

Holistic Management of the GPGPU Memory Hierarchy to Manage Warp-level Latency Tolerance

Rachata Ausavarungnirun¹ Saugata Ghose¹ Onur Kayiran^{2,3}
Gabriel H. Loh² Chita R. Das³ Mahmut T. Kandemir³ Onur Mutlu^{4,1}

¹Carnegie Mellon University ²AMD Research ³Pennsylvania State University ⁴ETH Zürich

This paper summarizes the idea of Memory Divergence Correction (MeDiC), which was published at PACT 2015 [6], and examines the work’s significance and future potential. In a modern GPU architecture, all threads within a warp execute the same instruction in lockstep. For a memory instruction, this can lead to memory divergence: the memory requests for some threads are serviced early, while the remaining requests incur long latencies. This divergence stalls the warp, as it cannot execute the next instruction until all requests from the current instruction complete.

In this work, we make three new observations. First, GPGPU warps exhibit heterogeneous memory divergence behavior at the shared cache: some warps have most of their requests hit in the cache (high cache utility), while other warps see most of their request miss (low cache utility). Second, a warp retains the same divergence behavior for long periods of execution. Third, due to high memory level parallelism, requests going to the shared cache can incur queuing delays as large as hundreds of cycles, exacerbating the effects of memory divergence.

We propose a set of techniques, collectively called Memory Divergence Correction (MeDiC), that reduce the negative performance impact of memory divergence and cache queuing. MeDiC uses online warp divergence characterization to guide three components: (1) a cache bypassing mechanism that exploits the latency tolerance of low cache utility warps to both alleviate queuing delay and increase the hit rate for high cache utility warps, (2) a cache insertion policy that prevents data from high cache utility warps from being prematurely evicted, and (3) a memory controller that prioritizes the few requests received from high cache utility warps to minimize stall time. We compare MeDiC to four cache management techniques, and find that it delivers an average speedup of 21.8%, and 20.1% higher energy efficiency, over a state-of-the-art GPU cache management mechanism across 15 different GPGPU applications.

1. Introduction

Graphics Processing Units (GPUs) have enormous parallel processing power to leverage thread-level parallelism. GPU applications are usually broken down into thousands of threads, allowing GPUs to use *fine-grained multithreading* [128, 136] to prevent GPU cores from stalling due to dependencies and long memory latencies. Ideally, there should always be available threads for GPU cores to continue execution, preventing stalls within the core. GPUs also take advantage of the *SIMD* (Single Instruction, Multiple Data)

execution model [30]. The thousands of threads within a GPU application are clustered into *thread blocks*, each of which contains multiple smaller bundles (*warps*) of threads that run concurrently. Each thread in a warp executes the same instruction on a different piece of data. A warp completes an instruction when all threads in the warp complete the instruction.

While many GPGPU applications can tolerate a significant amount of memory latency due to their parallelism and the use of fine-grained multithreading, *memory divergence* (where the threads of a warp reach a memory instruction, and some of the threads’ memory requests take longer to service than others) can significantly increase the stall time of a warp [51, 52, 63, 75, 89, 101, 116, 117, 155]. Because all threads within a warp operate in lockstep due to the SIMD execution model, the warp cannot proceed to the next instruction until the *slowest* request within the warp completes. Figures 1a and 1b show examples of memory divergence within a warp. Figure 1a shows a *mostly-hit warp*, where most of the warp’s memory accesses hit in the cache (❶). Only a single access misses in the cache and must go to main memory (❷). As a result, the *entire warp* is stalled until the much longer cache miss completes. Figure 1b shows a *mostly-miss warp*, where most of the warp’s memory requests miss in the cache (❸), resulting in many accesses to main memory. Even though some requests are cache hits (❹), these do not benefit the execution time of the warp since the execution of the warp ends when the slowest thread finishes the instruction.

Based on these three observations, we aim to devise a mechanism that has two major goals: (1) convert mostly-hit warps into *all-hit warps* (warps where *all* requests hit in the cache, as shown in Figure 1c), and (2) convert mostly-miss warps into *all-miss warps* (warps where *none* of the requests hit in the cache, as shown in Figure 1d). As we can see in Figure 1a, the stall time due to memory divergence for the mostly-hit warp can be eliminated by converting only the single cache miss (❷) into a hit. Doing so requires additional cache space. If we convert the two cache hits of the mostly-miss warp (Figure 1b, ❹) into cache misses, we can allocate the cache space previously used by these hits to the mostly-hit warp, thus converting the mostly-hit warp into an all-hit warp. Though the mostly-miss warp is now an all-miss warp (Figure 1d), it incurs no extra stall penalty, as the warp was already waiting on the other six cache misses to complete.

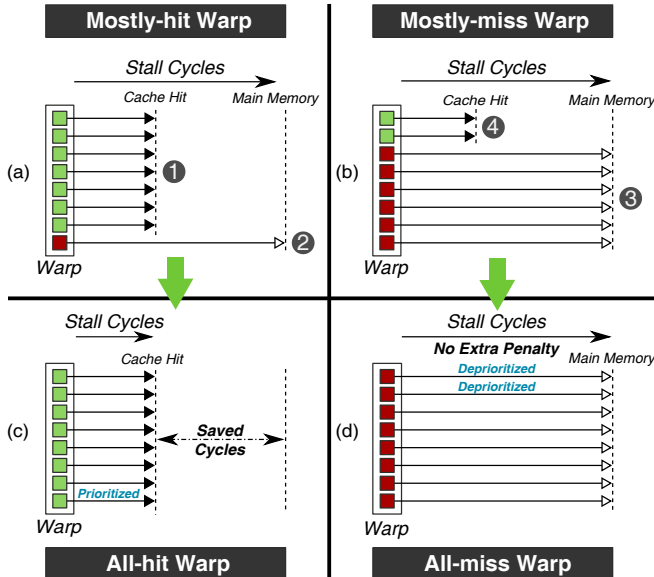


Figure 1: Memory divergence within a warp. (a) and (b) show the heterogeneity between *mostly-hit* and *mostly-miss* warps, respectively. (c) and (d) show the change in stall time from converting *mostly-hit* warps into *all-hit* warps, and *mostly-miss* warps into *all-miss* warps, respectively. Reproduced from [6].

Moreover, now that it is an all-miss warp, we can predict that its future memory requests will also not be in the L2 cache. Based on this prediction, we can simply have these requests *bypass the cache*. By doing so, the requests from the all-miss warp can completely avoid unnecessary L2 access and queuing delays, and enable the use of L2 cache bandwidth and buffer space by warps that benefit from the L2 cache. This decreases the total number of requests going to the L2 cache, thus reducing the queuing latencies for requests from *mostly-hit* and *all-hit* warps, as there is less contention.

2. Observation on GPU Memory Divergence

We make three new key observations about memory divergence (at the shared L2 cache). First, we observe that the degree of memory divergence can differ across warps (as illustrated in Figure 1). This inter-warp heterogeneity affects how well each warp takes advantage of the shared cache. Second, we observe that a warp’s memory divergence behavior tends to remain stable for long periods of execution, making it predictable. Third, we observe that requests to the shared cache experience long queuing delays due to the large amount of parallelism in GPGPU programs, which exacerbates the memory divergence problem and slows down GPU execution. Next, we describe each of these observations in detail and motivate our solutions.

2.1. Memory Divergence Heterogeneity

There is *heterogeneity across warps* in the degree of memory divergence experienced by each warp at the shared L2 cache. Figures 1a and 1b show examples of two different *types of warps* that exhibit different degrees of memory divergence.

We observe that different warps have different amounts of sensitivity to memory latency and cache utilization. We study the cache utilization of a warp by determining its *hit ratio*, the percentage of memory requests that hit in the cache when the warp issues a single memory instruction. As Figure 2 shows, the warps from each of our three representative GPGPU applications are distributed across all possible ranges of *hit ratio*, exhibiting significant heterogeneity. To better characterize warp behavior, we break the warps down into the five types shown in Figure 3 based on their hit ratios: *all-hit*, *mostly-hit*, *balanced*, *mostly-miss*, and *all-miss*.




Figure 2: L2 cache hit ratio of different warps in three representative GPGPU applications. Reproduced from [6].



Figure 3: Warp type categorization based on the shared cache hit ratios. Hit ratio values are empirically chosen. Reproduced from [6].

MeDiC provide two mechanisms, warp-type-aware cache bypassing and warp-type-aware cache insertion policy, in order to convert *mostly-hit* warps into *all-hit* warps, where all requests in the warp hit in the cache, thereby reducing the stall time of *mostly-hit* warp significantly. This is done at the cost of converting the *mostly-miss* warps into *all-miss* warps, but doing so does not increase the stall time of such warps. To speed up uncacheable cache misses from *mostly-hit* warps, the warp-type-aware memory scheduling policy in MeDiC prioritizes memory requests from *mostly-hit* warps over memory requests from *mostly-miss* warps.

2.2. Memory Divergence Stability Over Time

A warp tends to retain its memory divergence behavior (e.g., whether or not it is *mostly-hit* or *mostly-miss*) for long periods of execution, and is thus predictable. This is due to the spatial and temporal locality of each thread within the warp. Figure 4 shows a sample of warps from a representative application (i.e., BFS [10]) that shows this predictability. This predictability enables us to perform history-based warp divergence characterization.




Figure 4: Hit ratio of randomly selected warps over time from BFS. Reproduced from [6].

2.3. High Queuing Latencies at the Shared Cache

Due to the amount of thread parallelism within a GPU, a large number of memory requests can arrive at the L2 cache in a small window of execution time, leading to significant queuing delays. Prior work observes high access latencies for the shared L2 cache within a GPU [126, 127, 142], but does not identify *why* these latencies are so high. We show that when a large number of requests arrive at the L2 cache, both the limited number of read/write ports and backpressure from cache bank conflicts force many of these requests to queue up for long periods of time. We observe that this queuing latency can sometimes add *hundreds* of cycles to the cache access latency, and that non-uniform queuing across the different cache banks exacerbates memory divergence. Figure 5 quantifies the magnitude of this queue contention if we set the cache lookup latency at one cycle, for one application, BFS [10]. As shown, a significant number of requests experience tens to hundreds of cycles of queuing delay.




Figure 5: Distribution of per-request queuing latencies for L2 cache requests from BFS. Reproduced from [6].

The warp-type-aware bypassing logic in MeDiC helps to alleviate these L2 queuing latencies. By preventing mostly-miss and all-miss warps from accessing the cache, which yields little or no benefit to them, we reduce the access latencies for requests from (1) mostly-hit and all-hit warps, which benefit from the cache much more; and also (2) mostly-miss and all-miss warps themselves; thereby improving the overall performance of all warps and the system.

3. MeDiC: Memory Divergence Correction

Based on these three new observations we made, we define three major goals for our new mechanism. We would like to devise a mechanism that (1) converts mostly-hit warps into *all-hit warps* (warps where *all* requests hit in the cache, as shown in Figure 1c), (2) converts mostly-miss warps into *all-miss warps* (warps where *none* of the requests hit in the cache, as shown in Figure 1d) and (3) reduces L2 cache queuing delay for all warp types. As we can see in Figure 1a, the stall time due to memory divergence for the mostly-hit warp can be eliminated by converting only the single cache miss (Figure 1a, ②) into a cache hit.

To this end, we introduce *Memory Divergence Correction* (MeDiC), a GPU-specific mechanism that exploits *memory divergence heterogeneity* across warps at the shared cache and at main memory to improve the overall performance of GPGPU applications. MeDiC consists of three different components, which work together to achieve our three goals: (1) a warp-type-aware *cache bypassing mechanism*, which prevents requests from mostly-miss and all-miss warps from accessing the shared L2 cache; (2) a warp-type-aware *cache insertion policy*, which prioritizes requests from mostly-hit and all-hit warps, in order to increase the likelihood that they all become cache hits; and (3) a warp-type-aware *memory scheduling mechanism*, which prioritizes requests from mostly-hit warps that were not successfully converted to all-hit warps, in order to minimize the stall time due to divergence. These three components are all driven by an online mechanism that can identify the expected memory divergence behavior of each warp.

Figure 6 shows the overall MeDiC mechanism. MeDiC consists of four different components: ① a *warp-type-identification mechanism* that classifies warps into one of the four warp types as described in Section 2.1; ② a *bypass mechanism* that bypasses requests from all-miss and mostly-miss warps, reducing the queuing delay in the L2 cache; ③ an *insertion policy* that prevent mostly-hit requests from being evicted from the cache; and ④ a *memory scheduler* that prioritizes requests from mostly-hit warps, which are more latency sensitive.

3.1. Warp Type Identification

In order to take advantage of the memory divergence heterogeneity across warps, we must first add hardware that can identify the divergence behavior of each warp. The key idea is to periodically sample the hit ratio of a warp, and to classify the warp’s divergence behavior as one of the five types in Figure 3 based on the observed hit ratio. This information can then be used to drive the warp-type-aware components of MeDiC. In general, warps tend to retain the same memory divergence behavior for long periods of execution. However, there can be some long-term shifts in warp divergence behavior, requiring periodic resampling of the hit ratio to potentially re-evaluate the warp type. Warp type




Figure 6: Overview of MeDiC: ① warp type identification logic, ② warp-type-aware cache bypassing, ③ warp-type-aware cache insertion policy, ④ warp-type-aware memory scheduler. Reproduced from [6].

identification through hit ratio sampling requires hardware within the cache to periodically count the number of hits and misses each warp incurs. We append two counters to the metadata stored for each warp, which represent the total number of cache hits and cache accesses for the warp during the sampling interval.

3.2. Warp-type-aware Shared Cache Bypassing

Once the warp type is known and a warp generates a request to the L2 cache, our mechanism first decides whether to bypass the cache based on the warp type. The key idea behind *warp-type-aware cache bypassing* is to convert mostly-miss warps into all-miss warps, as they do not benefit greatly from the few cache hits that they get. By bypassing these requests, we achieve three benefits: (1) bypassed requests can avoid L2 queuing latencies entirely, (2) other requests that do hit in the L2 cache experience shorter queuing delays due to the reduced contention, and (3) space is created in the L2 cache for mostly-hit warps.

The cache bypassing logic must make a simple decision: if an incoming memory request was generated by a mostly-miss or all-miss warp, the request is bypassed directly to DRAM. This is determined by reading the warp type stored in the warp metadata from the warp type identification mechanism. A simple 2-bit demultiplexer can be used to determine whether a request is sent to the L2 bank arbiter, or directly to the DRAM request queue.

3.3. Warp-type-aware Cache Insertion Policy

Our cache bypassing mechanism frees up space within the L2 cache, which we want to use for the cache misses from mostly-hit warps (to convert the cache miss memory requests into cache hits). However, even with the new bypassing mechanism, other warps (e.g., balanced, mostly-miss) still insert some data into the cache. In order to aid the conversion of mostly-hit warps into all-hit warps, we develop a *warp-type-aware cache insertion policy*, whose key idea is to ensure that in a given cache set, data from mostly-miss warps are

evicted first, while data from mostly-hit warps and all-hit warps are evicted last.

To ensure that a cache block from a mostly-hit warp stays in the cache for as long as possible, we insert the block closer to the MRU position. A cache block requested by a mostly-miss warp is inserted closer to the LRU position, making it more likely to be evicted. To track the warp type associated with these cache blocks, we add two bits of metadata to each cache block, indicating the warp type. These bits are then appended to the replacement policy bits. The bits modify the replacement policy behavior, such that a cache block from a mostly-miss warp is more likely to get evicted than a block from a balanced warp. Similarly, a cache block from a balanced warp is more likely to be evicted than a block from a mostly-hit or all-hit warp.

3.4. Warp-type-aware Memory Scheduler

Our cache bypassing mechanism and cache insertion policy work to increase the likelihood that *all* requests from a mostly-hit warp become cache hits, converting the warp into an all-hit warp. However, due to cache conflicts, or due to poor locality, there may still be cases when a mostly-hit warp cannot be fully converted into an all-hit warp, and is therefore unable to avoid stalling due to memory divergence as at least one of its requests has to go to DRAM. In such a case, we want to minimize the amount of time that this warp stalls. To this end, we propose a *warp-type-aware memory scheduler* that prioritizes the occasional DRAM requests from mostly-hit warps.

The design of our memory scheduler is very simple. Each memory request is tagged with a single bit, which is set if the memory request comes from a mostly-hit warp (or an all-hit warp, in case the warp was mischaracterized). We modify the request queue at the memory controller to contain two different queues, where a *high-priority queue* contains all requests that have their mostly-hit bit set to one. The *low-priority queue* contains all other requests, whose mostly-hit bits are set to zero. Each queue uses FR-FCFS [115, 156] as

the scheduling policy; however, the scheduler always selects requests from the high priority queue over requests in the low priority queue.¹

We describe each component of MeDiC in more detail in Sections 4.1, 4.2, 4.3 and 4.4 of our PACT 2015 paper [6].

4. Methodology

We model our mechanism using GPGPU-Sim 3.2.1 [9]. We modified GPGPU-Sim to accurately model cache bank conflicts, and added the cache bypassing, cache insertion, and memory scheduling mechanisms needed to support MeDiC. We use GPUWattch [76] to evaluate power consumption. We have open sourced our simulator source code at [118]. We evaluate our system across multiple GPGPU applications from the CUDA SDK [102], Rodinia [19], MARS [39], and Lonestar [10] benchmark suites.

We report performance results using the harmonic average of the IPC speedup (over the baseline GPU) of each kernel of each application.² Harmonic speedup [28, 85] was shown to reflect the average normalized execution time in multi-programmed workloads. We calculate energy efficiency for each workload by dividing the IPC by the energy consumed. Section 5 of our PACT 2015 paper [6] provides more detail on our experimental methodology.

5. Evaluation

Figure 7 shows the performance of MeDiC compared to four GPU cache management mechanisms: the Evicted Address Filter insertion policy [123] (**EAF**), PCAL bypassing policy [79] (**PCAL**), PC-based cache bypassing policy (**PC-ByP**) and an idealized random bypassing policy (**Rand**) over 15 different GPGPU applications from 4 benchmark suites. We also show the performance of each individual component of MeDiC: our warp-type-aware insertion policy (**WIP**), our warp-type-aware memory scheduling policy (**WMS**) and our warp-type-aware bypassing policy (**WByP**).

We found that each component of MeDiC individually provides significant performance improvement: WIP (32.5%), WMS (30.2%), and WByP (33.6%). MeDiC, which combines all three mechanisms, provides a 41.5% performance improvement over Baseline, on average. MeDiC matches or outperforms its individual components for all benchmarks except BP, where MeDiC has a higher L2 miss rate and lower row buffer locality than WMS and WByP.

Our insertion policy, WIP, outperforms EAF [123] by 12.2%. We observe that the key benefit of WIP is that cache blocks

¹Using two queues ensures that high-priority requests are not blocked by low-priority requests even when the low-priority queue is full. Two-queue priority also uses simpler logic design than comparator-based priority [5, 132, 133].

²We confirm that for each application, all kernels have similar speedup values, and that aside from SS and PVC, there are no outliers (i.e., no kernel has a much higher speedup than the other kernels). To verify that harmonic speedup is not swayed greatly by these few outliers, we recompute it for SS and PVC *without* these outliers, and find that the outlier-free speedup is within 1% of the harmonic speedup we report in the paper.

from mostly-miss warps are much more likely to be evicted. In addition, WIP reduces the cache miss rate of several applications. Our memory scheduler, WMS, provides significant performance gains (30.2%) over Baseline, because the memory scheduler prioritizes requests from warps that have a high hit ratio, allowing these warps to become active much sooner than they do in Baseline. Our bypassing mechanism, WByP provides an average 33.6% performance improvement over Baseline, because it is effective at reducing the L2 queuing latency.

Compared to PCAL [79], WByP provides 12.8% better performance, and full MeDiC provides 21.8% better performance. We observe that while PCAL reduces the amount of cache thrashing, the reduction in thrashing does not directly translate into better performance. We observe that warps in the mostly-miss category sometimes have high reuse, and acquire tokens to access the cache. This causes less cache space to become available for mostly-hit warps, limiting how many of these warps become all-hit. However, when high-reuse warps that possess tokens are mainly in the mostly-hit category (PVC, PVR, SS, and BH), we find that PCAL performs better than WByP.

Compared to Rand,³ MeDiC performs 6.8% better, because MeDiC is able to make bypassing decisions that do not increase the miss rate significantly. This leads to lower off-chip bandwidth usage under MeDiC than under Rand. Rand increases the cache miss rate by 10% for the kernels of several applications (BP, PVC, PVR, BFS, and MST). We observe that in many cases, MeDiC improves the performance of applications that tend to generate a large number of memory requests, and thus experience substantial queuing latencies.

Compared to PC-ByP, MeDiC performs 12.4% better. We observe that the overhead of tracking the PC becomes significant, and that thrashing occurs as two PCs can hash to the same index, leading to inaccuracies in the bypassing decisions.

We conclude that each component of MeDiC, and the full MeDiC framework, are effective. Note that each component of MeDiC addresses the same problem (i.e., memory divergence of threads within a warp) using different techniques on different parts of the memory hierarchy. For the majority of workloads, one optimization is enough. However, we see that for certain high-intensity workloads (BFS and SSSP), the congestion is so high that we need to attack divergence on multiple fronts. Thus, MeDiC provides better average performance than all of its individual components, especially for such memory-intensive workloads.

We provide the following other evaluation results in our PACT 2015 paper [6]:

- Impact of MeDiC on cache miss rate.

³Note that our evaluation uses an ideal random bypassing mechanism, where we manually select the best individual percentage of requests to bypass the cache for each workload. As a result, the performance shown for Rand is better than can be practically realized.




Figure 7: Performance of MeDiC. Adapted from [6].




Figure 8: Energy efficiency of MeDiC. Adapted from [6].

- Impact of MeDiC on queuing latency.
- Impact of MeDiC on row buffer locality.
- Analysis of reuse in GPGPU applications.
- Hardware cost of MeDiC.

6. Related Work

To our knowledge, MeDiC is the first work that identifies inter-warp memory divergence heterogeneity and exploits it to achieve better system performance in GPGPU applications. Our new mechanism consists of warp-type-aware components for cache bypassing, cache insertion, and memory scheduling. We have already provided extensive quantitative and qualitative comparisons to state-of-the-art mechanisms in GPU cache bypassing [79], cache insertion [123], and memory scheduling [115, 156]. In this section, we discuss other related work in these areas.

Hardware-based Cache Bypassing. PCAL is a bypassing mechanism that addresses the cache thrashing problem by throttling the number of threads that time-share the cache at any given time [79]. The key idea of PCAL is to limit the number of threads that get to access the cache. Concurrent work by Li et al. [78] proposes a cache bypassing mechanism that allows only threads with high reuse to utilize the cache. The key idea is to use locality filtering based on the reuse characteristics of GPGPU applications, with only high reuse threads having access to the cache. Xie et al. [146] propose a bypassing mechanism at the thread block level. In their mechanism, the compiler statically marks whether thread blocks prefer caching or bypassing. At runtime, the mechanism dynamically selects a subset of thread blocks to use the cache, to increase cache utilization.

Chen et al. [20, 21] propose a combined warp throttling and bypassing mechanism for the L1 cache based on the cache-conscious warp scheduler [116]. The key idea is to bypass the cache when resource contention is detected. This is done by embedding history information into the L2 tag arrays. The L1 cache uses this information to perform bypassing decisions, and only warps with high reuse are allowed to access the L1 cache. Jia et al. propose an L1 bypassing mechanism [48], whose key idea is to bypass requests when there is an associativity stall. Dai et al. propose a mechanism to bypass cache based on a model of a cache miss rate [23].

MeDiC differs from these prior cache bypassing works because it uses warp memory divergence heterogeneity for bypassing decisions. We also show (in Section 6.4 of our PACT 2015 paper [6]) that our mechanism implicitly takes reuse information into account.

Software-based Cache Bypassing. Concurrent work by Li et al. [77] proposes a compiler-based technique that performs cache bypassing using a method similar to PCAL [79]. Xie et al. [145] propose a mechanism that allows the compiler to perform cache bypassing for global load instructions. Both of these mechanisms are different from MeDiC in that MeDiC applies bypassing to *all* loads and stores that utilize the shared cache, without requiring additional characterization at the compiler level. Mekkat et al. [87] propose a bypassing mechanism for when a CPU and a GPU share the last level cache. Their key idea is to bypass GPU cache accesses when CPU applications are cache sensitive, which is not applicable to GPU-only execution.

CPU Cache Bypassing. There are also several other CPU-based cache bypassing techniques. These techniques include using additional buffers track cache statistics to predict cache

blocks that have high utility based on reuse count [18, 27, 32, 50, 55, 81, 144, 152], reuse distance [18, 24, 29, 31, 34, 104, 143, 149], behavior of the cache block [46] or miss rate [22, 88, 120, 137]. As they do not operate on SIMD systems, these mechanisms do not (need to) account for memory divergence heterogeneity when performing bypassing decisions.

Cache Insertion and Replacement Policies. Many works propose different insertion policies for CPU systems (e.g., [44, 45, 54, 110, 112, 123]). We compare our insertion policy against the Evicted-Address Filter (EAF) [123], and show that our policy, which takes advantage of inter-warp divergence heterogeneity, outperforms EAF. Dynamic Insertion Policy (DIP) [44] and Dynamic Re-Reference Interval Prediction (DRRIP) [45] are insertion policies that account for cache thrashing. The downside of these two policies is that they are unable to distinguish between high-reuse and low-reuse blocks in the same thread [123]. The Bi-modal Insertion Policy [110] dynamically characterizes the cache blocks being inserted. None of these works take warp type characteristics or memory divergence behavior into account. Other work proposed prefetch-aware insertion and replacement policies [25, 124, 130]. MeDiC can be combined with such policies.

Memory Scheduling. Yuan et al. propose a GPU interconnect design that rearrange the sequence of memory requests that arrive at each memory channel to reduce the complexity of GPU memory scheduler [151]. Chatterjee et al. propose a GPU memory scheduler that allows requests from the same warp to be grouped together, in order to reduce the memory divergence across different memory requests within the same warp [17]. Jog et al. propose a GPU memory scheduler that exploit the criticality information of warps in the GPU cores in order to improve the performance of GPGPU applications [49]. Principles of MeDiC can be incorporated into these schedulers.

There are several memory scheduler designs that target systems with CPUs [26, 33, 43, 56, 57, 59, 60, 67, 68, 69, 82, 93, 94, 95, 96, 98, 99, 115, 131, 132, 133, 134, 135, 147, 153], and heterogeneous compute elements [5, 47, 138]. Memory schedulers for CPUs and heterogeneous systems generally aim to reduce interference across different applications.

Improving DRAM. An alternative approach to mitigate memory divergence is to improve the performance of the main memory itself. Previous works propose new DRAM designs that are capable of reducing memory latency in conventional DRAM [1, 2, 3, 4, 11, 12, 13, 13, 14, 14, 15, 16, 35, 36, 37, 38, 40, 41, 42, 53, 58, 61, 70, 71, 72, 73, 74, 83, 86, 92, 97, 100, 103, 109, 119, 121, 122, 125, 129, 141, 154] as well as non-volatile memory [62, 64, 65, 66, 80, 84, 90, 91, 111, 113, 114, 148, 150]. Data compression techniques can increase the effective DRAM bandwidth [105, 106, 107, 108, 140]. All these techniques are orthogonal to MeDiC and can be used to further improve the performance of GPGPU applications.

Other Ways to Handle Memory Divergence. In addition to cache bypassing, cache insertion policy, and memory scheduling, other works propose different methods of decreasing memory divergence [51, 52, 63, 75, 89, 101, 116, 117, 155]. These methods range from thread throttling [51, 52, 63, 116] to warp scheduling [75, 89, 101, 116, 117, 155]. While these methods share our goal of reducing memory divergence, none of them exploit inter-warp heterogeneity and, as a result, are either orthogonal or complementary to our proposal. Our work also makes new observations about memory divergence not covered by these works.

7. Potential Impact

While the problem that MeDiC is trying to solve, which is memory divergence, is not new, key findings in this work provide novelty and create potential research topics for the future. We discuss at least three such opportunities and future directions.

Taking Advantage of Memory Divergence Heterogeneity. MeDiC modifies the memory hierarchy to introduce awareness of the memory divergence heterogeneity between different types of warps. There are many other applications that can exploit warp type information. Other resources within the GPU (e.g., GPU cores, warp scheduler) can exploit the memory divergence heterogeneity across different warps to further improve the performance of GPGPU applications. For example, the warp type information can be used by the warp scheduler and thread block scheduler to ensure that they do not schedule warps of the same type to execute at the same time, to limit the amount of cache contention that occurs. Incorporating the warp type information with other techniques, such as assist warps to relieve execution bottlenecks [140], can enable GPUs to utilize resources based on the type of warps the GPU is executing. For example, mostly-hit warps favor a mechanism that provides low memory latency, while mostly-miss warps might favor a mechanism that provides higher off-chip bandwidth. Memory divergence heterogeneity can also be used to assist GPU resource virtualization [139], as virtual resource allocation can take into account the utilization of shared memory resources to determine how much of a particular memory resource to allocate to each thread block.

Warp type information can be used to improve the performance of GPU address translation. Prior works [7, 8] show that address translations that do not hit in a TLB can incur long-latency page table walks, which can affect hundreds of application threads at once. Such long-latency address translations might have a greater impact on warps that are latency sensitive (e.g., *mostly-hit* and *all-hit* warps). Thus, warp-type information can be combined with previously-proposed techniques that aim to reduce the overhead of GPU address translation [7, 8] to provide synergistic performance benefits.

We believe the idea of warp-type heterogeneity enables many different mechanisms to customize execution on a GPU

to achieve higher performance and energy efficiency. Hence, our PACT 2015 paper [6] paves the way for fine-grained customization of a GPU.

Identifying Long-Latency Threads in a Warp. Our PACT 2015 paper [6] shows how to intelligently reduce the memory latency of threads within a warp in order to reduce the memory divergence problem. However, MeDiC focuses on reducing the stall time of mostly-hit warps. Long-latency threads can still exist in the mostly-hit warps due to other problems such as load balancing at the memory partitions. Additional work on (1) how to identify latency-critical threads within a warp and (2) how to accelerate these specific threads can further improve the performance and energy efficiency of GPGPU applications.

Reducing High Queuing Delays and Memory Contention in the GPU Memory Hierarchy. As shown in our PACT 2015 paper [6], the queuing delay of throughput processors such as GPUs can become a performance bottleneck, as the delay increases the stall time of warps of *all* types. While the proposed warp-type-aware cache bypassing mechanism in MeDiC aims to reduce the queuing delay, non-uniform memory access patterns can still cause contention at a few L2 cache banks and memory partitions. In future systems, the parallelism of throughput processors is likely to increase further. For example, future GPUs will likely come with a higher number of GPU cores and larger SIMD widths. This is expected to greatly increase the amount of contention and, thus, queuing delay, for many resources. The different components of MeDiC can serve as a starting point for future research on alleviating cache and memory contention in future systems, and can ultimately enable a larger amount of thread-level parallelism. We believe studying the mitigation of high cache and memory contention is very promising for future parallel throughput processors and encourage future work in this area.

8. Conclusion

Warps from GPGPU applications exhibit heterogeneity in their memory divergence behavior at the shared L2 cache within the GPU. We find that (1) some warps benefit significantly from the cache, while others make poor use of it; (2) such divergence behavior for a warp tends to remain stable for long periods of the warp’s execution; and (3) the impact of memory divergence can be amplified by the high queuing latencies at the L2 cache.

We propose *Memory Divergence Correction* (MeDiC), whose key idea is to identify memory divergence heterogeneity online in hardware and use this information to drive cache management and memory scheduling, by prioritizing warps that take the greatest advantage of the shared cache. To achieve this, MeDiC consists of three *warp-type-aware* components for (1) cache bypassing, (2) cache insertion, and (3) memory scheduling. MeDiC delivers significant performance and energy improvements over multiple previously proposed

policies, and over a state-of-the-art GPU cache management technique. We conclude that exploiting inter-warp heterogeneity is effective, and hope future works explore other ways of improving systems based on this key observation of our work.

Acknowledgments

We thank the anonymous reviewers and SAFARI group members for their feedback. Special thanks to Mattan Erez for his valuable feedback on our PACT 2015 paper. We acknowledge the support of our industrial partners: Facebook, Google, IBM, Intel, Microsoft, NVIDIA, Qualcomm, VMware, and Samsung. This research was partially supported by the NSF (grants 0953246, 1065112, 1205618, 1212962, 1213052, 1302225, 1302557, 1317560, 1320478, 1320531, 1409095, 1409723, 1423172, 1439021, and 1439057), the Intel Science and Technology Center for Cloud Computing, and the Semiconductor Research Corporation.

References

- [1] J. Ahn, S. Hong, S. Yoo, O. Mutlu, and K. Choi, “A Scalable Processing-in-memory Accelerator for Parallel Graph Processing,” in *ISCA*, 2015.
- [2] J. Ahn, S. Yoo, O. Mutlu, and K. Choi, “PIM-enabled Instructions: A Low-overhead, Locality-aware Processing-in-memory Architecture,” in *ISCA*, 2015.
- [3] J. H. Ahn, J. Leverich, R. Schreiber, and N. P. Jouppi, “Multicore DIMM: an Energy Efficient Memory Module with Independently Controlled DRAMs,” *IEEE CAL*, 2009.
- [4] J. H. Ahn, N. P. Jouppi, C. Kozyrakis, J. Leverich, and R. S. Schreiber, “Improving System Energy Efficiency with Memory Rank Subsetting,” *ACM TACO*, vol. 9, no. 1, pp. 4:1–4:28, 2012.
- [5] R. Ausavarungnirun, K. K. Chang, L. Subramanian, G. Loh, and O. Mutlu, “Staged Memory Scheduling: Achieving High Performance and Scalability in Heterogeneous Systems,” in *ISCA*, 2012.
- [6] R. Ausavarungnirun, S. Ghose, O. Kayiran, G. H. Loh, C. R. Das, M. T. Kandemir, and O. Mutlu, “Exploiting Inter-Warp Heterogeneity to Improve GPGPU Performance,” in *PACT*, 2015.
- [7] R. Ausavarungnirun, J. Landgraf, V. Miller, S. Ghose, J. Gandhi, C. J. Rossbach, and O. Mutlu, “Mosaic: A GPU Memory Manager with Application-Transparent Support for Multiple Page Sizes,” in *MICRO*, 2017.
- [8] R. Ausavarungnirun, V. Miller, J. Landgraf, S. Ghose, J. Gandhi, A. Jog, C. Rossbach, and O. Mutlu, “MASK: Redesigning the GPU Memory Hierarchy to Support Multi-Application Concurrency,” in *ASPLOS*, 2018.
- [9] A. Bakhoda, G. Yuan, W. Fung, H. Wong, and T. Aamodt, “Analyzing CUDA Workloads Using a Detailed GPU Simulator,” in *ISPASS*, 2009.
- [10] M. Burtscher, R. Nasre, and K. Pingali, “A Quantitative Study of Irregular Programs on GPUs,” in *IISWC*, 2012.
- [11] K. Chandrasekar, S. Goossens, C. Weis, M. Koedam, B. Akesson, N. Wehn, and K. Goossens, “Exploiting Expendable Process-Margins in DRAMs for Run-Time Performance Optimization,” in *DATE*, 2014.
- [12] K. K. Chang, A. Kashyap, H. Hassan, S. Ghose, K. Hsieh, D. Lee, T. Li, G. Pekhimenko, S. Khan, and O. Mutlu, “Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization,” in *SIGMETRICS*, 2016.
- [13] K. K. Chang, D. Lee, Z. Chishti, A. Alameldeen, C. Wilkerson, Y. Kim, and O. Mutlu, “Improving DRAM Performance by Parallelizing Refreshes with Accesses,” in *HPCA*, 2014.
- [14] K. K. Chang, P. J. Nair, D. Lee, S. Ghose, M. K. Qureshi, and O. Mutlu, “Low-Cost Inter-Linked Subarrays (LISA): Enabling Fast Inter-Subarray Data Movement in DRAM,” in *HPCA*, 2016.
- [15] K. K. Chang, A. G. Yaglikci, S. Ghose, A. Agrawal, N. Chatterjee, A. Kashyap, D. Lee, M. O’Connor, H. Hassan, and O. Mutlu, “Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms,” in *SIGMETRICS*, 2017.
- [16] N. Chatterjee, M. Shevgoor, R. Balasubramonian, A. Davis, Z. Fang, R. Illikkal, and R. Iyer, “Leveraging Heterogeneity in DRAM Main Memories to Accelerate Critical Word Access,” in *MICRO*, 2012.
- [17] N. Chatterjee, M. O’Connor, G. H. Loh, N. Jayasena, and R. Balasubramonian, “Managing DRAM Latency Divergence in Irregular GPGPU Applications,” in *SC*, 2014.
- [18] M. Chaudhuri, J. Gaur, N. Bashyam, S. Subramoney, and J. Nuzman, “Introducing Hierarchy-awareness in Replacement and Bypass Algorithms for Last-level Caches,” in *PACT*, 2012.

- [19] S. Che, M. Boyer, J. Meng, D. Tarjan, J. Sheaffer, S.-H. Lee, and K. Skadron, "Rodinia: A Benchmark Suite for Heterogeneous Computing," in *IISWC*, 2009.
- [20] X. Chen, L.-W. Chang, C. I. Rodrigues, J. Lv, Z. Wang, and W. W. Hwu, "Adaptive Cache Management for Energy-Efficient GPU Computing," in *MICRO*, 2014.
- [21] X. Chen, S. Wu, L.-W. Chang, W.-S. Huang, C. Pearson, Z. Wang, and W. W. Hwu, "Adaptive Cache Bypass and Insertion for Many-Core Accelerators," in *MES*, 2014.
- [22] J. D. Collins and D. M. Tullsen, "Hardware Identification of Cache Conflict Misses," in *MICRO*, 1999.
- [23] H. Dai, C. Li, H. Zhou, S. Gupta, C. Kartsaklis, and M. Mantor, "A Model-driven Approach to Warp/thread-block Level GPU Cache Bypassing," in *DAC*, 2016.
- [24] N. Duong, D. Zhao, T. Kim, R. Cammarota, M. Valero, and A. V. Veidenbaum, "Improving Cache Management Policies Using Dynamic Reuse Distances," in *MICRO*, 2012.
- [25] E. Ebrahimi, C. J. Lee, O. Mutlu, and Y. N. Patt, "Prefetch-aware Shared Resource Management for Multi-core Systems," in *ISCA*, 2011.
- [26] E. Ebrahimi, R. Miftakhutdinov, C. Fallin, C. J. Lee, J. A. Joao, O. Mutlu, and Y. N. Patt, "Parallel Application Memory Scheduling," in *MICRO*, 2011.
- [27] Y. Etsion and D. G. Feitelson, "Exploiting Core Working Sets to Filter the L1 Cache with Random Sampling," *IEEE TC*, vol. 61, no. 11, pp. 1535–1550, 2012.
- [28] S. Eyerhan and L. Eeckhout, "System-Level Performance Metrics for Multiprogram Workloads," *IEEE Micro*, 2008.
- [29] M. Feng, C. Tian, and R. Gupta, "Enhancing LRU Replacement via Phantom Associativity," in *INTERACT*, Feb 2012.
- [30] M. Flynn, "Very High-Speed Computing Systems," *Proc. of the IEEE*, vol. 54, no. 2, 1966.
- [31] H. Gao and C. Wilkerson, "A Dueling Segmented LRU Replacement Algorithm with Adaptive Bypassing," in *JWAC*, 2010.
- [32] J. Gaur, M. Chaudhuri, and S. Subramoney, "Bypass and Insertion Algorithms for Exclusive Last-Level Caches," in *ISCA*, 2011.
- [33] S. Ghose, H. Lee, and J. F. Martinez, "Improving Memory Scheduling via Processor-side Load Criticality Information," in *ISCA*, 2013.
- [34] S. Gupta, H. Gao, and H. Zhou, "Adaptive Cache Bypassing for Inclusive Last Level Caches," in *IPDPS*, 2013.
- [35] C. A. Hart, "CDRAM in a Unified Memory Architecture," in *Intl. Computer Conference*, 1994.
- [36] M. Hashemi, O. Mutlu, and Y. N. Patt, "Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads," in *MICRO*, 2016.
- [37] H. Hassan, G. Pekhimenko, N. Vijaykumar, V. Seshadri, D. Lee, O. Ergin, and O. Mutlu, "ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality," in *HPCA*, 2016.
- [38] H. Hassan, N. Vijaykumar, S. Khan, S. Ghose, K. K. Chang, G. Pekhimenko, D. Lee, O. Ergin, and O. Mutlu, "SoftMC: A Flexible and Practical Open-source Infrastructure for Enabling Experimental DRAM Studies," in *HPCA*, 2017.
- [39] B. He, W. Fang, Q. Luo, N. K. Govindaraju, and T. Wang, "Mars: A MapReduce Framework on Graphics Processors," in *PACT*, 2008.
- [40] H. Hidaka, Y. Matsuda, M. Asakura, and K. Fujishima, "The Cache DRAM Architecture," *IEEE Micro*, 1990.
- [41] K. Hsieh, S. Khan, N. Vijaykumar, K. K. Chang, A. Boroumand, S. Ghose, and O. Mutlu, "Accelerating Pointer Chasing in 3D-stacked Memory: Challenges, Mechanisms, Evaluation," in *ICCD*, 2016.
- [42] W.-C. Hsu and J. E. Smith, "Performance of Cached DRAM Organizations in Vector Supercomputers," in *ISCA*, 1993.
- [43] E. Ipek, O. Mutlu, J. F. Martinez, and R. Caruana, "Self-optimizing memory controllers: A reinforcement learning approach," in *ISCA*, 2008.
- [44] A. Jaleel, W. Hasenplaugh, M. Qureshi, J. Sebot, S. Steely, Jr., and J. Emer, "Adaptive Insertion Policies for Managing Shared Caches," in *PACT*, 2008.
- [45] A. Jaleel, K. B. Theobald, S. C. Steely, Jr., and J. Emer, "High Performance Cache Replacement Using Re-reference Interval Prediction (RRIP)," in *ISCA*, 2010.
- [46] J. Jalminger and P. Stenstrom, "A Novel Approach to Cache Block Reuse Predictions," in *ICPP*, 2003.
- [47] M. K. Jeong, M. Erez, C. Sudanthi, and N. Paver, "A QoS-Aware Memory Controller for Dynamically Balancing GPU and CPU Bandwidth Use in an MPSoC," in *DAC*, 2012.
- [48] W. Jia, K. A. Shaw, and M. Martonosi, "MRPB: Memory Request Prioritization for Massively Parallel Processors," in *HPCA*, 2014.
- [49] A. Jog, O. Kayiran, A. Pattnaik, M. T. Kandemir, O. Mutlu, R. Iyer, and C. R. Das, "Exploiting Core Criticality for Enhanced GPU Performance," in *SIGMETRICS*, 2016.
- [50] L. K. John and A. Subramanian, "Design and Performance Evaluation of A Cache Assist to Implement Selective Caching," in *ICCD*, 1997.
- [51] O. Kayiran, A. Jog, M. T. Kandemir, and C. R. Das, "Neither More Nor Less: Optimizing Thread-Level Parallelism for GPGPUs," in *PACT*, 2013.
- [52] O. Kayiran, N. C. Nachiappan, A. Jog, R. Ausavarungnirun, M. T. Kandemir, G. H. Loh, O. Mutlu, and C. R. Das, "Managing GPU Concurrency in Heterogeneous Architectures," in *MICRO*, 2014.
- [53] G. Kedem and R. P. Goganti, "WCDRAM: A Fully Associative Integrated Cached-DRAM with Wide Cache Lines," *CS-1997-03, Duke*, 1997.
- [54] S. Khan, A. R. Alameldeen, C. Wilkerson, O. Mutlu, and D. A. Jimenez, "Improving Cache Performance using Read-write Partitioning," in *HPCA*, 2014.
- [55] M. Kharbutli and Y. Solihin, "Counter-Based Cache Replacement and Bypassing Algorithms," *IEEE TC*, vol. 57, no. 4, pp. 433–447, Apr. 2008.
- [56] H. Kim, D. de Niz, B. Andersson, M. Klein, O. Mutlu, and R. Rajkumar, "Bounding Memory Interference Delay in COTS-based Multi-core Systems," in *RTAS*, 2014.
- [57] H. Kim, D. de Niz, B. Andersson, M. Klein, O. Mutlu, and R. Rajkumar, "Bounding and Reducing Memory Interference in COTS-based Multi-core Systems," *Real-Time Systems*, vol. 52, no. 3, May 2016.
- [58] Y. Kim, W. Yang, and O. Mutlu, "Ramulator: A Fast and Extensible DRAM Simulator," *CAL*, 2015.
- [59] Y. Kim, D. Han, O. Mutlu, and M. Harchol-Balter, "ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers," in *HPCA*, 2010.
- [60] Y. Kim, M. Papamichael, O. Mutlu, and M. Harchol-Balter, "Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior," in *MICRO*, 2010.
- [61] Y. Kim, V. Seshadri, D. Lee, J. Liu, and O. Mutlu, "A Case for Exploiting Subarray-Level Parallelism (SALP) in DRAM," in *ISCA*, 2012.
- [62] E. Kultursay, M. Kandemir, A. Sivasubramanian, and O. Mutlu, "Evaluating STT-DRAM as an energy-efficient main memory alternative," in *ISPASS*, 2013.
- [63] H.-K. Kuo, B. C. Lai, and J.-Y. Jou, "Reducing Contention in Shared Last-Level Cache for Throughput Processors," *ACM TODAES*, vol. 20, no. 1, 2014.
- [64] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Architecting Phase Change Memory as a Scalable DRAM Alternative," in *ISCA*, 2009.
- [65] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Phase Change Memory Architecture and the Quest for Scalability," *CACM*, vol. 53, no. 7, pp. 99–106, 2010.
- [66] B. C. Lee, P. Zhou, J. Yang, Y. Zhang, B. Zhao, E. Ipek, O. Mutlu, and D. Burger, "Phase-Change Technology and the Future of Main Memory," *IEEE Micro*, vol. 30, no. 1, pp. 143–143, 2010.
- [67] C. J. Lee, V. Narasiman, O. Mutlu, and Y. N. Patt, "Improving Memory Bank-Level Parallelism in the Presence of Prefetching," in *MICRO*, 2009.
- [68] C. J. Lee, E. Ebrahimi, V. Narasiman, O. Mutlu, and Y. N. Patt, "DRAM-Aware Last-Level Cache Writeback: Reducing Write-Caused Interference in Memory Systems," Univ. of Texas at Austin, High Performance Systems Group, Tech. Rep. TR-HPS-2010-002, 2010.
- [69] C. J. Lee, O. Mutlu, V. Narasiman, and Y. N. Patt, "Prefetch-aware DRAM Controllers," in *MICRO*, 2008.
- [70] D. Lee, S. Ghose, G. Pekhimenko, S. Khan, and O. Mutlu, "Simultaneous Multi-layer Access: Improving 3D-stacked Memory Bandwidth at Low Cost," *ACM TACO*, vol. 12, no. 4, p. 63, 2016.
- [71] D. Lee, S. Khan, L. Subramanian, S. Ghose, R. Ausavarungnirun, G. Pekhimenko, V. Seshadri, and O. Mutlu, "Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms," in *SIGMETRICS*, 2017.
- [72] D. Lee, Y. Kim, G. Pekhimenko, S. Khan, V. Seshadri, K. K. Chang, and O. Mutlu, "Adaptive-latency DRAM: Optimizing DRAM Timing for the Common-case," in *HPCA*, 2015.
- [73] D. Lee, Y. Kim, V. Seshadri, J. Liu, L. Subramanian, and O. Mutlu, "Tiered-latency DRAM: A Low Latency and Low Cost DRAM Architecture," in *HPCA*, 2013.
- [74] D. Lee, L. Subramanian, R. Ausavarungnirun, J. Choi, and O. Mutlu, "Decoupled Direct Memory Access: Isolating CPU and IO Traffic by Leveraging a Dual-Data-Port DRAM," in *PACT*, 2015.
- [75] S.-Y. Lee and C.-J. Wu, "CAWS: Criticality-Aware Warp Scheduling for GPGPU Workloads," in *PACT*, 2014.
- [76] J. Leng, T. Hetherington, A. ElTantawy, S. Gilani, N. S. Kim, T. M. Aamodt, and V. J. Reddi, "GPUWattch: Enabling Energy Optimizations in GPGPUs," in *ISCA*, 2013.
- [77] A. Li, G.-J. van den Braak, A. Kumar, and H. Corporaal, "Adaptive and Transparent Cache Bypassing for GPUs," in *SC*, 2015.
- [78] C. Li, S. L. Song, H. Dai, A. Sidelnik, S. K. S. Hari, and H. Zhou, "Locality-Driven Dynamic GPU Cache Bypassing," in *ICS*, 2015.
- [79] D. Li, M. Rhu, D. Johnson, M. O'Connor, M. Erez, D. Burger, D. Fussell, and S. Redder, "Priority-Based Cache Allocation in Throughput Processors," in *HPCA*, 2015.
- [80] Y. Li, S. Ghose, J. Choi, J. Sun, H. Wang, and O. Mutlu, "Utility-Based Hybrid Memory Management," in *CLUSTER*, 2017.
- [81] H. Liu, M. Ferdman, J. Huh, and D. Burger, "Cache Bursts: A New Approach for Eliminating Dead Blocks and Increasing Cache Efficiency," in *MICRO*, 2008.
- [82] W. Liu, P. Huang, T. Kun, T. Lu, K. Zhou, C. Li, and X. He, "LAMs: A Latency-aware Memory Scheduling Policy for Modern DRAM Systems," in *IPCCC*, 2016.
- [83] G. H. Loh, "3D-stacked Memory Architectures for Multi-core Processors," in *ISCA*, 2008.
- [84] Y. Lu, J. Shu, L. Sun, and O. Mutlu, "Loose-Ordering Consistency for Persistent Memory," in *ICCD*, 2014.
- [85] K. Luo, J. Gummaraju, and M. Franklin, "Balancing Throughput and Fairness in SMT Processors," in *ISPASS*, 2001.
- [86] Y. Luo, S. Govindan, B. Sharma, M. Santaniello, J. Meza, A. Kansal, J. Liu, B. Kheshib, K. Vaid, and O. Mutlu, "Characterizing Application Memory Error Vulnerability to Optimize Datacenter Cost via Heterogeneous-Reliability Memory," in *DSN*, 2014.
- [87] V. Mekkat, A. Holey, P.-C. Yew, and A. Zhai, "Managing Shared Last-Level Cache in a Heterogeneous Multicore Processor," in *PACT*, 2013.

- [88] G. Memik, G. Reinman, and W. H. Mangione-Smith, "Just Say No: Benefits of Early Cache Miss Determination," in *HPCA*, 2003.
- [89] J. Meng, D. Tarjan, and K. Skadron, "Dynamic Warp Subdivision for Integrated Branch and Memory Divergence Tolerance," in *ISCA*, 2010.
- [90] J. Meza, Y. Luo, S. Khan, J. Zhao, Y. Xie, and O. Mutlu, "A Case for Efficient Hardware/Software Cooperative Management of Storage and Memory," in *WEED*, 2013.
- [91] J. Meza, J. Chang, H. Yoon, O. Mutlu, and P. Ranganathan, "Enabling Efficient and Scalable Hybrid Memories Using Fine-Granularity DRAM Cache Management," *IEEE CAL*, 2012.
- [92] Micron Technology, Inc., "576Mb: x18, x36 RLLDRAM3," 2011.
- [93] T. Moscibroda and O. Mutlu, "Memory Performance Attacks: Denial of Memory Service in Multi-core Systems," in *USENIX Security*, 2007.
- [94] T. Moscibroda and O. Mutlu, "Distributed Order Scheduling and Its Application to Multi-core DRAM Controllers," in *PODC*, 2008.
- [95] J. Mukundan and J. F. Martinez, "MORSE: Multi-objective Reconfigurable Self-optimizing Memory Scheduler," in *HPCA*, 2012.
- [96] S. P. Muralidhara, L. Subramanian, O. Mutlu, M. Kandemir, and T. Moscibroda, "Reducing Memory Interference in Multicore Systems via Application-Aware Memory Channel Partitioning," in *MICRO*, 2011.
- [97] O. Mutlu, "Memory Scaling: A Systems Architecture Perspective," in *IMW*, 2013.
- [98] O. Mutlu and T. Moscibroda, "Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors," in *MICRO*, 2007.
- [99] O. Mutlu and T. Moscibroda, "Parallelism-Aware Batch Scheduling: Enhancing Both Performance and Fairness of Shared DRAM Systems," in *ISCA*, 2008.
- [100] O. Mutlu and L. Subramanian, "Research Problems and Opportunities in Memory Systems," *SUPERFRI*, 2014.
- [101] V. Narasiman, M. Shebanow, C. J. Lee, R. Miftakhutdinov, O. Mutlu, and Y. N. Patt, "Improving GPU Performance via Large Warps and Two-Level Warp Scheduling," in *MICRO*, 2011.
- [102] NVIDIA Corp., "CUDA C/C++ SDK Code Samples," <http://developer.nvidia.com/cuda-cc-sdk-code-samples>, 2011.
- [103] S. O. Y. H. Son, N. S. Kim, and J. H. Ahn, "Row-Buffer Decoupling: A Case for Low-Latency DRAM Microarchitecture," in *ISCA*, 2014.
- [104] J. Park, R. M. Yoo, D. S. Khudia, C. J. Hughes, and D. Kim, "Location-aware Cache Management for Many-core Processors with Deep Cache Hierarchy," in *SC*, 2013.
- [105] G. Pekhimenko, E. Bolotin, N. Vijaykumar, O. Mutlu, T. C. Mowry, and S. W. Keckler, "A Case for Toggle-aware Compression for GPU Systems," in *HPCA*, 2016.
- [106] G. Pekhimenko, T. Huberty, R. Cai, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "Exploiting Compressed Block Size as an Indicator of Future Reuse," in *HPCA*, 2015.
- [107] G. Pekhimenko, V. Seshadri, Y. Kim, H. Xin, O. Mutlu, M. A. Kozuch, P. B. Gibbons, and T. C. Mowry, "Linearly Compressed Pages: A Main Memory Compression Framework with Low Complexity and Low Latency," in *MICRO*, 2013.
- [108] G. Pekhimenko, V. Seshadri, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "Base-delta-immediate Compression: Practical Data Compression for On-chip Caches," in *PACT*, 2012.
- [109] S. Phadke and S. Narayanasamy, "MLP Aware Heterogeneous Memory System," in *DATE*, 2011.
- [110] M. K. Qureshi, A. Jaleel, Y. N. Patt, S. C. Steely, and J. Emer, "Adaptive Insertion Policies for High Performance Caching," in *ISCA*, 2007.
- [111] M. K. Qureshi, J. Karidis, M. Franceschini, V. Srinivasan, L. Lastras, and B. Abali, "Enhancing Lifetime and Security of PCM-based Main Memory with Start-gap Wear Leveling," in *MICRO*, 2009.
- [112] M. K. Qureshi, D. N. Lynch, O. Mutlu, and Y. N. Patt, "A Case for MLP-Aware Cache Replacement," in *ISCA*, 2006.
- [113] M. K. Qureshi, V. Srinivasan, and J. A. Rivers, "Scalable High Performance Main Memory System Using Phase-change Memory Technology," in *ISCA*, 2009.
- [114] J. Ren, J. Zhao, S. Khan, J. Choi, Y. Wu, and O. Mutlu, "ThyNVM: Enabling software-transparent crash consistency in persistent memory systems," in *MICRO*, 2015.
- [115] S. Rixner, W. J. Dally, U. J. Kapasi, P. Mattson, and J. D. Owens, "Memory Access Scheduling," in *ISCA*, 2000.
- [116] T. G. Rogers, M. O'Connor, and T. M. Aamodt, "Cache-Conscious Wavefront Scheduling," in *MICRO*, 2012.
- [117] T. G. Rogers, M. O'Connor, and T. M. Aamodt, "Divergence-Aware Warp Scheduling," in *MICRO*, 2013.
- [118] SAFARI Research Group, "MeDiC - GitHub Repository," <https://github.com/CMU-SAFARI/MeDiC>.
- [119] Y. Sato *et al.*, "Fast cycle RAM (FCRAM): A 20-ns Random Row Access, Pipelined Operating DRAM," in *VLSIC*, 1998.
- [120] V. Seshadri, A. Bhowmick, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "The Dirty-Block Index," in *ISCA*, 2014.
- [121] V. Seshadri *et al.*, "RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization," in *ISCA*, 2013.
- [122] V. Seshadri, D. Lee, T. Mullins, H. Hassan, A. Boroumand, J. Kim, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry, "Ambit: In-memory Accelerator for Bulk Bitwise Operations using Commodity DRAM Technology," in *MICRO*, 2017.
- [123] V. Seshadri, O. Mutlu, M. A. Kozuch, and T. C. Mowry, "The Evicted-Address Filter: A Unified Mechanism to Address Both Cache Pollution and Thrashing," in *PACT*, 2012.
- [124] V. Seshadri, S. Yedkar, H. Xin, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "Mitigating Prefetcher-Caused Pollution Using Informed Caching Policies for Prefetched Blocks," *ACM TACO*, vol. 11, no. 4, pp. 51:1–51:22, 2015.
- [125] W. Shin, J. Yang, J. Choi, and L.-S. Kim, "NUAT: A Non-Uniform Access Time Memory Controller," in *HPCA*, 2014.
- [126] I. Singh, A. Shriraman, W. W. L. Fung, M. O'Connor, and T. M. Aamodt, "Cache Coherence for GPU Architectures," in *HPCA*, 2013.
- [127] SiSoftware, "Benchmarks : Measuring GP (GPU/APU) Cache and Memory Latencies," <http://www.sisoftware.net>, 2014.
- [128] B. J. Smith, "A Pipelined, Shared Resource MIMD Computer," in *ICPP*, 1978.
- [129] Y. H. Son, S. O. Y. Ro, J. W. Lee, and J. H. Ahn, "Reducing Memory Access Latency with Asymmetric DRAM Bank Organizations," in *ISCA*, 2013.
- [130] S. Srinath, O. Mutlu, H. Kim, and Y. Patt, "Feedback Directed Prefetching: Improving the Performance and Bandwidth-Efficiency of Hardware Prefetchers," in *HPCA*, 2007.
- [131] J. Stuecheli, D. Kaseridis, D. Daly, H. C. Hunter, and L. K. John, "The Virtual Write Queue: Coordinating DRAM and Last-level Cache Policies," in *ISCA*, 2010.
- [132] L. Subramanian, D. Lee, V. Seshadri, H. Rastogi, and O. Mutlu, "BLISS: Balancing Performance, Fairness and Complexity in Memory Access Scheduling," in *IEEE TPDS*, 2016.
- [133] L. Subramanian, D. Lee, V. Seshadri, H. Rastogi, and O. Mutlu, "The Blacklisting Memory Scheduler: Achieving high performance and fairness at low cost," in *ICCD*, 2014.
- [134] L. Subramanian, V. Seshadri, A. Ghosh, S. Khan, and O. Mutlu, "The Application Slowdown Model: Quantifying and Controlling the Impact of Inter-application Interference at Shared Caches and Main Memory," in *MICRO*, 2015.
- [135] L. Subramanian, V. Seshadri, Y. Kim, B. Jaiyen, and O. Mutlu, "MISE: Providing Performance Predictability and Improving Fairness in Shared Main Memory Systems," in *HPCA*, 2013.
- [136] J. E. Thornton, "Parallel Operation in the Control Data 6600," in *AFIPS FJCC*, 1964.
- [137] G. Tyson, M. Farrens, J. Matthews, and A. R. Pleszkun, "A Modified Approach to Data Cache Management," in *MICRO*, 1995.
- [138] H. Usui *et al.*, "DASH: Deadline-Aware High-Performance Memory Scheduler for Heterogeneous Systems with Hardware Accelerators," *ACM TACO*, 2016.
- [139] N. Vijaykumar, K. Hsieh, G. Pekhimenko, S. Khan, A. Shrestha, S. Ghose, A. Jog, P. B. Gibbons, and O. Mutlu, "Zorua: A Holistic Approach to Resource Virtualization in GPUs," in *MICRO*, 2016.
- [140] N. Vijaykumar, G. Pekhimenko, A. Jog, A. Bhowmick, R. Ausavarungnirun, C. Das, M. Kandemir, T. C. Mowry, and O. Mutlu, "A Case for Core-Assisted Bottleneck Acceleration in GPUs: Enabling Flexible Data Compression with Assist Warps," in *ISCA*, 2015.
- [141] F. A. Ware and C. Hampel, "Improving Power and Data Efficiency with Threaded Memory Modules," in *ICCD*, 2006.
- [142] H. Wong, M.-M. Papadopoulou, M. Sadooghi-Alvandi, and A. Moshovos, "Demystifying GPU Microarchitecture Through Microbenchmarking," in *ISPASS*, 2010.
- [143] Y. Wu, R. Rakvic, L.-L. Chen, C.-C. Miao, G. Chrysos, and J. Fang, "Compiler Managed Micro-cache Bypassing for High Performance EPIC Processors," in *MICRO*, 2002.
- [144] L. Xiang, T. Chen, Q. Shi, and W. Hu, "Less Reused Filter: Improving L2 Cache Performance via Filtering Less Reused Lines," in *ICS*, 2009.
- [145] X. Xie, Y. Liang, G. Sun, and D. Chen, "An Efficient Compiler Framework for Cache Bypassing on GPUs," in *ICCAD*, 2013.
- [146] X. Xie, Y. Liang, Y. Wang, G. Sun, and T. Wang, "Coordinated Static and Dynamic Cache Bypassing for GPUs," in *HPCA*, 2015.
- [147] D. Xiong, K. Huang, X. Jiang, and X. Yan, "Memory Access Scheduling Based on Dynamic Multilevel Priority in Shared DRAM Systems," *ACM TACO*, vol. 13, no. 4, Dec. 2016.
- [148] H. Yoon, J. Meza, R. Ausavarungnirun, R. Harding, and O. Mutlu, "Row Buffer Locality Aware Caching Policies for Hybrid Memories," in *ICCD*, 2012.
- [149] B. Yu, J. Ma, T. Chen, and M. Wu, "Global Priority Table for Last-Level Caches," in *DASC*, 2011.
- [150] X. Yu, C. J. Hughes, N. Satish, O. Mutlu, and S. Devadas, "Banshee: Bandwidth-efficient DRAM Caching via Software/Hardware Cooperation," in *MICRO*, 2017.
- [151] G. Yuan, A. Bakhoda, and T. Aamodt, "Complexity Effective Memory Access Scheduling for Many-Core Accelerator Architectures," in *MICRO*, 2009.
- [152] C. Zhang, G. Sun, P. Li, T. Wang, D. Niu, and Y. Chen, "SBAC: A Statistics Based Cache Bypassing Method for Asymmetric-access Caches," in *ISPLED*, 2014.
- [153] J. Zhao, O. Mutlu, and Y. Xie, "FIRM: Fair and High-Performance Memory Control for Persistent Memory Systems," in *MICRO*, 2014.
- [154] H. Zheng, J. Lin, Z. Zhang, E. Gorbatov, H. David, and Z. Zhu, "Mini-rank: Adaptive DRAM Architecture for Improving Memory Power Efficiency," in *MICRO*, 2008.
- [155] Z. Zheng, Z. Wang, and M. Lipasti, "Adaptive Cache and Concurrency Allocation on GPGPUs," *IEEE CAL*, 2014.
- [156] W. K. Zuravleff and T. Robinson, "Controller for a Synchronous DRAM That Maximizes Throughput by Allowing Memory Requests and Commands to Be Issued Out of Order," US Patent No. 5,630,096, 1997.