. Feedback Directed Prefetching: Improving the Performance and Bandwidth-Efficiency of Hardware Prefetchers. In *HPCA*, 2007.
- [109] S. Srinath et al. Feedback Directed Prefetching: Improving the Performance and Bandwidth-Efficiency of Hardware Prefetchers. In *HPCA*, 2007.
- [110] M. Stephenson et al. Flexible software profiling of GPU architectures. In *ISCA*, 2015.
- [111] K. Sundaramoorthy et al. Slipstream Processors: Improving Both Performance and Fault Tolerance. In *ASPLOS*, 2000.
- [112] J. E. Thornton. Parallel Operation in the Control Data 6600. *Proceedings of the AFIPS FJCC*, 1964.
- [113] S. Thoziyoor et al. CACTI 5.1. Technical Report HPL-2008-20, HP Laboratories, 2008.
- [114] M. Thuresson et al. Memory-Link Compression Schemes: A Value Locality Perspective. *IEEE Trans. Comput.*, 2008.
- [115] N. Vijaykumar et al. A Case for Core-Assisted Bottleneck Acceleration in GPUs: Enabling Flexible Data Compression with Assist Warps. In *ISCA*, 2015.
- [116] J. Wadden et al. Real-world design and evaluation of compiler-managed gpu redundant multithreading. *ISCA '14*, 2014.
- [117] Wang et al. Compiler-managed software-based redundant multithreading for transient fault detection. *CGO*, 2007.
- [118] J. Yang et al. Frequent Value Compression in Data Caches. In *MICRO*, 2000.
- [119] A. Yazdanbakhsh et al. Mitigating the memory bottleneck with approximate load value prediction. *IEEE Date and Test*, 2016.
- [120] W. Zhang et al. Accelerating and adapting precomputation threads for efficient prefetching. In *HPCA*, 2007.
- [121] C. Zilles et al. The use of multithreading for exception handling. *MICRO*, 1999.
- [122] C. Zilles et al. Execution-based Prediction Using Speculative Slices. In *ISCA*, 2001.