

DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks

Geraldo F. Oliveira¹ Juan Gómez-Luna¹ Lois Orosa¹ Saugata Ghose²
Nandita Vijaykumar³ Ivan Fernandez^{1,4} Mohammad Sadrosadati¹ Onur Mutlu²

¹ETH Zürich ²University of Illinois Urbana-Champaign ³University of Toronto ⁴University of Malaga

Abstract

Data movement between the CPU and main memory is a first-order obstacle against improving performance, scalability, and energy efficiency in modern systems. Computer systems employ a range of techniques to reduce overheads tied to data movement, spanning from traditional mechanisms (e.g., deep multi-level cache hierarchies, aggressive hardware prefetchers) to emerging techniques such as Near-Data Processing (NDP), where some computation is moved close to memory. Prior NDP works investigate the root causes of data movement bottlenecks using different profiling methodologies and tools. However, there is still a lack of understanding about the key metrics that can identify different data movement bottlenecks and their relation to traditional and emerging data movement mitigation mechanisms. Our goal is to methodically identify potential sources of data movement over a broad set of applications and to comprehensively compare traditional compute-centric data movement mitigation techniques (e.g., caching and prefetching) to more memory-centric techniques (e.g., NDP), thereby developing a rigorous understanding of the best techniques to mitigate each source of data movement.

With this goal in mind, we perform the first large-scale characterization of a wide variety of applications, across a wide range of application domains, to identify fundamental program properties that lead to data movement to/from main memory. We develop the first systematic methodology to classify applications based on the sources contributing to data movement bottlenecks. From our large-scale characterization of 77K functions across 345 applications, we select 144 functions to form the first open-source benchmark suite (DAMOV) for main memory data movement studies. We select a diverse range of functions that (1) represent different types of data movement bottlenecks, and (2) come from a wide range of application domains. Using NDP as a case study, we identify new insights about the different data movement bottlenecks and use these insights to determine the most suitable data movement mitigation mechanism for a particular application. We open-source DAMOV and the complete source code for our new characterization methodology at <https://github.com/CMU-SAFARI/DAMOV>.

Keywords

benchmarking, data movement, energy, memory systems, near-data processing, performance, processing-in-memory, workload characterization, 3D-stacked memory

1 Introduction

Today’s computing systems require moving data from main memory (consisting of DRAM) to the CPU cores so that computation can take place on the data. Unfortunately, this *data movement* is a major bottleneck for system performance and energy consumption [1, 2].

DRAM technology scaling is failing to keep up with the increasing memory demand from applications [2–29], resulting in significant latency and energy costs due to data movement [1–3, 5, 6, 30–49]. High-performance systems have evolved to include mechanisms that aim to alleviate data movement’s impact on system performance and energy consumption, such as deep cache hierarchies and aggressive prefetchers. However, such mechanisms not only come with significant hardware cost and complexity, but they also often fail to hide the latency and energy costs of accessing DRAM in many modern and emerging applications [1, 5, 50–52]. These applications’ memory behavior can differ significantly from more traditional applications since modern applications often have lower memory locality, more irregular access patterns, and larger working sets [36, 45, 46, 53–61]. One promising technique that aims to alleviate the data movement bottleneck in modern and emerging applications is Near-Data Processing (NDP) [1, 33, 34, 46–48, 54, 55, 59–118],¹ where the cost of data movement to/from main memory is reduced by placing computation capability close to memory. In NDP, the computational logic close to memory has access to data that resides in main memory with significantly higher memory bandwidth, lower latency, and lower energy consumption than the CPU has in existing systems. There is very high bandwidth available to the cores in the logic layer of 3D-stacked memories, as demonstrated by many past works (e.g., [1, 46, 59, 60, 62–64, 67–69, 74, 76, 99, 119]). To illustrate this, we use the STREAM Copy [120] workload to measure the peak memory bandwidth the host CPU and an NDP architecture with processing elements in the logic layer of a single 3D-stacked memory (e.g., Hybrid Memory Cube [73]) can leverage.² We observe that the peak memory bandwidth that the NDP logic can leverage (431 GB/s) is 3.7× the peak memory bandwidth that the host CPU can exploit (115 GB/s). This happens since the external memory bandwidth is bounded by the limited number of I/O pins available in the DRAM device [121].

Many recent works explore how NDP can benefit various application domains, such as graph processing [46, 47, 54, 63, 74, 93, 122–126], machine learning [1, 61, 69, 70, 84, 85, 103], bioinformatics [59, 60, 68], databases [55, 61, 63, 66, 67, 74, 86, 102], security [71, 105, 106], data manipulation [49, 86, 88, 89, 127–130], and mobile workloads [1, 61]. These works demonstrate that simple metrics such as last-level CPU cache Misses per Kilo-Instruction (MPKI) and Arithmetic Intensity (AI) are useful metrics that serve as a proxy for the amount of data movement an application experiences. These metrics can be used as a potential guide for choosing when to apply data movement mitigation mechanisms such as NDP. However, such metrics (and the corresponding insights) are

¹We use the term NDP to refer to *any* type of Processing-in-Memory [37].

²See Section 2 for our experimental evaluation methodology.

often extracted from a small set of applications, with similar or not-rigorously-analyzed data movement characteristics. Therefore, it is difficult to generalize the metrics and insights these works provide to a broader set of applications, making it unclear what different metrics can reveal about a new (i.e., previously uncharacterized) application’s data movement behavior (and how to mitigate its associated data movement costs).

We illustrate this issue by highlighting the limitations of two different methodologies commonly used to identify memory bottlenecks and often used as a guide to justify the use of NDP architectures for an application: (a) analyzing a roofline model [131] of the application, and (b) using last-level CPU cache MPKI as an indicator of NDP suitability of the application. The roofline model correlates the computation requirements of an application with its memory requirements under a given system. The model contains two roofs: (1) a diagonal line ($y = \text{Peak Memory Bandwidth} \times \text{Arithmetic Intensity}$) called the *memory roof*, and (2) a horizontal line ($y = \text{Peak System Throughput}$) called the *compute roof* [131]. If an application lies under the memory roof, the application is classified as *memory-bound*; if an application lies under the compute roof, it is classified as *compute-bound*. Many prior works [99, 103, 132–144] employ this roofline model to identify memory-bound applications that can benefit from NDP architectures. Likewise, many prior works [1, 36, 51, 54, 55, 145–150] observe that applications with high last-level cache MPKI³ are good candidates for NDP.

Figure 1 shows the roofline model (left) and a plot of MPKI vs. speedup (right) of a system with general-purpose NDP support over a baseline system without NDP for a diverse set of 44 applications (see Table 8). In the MPKI vs. speedup plot, the MPKI corresponds to a baseline host CPU system. The speedup represents the performance improvement of a general-purpose NDP system over the baseline (see Section 2.4 for our methodology). We make the following observations. First, analyzing the roofline model (Figure 1, left), we observe that most of the memory-bound applications (yellow dots) benefit from NDP, as foreseen by prior works. We later observe (in Section 3.3.1) that such applications are DRAM bandwidth-bound and are a natural fit for NDP. However, the roofline model does *not* accurately account for the NDP suitability of memory-bound applications that (i) benefit from NDP only under particular microarchitectural configurations, e.g., either at low or high core counts (green dots, which are applications that are either bottlenecked by DRAM latency or suffer from L3 cache contention; see Sections 3.3.3 and 3.3.4); or (ii) experience performance degradation when executed using NDP (blue dots, which are applications that suffer from the lack of a deep cache hierarchy in NDP architectures; see Section 3.3.6). Second, analyzing the MPKI vs. speedup plot (Figure 1, right), we observe that while all applications with high MPKI benefit from NDP (yellow dots with MPKI higher than 10), some applications with *low* MPKI can *also* benefit from NDP in all of the NDP microarchitecture configurations we evaluate (yellow dots with MPKI lower than 10) or under specific NDP microarchitecture configurations (green dots with MPKI lower than 10). Thus, even though both the roofline model

and MPKI can identify some specific sources of memory bottlenecks and can sometimes be used as a proxy for NDP suitability, they alone cannot definitively determine NDP suitability because they cannot comprehensively identify different possible sources of memory bottlenecks in a system.

Our *goal* in this work is (1) to understand the major sources of inefficiency that lead to data movement bottlenecks by observing and identifying relevant metrics and (2) to develop a benchmark suite for data movement that captures each of these sources. To this end, we develop a new three-step methodology to correlate application characteristics with the *primary* sources of data movement bottlenecks and to determine the potential benefits of three example data movement mitigation mechanisms: (1) a deep cache hierarchy, (2) a hardware prefetcher, and (3) a general-purpose NDP architecture.⁴ We use two main profiling strategies to gather key metrics from applications: (i) an architecture-independent profiling tool and (ii) an architecture-dependent profiling tool. The architecture-independent profiling tool provides metrics that characterize the application memory behavior independently of the underlying hardware. In contrast, the architecture-dependent profiling tool evaluates the impact of the system configuration (e.g., cache hierarchy) on the memory behavior. Our methodology has three steps. In *Step 1*, we use a hardware profiling tool to identify memory-bound functions across many applications. This step allows for a quick first-level identification of many applications that suffer from memory bottlenecks and functions that cause these bottlenecks. In *Step 2*, we use the architecture-independent profiling tool to collect metrics that provide insights about the memory access behavior of the memory-bottlenecked functions. In *Step 3*, we collect architecture-dependent metrics and analyze the performance and energy of each function in an application when each of our three candidate data movement mitigation mechanisms is applied to the system. By combining the data obtained from all three steps, we can systematically classify the leading causes of data movement bottlenecks in an application or function into different bottleneck classes.

Using this new methodology, we characterize a large, heterogeneous set of applications (345 applications from 37 different workload suites) across a wide range of domains. Within these applications, we analyze 77K functions and find a subset of 144 functions from 74 different applications that are memory-bound (and that consume a significant fraction of the overall execution time). We fully characterize this set of 144 representative functions to serve as a core set of application kernel benchmarks, which we release as the open-source DAMOV (Data MOVement) Benchmark Suite [158]. Our analyses reveal six new insights about the sources of memory bottlenecks and their relation to NDP:

- (1) Applications with high last-level cache MPKI and low temporal locality are *DRAM bandwidth-bound*. These applications benefit from the large memory bandwidth available to the NDP system (Section 3.3.1).
- (2) Applications with low last-level cache MPKI and low temporal locality are *DRAM latency-bound*. These applications do *not*

³Typically, an MPKI value greater than 10 is considered *high* by prior works [151–157].

⁴We focus on these three data movement mitigation mechanisms for two different reasons: (1) deep cache hierarchies and hardware prefetchers are standard mechanisms in almost all modern systems, and (2) NDP represents a promising paradigm shift for many modern data-intensive applications.




Figure 1: Roofline (left) and last-level cache MPKI vs. NDP speedup (right) for 44 memory-bound applications. Applications are classified into four categories: (1) those that experience performance degradation due to NDP (blue; Faster on CPU), (2) those that experience performance improvement due to NDP (yellow; Faster on NDP), (3) those where the host CPU and NDP performance are similar (red; Similar on CPU/NDP), (4) those that experience either performance degradation or performance improvement due to NDP depending on the microarchitectural configuration (green; Depends).

benefit from L2/L3 caches. The NDP system improves performance and energy efficiency by sending L1 misses directly to DRAM (Section 3.3.2).

- (3) A second group of applications with low LLC MPKI and low temporal locality are *bottlenecked by L1/L2 cache capacity*. These applications benefit from the NDP system at low core counts. However, at high core counts (and thus larger L1/L2 cache space), the caches capture most of the data locality in these applications, decreasing the benefits the NDP system provides (Section 3.3.3). We make this observation using a *new* metric that we develop, called *last-to-first miss-ratio (LFMR)*, which we define as the ratio between the number of LLC misses and the total number of L1 cache misses. We find that this metric accurately identifies how efficient the cache hierarchy is in reducing data movement.
- (4) Applications with high temporal locality and low LLC MPKI are *bottlenecked by L3 cache contention* at high core counts. In such cases, the NDP system provides a cost-effective way to alleviate cache contention over increasing the L3 cache capacity (Section 3.3.4).
- (5) Applications with high temporal locality, low LLC MPKI, and low AI are bottlenecked by the *L1 cache capacity*. The three candidate data movement mitigation mechanisms achieve similar performance and energy consumption for these applications (Section 3.3.5).
- (6) Applications with high temporal locality, low LLC MPKI, and high AI are *compute-bound*. These applications benefit from a deep cache hierarchy and hardware prefetchers, but the NDP system degrades their performance (Section 3.3.6).

We publicly release our 144 representative data movement bottlenecked functions from 74 applications as the first open-source benchmark suite for data movement, called DAMOV Benchmark Suite, along with the complete source code for our new characterization methodology [158].

This work makes the following key contributions:

- We propose the first methodology to characterize data-intensive workloads based on the source of their data movement bottlenecks. This methodology is driven by insights obtained from a large-scale experimental characterization of 345 applications from 37 different benchmark suites and an evaluation of the performance of memory-bound functions from these applications with three data-movement mitigation mechanisms.
- We release DAMOV, the first open-source benchmark suite for main memory data movement-related studies, based on our systematic characterization methodology. This suite consists of 144 functions representing different sources of data movement bottlenecks and can be used as a baseline benchmark set for future data-movement mitigation research.
- We show how our DAMOV benchmark suite can aid the study of open research problems for NDP architectures, via four case studies. In particular, we evaluate (i) the impact of load balance and inter-vault communication in NDP systems, (ii) the impact of NDP accelerators on our memory bottleneck analysis, (iii) the impact of different core models on NDP architectures, and (iv) the potential benefits of identifying simple NDP instructions. We conclude that our benchmark suite and methodology can be employed to address many different open research and development questions on data movement mitigation mechanisms, particularly topics related to NDP systems and architectures.

2 Methodology Overview

We develop a new workload characterization methodology to analyze data movement bottlenecks and the suitability of different data movement mitigation mechanisms for these bottlenecks, with a focus on Near-Data Processing (NDP). Our methodology consists of three main steps, as Figure 2 depicts: (1) *memory-bound function identification* using application profiling; (2) *locality-based clustering* to analyze spatial and temporal locality in an architecture-independent manner; and (3) *memory bottleneck classification* using

a scalability analysis to nail down the sources of memory boundedness, including architecture-dependent characterization. Our methodology takes as input an application’s source code and its input datasets, and produces as output a classification of the primary source of memory bottleneck of important functions in an application (i.e., bottleneck class of each key application function). We illustrate the applicability of this methodology with a detailed characterization of 144 functions that we select from among 77K analyzed functions of 345 characterized applications. In this section, we give an overview of our workload characterization methodology. We use this methodology to drive the analyses we perform in Section 3.

2.1 Experimental Evaluation Framework

As our scalability analysis depends on the hardware architecture, we need a hardware platform that can allow us to replicate and control all of our configuration parameters. Unfortunately, such an analysis cannot be performed practically using real hardware, as (1) there are very few available NDP hardware platforms, and the ones that currently exist do not allow us to comprehensively analyze our general-purpose NDP configuration in a controllable way (as existing platforms are specialized and non-configurable); and (2) the configurations of real CPUs can vary significantly across the range of core counts that we want to analyze, eliminating the possibility of a carefully controlled study. As a result, we must rely on accurate simulation platforms to perform an accurate comparison across different configurations. To this end, we build a framework that integrates the ZSim CPU simulator [159] with the Ramulator memory simulator [160] to produce a fast, scalable, and cycle-accurate open-source simulator called DAMOV-SIM [158]. We use ZSim to simulate the core microarchitecture, cache hierarchy, coherence protocol, and prefetchers. We use Ramulator to simulate the DRAM architecture, memory controllers, and memory accesses. To compute spatial and temporal locality, we modify ZSim to generate a single-thread memory trace for each application, which we use as input for the locality analysis algorithm described in Section 2.3 (which statically computes the temporal and spatial locality at word-level granularity).

2.2 Step 1: Memory-Bound Function Identification

The first step (labeled ❶ in Figure 2) aims to identify the functions of an application that are *memory-bound* (i.e., functions that suffer from data movement bottlenecks). These bottlenecks might be caused at any level of the memory hierarchy. There are various potential sources of memory boundedness, such as cache misses, cache coherence traffic, and long queuing latencies. Therefore, we need to take all such potential causes into account. This step is optional if the application’s memory-bound functions (i.e., regions of interest, *roi*, in Figure 2) are already known *a priori*.

Hardware profiling tools, both open-source and proprietary, are available to obtain hardware counters and metrics that characterize the application behavior on a computing system. In this work, we use the Intel VTune Profiler [161], which implements the well-known *top-down analysis* [162]. Top-down analysis uses the available CPU hardware counters to hierarchically identify different sources of CPU system bottlenecks for an application. Among the

various metrics measured by top-down analysis, there is a relevant one called *Memory Bound* [163] that measures the percentage of CPU pipeline slots that are *not* utilized due to any issue related to data access. We employ this metric to identify functions that suffer from data movement bottlenecks (which we define as functions where *Memory Bound* is greater than 30%).

2.3 Step 2: Locality-Based Clustering

Two key properties of an application’s memory access pattern are its inherent spatial locality (i.e., the likelihood of accessing nearby memory locations in the near future) and temporal locality (i.e., the likelihood of accessing a memory location again in the near future). These properties are closely related to how well the application can exploit the memory hierarchy in computing systems and how accurate hardware prefetchers can be. Therefore, to understand the sources of memory bottlenecks for an application, we should analyze how much spatial and temporal locality its memory accesses inherently exhibit. However, we should isolate these properties from particular configurations of the memory subsystem. Otherwise, it would be unclear if memory bottlenecks are due to the nature of the memory accesses or due to the characteristics and limitations of the memory subsystem (e.g., limited cache size, too simple or inaccurate prefetching policies). As a result, in this step (labeled ❷ in Figure 2), we use *architecture-independent* static analysis to obtain spatial and temporal locality metrics for the functions selected in the previous step (Section 2.2). Past works [164–173] propose different ways of analyzing spatial and temporal locality in an architecture-independent manner. In this work, we use the definition of spatial and temporal metrics presented in [166, 167].

The spatial locality metric is calculated for a window of memory references⁵ of length W using Equation 1. First, for every W memory references, we calculate the minimum distance between any two addresses (*stride*). Second, we create a histogram called the *stride profile*, where each bin i stores how many times each *stride* appears. Third, to calculate the spatial locality, we divide the *percentage* of times *stride* i is referenced (*stride profile*(i)) by the *stride length* i and sum the resulting value across all instances of i .

$$\text{Spatial Locality} = \sum_{i=1}^{\text{\#bins}} \frac{\text{stride profile}(i)}{i} \quad (1)$$

A spatial locality value close to 0 is caused by large *stride* values (e.g., regular accesses with large strides) or random accesses, while a value equal to 1 is caused by a completely sequential access pattern.

The temporal locality metric is calculated by using a histogram of reused addresses. First, we count the number of times each memory address is repeated in a window of L memory references. Second, we create a histogram called *reuse profile*, where each bin i represents the number of times a memory address is reused, expressed as a power of 2. For each memory address, we increment the bin that represents the corresponding number of repetitions. For example, *reuse profile*(0) represents memory addresses that are reused only once. *reuse profile*(1) represents memory addresses that are reused twice. Thus, if a memory address is reused N times, we increment

⁵We compute both the spatial and temporal locality metrics at the word granularity. In this way, we keep our locality analysis architecture-independent, using *only* properties of the application under study.

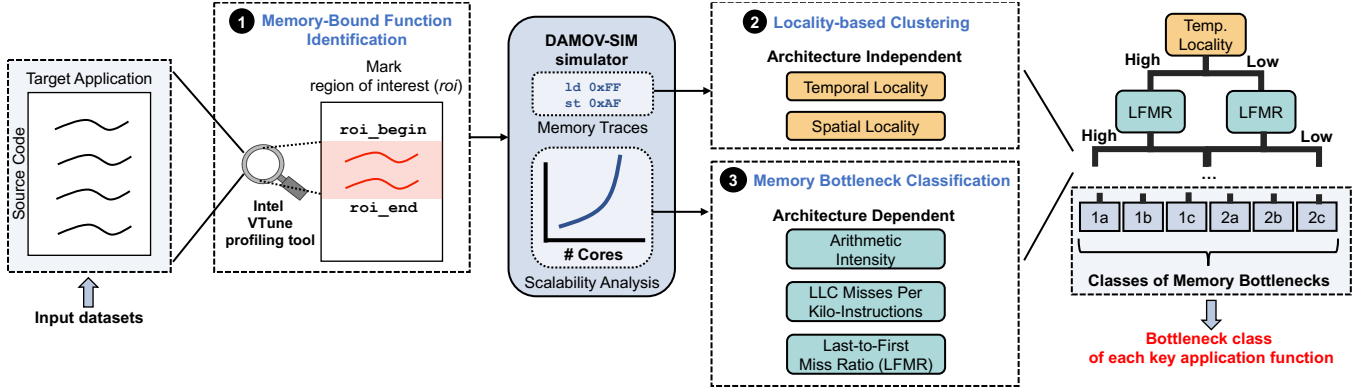


Figure 2: Overview of our three-step workload characterization methodology.

$reuse\ profile(\lfloor \log_2 N \rfloor)$ by one. Third, we obtain the temporal locality metric with Equation 2.

$$Temporal\ Locality = \sum_{i=0}^{\#bins} \frac{2^i \times reuse\ profile(i)}{total\ memory\ accesses} \quad (2)$$

A temporal locality value of 0 indicates no data reuse, while a value close to 1 indicates very high data reuse (i.e., a value equal to 1 means that the application accesses a single memory address continuously).

To calculate these metrics, we empirically select window lengths W and L to 32. We find that different values chosen for W and L do not significantly change the conclusions of our analysis. We observe that our conclusions remain the same when we set both values to 8, 16, 32, 64, and 128.

2.4 Step 3: Bottleneck Classification

While Step 2 allows us to understand inherent application sources for memory boundedness, it is important to understand how hardware architectural features can also result in memory bottlenecks. As a result, in our third step (3 in Figure 2), we perform a scalability analysis of the functions selected in Step 1, where we evaluate performance and energy scaling for three different system configurations. The scalability analysis makes use of three *architecture-dependent* metrics: (1) *Arithmetic Intensity (AI)*, (2) *Misses per Kilo-Instruction (MPKI)*, and (3) a new metric called *Last-to-First Miss-Ratio (LFMR)*. We select these metrics for the following reasons. First, AI can measure the compute intensity of an application. Intuitively, we expect an application with high compute intensity to not suffer from severe data movement bottlenecks, as demonstrated by prior work [174]. Second, MPKI serves as a proxy for the memory intensity of an application. It can also indicate the memory pressure experienced by the main memory system [45, 47, 48, 58, 151, 153, 156, 175–177]. Third, LFMR, a new metric we introduce and is described in detail later in this subsection, indicates how efficient the cache hierarchy is in reducing data movement.

As part of our methodology development, we evaluate other metrics related to data movement, including raw cache misses, coherence traffic, and DRAM row misses/hits/conflicts. We observe that even though such metrics are useful for further characterizing an application (as we do in some of our later analyses in Section 3.3), they do not necessarily characterize a specific type of

data movement bottleneck. We show in Section 4.1 that the three architecture-dependent and two architecture-independent metrics we select for our classification are enough to accurately characterize and cluster the different types of data movement bottlenecks in a wide variety of applications.

2.4.1 Definition of Metrics. We define Arithmetic Intensity (AI) as the number of arithmetic and logic operations performed per L1 cache line accessed.⁶ This metric indicates how much computation there is per memory request. Intuitively, applications with high AI are likely to be computationally intensive, while applications with low AI tend to be memory intensive. We use MPKI at the last-level cache (LLC), i.e., the number of LLC misses per one thousand instructions. This metric is considered to be a good indicator of NDP suitability by several prior works [1, 36, 51, 54, 55, 145–149]. We define the LFMR of an application as the ratio between the number of LLC misses and the total number of L1 cache misses. We find that this metric accurately identifies how much an application benefits from the deep cache hierarchy of a contemporary CPU. An LFMR value close to 0 means that the number of LLC misses is very small compared to the number of L1 misses, i.e., the L1 misses are likely to hit in the L2 or L3 caches. However, an LFMR value close to 1 means that very few L1 misses hit in L2 or L3 caches, i.e., the application does not benefit much from the deep cache hierarchy, and most L1 misses need to be serviced by main memory.

2.4.2 Scalability Analysis and System Configuration. The goal of the scalability analysis we perform is to nail down the specific sources of data movement bottlenecks in the application. In this analysis, we (i) evaluate the performance and energy scaling of an application in three different system configurations; and (ii) collect the key metrics for our bottleneck classification (i.e., AI, MPKI, and LFMR). During scalability analysis, we simulate three system configurations of a general-purpose multicore processor:

- A host CPU with a deep cache hierarchy (i.e., private L1 (32 kB) and L2 (256 kB) caches, and a shared L3 (8 MB) cache with 16 banks). We call this configuration *Host CPU*.
- A host CPU with a deep cache hierarchy (same cache configurations as in *Host CPU*), augmented with a stream prefetcher [178]. We call this configuration *Host CPU with prefetcher*.

⁶We consider AI to be architecture-dependent since we consider the number of cache lines accessed by the application (and hence the hardware cache block size) to compute the metric. This is the same definition of AI used by the hardware profiling tool we employ in Step 1 (i.e., the Intel VTune Profiler [161]).

- An NDP CPU with a single level of cache (only a private read-only⁷ L1 cache (32 kB), as assumed in many prior NDP works [1, 46, 51, 63, 66, 74, 99, 101, 119, 179]) and no hardware prefetcher. We call this configuration *NDP*.

The remaining components of the processor configuration are kept the same (e.g., number of cores, instruction window size, branch predictor) to isolate the impact of only the caches, prefetchers, and NDP. This way, we expect that the performance and energy differences between the three configurations to come *exclusively* from the different data movement requirements. For the three configurations, we sweep the number of CPU cores in our analysis from 1 to 256, as previous works [46, 66, 180] show that large core counts are necessary to saturate the bandwidth provided by modern high-bandwidth memories, and because modern CPUs and NDP proposals can have varying core counts. The core count sweep allows us to observe (1) how an application’s performance changes when increasing the pressure on the memory subsystem, (2) how much Memory-Level Parallelism (MLP) [176, 181–184] the application has, and (3) how much the cores leverage the cache hierarchy and the available memory bandwidth. We proportionally increase the size of the CPU’s private L1 and L2 caches when increasing the number of CPU cores in our analysis (e.g., when scaling the CPU core count from 1 to 4, we also scale the aggregated L1/L2 cache size by a factor of 4). We use out-of-order and in-order CPU cores in our analysis for all three configurations. In this way, we build confidence that our trends and findings are independent of a specific underlying general-purpose core microarchitecture. We simulate a memory architecture similar to the Hybrid Memory Cube (HMC) [73], where (1) the host CPU accesses memory through a high-speed off-chip link, and (2) the NDP logic resides in the logic layer of the memory chip and has direct access to the DRAM banks (thus taking advantage of higher memory bandwidth and lower memory latency). Table 1 lists the parameters of our host CPU, host CPU with prefetcher, and NDP baseline configurations.

2.4.3 Choosing an NDP Architecture. We note that across the proposed NDP architectures in literature, there is a lack of consensus on whether the architectures should make use of general-purpose NDP cores or specialized NDP accelerators [36, 37]. In this work, we focus on general-purpose NDP cores for two major reasons. First, many prior works (e.g., [1, 46, 51, 63, 66, 76, 99, 101, 119, 147, 179, 190, 192–194]) suggest that general-purpose cores (especially simple in-order cores) can successfully accelerate memory-bound applications in NDP architectures. In fact, UPMEM [83], a start-up building some of the first commercial in-DRAM NDP systems, utilizes simple in-order cores in their NDP units inside DRAM chips [83, 140]. Therefore, we believe that general-purpose NDP cores are a promising candidate for future NDP architectures. Second, the goal of our work is not to perform a design space exploration of different NDP architectures, but rather to understand the key properties of applications that lead to memory bottlenecks that can be mitigated by a simple NDP engine. While we expect that each application could potentially benefit further from an NDP accelerator tailored to its computational and memory requirements, such

customized architectures open many challenges for a methodical characterization, such as the need for significant code refactoring, changes in data mapping, and code partitioning between NDP accelerators and host CPUs.^{8,9}

3 Characterizing Memory Bottlenecks

In this section, we apply our three-step workload characterization methodology to characterize the sources of memory bottlenecks across a wide range of applications. First, we apply *Step 1* to identify memory-bound functions within an application (Section 3.1). Second, we apply *Step 2* and cluster the identified functions using two architecture-independent metrics (spatial and temporal locality) (Section 3.2). Third, we apply *Step 3* and combine the architecture-dependent and architecture-independent metrics to classify the different sources of memory bottlenecks we observe (Section 3.3).

We also evaluate various other aspects of our three-step workload characterization methodology. We investigate the effect of increasing the last-level cache on our memory bottleneck classification in Section 3.4. We provide a validation of our memory bottleneck classification in Section 3.5. We discuss the limitations of our proposed methodology in Section 3.6.

3.1 Step 1: Memory-Bound Function Identification

We first apply Step 1 of our methodology across 345 applications (listed in Appendix C) to identify functions whose performance is significantly affected by data movement. We use the previously-proposed top-down analysis methodology [162] that has been used by several recent workload characterization studies [5, 195, 196]. As discussed in Section 2.2, we use the Intel VTune Profiler [161], which we run on an Intel Xeon E3-1240 processor [197] with four cores. We disable hyper-threading for more accurate profiling results, as recommended by the VTune documentation [198]. For the applications that we analyze, we select functions (1) that take at least 3% of the clock cycles, and (2) that have a Memory Bound percentage that is greater than 30%. We choose 30% as the threshold for this metric because, in preliminary simulation experiments, we do not observe significant performance improvement or energy savings with data movement mitigation mechanisms for functions whose Memory Bound percentage is less than 30%.

The applications we analyze come from a variety of sources, such as popular workload suites (Chai [199], CORAL [200], Parboil [201], PARSEC [202], Rodinia [203], SD-VBS [204], SPLASH-2 [205]), benchmarking (STREAM [120], HPCC [206], HPCG [207]), bioinformatics [208], databases [209, 210], graph processing frameworks (GraphMat [211], Ligra [212]), a map-reduce framework (Phoenix [213]), and neural networks (AlexNet [214], Darknet [215]). We explore different input dataset sizes for the applications and choose real-world input datasets that impose high pressure on the memory subsystem (as we expect that such real-world inputs are best suited for stressing the memory hierarchy). We also use different inputs for applications whose performance

⁷We use read-only L1 caches to simplify the cache coherence model of the NDP system. Enabling efficient synchronization and cache coherence in NDP architectures is an open-research problem, as we discuss in Section 3.6.

⁸We show in Section 5.2 that our DAMOV benchmark suite is useful to rigorously study NDP accelerators.

⁹The development of a new methodology or extension of our methodology to perform analysis targeting function-specific, customized, or reconfigurable NDP accelerators is a good direction for future work.

Table 1: Evaluated Host CPU and NDP system configurations.

Host CPU Configuration	
Host CPU Processor	1, 4, 16, 64, and 256 cores @2.4 GHz, 32 nm; 4-wide out-of-order 1, 4, 16, 64, and 256 cores @2.4 GHz, 32 nm; 4-wide in-order Buffers: 128-entry ROB; 32-entry LSQ (each) Branch predictor: Two-level GAs [185]. 2,048 entry BTB; 1 branch per fetch
Private L1 Cache	32 KB, 8-way, 4-cycle; 64 B line; LRU policy Energy: 15/33 pJ per hit/miss [51, 186]
Private L2 Cache	256 KB, 8-way, 7-cycle; 64 B line; LRU policy MSHR size: 20-request, 20-write, 10-eviction Energy: 46/93 pJ per hit/miss [51, 186]
Shared L3 Cache	8 MB (16-banks), 0.5 MB per bank, 16-way, 27-cycle 64 B line; LRU policy; Bi-directional ring [187]; Inclusive; MESI protocol [188] MSHR size: 64-request, 64-write, 64-eviction Energy: 945/1904 pJ per hit/miss [51, 186]
Host CPU with Prefetcher Configuration	
Processor, Private L1 Cache, Private L2 Cache, and Share L3 Cache	Same as in Host CPU Configuration
L2 Cache Prefetcher	Stream prefetcher [178, 189]; 2-degree; 16 stream buffers; 64 entries
NDP Configuration	
NDP CPU Processor	1, 4, 16, 64, and 256 cores @2.4 GHz, 32 nm; 4-wide out-of-order 1, 4, 16, 64, and 256 cores @ 2.4 GHz, 32 nm; 4-wide in-order Buffers: 128-entry ROB; 32-entry LSQ (each) Branch predictor: Two-level GAs [185]. 2,048 entry BTB; 1 branch per fetch
Private L1 Cache	32 KB, 8-way, 4-cycle; 64 B line; LRU policy; Read-only Data Cache Energy: 15/33 pJ per hit/miss [51, 186]
Common	
Main Memory	HMC v2.0 Module [73] 32 vaults, 8 DRAM banks/vault, 256 B row buffer 8 GB total size; DRAM@166 MHz; 4-links@8 GHz 8 B burst width at 2:1 core-to-bus freq. ratio; Open-page policy; HMC default interleaving [45, 73] ¹⁰ Energy: 2 pJ/bit internal, 8 pJ/bit logic layer [51, 64, 190], 2 pJ/bit links [51, 76, 191]

is tightly related to the input dataset properties. For example, we use two different graphs with varying connectivity degrees (rMat [216] and USA [217]) to evaluate graph processing applications and two different read sequences to evaluate read alignment algorithms[60, 218, 219].

In total, our application analysis covers more than 77K functions. To date, this is the most extensive analysis of data movement bottlenecks in real-world applications. We find a set of 144 functions that take at least 3% of the total clock cycles and have a value of the Memory Bound metric greater or equal to 30%, which forms the basis of DAMOV, our new data movement benchmark suite. We provide a list of all 144 functions selected based on our analysis and their major characteristics in Appendix A.

After identifying memory-bound functions over a wide range of applications, we apply Steps 2 and 3 of our methodology to classify the primary sources of memory bottlenecks for our selected functions. We evaluate a total of 144 functions out of the 77K functions we analyze in Step 1. These functions span across 74 different applications, belonging to 16 different widely-used benchmark suites or frameworks.

From the 144 functions that we analyze further, we select a subset of 44 representative functions to explore in-depth in Sections 3.2 and 3.3 and to drive our bottleneck classification analysis. We use the 44 representative functions to ease our explanations and make figures more easily readable. Table 8 in Appendix A lists the 44 representative functions that we select. The table includes one column that indicates the class of data movement bottleneck experienced by each function (we discuss the classes in Section 3.3), and another column representing the percentage of clock cycles of the selected function in the whole application. We select representative functions that belong to a variety of domains: benchmarking, bioinformatics, data analytics, databases, data mining, data reorganization, graph processing, neural networks, physics, and signal processing. In Section 3.5, we validate our classification using the remaining 100 functions and provide a summary of the results of our methodology when applied to all 144 functions.

3.2 Step 2: Locality-Based Clustering

We cluster the 44 representative functions across both spatial and temporal locality using the K-means clustering algorithm [221]. Figure 3 shows how each function is grouped. We find that two groups emerge from the clustering: (1) low temporal locality functions (orange boxes in Figure 3), and (2) high temporal locality functions (blue boxes in Figure 3). Intuitively, the closer a function

¹⁰The default HMC interleaving scheme (Row:Column:Bank:Vault [73]) interleaves consecutive cache lines across vaults, and then across banks [220].




Figure 3: Locality-based clustering of 44 representative functions.

is to the bottom-left corner of the figure, the less likely it is to take advantage of a multi-level cache hierarchy. These functions are more likely to be good candidates for NDP. However, as we see in Section 3.3, the NDP suitability of a function also depends on a number of other factors.

3.3 Step 3: Bottleneck Classification

Within the two groups of functions identified in Section 3.2, we use three key metrics (AI, MPKI, and LFMR) to classify the memory bottlenecks. We observe that the AI of the analyzed low temporal locality functions is low (i.e., always less than 2.2 ops/cache line, with an average of 1.3 ops/cache line). Among the high temporal locality functions, there are some with low AI (minimum of 0.3 ops/cache line) and others with high AI (maximum of 44 ops/cache line). LFMR indicates whether a function benefits from a deeper cache hierarchy. When LFMR is low (i.e., less than 0.1), then a function benefits significantly from a deeper cache hierarchy, as most misses from the L1 cache hit in either the L2 or L3 caches. When LFMR is high (i.e., greater than 0.7), then most L1 misses are not serviced by the L2 or L3 caches, and must go to memory. A medium LFMR (0.1–0.7) indicates that a deeper cache hierarchy can mitigate some, but not a very large fraction of L1 cache misses. MPKI indicates the memory intensity of a function (i.e., the rate at which requests are issued to DRAM). We say that a function is memory-intensive (i.e., it has a high MPKI) when the MPKI is greater than 10, which is the same threshold used by prior works [151–157].

We find that six classes of functions emerge, based on their temporal locality, AI, MPKI, and LFMR values, as we observe from

Figures 3 and 4. We observe that spatial locality is not a key metric for our classification (i.e., it does not define a bottleneck class) because the L1 cache, which is present in both host CPU and NDP system configurations, can capture most of the spatial locality for a function. Figure 4 shows the LFMR and MPKI values for each class. Note that we do not have classes of functions for all possible combinations of metrics. In our analysis, we obtain the temporal locality, AI, MPKI, and LFMR values and their combinations empirically. Fundamentally, not all value combinations of different metrics are possible. We list some of the combinations we do *not* observe in our analysis of 144 functions:

- A function with high LLC MPKI does *not* display low LFMR. This is because a low LFMR happens when most L1 misses hit the L2/L3 caches. Thus, it becomes highly unlikely for the L3 cache to suffer many misses when the L2/L3 caches do a good job in fulfilling L1 cache misses.
- A function with high temporal locality does *not* display both high LFMR and high MPKI. This is because a function with high temporal locality will likely issue repeated memory requests to few memory addresses, which will likely be serviced by the cache hierarchy.
- A function with low temporal locality does *not* display low LFMR since there is little data locality to be captured by the cache hierarchy.

We discuss each class in detail below, identifying the memory bottlenecks for each class and whether the NDP system can alleviate these bottlenecks. To simplify our explanations, we focus on a

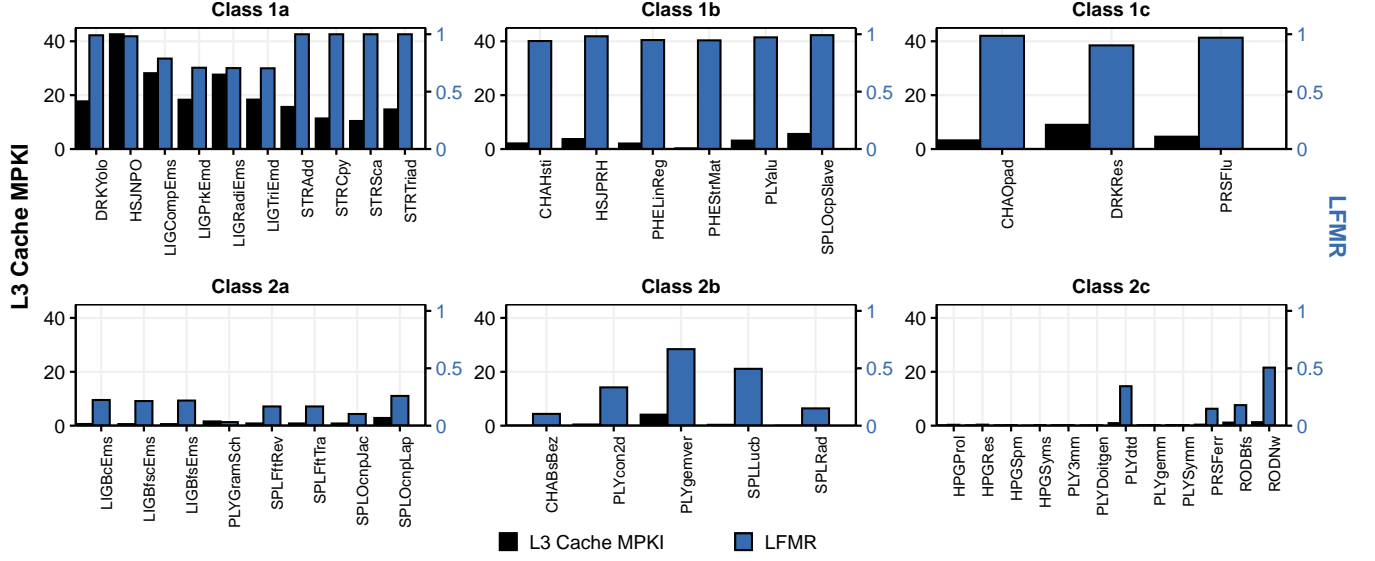


Figure 4: L3 Cache MPKI and Last-to-First Miss Ratio (LFMR) for 44 representative functions.

smaller set of 12 representative functions (out of the 44 representative functions) for this part of the analysis. Figure 5 shows how each of the 12 functions scales in terms of performance for the *host CPU*, *host CPU with prefetcher*, and *NDP system* configurations.

3.3.1 Class 1a: Low Temporal Locality, Low AI, High LFMR, and High MPKI (DRAM Bandwidth-Bound Functions) Functions in this class exert high main memory pressure since they are highly memory intensive and have low data reuse. To understand how this affects a function’s suitability for NDP, we study how performance scales as we increase the number of cores available to a function, for the host CPU, host CPU with prefetcher, and NDP system configurations. Figure 5(a) depicts performance¹¹ as we increase the core count, normalized to the performance of one host CPU core, for two representative functions from Class 1a (HSJNPO and LIGPrkEmd; we see similar trends for all functions in the class).

We make three observations from the figure. First, as the number of host CPU cores increases, performance eventually stops increasing significantly. For HSJNPO, host CPU performance increases by 27.5× going from 1 to 64 host CPU cores but only 27% going from 64 host CPU cores to 256 host CPU cores. For LIGPrkEmd, host CPU performance increases by 33× going from 1 to 64 host CPU cores but *decreases* by 20% going from 64 to 256 host CPU cores. We find that the lack of performance improvement at large host CPU core counts is due to main memory bandwidth saturation, as shown in Figure 6. Given the limited DRAM bandwidth available across the off-chip memory channel, we find that Class 1a functions saturate the DRAM bandwidth once enough host CPU cores (e.g., 64) are used, and thus these functions are *bottlenecked by the DRAM bandwidth*. Second, the host CPU system with prefetcher slows down the execution of the HSJNPO (LIGPrkEmd) function compared with the host CPU system without prefetcher by 43% (38%), on average across all core counts. The prefetcher is ineffective since these functions have low temporal and spatial locality. Third, when

running on the NDP system, the functions see continued performance improvements as the number of NDP cores increases. By providing the functions with access to the much higher bandwidth available inside memory, the NDP system can greatly outperform the host CPU system at a high enough core count. For example, at 64/256 cores, the NDP system outperforms the host CPU system by 1.7×/4.8× for HSJNPO, and by 1.5×/4.1× for LIGPrkEmd.

Figure 7 depicts the energy breakdown for our two representative functions. We make two observations from the figure. First, for HSJNPO, the energy spent on DRAM for both host CPU system and NDP system are similar. This is due to the function’s poor locality, as 98% of its memory requests miss in the L1 cache. Since LFMR is near 1, L1 miss requests almost always miss in the L2 and L3 caches and go to DRAM in the host CPU system for all core counts we evaluate, which requires significant energy to query the large caches and then to perform off-chip data transfers. The NDP system does not access L2, L3, and off-chip links, leading to large system energy reduction. Second, for LIGPrkEmd, the DRAM energy is higher in the NDP system than in the host CPU system. Since the function’s LFMR is 0.7, some memory requests that would be cache hits in the host CPU’s L2 and L3 caches are instead sent directly to DRAM in the NDP system. However, the total energy consumption on the host CPU system is still larger than that on the NDP system, again because the NDP system eliminates the L2, L3 and off-chip link energy.

DRAM bandwidth-bound applications such as those in Class 1a have been the primary focus of a large number of proposed NDP architectures (e.g., [1, 46, 54, 69, 76, 132, 133, 192, 222, 223]), as they benefit from increased main memory bandwidth and do not have high AI (and, thus, do not benefit from complex cores on the host CPU system). An NDP architecture for a function in Class 1a needs to extract enough MLP [57, 176, 181–184, 224–229] to maximize the usage of the available internal memory bandwidth. However, prior work has shown that this can be challenging due to the area and power constraints in the logic layer of a 3D-stacked DRAM [1, 46]. To exploit the high memory bandwidth while satisfying these

¹¹Performance is the inverse of application execution time.

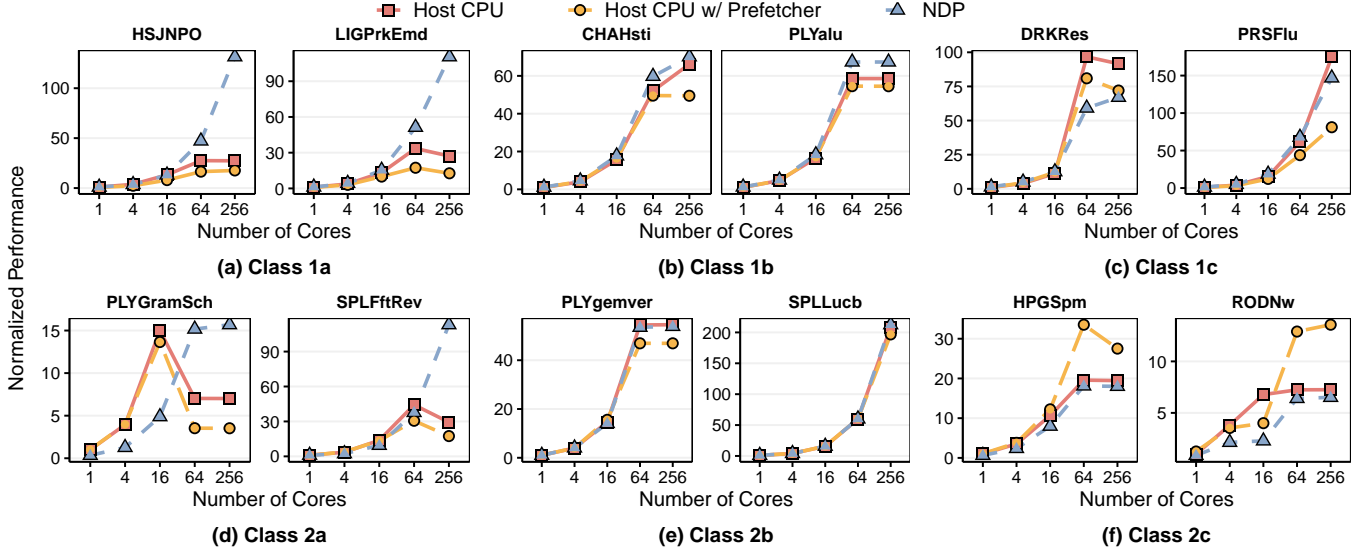


Figure 5: Performance of 12 representative functions on three systems: host CPU, host CPU with prefetcher, and NDP, normalized to one host CPU core.




Figure 6: Host CPU system IPC vs. utilized DRAM Bandwidth for representative Class 1a functions.




Figure 7: Cache and DRAM energy breakdown for representative Class 1a functions at 1, 4, 16, 64, and 256 cores.

area and power constraints, the NDP architecture should leverage application memory access patterns to efficiently maximize main memory bandwidth utilization.

We find that there are two dominant types of memory access patterns among our Class 1a functions. First, functions with regular access patterns (DRKYoIo, STRAdd, STRCpy, STRSca, STRTriad) can take advantage of specialized accelerators or Single Instruction Multiple Data (SIMD) architectures [1, 66], which can exploit the regular access patterns to issue many memory requests concurrently. Such accelerators or SIMD architectures have hardware area and thermal dissipation that fall well within the constraints of 3D-stacked DRAM [1, 46, 64, 230]. Second, functions with irregular

access patterns (HSJNPO, LIGCompEms, LIGPrkEmd, LIGRadiEms) require techniques to extract MLP while still fitting within the design constraints. This requires techniques that cater to the irregular memory access patterns, such as prefetching algorithms designed for graph processing [46, 231–235], pre-execution of difficult access patterns [57, 58, 151, 183, 184, 236–243] or hardware accelerators for pointer chasing [55, 56, 149, 193, 244–246].

3.3.2 Class 1b: Low Temporal Locality, Low AI, High LFM, and Low MPKI (DRAM Latency-Bound Functions) While functions in this class do not effectively use the host CPU caches, they do *not* exert high pressure on the main memory due to their low MPKI. Across all Class 1b functions, the average DRAM bandwidth consumption is only 0.5 GB/s. However, all the functions have very high LFM values (the minimum is 0.94 for CHAHsti), indicating that the host CPU L2 and L3 caches are ineffective. Because the functions cannot exploit significant MLP but still incur long-latency requests to DRAM, the DRAM requests fall on the critical path of execution and stall forward progress [57, 58, 151, 176, 247]. Thus, Class 1b functions are *bottlenecked by DRAM latency*. Figure 5(b) shows performance of both the host CPU system and the NDP system for two representative functions from Class 1b (CHAHsti and PLYalu). We observe that while performance of both the host CPU system and the NDP system scale well as the core count increases, NDP system performance is always higher than the host CPU system performance for the same core count. The maximum (average) speedup with NDP over host CPU at the same core count is $1.15\times$ ($1.12\times$) for CHAHsti and $1.23\times$ ($1.13\times$) for PLYalu.

We find that the NDP system’s improved performance is due to a reduction in the Average Memory Access Time (AMAT) [248]. Figure 8 shows the AMAT for our two representative functions. Memory accesses take significantly longer in the host CPU system than in the NDP system due to the additional latency of looking up requests in the L2 and L3 caches, even though data is rarely present in those caches, and going through the off-chip links.




Figure 8: Average Memory Access Time (AMAT) for representative Class 1b functions.

Figure 9 shows the energy breakdown for Class 1b representative functions. Similar to Class 1a, we observe that the L2/L3 caches and off-chip links are a large source of energy usage in the host CPU system. While DRAM energy increases in the NDP system, as L2/L3 hits in the host CPU system become DRAM lookups with NDP, the overall energy consumption in the NDP system is greatly smaller (by 69% maximum and 39% on average) due to the lack of L2 and L3 caches.




Figure 9: Energy breakdown for representative Class 1b functions.

Class 1b functions benefit from the NDP system, but primarily because of the lower memory access latency (and energy) that the NDP system provides for memory requests that need to be serviced by DRAM. These functions could benefit from other latency and energy reduction techniques, such as L2/L3 cache bypassing [51, 249–260], low-latency DRAM [15, 22–26, 89, 127, 261–276], and better memory access scheduling [153–157, 175–177, 247, 277–290]. However, they generally do *not* benefit significantly from prefetching (as seen in Figure 5(b)), since infrequent memory requests make it difficult for the prefetcher to successfully train on an access pattern.

3.3.3 Class 1c: Low Temporal Locality, Low AI, Decreasing LFMR with Core Count, and Low MPKI (L1/L2 Cache Capacity Bottlenecked Functions) We find that the behavior of functions in this class depends on the number of cores they are using. Figure 5(c) shows the host CPU system and the NDP system performance as we increase the core count for two representative functions (DRKRes and PRSFlu). We make two observations from the figure. First, at low core counts, the NDP system outperforms the host CPU system. With a low number of cores, the functions have medium to high LFMR (0.5 for DRKRes at 1 and 4 host CPU cores; 0.97 at 1 host CPU core and 0.91 at 4 host CPU cores for PRSFlu), and behave like Class 1b functions, where they are DRAM latency-sensitive. Second, as the core count increases, the host CPU system begins to outperform the NDP system. For example, beyond 16 (64) cores, the host CPU system outperforms the NDP system for DRKRes (PRSFlu). This is because as the core count increases, the

aggregate L1 and L2 cache size available at the host CPU system grows, which reduces the miss rates of both L2 and L3 caches. As a result, the LFMR decreases significantly (e.g., at 256 cores, LFMR is 0.09 for DRKRes and 0.35 for PRSFlu). This indicates that the *available L1/L2 cache capacity* bottlenecks Class 1c functions.

Figure 10 shows the energy breakdown for Class 1c functions. We make three observations from the figure. First, for functions with larger LFMR values (PRSFlu), the NDP system provides energy savings over the host CPU system at lower core counts, since the NDP system eliminates the energy consumed due to L3 and off-chip link accesses. Second, for functions with smaller LFMR values (DRKRes), the NDP system does not provide energy savings even for low core counts. Due to the medium LFMR, enough requests still hit in the host CPU system L2/L3 caches, and these cache hits become DRAM accesses in the NDP system, which consume more energy than the cache hits. Third, at high-enough core counts, the NDP system consumes more energy than the host CPU system for all Class 1c functions. As the LFMR decreases, the functions effectively utilize the caches in the host CPU system, reducing the off-chip traffic and, consequently, the energy Class 1c functions spend on accessing DRAM. The NDP system, which does not have L2 and L3 caches, pays the larger energy cost of a DRAM access for all L2/L3 hits in the host CPU system.




Figure 10: Energy breakdown for representative Class 1c functions.

We find that the primary source of the memory bottleneck in Class 1c functions is limited L1/L2 cache capacity. Therefore, while the NDP system improves performance and energy of some Class 1c functions at low core counts (with lower associated L1/L2 cache capacity), the NDP system does not provide performance and energy benefits across all core counts for Class 1c functions.

3.3.4 Class 2a: High Temporal Locality, Low AI, Increasing LFMR with Core Count, and Low MPKI (L3 Cache Contention Bottlenecked Functions) Like Class 1c functions, the behavior of the functions in this class depends on the number of cores that they use. Figure 5(d) shows the host CPU system and the NDP system performance as we increase the core count for two representative functions (PLYGramSch and SPLFFtRev). We make two observations from the figure. First, at low core counts, the functions do *not* benefit from the NDP system. In fact, for a single core (16 cores), PLYGramSch *slows down* by 67% (3×) when running on the NDP system, compared to running on the host CPU system. This is because, at low core counts, these functions make reasonably good use of the cache hierarchy, with LFMR values of 0.03 for PLYGramSch and lower than 0.44 for SPLFFtRev until 16 host CPU cores. We confirm this in Figure 11, where we see that very few memory requests for PLYGramSch and SPLFFtRev go to DRAM (5% for PLYGramSch, and

at most 13% for SPLFFtRev) at core counts lower than 16. Second, at high core counts (i.e., 64 for PLYGramSch and 256 for SPLFFtRev), the host CPU system performance starts to *decrease*. This is because Class 2a functions are *bottlenecked by cache contention*. At 256 cores, this contention undermines the cache effectiveness and causes the LFRM to increase to 0.97 for PLYGramSch and 0.93 for SPLFFtRev. With the last-level cache rendered essentially ineffective, the NDP system greatly improves performance over the host CPU system: by 2.23 \times for PLYGramSch and 3.85 \times for SPLFFtRev at 256 cores.




Figure 11: Memory request breakdown for representative Class 2a functions.

One impact of the increased cache contention is that it converts these high-temporal-locality functions into memory latency-bound functions. We find that with the increased number of requests going to DRAM due to cache contention, the AMAT increases significantly, in large part due to queuing at the memory controller. At 256 cores, the queuing becomes so severe that a large fraction of requests (24% for PLYGramSch and 67% for SPLFFtRev) must be reissued because the memory controller queues are full. The increased main memory bandwidth available to the NDP cores allows the NDP system to issue many more requests concurrently, which reduces the average length of the queue and, thus, the main memory latency. The NDP system also reduces memory access latency by getting rid of L2/L3 cache lookup and interconnect latencies.

Figure 12 shows the energy breakdown for the two representative Class 2a functions. We make two observations. First, the host CPU system is more energy-efficient than the NDP system at low core counts, as most of the memory requests are served by on-chip caches in the host CPU system. Second, the NDP system provides large energy savings over the host CPU system at high core counts. This is due to the increased cache contention, which increases the number of off-chip requests that the host CPU system must make, increasing the L3 and off-chip link energy.




Figure 12: Energy breakdown for representative Class 2a functions.

We conclude that cache contention is the primary scalability bottleneck for Class 2a functions, and the NDP system can provide an effective way of mitigating this cache contention bottleneck

without incurring the high area and energy overheads of providing additional cache capacity in the host CPU system, thereby improving the scalability of these applications to high core counts.

3.3.5 Class 2b: High Temporal Locality, Low AI, Low/Medium LFRM, and Low MPKI (L1 Cache Capacity Bottlenecked Functions) Figure 5(e) shows the host CPU system and the NDP system performance for PLYgemver and SPLLucb. We make two observations from the figure. First, as the number of cores increases, performance of the host CPU system and the NDP system scale in a very similar fashion. The NDP system and the host CPU system perform essentially on par with (i.e., within 1% of) each other at all core counts. Second, even though the NDP system does not provide any performance improvement for Class 2b functions, it also does not hurt performance. Figure 13 shows the AMAT for our two representative functions. When PLYgemver executes on the host CPU system, up to 77% of the memory latency comes from accessing L3 and DRAM, which can be explained by the function's medium LFRM (0.5). For SPLLucb, even though up to 73% of memory latency comes from L1 accesses, some requests still hit in the L3 cache (its LFRM is 0.2), translating to around 10% of the memory latency. However, the latency that comes from L3 + DRAM for the host CPU system is similar to the latency to access DRAM in the NDP system, resulting in similar performance between the host CPU system and the NDP system.




Figure 13: AMAT for representative Class 2b functions.

We make a similar observation for the energy consumption for the host CPU system and the NDP system (Figure 14). Even though a small number of memory requests hit in L3, the total energy consumption for both the host CPU system and the NDP system is similar due to L3 and off-chip link energy. For some functions in Class 2b, we observe that the NDP system slightly reduces energy consumption compared to the host CPU system. For example, the NDP system provides an 12% average reduction in energy consumption, across all core counts, compared to the host CPU system for PLYgemver.




Figure 14: Energy breakdown for representative Class 2b functions.

We conclude that while the NDP system does not solve any memory bottlenecks for Class 2b functions, it can be used to reduce the overall SRAM area in the system without any performance or energy penalty (and sometimes with energy savings).

3.3.6 Class 2c: High Temporal Locality, High AI, Low LFMR, and Low MPKI (Compute-Bound Functions). Aside from one exception (PLYSymm), all of the 11 functions in this class exhibit high temporal locality. When combined with the high AI and low memory intensity, we find that these characteristics significantly impact how the NDP system performance scales for this class. Figure 5(f) shows the host CPU system and the NDP system performance for HPGSpm and RODNw, two representative functions from the class. We make two observations from the figure. First, the host CPU system performance is *always* greater than the NDP system performance (by 44% for HPGSpm and 54% for RODNw, on average). The high AI (more than 12 ops per cache line), combined with the high temporal locality and low MPKI, enables these functions to make excellent use of the host CPU system resources. Second, both of the functions benefit greatly from prefetching in the host CPU system. This is a direct result of these functions' high spatial locality, which allows the prefetcher to be highly accurate and effective in predicting which lines to retrieve from main memory.

Figure 15 shows the energy breakdown consumption for the two representative Class 2c functions. We make two observations. First, the host CPU system is 77% more energy-efficient than the NDP system for HPGSpm, on average across all core counts. Second, the NDP system provides energy savings over the host CPU system at high core counts for RODNw (up to 65% at 256 cores). When the core count increases, the aggregate L1 cache capacity across all cores increases as well, which in turn decreases the number of L1 cache misses. Compared to executing on a single core, executing on 256 cores decreases the L1 cache miss count by 43%, reducing the memory subsystem energy consumption by 40%. However, due to RODNw's medium LFMR of 0.5, the host CPU system still suffers from L2 and L3 cache misses at high core counts, which require the large L3 and off-chip link energy. In contrast, the NDP system eliminates the energy of accessing the L3 cache and the off-chip link energy by directly sending L1 cache misses to DRAM, which, at high core counts, leads to lower energy consumption than the host CPU system.




Figure 15: Energy breakdown for representative Class 2c functions.

We conclude that Class 2c functions do not experience large memory bottlenecks and are not a good fit for the NDP system in terms of performance. However, the NDP system can sometimes provide energy savings for functions that experience medium LFMR.

3.4 Effect of the Last-Level Cache Size

The bottleneck classification we present in Section 3.3 depends on two key architecture-dependent metrics (LFMR and MPKI) that are directly affected by the parameters and the organization of the cache hierarchy. Our analysis in Section 3.3 partially evaluates the effect of caching by scaling the aggregated size of the private (L1/L2) caches with the number of cores in the system while maintaining the size of the L3 cache fixed at 8 MB for the host CPU system. However, we also need to understand the impact of the L3 cache size on our bottleneck classification analysis. To this end, this section evaluates the effects on our bottleneck classification analysis of using an alternative cache hierarchy configuration, where we employ a Non-Uniform Cache Architecture (NUCA) [291] model to scale the size of the L3 cache with the number of cores in the host CPU system.

In this configuration, we maintain the sizes of the private L1 and L2 caches (32 kB and 256 kB per core, respectively) while increasing the shared L3 cache size with the core count (we use 2 MB/core) in the host CPU system. The cores, shared L3 caches, and DRAM memory controller are interconnected using a 2D-mesh Network-on-Chip (NoC) [292–299] of size $(n+1) \times (n+1)$ (an extra interconnection dimension is added to place the DRAM memory controllers). To faithfully simulate the NUCA model (e.g., including network contention in our simulations), we integrate the M/D/1 network model proposed by ZSim++ [300] in our DAMOV simulator [158]. We use a latency of 3 cycles per hop in our analysis, as suggested by prior work [301]. We adapt our energy model to account for the energy consumption of the NoC in the NUCA system. We consider router energy consumption of 63 pJ per request and energy consumed per link traversal of 71 pJ, same as previous work [251].

Figure 16 shows the performance scalability curves for representative functions from each one of our bottleneck classes presented in Section 3.3 for the baseline host CPU system (*Host with 8MB Fixed LLC*), the host CPU NUCA system (*Host with NUCA 2MB/Core LLC*), and the NDP system. We make two observations. First, the observations we make for our bottleneck classification (Section 3.3) are *not* affected by increasing the L3 cache size for Classes 1a, 1b, 1c, 2b, and 2c. We observe that Class 1a functions benefit from a large L3 cache size (by up to $1.9 \times / 2.3 \times$ for HSNPO/LIGPrkEmd at 256 cores). However, the NDP system still provides performance benefits compared to the host CPU NUCA system. We observe that increasing the L3 size reduces some of the pressure on main memory but cannot fully reduce the DRAM bandwidth bottleneck for Class 1a functions. Functions in Class 1b do *not* benefit from extra L3 capacity (we do not observe a decrease in LFMR or MPKI). Functions in Class 1c do *not* benefit from extra L3 cache capacity. We observe that the private L1 and L2 caches capture most of their data locality, as mentioned in Section 3.3.3, and thus, these functions do *not* benefit from increasing the L3 size. Functions in Class 2b do *not* benefit from extra L3 cache capacity, which can even lead to a decrease in performance at high core counts for the host CPU NUCA system in some Class 2b functions due to long NUCA L3 access latencies. For example, we observe that PLYgemver's performance drops 18% when increasing the core count from 64 to 256 in the host CPU NUCA system. We do *not* observe such a performance drop for the host CPU system with fixed LLC size. The performance

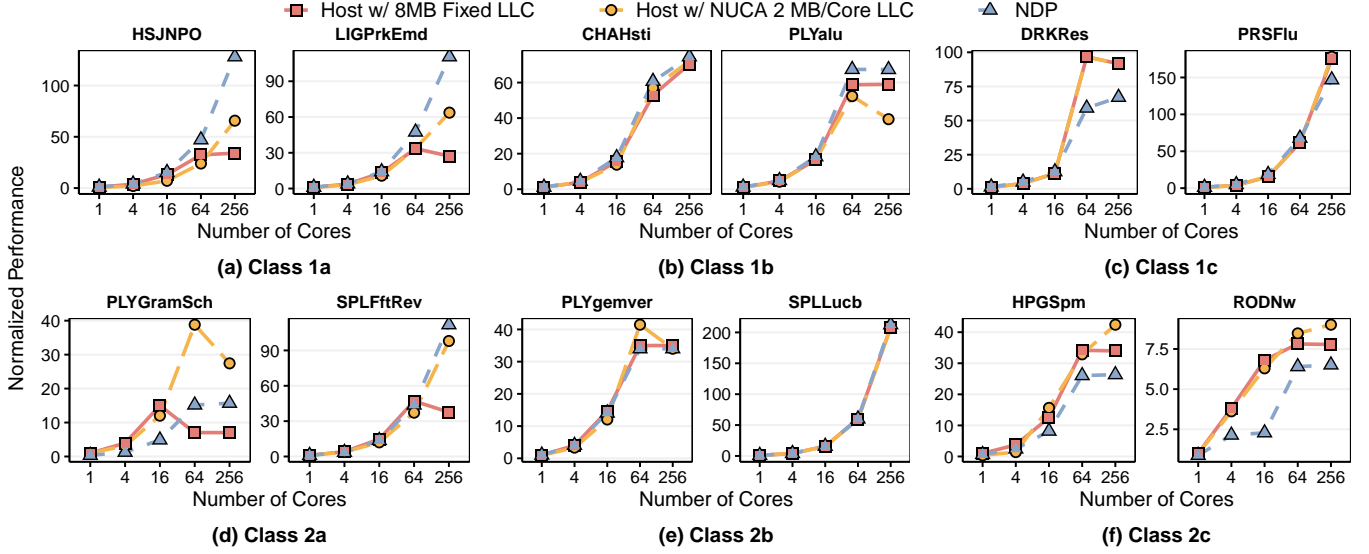


Figure 16: Performance of the host and the NDP system as we vary the LLC size, normalized to one host core with a fixed 8MB LLC size.

drop in the host CPU NUCA system is due to the increase in the number of hops that L3 requests need to travel in the NoC at high core counts, which increase the function’s AMAT. Class 2c functions benefit from a larger last-level cache. We observe that their performance improves by $1.3\times/1.2\times$ for HPGSpm/RODNw compared to the host CPU system with 8MB fixed LLC at 256 cores.

Second, we observe two different types of behavior for functions in Class 2a. Since cache conflicts are the major bottleneck for functions in this class, we observe that increasing the L3 cache size can mitigate this bottleneck. In Figure 16, we observe that for both PLYGramSch and SPLFftRev, the host system with NUCA 2MB/Core LLC provides better performance than the host system with 8MB fixed LLC. However, the NDP system can still provide performance benefits in case of contention on the L3 NoC (e.g., in SPLFftRev). For example, the NDP system provides 14% performance improvement for SPLFftRev compared to the NUCA system (with 512 MB L3 cache) for 256 cores.

In summary, we conclude that the key takeaways and observations we present in our bottleneck classification in Section 3.3 are also valid for a host system with a shared last-level cache whose size scales with core count. In particular, different workload classes get affected by an increase in L3 cache size as expected by their characteristics distilled by our classification.

Figure 17 shows the energy consumption for representative functions from each one of our bottleneck classes presented in Section 3.3. We observe that the NDP system can provide substantial energy savings for functions in different bottleneck classes, even compared against a system with very large (e.g., 512 MB) cache sizes. We make the following observations for each bottleneck class:

- **Class 1a:** First, for both representative functions in this bottleneck class, the host CPU NUCA system and the NDP system reduce energy consumption compared to the baseline host CPU system. However, we observe that the NDP system provides larger energy savings than the host CPU NUCA system. On average, across all core counts, the NDP system and the host CPU NUCA system

reduce energy consumption compared to the host CPU system for HSJNPO/LIGPrkEmd by 46%/65% and 25%/22%, respectively. Second, at 256 cores, the host CPU NUCA system provides larger energy savings than the NDP system for both representative functions. This happens because at 256 cores, the large L3 cache (i.e., 512 MB) captures a large portion of the dataset for these functions, reducing costly DRAM traffic. The host CPU NUCA system reduces energy consumption compared to the host CPU system for HSJNPO/LIGPrkEmd at 256 cores by $2.0\times/2.2\times$ while the NDP system reduces energy consumption by $1.6\times/1.8\times$. The L3 cache capacity needed to make the host CPU NUCA system more energy efficient than the NDP system is *very large* (512 MB SRAM), which is likely not cost-effective.

- **Class 1b:** First, for CHAHsti, the host CPU NUCA system *increases* energy consumption compared to the host CPU system by 9%, on average across all core counts. In contrast, the NDP system *reduces* energy consumption by 57%. Due to its low spatial and temporal locality (Figure 3), this function does not benefit from a deep cache hierarchy. In the host CPU NUCA system, the extra energy from the large amount of NoC traffic further increases the cache hierarchy’s overall energy consumption. Second, for PLYalu, the host CPU NUCA system and the NDP system reduce energy consumption compared to the host CPU system by 76% and 23%, on average across all core counts. Even though the increase in LLC size does not translate to performance improvements, the large LLC sizes in the host CPU NUCA system aid to reduce DRAM traffic, thereby providing energy savings compared to the baseline host CPU system.
- **Class 1c:** First, for DRKRes, the host CPU NUCA system reduces energy consumption compared to the host CPU system by 15%, on average across all core counts. In contrast, the NDP system increases energy consumption by 30%, which is due to the function’s medium LFMR (Section 3.3.3). Second, for PRSFlu, we observe that the NDP system provides large energy savings than

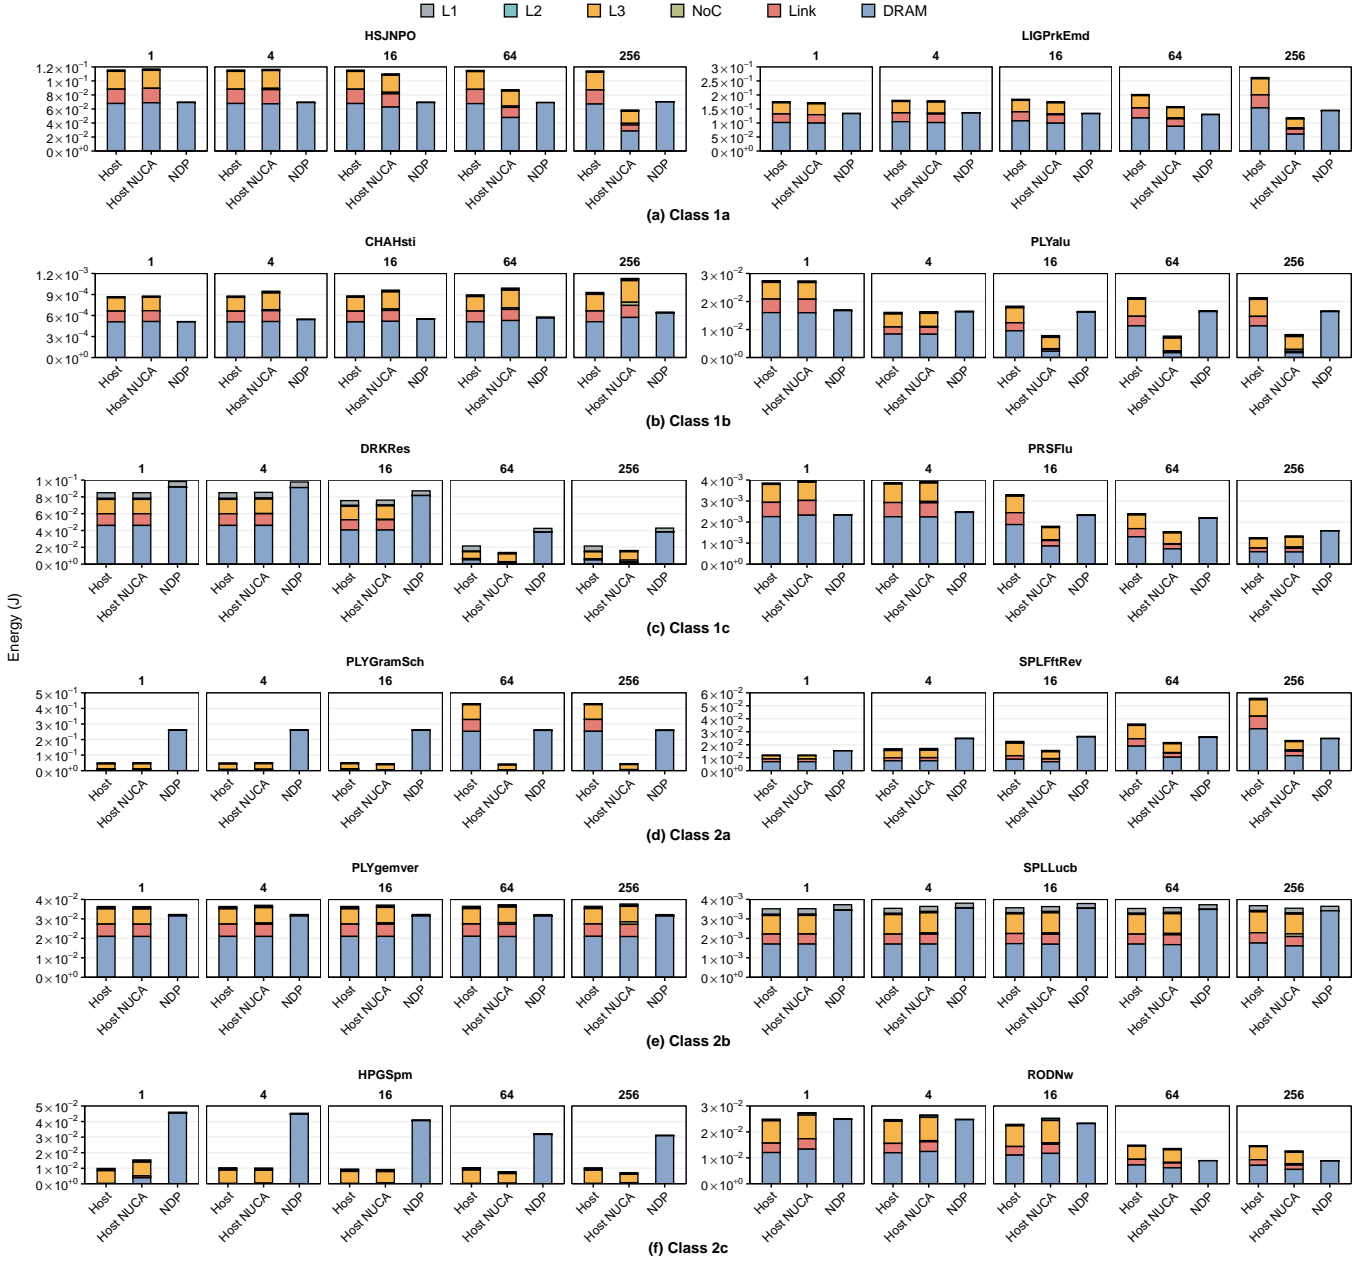


Figure 17: Energy of the host and the NDP system as we vary the LLC size. *Host* refers to the host system with a fixed 8MB LLC size; *Host NUCA* refers to the host system with 2MB/Core LLC.

the host CPU NUCA system. The host CPU NUCA system reduces energy consumption compared to the host CPU system by 21%, while the NDP system reduces energy consumption by 25%, on average across all core counts. However, the energy savings of both host CPU NUCA and NDP systems compared to the host CPU system reduces at high-enough core counts (the energy consumption of the host CPU NUCA system (NDP system) is $0.6\times$ ($0.9\times$) that of the host CPU system at 64 cores and $1.1\times$ ($1.3\times$) that of the host CPU system at 256 cores). This result is expected for Class 1c functions since the functions in this class have decreasing LFM, i.e., the functions effectively utilize the

private L1/L2 caches in the host CPU system at high-enough core counts.

- **Class 2a:** First, for PLYGramSch, compared to the host CPU system the host CPU NUCA system reduces energy consumption by $2.53\times$ and the NDP system increases energy consumption by 55%, on average across all core counts. Even though at high core counts (64 and 256 cores) the host CPU NUCA system provides larger energy savings than the NDP system compared to the host CPU system (the host CPU NUCA system and the NDP system reduce energy consumption compare to the host CPU system by $9\times$ and $65\times$ respectively, averaged across 64 and 256 cores), such

large energy savings come at the cost of very large (e.g., 512 MB) cache sizes. Second, for SPLFFtRev, the host CPU NUCA system and the NDP system reduce energy consumption compared to the host CPU system by 42% and 7%, on average across all core counts. The NDP system increases energy consumption compared to the host CPU system at low core counts (an increase of 33%, averaged across 1, 4, and 16 cores). However, it provides similar energy savings as the host CPU NUCA system for large core counts (99% and 75% energy reduction compare to the host CPU system for the host CPU NUCA system and the NDP system, respectively, averaged across 64 and 256 cores counts). Since the function suffers from high network contention, the increase in core count increases NoC traffic, which in turn increases energy consumption for the host CPU NUCA system. We conclude that the NDP system provides energy savings for Class 2a applications compared to the host CPU system at lower cost than the host CPU NUCA system.

- **Class 2b:** First, for PLYgemver, the host CPU NUCA system increases energy consumption compared to the host CPU system by 2%, on average across all core counts. In contrast, the NDP system reduces energy consumption by 13%. This function does not benefit from large L3 cache sizes since Class 2b functions are bottlenecked by L1 capacity. Thus, the NoC only adds extra static and dynamic energy consumption. Second, for SPLLucb, the host CPU NUCA system consumes the same energy as the host CPU system while the NDP system increases energy consumption by 5%, averaged across all core counts.
- **Class 2c:** For both representative functions in this class, the host CPU NUCA system reduces energy consumption compared to the host CPU system while the NDP system increases energy consumption. For HPGSpm/RODNw, the host CPU NUCA system reduces energy consumption by 6%/9% while the NDP system increases energy consumption by 74%/22%, averaged across all core counts. This result is expected since Class 2c functions are compute-bound and highly benefit from a deep cache hierarchy.

In conclusion, the NDP system can provide substantial energy savings for functions in different bottleneck classes, even compared against a system with very large (e.g., 512 MB) cache sizes.

3.5 Validation and Summary of Our Workload Characterization Methodology

In this section, we present the validation and a summary of our new workload characterization methodology. First, we use the remaining 100 memory-bound functions we obtain from *Step 1* (see Section 3.1) to validate our workload characterization methodology. To do so, we calculate the accuracy of our workload classification by using the remaining 100 memory-bound functions, which were not used to identify the six classes we found and described in Section 3.3. Second, we present a summary of the key metrics we obtain for all 144 memory-bound functions, including our analysis of the host CPU system and the NDP system using two types of cores (in-order and out-of-order).

3.5.1 Validation of Our Workload Characterization Methodology Our goal is to evaluate the accuracy of our workload characterization methodically on a large set of functions. To this end, we apply *Step 2* and *Step 3* of our memory bottleneck classification methodology (as described in Sections 2.3 and 2.4) to the remaining

100 memory-bound functions we obtain from *Step 1* (in Section 3.1). Then, we perform a two-phase validation to calculate the accuracy of our workload characterization.


In *phase 1* of our validation, we calculate the threshold values that define the low/high boundaries of each of the four metrics we use to cluster the initial 44 functions in the six memory bottleneck classes in Section 3.3 (i.e., temporal locality, LFMR, LLC MPKI, and AI). We also include the LFMR curve slope to indicate when the LFMR increases, decreases or stays constant as we scale the core count. We calculate the threshold values for a metric M by computing the middle point between (i) the average value of M across the memory bottleneck classes with *low* values of M and (ii) the average value of M across the memory bottleneck classes with *high* values of M values out of the 44 functions. In *phase 2* of our validation, we calculate the accuracy of our workload characterization by classifying the remaining 100 memory-bound functions using the threshold values obtained from *phase 1* and the LFMR curve slope. After *phase 2*, a function is considered to be *accurately* classified into a correct memory bottleneck class if and only if it (1) fits the definition of the assigned class using the threshold values obtained from *phase 1* and (2) follows the expected performance trends of the assigned class when the function is executed in the host CPU system and the NDP system. For example, a function is correctly classified into Class 1a *if and only if* it (1) displays low temporal locality, low AI, high LFMR, high MPKI and (2) the NDP system outperforms the host CPU system as we scale the core count when executing the function. The final *accuracy of our workload characterization methodology* is calculated by computing the percentage of the functions that are *accurately* classified into one of the six memory bottleneck classes.

First, by applying *phase 1* of our two-phase validation, we obtain that the threshold values are: 0.48 for *temporal locality*, 0.56 for *LFMR*, 11.0 for *MPKI*, and 8.5 for *AI*. Second, by applying *phase 2* of our two-phase validation, we find that we can accurately classify 97% of the 100 memory-bound functions into one of our six memory bottleneck classes (i.e., the accuracy of our workload characterization methodology is 97%). We observe that three functions (*Ligra:ConnectedComponents:compute:rMat*, *Ligra:MaximalIndependentSet:edgeMapDense:USA*, and *SPLASH-2:Oceanncp:relax*) could not be accurately classified into their correct memory bottleneck class (Class 1a). We observe that these functions have LLC MPKI values *lower* than the MPKI threshold expected for Class 1a functions. We expect that the accuracy of our methodology can be further improved by incorporating more workloads into our workload suite and fine-tuning each metric to encompass an even larger set of applications.

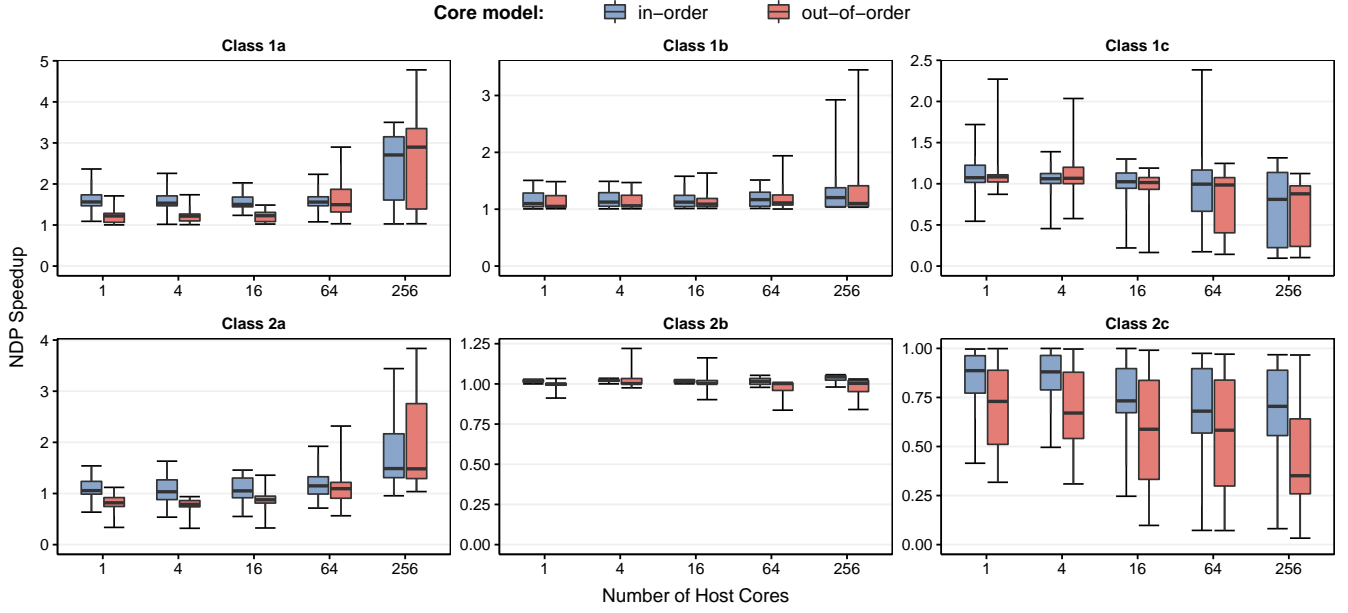
We conclude that our workload characterization methodology can accurately classify a given new application/function into its appropriate memory bottleneck class.

3.5.2 Summary of Our Workload Characterization Results.

Figure 18a summarizes the metrics we collect for all 144 functions across all core counts (i.e., from 1 to 256 cores) and different core microarchitectures (i.e., out-of-order and in-order cores). The figure shows the distribution of the key metrics we use during our workload characterization for each memory bottleneck class in Section 3.3, including architecture-independent metrics (i.e., temporal locality) and architecture-dependent metrics (i.e., AI, LFMR,



(a) Summary of the key metrics for each memory bottleneck class.



(b) Summary of NDP speedup for each memory bottleneck class at 1, 4, 16, 64, and 256 cores.

Figure 18: Summary of our characterization for all 144 memory-bound functions. Each box is lower-bounded by the first quartile and upper-bounded by the third quartile. The median falls within the box. The inter-quartile range (IQR) is the distance between the first and third quartiles (i.e., box size). Whiskers extend to the minimum and maximum data point values on either sides of the box.

and LLC MPKI). We report the architecture-dependent metrics for two core models: (i) in-order and (ii) out-of-order cores.¹² Together with the out-of-order core model that we use in Section 3.3, we incorporate an in-order core model to our analysis, so as to show that our memory bottleneck classification methodology focuses on data movement requirements and works independently of the core microarchitecture. Figure 18b shows the distribution of speedups we observe for when we offload the function to our general-purpose NDP cores, while employing the same core type as the host CPU system.

We make two key observations from Figure 18. First, we observe similar values for each architecture-dependent key metric (i.e., LFMR, MPKI, AI) regardless of core type for all 144 functions (in

Figure 18a). Second, we observe that the NDP system achieves similar speedups over the host CPU system, when using both in-order and out-of-order core configurations (in Figure 18b). The speedup provided by the NDP system compared to the host CPU system when both systems use out-of-order (in-order) cores for Classes 1a, 1b, 1c, 2a, 2b, and 2c is 1.59 (1.77), 1.22 (1.15), 0.96 (0.95), 1.04 (1.22), 0.94 (1.01), and 0.56 (0.76), respectively, on average across all core counts and functions within a memory bottleneck class. The NDP system greatly outperforms the host CPU system across *all core counts* for Class 1a and 1b functions, with a maximum speedup for the out-of-order (in-order) core model of 4.8 (3.5) and 3.4 (2.9), respectively. The NDP system greatly outperforms the host CPU system at *low core counts* for Class 1c functions and at *high core counts* for Class 2a functions, with a maximum speedup for the out-of-order (in-order) core model of 2.3 (2.4) and 3.8 (3.4), respectively. The NDP system provides a modest speedup compared to the host CPU system across *all core counts* for Class 2b functions and slowdown for Class 2c functions, with a maximum speedup

¹²In Section 3.3, we collect and report the values of the architecture-independent metrics and architecture-dependent metrics for a subset of 44 representative functions out of the 144 memory-bound functions we identify in Step 1 of our workload characterization methodology. In Section 3.5.2, we report values for the *complete set* of 144 memory-bound functions.

for the out-of-order (in-order) core model of 1.2 (1.1) and 1.0 (1.0), respectively. We observe that, averaged across all classes and core types, the average speedup provided by the NDP system using in-order cores is 11% higher than the average speedup offered by the NDP system using out-of-order cores. This is because the host CPU system with out-of-order cores can hide the performance impact of memory access latency to some degree (e.g., using dynamic instruction scheduling) [57, 58, 183, 184, 240, 302]. On the other hand, the host CPU system using in-order cores has little tolerance to hide memory access latency [57, 58, 183, 184, 240, 302].

We conclude that our methodology to classify memory bottlenecks of applications is *robust* and *effective* since we observe similar trends for the six memory bottleneck classes across a large range of (144) functions and two very different core models.

3.6 Limitations of Our Methodology

We identify three limitations to our workload characterization methodology. We discuss each limitation next.

NDP Architecture Design Space. Our methodology uses the same type and number of cores in the host CPU and the NDP system configurations for our scalability analysis (Section 3.3) because our main goal is to highlight the performance and energy differences between the host CPU system and the NDP system that are caused by data movement. We do not consider practical limitations related to area or thermal dissipation that could affect the type and the maximum number of cores in the NDP system, because our goal is **not** to propose NDP architectures but to characterize data movement and understand the different data movement bottlenecks in modern workloads. Proposing NDP architectures for the workload classes that our methodology identifies as suitable for NDP is a promising topic for future work.

Function-level Analysis. We choose to conduct our analysis at a function granularity rather than at the application granularity for two major reasons. First, general-purpose NDP architectures are typically leveraged as accelerators to which only *parts* of the application or specific functions are offloaded [1, 47, 48, 54, 59, 64, 65, 83, 86, 89, 92, 98, 100, 102, 133, 192, 193, 303–307], rather than the entire application. Functions typically form natural boundaries for parts of algorithms/applications that can potentially be offloaded. Second, it is well-known that applications go through distinct phases during execution. Each phase may have different characteristics (e.g., a phase might be more compute-bound, while another one might be more memory-bound) and thus fall into different classes in our analysis. A fine-grained analysis at the function level enables us to identify each of those phases and hence, identify more fine-grained opportunities for NDP offloading. However, the main drawback of function-level analysis is that it does not take into account data movement across function boundaries, which affects the performance and energy benefits the NDP system provides over the host CPU system. For example, the NDP system might hurt overall system performance and energy consumption when a large amount of data needs to be continuously moved between a function executing on the NDP cores and another executing on the host CPU cores [63, 74].

Overestimating NDP Potential. Offloading kernels to NDP cores incurs overheads that our analysis does not account for (e.g., maintaining coherence between the host CPU and the NDP cores [63, 74],

efficiently synchronizing computation across NDP cores [101, 140], providing virtual memory support for the NDP system [47, 55, 308], and dynamic offloading support for NDP-friendly functions [48]). Such overheads can impact the performance benefits NDP can provide when considering the end-to-end application. However, deciding how to and whether or not to offload computation to NDP is an open research topic, which involves several architecture-dependent components in the system, such as the following two examples. First, maintaining coherence between the host CPU and the NDP cores is a challenging task that recent works tackle [63, 74]. Second, enabling efficient synchronization across NDP cores is challenging due to the lack of shared caches and hardware cache coherence protocols in NDP systems. Recent works, such as [101, 309], provide solutions to the NDP synchronization problem. Therefore, to focus our analysis on the data movement characteristics of workloads and the broad benefits of NDP, we minimize our assumptions about our target NDP architecture, making our evaluation as broadly applicable as possible.

4 DAMOV: The Data Movement Benchmark Suite

In this section, we present DAMOV, the Data MOVement Benchmark Suite. DAMOV is the collection of the 144 functions we use to drive our memory bottleneck classification in Section 3. The benchmark suite is divided into each one of the six classes of memory bottlenecks presented in Section 3. DAMOV is the first benchmark suite that encompasses *real* applications from a diverse set of application domains tailored to stress different memory bottlenecks in a system. We present the complete description of the functions in DAMOV in Appendix A. We highlight the benchmark diversity of the functions in DAMOV in Section 4.1. We open source DAMOV [158] to facilitate further rigorous research in mitigating data movement bottlenecks, including in near data processing.

4.1 Benchmark Diversity

We perform a hierarchical clustering algorithm with the 44 representative functions we employ in Section 3.3.¹³ Our goal is to showcase our benchmark suite’s diversity and observe whether a clustering algorithm produces a noticeable difference from the application clustering presented Section 3. The hierarchical clustering algorithm [310] takes as input a dataset containing features that define each object in the dataset. The algorithm works by incrementally grouping objects in the dataset that are similar to each other in terms of some distance metric (called *linkage distance*), which is calculated based on the features’ values. Two objects with a short linkage distance have more affinity to each other than two objects with a large linkage distance. To apply the hierarchical clustering algorithm, we create a dataset where each object is one of the 44 representative functions from DAMOV. We use as features the same metrics we use for our analysis, i.e., temporal locality, MPKI, LFMR, and AI. We also include the LFMR curve slope to indicate when the LFMR increases, decreases or stays constant when scaling the core count. We use Euclidean distance [310] to calculate the linkage distance across features in our dataset. We evaluate other

¹³In Section 4.1, we use the same 44 representative functions that we use during our bottleneck classification instead of the entire set of 144 functions in DAMOV, in order to visualize better the clustering produced by the hierarchical clustering algorithm.




Figure 19: Hierarchical clustering of 44 representative functions.

linkage distance metrics (such as Manhattan distance [310]), and we observe similar clustering results.

Figure 19 shows the dendrogram that the hierarchical clustering algorithm produces for our 44 representative functions. We indicate in the figure the application class each function belongs to, according to our classification. We make three observations from the figure.

First, our benchmarks exhibit a wide range of behavior diversity, even among those belonging to the same class. For example, we observe that the functions from Class 1a are divided into two groups, with a linkage distance of 3. Intuitively, functions in the first group (HSJNPO, STRAdd, STRCpy, STRSca, STRTriad) have regular access patterns while functions in the second group (DRKYolo, LIGCompEms, LIGPrkEms, LIGRadEms) have irregular access patterns. We observe a similar clustering in Section 3.3.1.

Second, we observe that our application clustering (Section 3.3) matches the clustering that the hierarchical clustering algorithm provides (Figure 19). From the dendrogram root, we observe that the right part of the dendrogram consists of functions with high temporal locality (from Classes 2a, 2b, and 2c). Conversely, the left part of the dendrogram consists of functions with low temporal locality (from Classes 1a, 1b, and 1c). The functions in the right and left part of the dendrogram have a high linkage distance (higher than 15), which implies that the metrics we use for our clustering are significantly different from each other for these functions. Third, we observe that functions within the same class are clustered into groups with a linkage distance lower than 5. This grouping matches the six classes of data movement bottlenecks present in DAMOV. Therefore, we conclude that our methodology can successfully cluster functions into distinct classes, each one representing a different memory bottleneck.

We conclude that (i) DAMOV provides a heterogeneous and diverse set of functions to study data movement bottlenecks and (ii) our memory bottleneck clustering methodology matches the clustering provided by a hierarchical clustering algorithm (this section; Figure 19).

5 Case Studies

In this section, we demonstrate how our benchmark suite is useful to study open questions related to NDP system designs. We

provide four case studies. The first study analyzes the impact of load balance and communication on NDP execution. The second study assesses the impact of tailored NDP accelerators on our memory bottleneck analysis. The third study evaluates the effect of different core designs on NDP system performance. The fourth study analyzes the impact of fine-grained offloading (i.e., offloading small blocks of instructions to NDP cores) on performance.

5.1 Case Study 1: Impact of Load Balance and Inter-Vault Communication on NDP Systems

Communication between NDP cores is one of the key challenges for future NDP system designs, especially for NDP architectures based on 3D-stacked memories, where accessing a remote vault incurs extra latency overhead due to network traffic [46, 101, 311]. This case study aims to evaluate the load imbalance and inter-vault communication that the NDP cores experience when executing functions from the DAMOV benchmark suite. We statically map a function to an NDP core, and we assume that NDP cores are connected using a 6x6 2D-mesh Network-on-Chip (NoC), similar to previous works [66, 70, 312–314]. Figure 20 shows the performance overhead that the interconnection network imposes to NDP cores when running several functions from our benchmark suite. We report performance overheads of functions from different bottleneck classes (i.e., from Classes 1a, 1b, 2a, and 2b) that experience at least 5% of performance overhead due to the interconnection network. We calculate the interconnection network performance overhead by comparing performance with the 2D-mesh versus that with an ideal zero-latency interconnection network. We observe that the interconnection network performance overhead varies across functions, with a minimum overhead of 5% for SPL0cpSlave and a maximum overhead of 26% for SPLLucb.

We further characterize the traffic of memory requests injected into the interconnection network for these functions, aiming to understand the communication patterns across NDP cores. Figure 21 shows the distribution of all memory requests (y-axis) in terms of how many hops they need to travel in the NoC between NDP cores (x-axis) for each function. We make the following observations. First, we observe that, on average, 40% of all memory requests need




Figure 20: Interconnection network performance overhead in our NDP system.




Figure 21: Distribution of NoC hops traveled per memory request.

to travel 3 to 4 hops in the NoC, and less than 5% of all requests are issued to a local vault (0 hops). Even though the functions follow different memory access patterns, they all inject similar network traffic into the NoC.¹⁴ Therefore, we conclude that the NDP design can be further optimized by (i) employing more intelligent data mapping and scheduling mechanisms that can efficiently allocate data nearby the NDP core that accesses the data (thereby reducing inter-vault communication and improving data locality) and (ii) designing interconnection networks that can better fit the traffic patterns that NDP workloads produce. The DAMOV benchmark suite can be used to develop new ideas as well as evaluate existing ideas in both directions.

5.2 Case Study 2: Impact of NDP Accelerators on Our Memory Bottleneck Analysis

In our second case study, we aim to leverage our memory bottleneck classification to evaluate the benefits an NDP accelerator provides compared to the same accelerator accessing memory externally. We use the Aladdin accelerator simulator [315] to tailor an accelerator for an application function. Aladdin works by estimating the performance of a custom accelerator based on the data-flow graph of the application. The main difference between an NDP accelerator and a regular accelerator (i.e., compute-centric accelerator) is that the former is placed in the logic layer of a 3D-stacked memory device and thus can leverage larger memory bandwidth, shorter memory access latency, and lower memory access energy, compared to the compute-centric accelerator that is exemplary of existing compute-centric accelerator designs.

To evaluate the benefits of NDP accelerators, we select three functions from our benchmark suite for this case study: DRKYolo

(from Class 1a), PLYalu (from Class 1b), and PLY3mm (from Class 2c). We select these functions and memory bottleneck classes because we expect them to benefit the most (or to show no benefit) from the near-memory placement of an accelerator. According to our memory bottleneck analysis, we expect that the functions we select to (i) benefit from NDP due to its high DRAM bandwidth (Class 1a), (ii) benefit from NDP due to its shorter DRAM access latency (Class 1b), or (iii) do *not* benefit from NDP in any way (Class 2c).

Figure 22 shows the speedup that the NDP accelerator provides for the different functions compared to the compute-centric accelerator. We make four observations. First, as expected based on our classification, the NDP accelerator provides performance benefits compared to the compute-centric accelerator for functions in Classes 1a and 1b. It does not provide performance improvement for the function in Class 2c. Second, the NDP accelerator for DRKYolo shows the largest performance benefits (1.9× performance improvement compared to the compute-centric accelerator). Since this function is DRAM bandwidth-bound (Class 1a, Section 3.3.1), the NDP accelerator can leverage the larger memory bandwidth available in the logic layer of the 3D-stacked memory device. Third, we observe that the NDP accelerator also provides speedup (1.25×) for the PLYalu function compared to the compute-centric accelerator, since the NDP accelerator provides shorter memory access latency to the function, which is latency-bound (Class 1b, Section 3.3.2). Fourth, the NDP accelerator does not provide performance improvement for the PLY3mm function since this function is compute-bound (Class 2c, Section 3.3.6).




Figure 22: Speedup of the NDP Accelerators over the Compute-Centric Accelerators for three functions from Classes 1a, 1b, and 2c.

In conclusion, our observations for the performance of NDP accelerators are in line with the characteristics of the three memory bottleneck classes we evaluate in this case study. Therefore, our memory bottleneck classification can be applied to study other types of system configurations, e.g., the accelerators used in this section. However, since NDP accelerators are often employed under restricted area and power constraints (e.g., limited area available in the logic layer of a 3D-stacked memory [63, 74]), the core model of the compute-centric and NDP accelerators cannot always be the same. We leave a thorough analysis that takes area and power constraints in the study of NDP accelerators into consideration for future research.

5.3 Case Study 3: Impact of Different Core Models on NDP Architectures

This case study aims to analyze when a workload can benefit from different core models and numbers of cores while respecting the area and power envelope of the logic layer of a 3D-stacked

¹⁴We use the default HMC data interleaving scheme in our experiments (Table 1).

memory. Many prior works employ 3D-stacked memories as the substrate to implement NDP architectures [1, 46–48, 54, 55, 59–61, 63–70, 74–77, 79, 80, 99, 101–103, 137, 146, 192, 194, 305, 316–324]. However, 3D-stacked memories impose severe area and power restrictions on NDP architectures. For example, the area and power budget of the logic layer of a single HMC vault are 4.4 mm^2 and 312 mW , respectively [1, 63].

In the case study, we perform an iso-area and iso-power performance evaluation of three functions from our benchmark suite. We configure the host CPU system and the NDP system to guarantee an iso-area and iso-power evaluation, considering the area and power budget for a 32-vault HMC device [1, 63]. We use four out-of-order cores with a deep cache hierarchy for the host system configuration and two different NDP configurations: (1) one using six out-of-order NDP cores (*NDP+out-of-order*) and (2) using 128 in-order NDP cores (*NDP+in-order*), without a deep cache hierarchy. We choose functions from Classes 1a, 1b, and 2b for this case study since the major effects distinct microarchitectures have on the memory system are: (a) how much DRAM bandwidth they can sustain, and (b) how much DRAM latency they can hide. Classes 1a, 1b, and 2b are the most affected by memory bandwidth and access latency (as shown in Section 3). We choose two representative functions from each of these classes.

Figure 23 shows the speedup provided by the two NDP system configurations compared to the baseline host system. We make two observations. First, in all cases, the *NDP+in-order* system provides higher speedup than the *NDP+out-of-order* system, both compared to the host system. On average across all six functions, the *NDP+in-order* system provides 4 \times the speedup of the *NDP+out-of-order* system. The larger speedup the *NDP+in-order* system provides is due to the high number of NDP cores in the *NDP+in-order* system. We can fit 128 in-order cores in the logic layer of the 3D-stacked memory as opposed to only six out-of-order cores in the same area/power budget. Second, we observe that the speedup the *NDP+in-order* system provides compared to the *NDP+out-of-order* system does not scale with the number of cores. For example, the *NDP+in-order* system provides only 2 \times the performance of the *NDP+out-of-order* system for DRKYolo and PLYalu, even though the *NDP+in-order* system has 21 \times the number of NDP cores of the *NDP+out-of-order* system. This implies that even though the functions benefit from a large number of NDP cores available in the *NDP+in-order* system, static instruction scheduling limits performance on the *NDP+in-order* system.




Figure 23: Speedup of NDP architectures over 4 out-of-order host CPU cores for two NDP configurations: using 128 in-order NDP cores (*NDP+in-order*) and 6 out-of-order NDP cores (*NDP+out-of-order*) for representative functions from Classes 1a, 1b, and 2b.

We believe, and our previous observations suggest, that an efficient NDP architecture can be achieved by leveraging mechanisms that can exploit both dynamic instruction scheduling and many-core design while fitting in the area and power budget of 3D-stacked memories. For example, past works [57, 58, 183, 184, 224, 325–343] propose techniques that enable the benefits of simple and complex cores at the same time, via heterogeneous or adaptive architectures. These ideas can be examined to enable better core and system designs for NDP systems, and DAMOV can facilitate their proper design, exploration, and evaluation.

5.4 Case Study 4: Impact of Fine-Grained Offloading to NDP on Performance

Several prior works on NDP (e.g., [47, 54, 86, 89, 100, 146, 303, 306, 344–346]) propose to identify and offload to the NDP system simple primitives (e.g., instructions, atomic operations). We refer to this NDP offloading scheme as a *fine-grained NDP offloading*, in contrast to a *coarse-grained NDP offloading scheme* that offloads whole functions and applications to NDP systems. A fine-grained NDP offloading scheme provides two main benefits compared to a coarse-grained NDP offloading scheme. First, a fine-grained NDP offloading scheme allows for a reduction in the complexity of the processing elements used as NDP logic, since the NDP logic can consist of simple processing elements (e.g., arithmetic units, fixed function units) instead of entire in-order or out-of-order cores often utilized when employing a coarse-grained NDP offloading scheme. Second, a fine-grained NDP offloading scheme can help developing simple coherence mechanism needed to allow shared host and NDP execution [47]. However, identifying arbitrary NDP instructions can be a daunting task since there is no comprehensive methodology that indicates what types of instructions are good offloading candidates.

As the first step in this direction, we exploit the key insight provided by [151, 347] to identify potential regions of code that can be candidates for fine-grained NDP offloading. [151, 347, 348] show that few instructions are responsible for generating most of the cache misses during program execution in memory-intensive applications. Thus, these instructions are naturally good candidates for fine-grained NDP offloading. Figure 24 shows the distribution of unique basic blocks (x-axis) and the percentage of last-level cache misses (y-axis) the basic block produces for three representative functions from our benchmark suite. We select functions from Classes 1a (LIGCrEms), 1b (HSJPRH), and 1c (DRKRes) since functions in these classes have higher L3 MPKI than functions in Classes 2a, 2b, and 2c. We observe from the figure that 1% to 10% of the basic blocks in each function are responsible for up to 95.3% of the LLC misses. We call these basic blocks the *hottest* basic blocks.¹⁵ We investigate the data-flow of each basic block and observe that these basic blocks often execute simple read-modify-write operations, with few arithmetic operations. Therefore, we believe that such basic blocks are good candidates for fine-grained offloading. Figure 25 shows the speedup obtained by offloading (i) the hottest basic block we identified for the three representative functions and (ii) the entire function to the NDP system, compared to the host

¹⁵We observe for the 44 functions we evaluate in Section 3 that in many cases (for 65% of the evaluated workloads), a single basic block is responsible for 90% to 100% of the LLC misses during the function’s execution.

system. Our initial evaluations show that offloading the hottest basic block of each function to the NDP system can provide up to 1.25 \times speedup compared to the host CPU, which is half of the 1.5 \times speedup achieved when offloading the entire function. Therefore, we believe that methodically identifying simple NDP instructions can be a promising research direction for future NDP system designs, which our DAMOV Benchmark Suite can help with.




Figure 24: Distribution of unique basic blocks (x-axis) and the percentage of last-level cache misses they produce (y-axis) for three representative functions from Classes 1a (LIGKcrEms), 1b (HSJPRH), and 1c (DRKRes).




Figure 25: Speedup of offloading to NDP the *hottest* basic block in each function versus the entire function.

6 Key Takeaways

We summarize the key takeaways from our extensive characterization of 144 functions using our new three-step methodology to identify data movement bottlenecks. We also highlight when NDP is a good architectural choice to mitigate a particular memory bottleneck.

Figure 26 pictorially represents the key takeaways we obtain from our memory bottleneck classification. Based on four key metrics, we classify workloads into six classes of memory bottlenecks. We provide the following key takeaways:

- (1) Applications with low temporal locality, high LFMR, high MPKI, and low AI are DRAM *bandwidth-bound* (Class 1a, Section 3.3.1). They are bottlenecked by the limited off-chip memory bandwidth as they exert high pressure on main memory. We make three observations for Class 1a applications. First, these applications do benefit from prefetching since they display a low degree of spatial locality. Second, these applications highly benefit from NDP architectures because they take advantage of the high memory bandwidth available within the memory device. Third, NDP architectures significantly improve energy for these applications since they eliminate the off-chip I/O traffic between the CPU and the main memory.

- (2) Applications with low temporal locality, high LFMR, low MPKI, and low AI are DRAM *latency-bound* (Class 1b, Section 3.3.2). We make three observations for Class 1b applications. First, these applications do not significantly benefit from prefetching since infrequent memory requests make it difficult for the prefetcher to train successfully on an access pattern. Second, these applications benefit from NDP architectures since they take advantage of NDP's lower memory access latency and the elimination of deep L2/L3 cache hierarchies, which fail to capture data locality for these workloads. Third, NDP architectures significantly improve energy for these applications since they eliminate costly (and unnecessary) L3 cache look-ups and the off-chip I/O traffic between the CPU and the main memory.
- (3) Applications with low temporal locality, decreasing LFMR with core count, low MPKI, and low AI are *bottlenecked by the available L1/L2 cache capacity* (Class 1c, Section 3.3.3). We make three observations for Class 1c applications. First, these applications are DRAM latency-bound at low core counts, thus taking advantage of NDP architectures, both in terms of performance improvement and energy reduction. Second, NDP's benefits reduce when core count becomes larger, which consequently allows the working sets of such applications to fit inside the cache hierarchy at high core counts. Third, NDP architectures can be a good design choice for such workloads in systems with limited area budget since NDP architectures do not require large L2/L3 caches to outperform or perform similarly to the host CPU (in terms of both system throughput and energy) for these workloads.
- (4) Applications with high temporal locality, increasing LFMR with core count, low MPKI, and low AI are *bottlenecked by L3 cache contention* (Class 2a, Section 3.3.4). We make three observations for Class 2a applications. First, these applications benefit from a deep cache hierarchy and do not take advantage of NDP architectures at low core counts. Second, the number of cache conflicts increases when the number of cores in the system increases, leading to more pressure on main memory. We observe that NDP can effectively mitigate such cache contention for these applications without incurring the high area and energy overheads of providing additional cache capacity in the host. Third, NDP can improve energy for these workloads at high core counts, since it eliminates the costly data movement between the last-level cache and the main memory.
- (5) Applications with high temporal locality, low LFMR, low MPKI, and low AI are bottlenecked by *L1 cache capacity* (Class 2b, Section 3.3.5). We make two observations for Class 2b applications. First, NDP can provide similar performance and energy consumption than the host system by leveraging lower memory access latency and avoiding off-chip energy consumption for these applications. Second, NDP can be used to reduce the overall SRAM area (by eliminating L2/L3 caches) in the system without a performance or energy penalty.

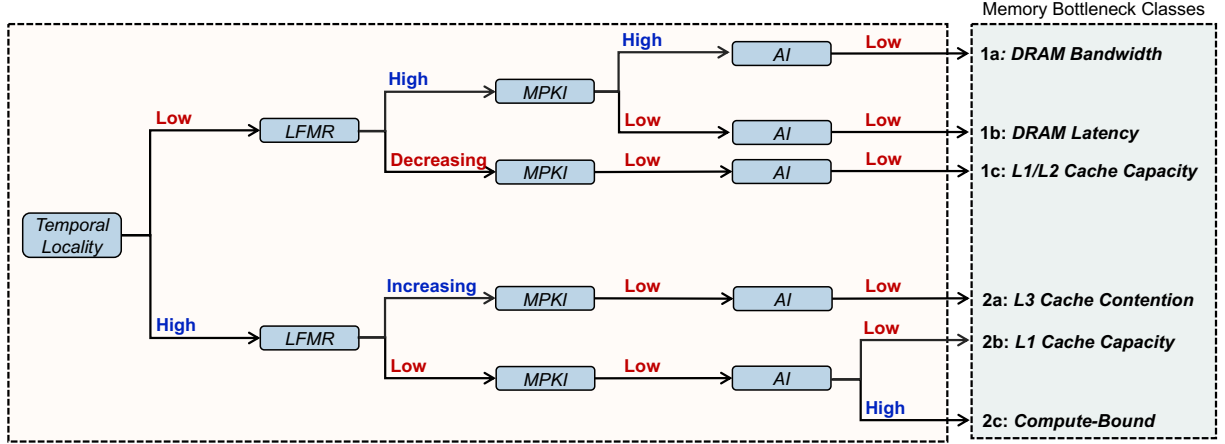


Figure 26: Summary of our memory bottleneck classification.

- (6) Applications with high temporal locality, low LFMR, low MPKI, and high AI are *compute-bound* (Class 2c, Section 3.3.6). We make three observations for Class 2c applications. First, these applications suffer performance and energy penalties due to the lack of a deep L2/L3 cache hierarchy when executed on the NDP architecture. Second, these applications highly benefit from prefetching due to their high temporal and spatial locality. Third, these applications are not good candidates to execute on NDP architectures.

6.1 Shaping Future Research with DAMOV

A key contribution of our work is DAMOV, the first benchmark suite for main memory data movement studies. DAMOV is the collection of 144 functions from 74 different applications, belonging to 16 different benchmark suites or frameworks, classified into six different classes of data movement bottlenecks.

We believe that DAMOV can be used to explore a wide range of research directions on the study of data movement bottlenecks, appropriate mitigation mechanisms, and open research topics on NDP architectures. We highlight DAMOV’s usability and potential benefits with four brief case studies, which we summarize below:

- In the first case study (Section 5.1), we use DAMOV to evaluate the interconnection network overheads that NDP cores placed in different vaults of a 3D-stacked memory suffer from. We observe that a large portion of the memory requests an NDP core issues go to remote vaults, which increases the memory access latency for the NDP core. We believe that DAMOV can be employed to study better data mapping techniques and interconnection network designs that aim to minimize (i) the number of remote memory accesses the NDP cores execute and (ii) the interconnection network latency overheads.
- In the second case study (Section 5.2), we evaluate the benefits that NDP accelerators can provide for three applications from our benchmark suite. We compare the performance improvements an NDP accelerator provides against the compute-centric version of the same accelerator. We observe that the NDP accelerator provides significant performance benefits compared to the compute-centric accelerator for applications in Classes 1a and 1b. At the same time,

it does not improve performance for an application in Class 2c. We believe that DAMOV can aid the design of NDP accelerators that target different memory bottlenecks in the system.

- In the third case study (Section 5.3), we perform an iso-area/-power performance evaluation to compare NDP systems using in-order and out-of-order cores. We observe that the in-order cores’ performance benefits for some applications are limited by the cores’ static instruction scheduling mechanism. We believe that better NDP systems can be built by leveraging techniques that enable dynamic instruction scheduling without incurring the large area and power overheads of out-of-order cores. DAMOV can help in the analysis and development of such NDP architectures.
- In the fourth case study (Section 5.4), we evaluate the benefits of offloading small portions of code (i.e., a basic block) to NDP, which simplifies the design of NDP systems. We observe that for many applications, a small percentage of basic blocks is responsible for most of the last-level cache misses. By offloading these basic blocks to an NDP core, we observe a performance improvement of up to 1.25×. We believe that DAMOV can be used to identify simple NDP instructions that enable building efficient NDP systems in the future.

7 Related Work

To our knowledge, this is the first work that methodically characterizes data movement bottlenecks and evaluates the benefits of different data movement mitigation mechanisms, with a focus on Near-Data Processing (NDP), for a broad range of applications. This is also the first work that provides an extensive open-source benchmark suite, with a diverse range of real world applications, tailored to stress different memory-related data movement bottlenecks in a system.

Many past works investigate how to reduce data movement cost using a range of different compute-centric (e.g., prefetchers [56, 189, 244, 349–369], speculative execution [57, 58, 183, 184, 349, 370], value-prediction [349, 356, 371–387], data compression [388–405], approximate computing [40, 371, 406, 407]) and memory-centric techniques [1, 35, 47, 54, 63–65, 81, 194, 222, 223, 250, 251, 335, 408–417]. These works evaluate the impact of data movement in

different systems, including mobile systems [1, 39, 418–420], data centers [5, 31, 355, 421–425], accelerators-based systems [1, 59, 60, 179, 219, 423, 426], and desktop computers [202, 427, 428]. They use very different profiling frameworks and methodologies to identify the root cause of data movement for a small set of applications. Thus, it is not possible to generalize prior works’ findings to other applications than the limited set they analyze.

We highlight two of these prior works, [426] and [1], since they also focus on characterizing applications for NDP architectures. In [426], the authors provide the first work that characterizes workloads for NDP. They analyze five applications (FFT, ray tracing, method of moments, image understanding, data management). The NDP organization [426] targets is similar to [429], where vector processing compute units are integrated into the DDRx memory modules. Even though [426] has a similar goal to our work, it understandably does not provide insights into modern data-intensive applications and NDP architectures as it dates from 2001. Also, [426] focuses its analysis only on a few workloads, whereas we conduct a broader workload analysis starting from 345 applications. Therefore, a new, more comprehensive and rigorous analysis methodology of data movement bottlenecks in modern workloads and modern NDP systems is necessary. A more recent work investigates the memory bottlenecks in widely-used consumer workloads from Google and how NDP can mitigate such bottlenecks [1]. This work focuses its analysis on a small number of consumer workloads. Our work presents a comprehensive analysis of a much broader set of applications (345 different applications, and a total of 77K application functions), which allows us to provide a general methodology, a comprehensive workload suite, and general takeaways and guidelines for future NDP research. With our comprehensive analysis, this work is the first to develop a rigorous methodology to classify applications into six groups, which have different characteristics with respect to how they benefit from NDP systems as well as other data movement bottleneck mitigation techniques.

8 Conclusion

This paper introduces the first rigorous methodology to characterize memory-related data movement bottlenecks in modern workloads and the first data movement benchmark suite, called DAMOV. We perform the first large-scale characterization of applications to develop a three-step workload characterization methodology that introduces and evaluates four key metrics to identify the sources of data movement bottlenecks in real applications. We use our new methodology to classify the primary sources of memory bottlenecks of a broad range of applications into six different classes of memory bottlenecks. We highlight the benefits of our benchmark suite with four case studies, which showcase how representative workloads in DAMOV can be used to explore open-research topics on NDP systems and reach architectural as well as workload-level insights and conclusions. We open-source our benchmark suite and our bottleneck analysis toolchain [158]. We hope that our work enables further studies and research on hardware and software solutions for data movement bottlenecks, including near-data processing.

Acknowledgments

We thank the SAFARI Research Group members for valuable feedback and the stimulating intellectual environment they provide. We acknowledge support from the SAFARI Research Group’s

industrial partners, especially ASML, Facebook, Google, Huawei, Intel, Microsoft, VMware, and the Semiconductor Research Corporation. An earlier version of this work was posted on arxiv.org (<https://arxiv.org/pdf/2105.03725.pdf>) on May 8, 2021. Talk videos for this work are available on YouTube, including a short talk video (<https://youtu.be/HkMYuYMuZog>), a long talk video (<https://youtu.be/GWideVyo0nM>), and a tutorial on the DAMOV framework and benchmarks (<https://youtu.be/GWideVyo0nM?t=8028>).

References

- [1] A. Boroumand, S. Ghose, Y. Kim, R. Ausavarungrun, E. Shiu, R. Thakur, D. Kim, A. Kuusela, A. Knies, P. Ranganathan *et al.*, “Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks,” in *ASPLOS*, 2018.
- [2] O. Mutlu, “Memory Scaling: A Systems Architecture Perspective,” in *IMW*, 2013.
- [3] U. Kang, H.-S. Yu, C. Park, H. Zheng, J. Halbert, K. Bains, S. Jang, and J. S. Choi, “Co-Architecting Controllers and DRAM to Enhance DRAM Process Scaling,” in *The Memory Forum*, 2014.
- [4] S. Hong, “Memory Technology Trend and Future Challenges,” in *IEDM*, 2010.
- [5] S. Kanev, J. P. Darago, K. Hazelwood, P. Ranganathan, T. Moseley, G.-Y. Wei, and D. Brooks, “Profiling a Warehouse-Scale Computer,” in *ISCA*, 2015.
- [6] O. Mutlu and L. Subramanian, “Research Problems and Opportunities in Memory Systems,” *SUPERFRI*, 2014.
- [7] O. Mutlu, “Main Memory Scaling: Challenges and Solution Directions,” in *More Than Moore Technologies for Next Generation Computer Design*. Springer-Verlag, 2015.
- [8] J. S. Kim, M. Patel, A. G. Yağlıkçı, H. Hassan, R. Azizi, L. Orosa, and O. Mutlu, “Revisiting RowHammer: An Experimental Analysis of Modern DRAM Devices and Mitigation Techniques,” in *ISCA*, 2020.
- [9] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu, “Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors,” in *ISCA*, 2014.
- [10] O. Mutlu, “The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser,” in *DATE*, 2017.
- [11] S. Ghose, A. G. Yağlıkçı, R. Gupta, D. Lee, K. Kudrolli, W. X. Liu, H. Hassan, K. K. Chang, N. Chatterjee, A. Agrawal *et al.*, “What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study,” in *SIGMETRICS*, 2018.
- [12] O. Mutlu and J. S. Kim, “RowHammer: A Retrospective,” *TCAD*, 2020.
- [13] P. Frigo, E. Vannacc, H. Hassan, V. Van Der Veen, O. Mutlu, C. Giuffrida, H. Bos, and K. Razavi, “TRRespass: Exploiting the Many Sides of Target Row Refresh,” in *SP*, 2020.
- [14] J. Liu, B. Jaiyen, Y. Kim, C. Wilkerson, O. Mutlu, J. Liu, B. Jaiyen, Y. Kim, C. Wilkerson, and O. Mutlu, “An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms,” in *ISCA*, 2013.
- [15] J. Liu, B. Jaiyen, R. Veras, and O. Mutlu, “RAIDR: Retention-Aware Intelligent DRAM Refresh,” in *ISCA*, 2012.
- [16] M. Patel, J. S. Kim, and O. Mutlu, “The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions,” in *ISCA*, 2017.
- [17] M. K. Qureshi, D. Kim, S. Khan, P. J. Nair, and O. Mutlu, “AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems,” in *DSN*, 2015.
- [18] J. A. Mandelman, R. H. Dennard, G. B. Bronner, J. K. DeBrosse, R. Divakaruni, Y. Li, and C. J. Radens, “Challenges and Future Directions for the Scaling of Dynamic Random-Access Memory (DRAM),” *IBM JRD*, 2002.
- [19] S. Khan, D. Lee, Y. Kim, A. R. Alameldeen, C. Wilkerson, and O. Mutlu, “The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study,” in *SIGMETRICS*, 2014.
- [20] S. Khan, D. Lee, and O. Mutlu, “PARBOR: An Efficient System-Level Technique to Detect Data-Dependent Failures in DRAM,” in *DSN*, 2016.
- [21] S. Khan, C. Wilkerson, Z. Wang, A. R. Alameldeen, D. Lee, and O. Mutlu, “Detecting and Mitigating Data-Dependent DRAM Failures by Exploiting Current Memory Content,” in *MICRO*, 2017.
- [22] D. Lee, Y. Kim, G. Pekhimenko, S. Khan, V. Seshadri, K. Chang, and O. Mutlu, “Adaptive-Latency DRAM: Optimizing DRAM Timing for the Common-Case,” in *HPCA*, 2015.
- [23] D. Lee, S. Khan, L. Subramanian, S. Ghose, R. Ausavarungrun, G. Pekhimenko, V. Seshadri, and O. Mutlu, “Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms,” in *SIGMETRICS*, 2017.
- [24] K. K. Chang, A. G. Yağlıkçı, 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.

- [25] 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.
- [26] K. K.-W. Chang, D. Lee, Z. Chishti, A. R. Alameldeen, C. Wilkerson, Y. Kim, and O. Mutlu, "Improving DRAM Performance by Parallelizing Refreshes with Accesses," in *HPCA*, 2014.
- [27] J. Meza, Q. Wu, S. Kumar, and O. Mutlu, "Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field," in *DSN*, 2015.
- [28] H. David, C. Fallin, E. Gorbato, U. R. Hanebutte, and O. Mutlu, "Memory Power Management via Dynamic Voltage/Frequency Scaling," in *ICAC*, 2011.
- [29] Q. Deng, D. Meisner, L. Ramos, T. F. Wenisch, and R. Bianchini, "MemScale: Active Low-Power Modes for Main Memory," in *ASPLOS*, 2011.
- [30] J. Dean and L. A. Barroso, "The Tail at Scale," *CACM*, 2013.
- [31] M. Ferdman, A. Adileh, O. Kocberber, S. Volos, M. Alisafae, D. Jevdjic, C. Kaynak, A. D. Popescu, A. Ailamaki, and B. Falsafi, "Clearing the Clouds: A Study of Emerging Scale-Out Workloads on Modern Hardware," in *ASPLOS*, 2012.
- [32] L. Wang, J. Zhan, C. Luo, Y. Zhu, Q. Yang, Y. He, W. Gao, Z. Jia, Y. Shi, S. Zhang *et al.*, "BigDataBench: A Big Data Benchmark Suite from Internet Services," in *HPCA*, 2014.
- [33] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "Enabling Practical Processing in and Near Memory for Data-Intensive Computing," in *DAC*, 2019.
- [34] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "Processing Data Where It Makes Sense: Enabling In-Memory Computation," *MicPro*, 2019.
- [35] O. Mutlu, "Intelligent Architectures for Intelligent Machines," in *VLSI-DAT*, 2020.
- [36] S. Ghose, A. Boroumand, J. S. Kim, J. Gómez-Luna, and O. Mutlu, "Processing-in-Memory: A Workload-Driven Perspective," *IBM JRD*, 2019.
- [37] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "A Modern Primer on Processing in Memory," in *Emerging Computing: From Devices to Systems – Looking Beyond Moore and Von Neumann*. Springer, 2021.
- [38] S. Wang and E. Ipek, "Reducing Data Movement Energy via Online Data Clustering and Encoding," in *MICRO*, 2016.
- [39] D. Pandiyan and C.-J. Wu, "Quantifying the Energy Cost of Data Movement for Emerging Smart Phone Workloads on Mobile Platforms," in *IISWC*, 2014.
- [40] S. Koppula, L. Orosa, A. G. Yağlıkcı, R. Azizi, T. Shahroodi, K. Kanellopoulos, and O. Mutlu, "EDEN: Enabling Energy-Efficient, High-Performance Deep Neural Network Inference Using Approximate DRAM," in *MICRO*, 2019.
- [41] S. A. McKee, "Reflections on the Memory Wall," in *CF*, 2004.
- [42] M. V. Wilkes, "The Memory Gap and the Future of High Performance Memories," *CAN*, 2001.
- [43] Y. Kim, V. Seshadri, D. Lee, J. Liu, and O. Mutlu, "A Case for Exploiting Subarray-Level Parallelism (SALP) in DRAM," in *ISCA*, 2012.
- [44] W. A. Wulf and S. A. McKee, "Hitting the Memory Wall: Implications of the Obvious," *CAN*, 1995.
- [45] S. Ghose, T. Li, N. Hajinazar, D. S. Cali, and O. Mutlu, "Demystifying Complex Workload-DRAM Interactions: An Experimental Study," in *SIGMETRICS*, 2020.
- [46] J. Ahn, S. Hong, S. Yoo, O. Mutlu, and K. Choi, "A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing," in *ISCA*, 2015.
- [47] J. Ahn, S. Yoo, O. Mutlu, and K. Choi, "PIM-Enabled Instructions: A Low-Overhead, Locality-Aware Processing-in-Memory Architecture," in *ISCA*, 2015.
- [48] K. Hsieh, E. Ebrahimi, G. Kim, N. Chatterjee, M. O'Connor, N. Vijaykumar, O. Mutlu, and S. W. Keckler, "Transparent Offloading and Mapping (TOM) Enabling Programmer-Transparent Near-Data Processing in GPU Systems," in *ISCA*, 2016.
- [49] Y. Wang, L. Orosa, X. Peng, Y. Guo, S. Ghose, M. Patel, J. S. Kim, J. G. Luna, M. Sadrosadati, N. M. Ghiasi *et al.*, "FIGARO: Improving System Performance via Fine-Grained In-DRAM Data Relocation and Caching," in *MICRO*, 2020.
- [50] Z. Jia, J. Zhan, L. Wang, C. Luo, W. Gao, Y. Jin, R. Han, and L. Zhang, "Understanding Big Data Analytics Workloads on Modern Processors," *TPDS*, 2016.
- [51] P.-A. Tsai, C. Chen, and D. Sanchez, "Adaptive Scheduling for Systems with Asymmetric Memory Hierarchies," in *MICRO*, 2018.
- [52] R. Sites, "It's the Memory, Stupid!" *MPR*, 1996.
- [53] V. Seshadri, T. Mullins, A. Boroumand, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "Gather-Scatter DRAM: In-DRAM Address Translation to Improve the Spatial Locality of Non-Unit Strided Accesses," in *MICRO*, 2015.
- [54] L. Nai, R. Hadidi, J. Sim, H. Kim, P. Kumar, and H. Kim, "GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks," in *HPCA*, 2017.
- [55] 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.
- [56] E. Ebrahimi, O. Mutlu, and Y. N. Patt, "Techniques for Bandwidth-Efficient Prefetching of Linked Data Structures in Hybrid Prefetching Systems," in *HPCA*, 2009.
- [57] O. Mutlu, J. Stark, C. Wilkerson, and Y. N. Patt, "Runahead Execution: An Alternative to Very Large Instruction Windows for Out-of-Order Processors," in *HPCA*, 2003.
- [58] M. Hashemi, O. Mutlu, and Y. N. Patt, "Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads," in *MICRO*, 2016.
- [59] J. S. Kim, D. S. Cali, H. Xin, D. Lee, S. Ghose, M. Alser, H. Hassan, O. Ergin, C. Alkan, and O. Mutlu, "GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies," *BMC Genomics*, 2018.
- [60] D. S. Cali, G. S. Kalsi, Z. Bingöl, C. Firtina, L. Subramanian, J. S. Kim, R. Ausavarungnirun, M. Alser, J. Gomez-Luna, A. Boroumand *et al.*, "GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis," in *MICRO*, 2020.
- [61] A. Boroumand, "Practical Mechanisms for Reducing Processor-Memory Data Movement in Modern Workloads," Ph.D. dissertation, Carnegie Mellon University, 2020.
- [62] D. Lee, S. Ghose, G. Pekhimenko, S. Khan, and O. Mutlu, "Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost," *TACO*, 2016.
- [63] A. Boroumand, S. Ghose, B. Lucia, K. Hsieh, K. Malladi, H. Zheng, and O. Mutlu, "LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory," *CAL*, 2017.
- [64] D. Zhang, N. Jayasena, A. Lyashevsky, J. L. Greathouse, L. Xu, and M. Ignatowski, "TOP-PIM: Throughput-Oriented Programmable Processing in Memory," in *HPDC*, 2014.
- [65] M. Gao and C. Kozyrakis, "HRL: Efficient and Flexible Reconfigurable Logic for Near-Data Processing," in *HPCA*, 2016.
- [66] M. Drumond, A. Daglis, N. Mirzadeh, D. Ustiugov, J. Picorel, B. Falsafi, B. Grot, and D. Pnevmatikatos, "The Mondrian Data Engine," in *ISCA*, 2017.
- [67] P. C. Santos, G. F. Oliveira, D. G. Tomé, M. A. Z. Alves, E. C. Almeida, and L. Carro, "Operand Size Reconfiguration for Big Data Processing in Memory," in *DATE*, 2017.
- [68] G. F. Oliveira, P. C. Santos, M. A. Alves, and L. Carro, "NIM: An HMC-Based Machine for Neuron Computation," in *ARC*, 2017.
- [69] M. Gao, J. Pu, X. Yang, M. Horowitz, and C. Kozyrakis, "TETRIS: Scalable and Efficient Neural Network Acceleration with 3D Memory," in *ASPLOS*, 2017.
- [70] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay, "Neurocube: A Programmable Digital Neuromorphic Architecture with High-Density 3D Memory," in *ISCA*, 2016.
- [71] P. Gu, S. Li, D. Stow, R. Barnes, L. Liu, Y. Xie, and E. Kursun, "Leveraging 3D Technologies for Hardware Security: Opportunities and Challenges," in *GLSVLSI*, 2016.
- [72] D. U. Lee, K. W. Kim, K. W. Kim, H. Kim, J. Y. Kim, Y. J. Park, J. H. Kim, D. S. Kim, H. B. Park, J. W. Shin *et al.*, "A 1.2V 8Gb 8-Channel 128GB/s High-Bandwidth Memory (HBM) Stacked DRAM with Effective Microbump I/O Test Methods Using 29nm Process and TSV," in *ISSCC*, 2014.
- [73] Hybrid Memory Cube Consortium, "Hybrid Memory Cube Specification Rev. 2.0," <http://www.hybridmemorycube.org/>.
- [74] A. Boroumand, S. Ghose, M. Patel, H. Hassan, B. Lucia, R. Ausavarungnirun, K. Hsieh, N. Hajinazar, K. T. Malladi, H. Zheng *et al.*, "CoNDA: Efficient Cache Coherence Support for Near-Data Accelerators," in *ISCA*, 2019.
- [75] Q. Zhu, T. Graf, H. E. Sumbul, L. Pileggi, and F. Franchetti, "Accelerating Sparse Matrix-Matrix Multiplication with 3D-Stacked Logic-in-Memory Hardware," in *HPEC*, 2013.
- [76] S. H. Pugsley, J. Jests, H. Zhang, R. Balasubramanian *et al.*, "NDC: Analyzing the Impact of 3D-Stacked Memory+Logic Devices on MapReduce Workloads," in *ISPASS*, 2014.
- [77] A. Farnahini-Farahani, J. H. Ahn, K. Morrow, and N. S. Kim, "NDA: Near-DRAM Acceleration Architecture Leveraging Commodity DRAM Devices and Standard Memory Modules," in *HPCA*, 2015.
- [78] G. H. Loh, N. Jayasena, M. Oskun, M. Nutter, D. Roberts, M. Meswani, D. P. Zhang, and M. Ignatowski, "A Processing in Memory Taxonomy and a Case for Studying Fixed-Function PIM," in *WoNDP*, 2013.
- [79] A. Pattnaik, X. Tang, A. Jog, O. Kayiran, A. K. Mishra, M. T. Kandemir, O. Mutlu, and C. R. Das, "Scheduling Techniques for GPU Architectures with Processing-in-Memory Capabilities," in *PACT*, 2016.
- [80] B. Akin, F. Franchetti, and J. C. Hoe, "Data Reorganization in Memory Using 3D-Stacked DRAM," in *ISCA*, 2015.
- [81] O. O. Babarinsa and S. Ideos, "JAFAR: Near-Data Processing for Databases," in *SIGMOD*, 2015.
- [82] J. H. Lee, J. Sim, and H. Kim, "BSSync: Processing Near Memory for Machine Learning Workloads with Bounded Staleness Consistency Models," in *PACT*, 2015.
- [83] F. Devaux, "The True Processing in Memory Accelerator," in *Hot Chips*, 2019.
- [84] P. Chi, S. Li, C. Xu, T. Zhang, J. Zhao, Y. Liu, Y. Wang, and Y. Xie, "PRIME: A Novel Processing-in-Memory Architecture for Neural Network Computation in ReRAM-Based Main Memory," in *ISCA*, 2016.
- [85] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramanian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar, "ISAAC: A Convolutional Neural Network Accelerator with In-Situ Analog Arithmetic in Crossbars," in *ISCA*, 2016.

- [86] 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.
- [87] V. Seshadri and O. Mutlu, "In-DRAM Bulk Bitwise Execution Engine," arXiv:1905.09822 [cs.AR], 2019.
- [88] S. Li, D. Niu, K. T. Malladi, H. Zheng, B. Brennan, and Y. Xie, "Drisa: A DRAM-Based Reconfigurable In-Situ Accelerator," in *MICRO*, 2017.
- [89] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch *et al.*, "RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization," in *MICRO*, 2013.
- [90] V. Seshadri and O. Mutlu, "The Processing Using Memory Paradigm: In-DRAM Bulk Copy, Initialization, Bitwise AND and OR," arXiv:1610.09603 [cs.AR], 2016.
- [91] Q. Deng, L. Jiang, Y. Zhang, M. Zhang, and J. Yang, "DrAcc: A DRAM Based Accelerator for Accurate CNN Inference," in *DAC*, 2018.
- [92] X. Xin, Y. Zhang, and J. Yang, "ELP2IM: Efficient and Low Power Bitwise Operation Processing in DRAM," in *HPCA*, 2020.
- [93] L. Song, Y. Zhuo, X. Qian, H. Li, and Y. Chen, "GraphR: Accelerating Graph Processing Using ReRAM," in *HPCA*, 2018.
- [94] L. Song, X. Qian, H. Li, and Y. Chen, "PipeLayer: A Pipelined ReRAM-Based Accelerator for Deep Learning," in *HPCA*, 2017.
- [95] F. Gao, G. Tziatzoulis, and D. Wentzlaff, "ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs," in *MICRO*, 2019.
- [96] C. Eckert, S. Wang, J. Wang, A. Subramanian, R. Iyer, D. Sylvester, D. Blaauw, and R. Das, "Neural Cache: Bit-Serial In-Cache Acceleration of Deep Neural Networks," in *ISCA*, 2018.
- [97] S. Aga, S. Jeloka, A. Subramanian, S. Narayanasamy, D. Blaauw, and R. Das, "Compute Caches," in *HPCA*, 2017.
- [98] D. Fujiki, S. Mahlke, and R. Das, "Duality Cache for Data Parallel Acceleration," in *ISCA*, 2019.
- [99] I. Fernandez, R. Quisilant, E. Gutiérrez, O. Plata, C. Giannoula, M. Alser, J. Gómez-Luna, and O. Mutlu, "NATSA: A Near-Data Processing Accelerator for Time Series Analysis," in *ICCD*, 2020.
- [100] N. Hajinazar, G. F. Oliveira, S. Gregorio, J. D. Ferreira, N. M. Ghiasi, M. Patel, M. Alser, S. Ghose, J. Gómez-Luna, and O. Mutlu, "SIMDRAM: A Framework for Bit-Serial SIMD Processing Using DRAM," in *ASPLOS*, 2021.
- [101] C. Giannoula, N. Vijaykumar, N. Papadopoulou, V. Karakostas, I. Fernandez, J. Gómez-Luna, L. Orosa, N. Koziris, G. Goumas, and O. Mutlu, "SynCron: Efficient Synchronization Support for Near-Data-Processing Architectures," in *HPCA*, 2021.
- [102] A. Boroumand, S. Ghose, G. F. Oliveira, and O. Mutlu, "Polynesia: Enabling Effective Hybrid Transactional/Analytical Databases with Specialized Hardware/Software Co-Design," arXiv:2103.00798 [cs.AR], 2021.
- [103] A. Boroumand, S. Ghose, B. Akin, R. Narayanaswami, G. F. Oliveira, X. Ma, E. Shiu, and O. Mutlu, "Mitigating Edge Machine Learning Inference Bottlenecks: An Empirical Study on Accelerating Google Edge Models," arXiv:2103.00768 [cs.AR], 2021.
- [104] G. F. Oliveira, P. C. Santos, M. A. Alves, and L. Carro, "A Generic Processing in Memory Cycle Accurate Simulator Under Hybrid Memory Cube Architecture," in *SAMOS*, 2017.
- [105] J. S. Kim, M. Patel, H. Hassan, L. Orosa, and O. Mutlu, "D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers With Low Latency and High Throughput," in *HPCA*, 2019.
- [106] J. S. Kim, M. Patel, H. Hassan, and O. Mutlu, "The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices," in *HPCA*, 2018.
- [107] M. Besta, R. Kanakagiri, G. Kwasniewski, R. Ausavarungnirun, J. Beránek, K. Kanellopoulos, K. Janda, Z. Vónarburg-Shmaria, L. Gianinazzi, I. Stefan *et al.*, "SISA: Set-Centric Instruction Set Architecture for Graph Mining on Processing-in-Memory Systems," arXiv:2104.07582 [cs.AR], 2021.
- [108] J. D. Ferreira, G. Falcao, J. Gómez-Luna, M. Alser, L. Orosa, M. Sadrosadati, J. S. Kim, G. F. Oliveira, T. Shahroodi, A. Nori *et al.*, "pLUTo: In-DRAM Lookup Tables to Enable Massively Parallel General-Purpose Computation," arXiv:2104.07699 [cs.AR], 2021.
- [109] V. Seshadri, D. Lee, T. Mullins, H. Hassan, A. Boroumand, J. Kim, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry, "Buddy-RAM: Improving the Performance and Efficiency of Bulk Bitwise Operations Using DRAM," arXiv:1611.09988 [cs.AR], 2016.
- [110] A. Boroumand, S. Ghose, M. Patel, H. Hassan, B. Lucia, N. Hajinazar, K. Hsieh, K. T. Malladi, H. Zheng, and O. Mutlu, "LazyPIM: Efficient Support for Cache Coherence in Processing-in-Memory Architectures," arXiv:1706.03162 [cs.AR], 2017.
- [111] J. S. Kim, D. Senol, H. Xin, D. Lee, S. Ghose, M. Alser, H. Hassan, O. Ergin, C. Alkan, and O. Mutlu, "GRIM-Filter: Fast Seed Filtering in Read Mapping using Emerging Memory Technologies," arXiv:1708.04329 [q-bio.GN], 2017.
- [112] S. Ghose, K. Hsieh, A. Boroumand, R. Ausavarungnirun, and O. Mutlu, "Enabling the Adoption of Processing-in-Memory: Challenges, Mechanisms, Future Research Directions," arXiv:1802.00320 [cs.AR], 2018.
- [113] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch *et al.*, "RowClone: Accelerating Data Movement and Initialization Using DRAM," arXiv:1805.03502 [cs.AR], 2018.
- [114] S. Ghose, A. Boroumand, J. S. Kim, J. Gómez-Luna, and O. Mutlu, "A Workload and Programming Ease Driven Perspective of Processing-in-Memory," arXiv:1907.12947 [cs.DC], 2019.
- [115] A. Olgun, M. Patel, A. G. Yaglikci, H. Luo, J. S. Kim, F. N. Bostanci, N. Vijaykumar, O. Ergin, and O. Mutlu, "QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAMs," in *ISCA*, 2021.
- [116] M. Gokhale, B. Holmes, and K. Iobst, "Processing in Memory: The Terasys Massively Parallel PIM Array," *Computer*, 1995.
- [117] R. Balasubramanian, J. Chang, T. Manning *et al.*, "Near-Data Processing: Insights from a MICRO-46 Workshop," *IEEE Micro*, 2014.
- [118] R. Nair, "Evolution of Memory Architecture," *Proceedings of the IEEE*, 2015.
- [119] E. Azarkhish, D. Rossi, I. Loi, and L. Benini, "A Case for Near Memory Computation Inside the Smart Memory Cube," in *EMS*, 2016.
- [120] J. D. McCalpin *et al.*, "Memory Bandwidth and Machine Balance in Current High Performance Computers," *IEEE TCCA Newsletter*, 1995.
- [121] 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.
- [122] M. Zhang, Y. Zhuo, C. Wang, M. Gao, Y. Wu, K. Chen, C. Kozyrakis, and X. Qian, "GraphP: Reducing Communication for PIM-Based Graph Processing with Efficient Data Partition," in *HPCA*, 2018.
- [123] S. Angizi, J. Sun, W. Zhang, and D. Fan, "GraphS: A Graph Processing Accelerator Leveraging SOT-MRAM," in *DATE*, 2019.
- [124] K. K. Matam, G. Koo, H. Zha, H.-W. Tseng, and M. Annavaram, "GraphSSD: Graph Semantics Aware SSD," in *ISCA*, 2019.
- [125] S. Angizi and D. Fan, "GraphiDe: A Graph Processing Accelerator Leveraging In-DRAM-Computing," in *GLSVLSI*, 2019.
- [126] Y. Zhuo, C. Wang, M. Zhang, R. Wang, D. Niu, Y. Wang, and X. Qian, "GraphQ: Scalable PIM-Based Graph Processing," in *MICRO*, 2019.
- [127] 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.
- [128] S. Li, C. Xu, Q. Zou, J. Zhao, Y. Lu, and Y. Xie, "Pinatubo: A Processing-in-Memory Architecture for Bulk Bitwise Operations in Emerging Non-Volatile Memories," in *DAC*, 2016.
- [129] S. H. S. Rezaei, M. Modarressi, R. Ausavarungnirun, M. Sadrosadati, O. Mutlu, and M. Daneshmand, "NoM: Network-on-Memory for Inter-Bank Data Transfer in Highly-Banked Memories," *CAL*, 2020.
- [130] V. Seshadri, K. Hsieh, A. Boroumand, D. Lee, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry, "Fast Bulk Bitwise AND and OR in DRAM," *CAL*, 2015.
- [131] S. Williams, A. Waterman, and D. Patterson, "Roofline: An Insightful Visual Performance Model for Multicore Architectures," *CACM*, 2009.
- [132] E. Azarkhish, D. Rossi, I. Loi, and L. Benini, "Neurostream: Scalable and Energy Efficient Deep Learning with Smart Memory Cubes," *TPDS*, 2018.
- [133] L. Ke, U. Gupta, B. Y. Cho, D. Brooks, V. Chandra, U. Diril, A. Firoozshahian, K. Hazelwood, B. Jia, H.-H. S. Lee *et al.*, "RecNMP: Accelerating Personalized Recommendation with Near-Memory Processing," in *ISCA*, 2020.
- [134] B. Asgari, S. Mukhopadhyay, and S. Yalamanchili, "MAHASIM: Machine-Learning Hardware Acceleration Using a Software-Defined Intelligent Memory System," *J. Signal Process Syst.*, 2020.
- [135] N. Kim, J. Ahn, S. Hong, H. Chafi, and K. Choi, "How Much Computation Power Do You Need for Near-Data Processing in Cloud?" in *ASBD*, 2017.
- [136] S. Liang, Y. Wang, C. Liu, H. Li, and X. Li, "InS-DLA: An In-SSD Deep Learning Accelerator for Near-Data Processing," in *FPL*, 2019.
- [137] A. O. Glova, I. Akgun, S. Li, X. Hu, and Y. Xie, "Near-Data Acceleration of Privacy-Preserving Biomarker Search with 3D-Stacked Memory," in *DATE*, 2019.
- [138] P. Gu, X. Xie, S. Li, D. Niu, H. Zheng, K. T. Malladi, and Y. Xie, "DLUX: A LUT-Based Near-Bank Accelerator for Data Center Deep Learning Training Workloads," *TCAD*, 2020.
- [139] G. Singh, D. Diamantopoulos, C. Hagleitner, J. Gomez-Luna, S. Stuijk, O. Mutlu, and H. Corporaal, "NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling," in *FPL*, 2020.
- [140] J. Gómez-Luna, I. E. Hajj, I. Fernández, C. Giannoula, G. F. Oliveira, and O. Mutlu, "Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture," arXiv:2105.03814 [cs.AR], 2021.
- [141] M. Radulovic, D. Zivanovic, D. Ruiz, B. R. de Supinski, S. A. McKee, P. Radojkovic, and E. Ayguadé, "Another Trip to the Wall: How Much Will Stacked DRAM Benefit HPC?" in *MEMSYS*, 2015.
- [142] L. Yavits, R. Kaplan, and R. Ginosar, "GIRAF: General Purpose In-Storage Resistive Associative Framework," *TPDS*, 2021.

- [143] J. M. Herruzo, I. Fernandez, S. González-Navarro, and O. Plata, "Enabling Fast and Energy-Efficient FM-Index Exact Matching Using Processing-Near-Memory," *The Journal of Supercomputing*, 2021.
- [144] B. Asgari, R. Hadidi, J. Cao, D. E. Shim, S.-K. Lim, and H. Kim, "FAFNIR: Accelerating Sparse Gathering by Using Efficient Near-Memory Intelligent Reduction," in *HPCA*, 2021.
- [145] H. Kim, H. Kim, S. Yalamanchili, and A. F. Rodrigues, "Understanding Energy Aspects of Processing-Near-Memory for HPC Workloads," in *MEMSYS*, 2015.
- [146] L. Nai and H. Kim, "Instruction Offloading with HMC 2.0 Standard: A Case Study for Graph Traversals," in *MEMSYS*, 2015.
- [147] H. Lim and G. Park, "Triple Engine Processor (TEP): A Heterogeneous Near-Memory Processor for Diverse Kernel Operations," *TACO*, 2017.
- [148] H. Kim, R. Hadidi, L. Nai, H. Kim, N. Jayasena, Y. Eckert, O. Kayiran, and G. Loh, "CODA: Enabling Co-Location of Computation and Data for Multiple GPU Systems," *TACO*, 2018.
- [149] B. Hong, G. Kim, J. H. Ahn, Y. Kwon, H. Kim, and J. Kim, "Accelerating Linked-List Traversal Through Near-Data Processing," in *PACT*, 2016.
- [150] M. Gries, P. Cabré, and J. Gago, "Performance Evaluation and Feasibility Study of Near-Data Processing on DRAM Modules (DIMM-NDP) for Scientific Applications," *Technical Report*, 2019.
- [151] M. Hashemi, E. Ebrahimi, O. Mutlu, Y. N. Patt *et al.*, "Accelerating Dependent Cache Misses with an Enhanced Memory Controller," in *ISCA*, 2016.
- [152] C. Chou, P. Nair, and M. K. Qureshi, "Reducing Refresh Power in Mobile Devices with Morphable ECC," in *DSN*, 2015.
- [153] 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.
- [154] Y. Kim, M. Papamichael, O. Mutlu, and M. Harchol-Balter, "Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior," in *MICRO*, 2010.
- [155] 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.
- [156] L. Subramanian, D. Lee, V. Seshadri, H. Rastogi, and O. Mutlu, "BLISS: Balancing Performance, Fairness and Complexity in Memory Access Scheduling," *TPDS*, 2016.
- [157] H. Usui, L. Subramanian, K. K.-W. Chang, and O. Mutlu, "DASH: Deadline-Aware High-Performance Memory Scheduler for Heterogeneous Systems With Hardware Accelerators," *TACO*, 2016.
- [158] SAFARI Research Group, "DAMOV Benchmark Suite and Simulation Framework," <https://github.com/CMU-SAFARI/DAMOV>.
- [159] D. Sanchez and C. Kozyrakis, "ZSim: Fast and Accurate Microarchitectural Simulation of Thousand-Core Systems," in *ISCA*, 2013.
- [160] Y. Kim, W. Yang, and O. Mutlu, "Ramulator: A Fast and Extensible DRAM Simulator," *CAL*, 2016.
- [161] Intel Corp., "Intel VTune Profiler User Guide," <https://software.intel.com/content/www/us/en/develop/documentation/vtune-help/>, 2021.
- [162] A. Yasin, "A Top-Down Method for Performance Analysis and Counters Architecture," in *ISPASS*, 2014.
- [163] Intel Corp., "Intel VTune Profiler User Guide," <https://software.intel.com/content/www/us/en/develop/documentation/vtune-help/top/reference/cpumetrics-reference/memory-bound.html>.
- [164] T. M. Conte and W.-m. W. Hwu, "Benchmark Characterization," *IEEE Computer*, 1991.
- [165] M. S. Johnstone and P. R. Wilson, "The Memory Fragmentation Problem: Solved?" in *ISMM*, 1998.
- [166] J. Weinberg, M. O. McCracken, E. Strohmaier, and A. Snavey, "Quantifying Locality in the Memory Access Patterns of HPC Applications," in *SC*, 2005.
- [167] Y. S. Shao and D. Brooks, "ISA-Independent Workload Characterization and Its Implications for Specialized Architectures," in *ISPASS*, 2013.
- [168] Y. Zhong, X. Shen, and C. Ding, "Program Locality Analysis Using Reuse Distance," *TOPLAS*, 2009.
- [169] X. Gu, I. Christopher, T. Bai, C. Zhang, and C. Ding, "A Component Model of Spatial Locality," in *ISMM*, 2009.
- [170] J. C. Beard and J. Randall, "Eliminating Dark Bandwidth: A Data-Centric View of Scalable, Efficient Performance, Post-Moore," in *HiPC*, 2017.
- [171] S. Lloyd and M. Gokhale, "In-Memory Data Rearrangement for Irregular, Data-Intensive Computing," *IEEE Computer*, 2015.
- [172] T. M. Conte and W.-m. W. Hwu, "A Brief Survey of Benchmark Usage in the Architecture Community," *CAN*, 1991.
- [173] T. M. Conte and W.-m. W. Hwu, "Advances in Benchmarking Techniques: New Standards and Quantitative Metrics," in *Advances in Computers*, 1995.
- [174] D. Doerfler, J. Deslippe, S. Williams, L. Oliker, B. Cook, T. Kurth, M. Lobet, T. Malas, J.-L. Vay, and H. Vincenti, "Applying the Roofline Performance Model to the Intel Xeon Phi Knights Landing Processor," in *HiPC*, 2016.
- [175] O. Mutlu and T. Moscibroda, "Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors," in *MICRO*, 2007.
- [176] O. Mutlu and T. Moscibroda, "Parallelism-Aware Batch Scheduling: Enhancing Both Performance and Fairness of Shared DRAM Systems," in *ISCA*, 2008.
- [177] E. Ebrahimi, C. J. Lee, O. Mutlu, and Y. N. Patt, "Fairness via Source Throttling: A Configurable and High-Performance Fairness Substrate for Multi-Core Memory Systems," in *ASPLOS*, 2010.
- [178] S. Palacharla and R. E. Kessler, "Evaluating Stream Buffers as a Secondary Cache Replacement," in *ISCA*, 1994.
- [179] G. Singh, G. F. Oliveira, S. Corda, S. Stuijk, O. Mutlu, and H. Corporaal, "NAPEL: Near-Memory Computing Application Performance Prediction via Ensemble Learning," in *DAC*, 2019.
- [180] P. C. Santos, M. A. Alves, M. Diener, L. Carro, and P. O. Navaux, "Exploring Cache Size and Core Count Tradeoffs in Systems with Reduced Memory Access Latency," in *PDP*, 2016.
- [181] A. Glew, "MLP Yes! ILP No!" in *ASPLOS WACI*, 1998.
- [182] M. K. Qureshi, D. N. Lynch, O. Mutlu, and Y. N. Patt, "A Case for MLP-Aware Cache Replacement," in *ISCA*, 2006.
- [183] O. Mutlu, H. Kim, and Y. N. Patt, "Efficient Runahead Execution: Power-Efficient Memory Latency Tolerance," *IEEE Micro*, 2006.
- [184] O. Mutlu, H. Kim, and Y. N. Patt, "Techniques for Efficient Processing in Runahead Execution Engines," in *ISCA*, 2005.
- [185] T.-Y. Yeh and Y. N. Patt, "Two-Level Adaptive Training Branch Prediction," in *MICRO*, 1991.
- [186] N. Muralimanohar, R. Balasubramanian, and N. Jouppi, "Optimizing NUCA Organizations and Wiring Alternatives for Large Caches with CACTI 6.0," in *MICRO*, 2007.
- [187] R. Ausavarungnirun, C. Fallin, X. Yu, K. K.-W. Chang, G. Nazario, R. Das, G. H. Loh, and O. Mutlu, "Design and Evaluation of Hierarchical Rings with Deflection Routing," in *SBAC-PAD*, 2014.
- [188] M. S. Papamarcos and J. H. Patel, "A Low-Overhead Coherence Solution for Multiprocessors with Private Cache Memories," in *ISCA*, 1984.
- [189] S. Srinath, O. Mutlu, H. Kim, and Y. N. Patt, "Feedback Directed Prefetching: Improving the Performance and Bandwidth-Efficiency of Hardware Prefetchers," in *HPCA*, 2007.
- [190] M. Gao, G. Ayers, and C. Kozyrakis, "Practical Near-Data Processing for in-Memory Analytics Frameworks," in *PACT*, 2015.
- [191] G. Kim, J. Kim, J. H. Ahn, and J. Kim, "Memory-Centric System Interconnect Design with Hybrid Memory Cubes," in *PACT*, 2013.
- [192] R. Nair, S. F. Antao, C. Bertolli, P. Bose *et al.*, "Active Memory Cube: A Processing-in-Memory Architecture for Exascale Systems," *IBM JRD*, 2015.
- [193] E. Lockerman, A. Feldmann, M. Bakhshalipour, A. Stanescu, S. Gupta, D. Sanchez, and N. Beckmann, "Livia: Data-Centric Computing Throughout the Memory Hierarchy," in *ASPLOS*, 2020.
- [194] J. P. C. de Lima, P. C. Santos, M. A. Alves, A. Beck, and L. Carro, "Design Space Exploration for PIM Architectures in 3D-Stacked Memories," in *CF*, 2018.
- [195] U. Sirin, A. Yasin, and A. Ailamaki, "A Methodology for OLTP Micro-Architectural Analysis," in *DAMON*, 2017.
- [196] R. Appuswamy, J. Fellay, and N. Chaturvedi, "Sequence Alignment Through the Looking Glass," in *IPDPSW*, 2018.
- [197] Intel Corp., "Intel Xeon Processor E3-1240," <http://ark.intel.com/products/52273/>, 2011.
- [198] Intel Corp., "Understanding How General Exploration Works in Intel VTune Amplifier," <https://software.intel.com/en-us/articles/understanding-how-general-exploration-works-in-intel-vtune-amplifier-xe>, 2018.
- [199] J. Gómez-Luna, I. El Hajj, V. Chang, Li-Wen Garcia-Flores, S. Garcia de Gonzalo, T. Jablin, A. J. Pena, and W.-m. Hwu, "Chai: Collaborative Heterogeneous Applications for Integrated Architectures," in *ISPASS*, 2017.
- [200] United States Department of Energy, "CORAL Benchmark Codes," <https://asc.llnl.gov/CORAL-benchmarks/>, 2014.
- [201] J. A. Stratton, C. Rodrigues, I.-J. Sung, N. Obeid, L.-W. Chang, N. Anssari, G. D. Liu, and W.-m. W. Hwu, "Parboil: A Revised Benchmark Suite for Scientific and Commercial Throughput Computing," Univ. of Illinois at Urbana-Champaign, IMPACT Research Group, Tech. Rep. IMPACT-12-01, 2012.
- [202] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC Benchmark Suite: Characterization and Architectural Implications," in *PACT*, 2008.
- [203] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S. Lee, and K. Skadron, "Rodinia: A Benchmark Suite for Heterogeneous Computing," in *IISWC*, 2009.
- [204] S. K. Venkata, I. Ahn, D. Jeon, A. Gupta, C. Louie, S. Garcia, S. Belongie, and M. B. Taylor, "SD-VBS: The San Diego Vision Benchmark Suite," in *IISWC*, 2009.
- [205] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The SPLASH-2 Programs: Characterization and Methodological Considerations," in *ISCA*, 1995.
- [206] P. R. Luszczyk, D. H. Bailey, J. J. Dongarra, J. Kepner, R. F. Lucas, R. Rabenseifner, and D. Takahashi, "The HPC Challenge (HPCC) Benchmark Suite," in *SC*, 2006.
- [207] J. Dongarra, M. A. Heroux, and P. Luszczyk, "HPCG Benchmark: A New Metric for Ranking High Performance Computing Systems," Univ. of Tennessee Dept. of Electrical Engg. and Computer Sci., Tech. Rep. UT-EECS-15-736, 2015.
- [208] N. Ahmed, K. Bertels, and Z. Al-Ars, "A Comparison of Seed-and-Extend Techniques in Modern DNA Read Alignment Algorithms," in *BIBM*, 2016.
- [209] C. Balkesen, J. Teubner, G. Alonso, and M. T. Ozsu, "Main-Memory Hash Joins on Modern Processor Architectures," *TKDE*, 2015.

- [210] J. Gómez-Luna, L. Chang, I. Sung, W. Hwu, and N. Guil, "In-Place Data Sliding Algorithms for Many-Core Architectures," in *ICPP*, 2015.
- [211] N. Sundaram, N. Satish, M. M. A. Patwary, S. R. Dulloor, M. J. Anderson, S. G. Vadlamudi, D. Das, and P. Dubey, "GraphMat: High Performance Graph Analytics Made Productive," *PVLDB*, 2015.
- [212] J. Shun and G. E. Blelloch, "Ligra: A Lightweight Graph Processing Framework for Shared Memory," in *PPoPP*, 2013.
- [213] R. M. Yoo, A. Romano, and C. Kozyrakis, "Phoenix Rebirth: Scalable MapReduce on a Large-Scale Shared-Memory System," in *IISWC*, 2009.
- [214] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *NIPS*, 2012.
- [215] J. Redmon, "Darknet: Open Source Neural Networks in C," <http://pjreddie.com/darknet>.
- [216] D. Chakrabarti, Y. Zhan, and C. Faloutsos, "R-MAT: A Recursive Model for Graph Mining," in *SDM*, 2004.
- [217] Center for Discrete Mathematics & Theoretical Computer Science, "9th DIMACS Implementation Challenge," <http://www.dis.uniroma1.it/challenge9/>, 2014.
- [218] M. Alser, H. Hassan, H. Xin, O. Ergin, O. Mutlu, and C. Alkan, "GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping," *Bioinformatics*, 2017.
- [219] M. Alser, Z. Bingöl, D. S. Cali, J. Kim, S. Ghose, C. Alkan, and O. Mutlu, "Accelerating Genome Analysis: A Primer on an Ongoing Journey," *IEEE Micro*, 2020.
- [220] R. Hadidi, B. Asgari, B. A. Mudassar, S. Mukhopadhyay, S. Yalamanchili, and H. Kim, "Demystifying the Characteristics of 3D-Stacked Memories: A Case Study for Hybrid Memory Cube," in *IISWC*, 2017.
- [221] J. A. Hartigan and M. A. Wong, "Algorithm AS 136: A K-Means Clustering Algorithm," *J. R. Stat. Soc. C-Appl.*, 1979.
- [222] E. Azarkhish, "Memory Hierarchy Design for Next Generation Scalable Many-Core Platforms," Ph.D. dissertation, Università di Bologna, 2016.
- [223] S. Pugsley, R. Balasubramanian, J. Jestes, F. Li, A. Davis, V. Srinivasan, and A. Buyuktosunoglu, "Comparing Different Implementations of Near Data Computing with In-Memory MapReduce Workloads," *IEEE Micro*, 2014.
- [224] Y. Chou, B. Fahs, and S. Abraham, "Microarchitecture Optimizations for Exploiting Memory-Level Parallelism," in *ISCA*, 2004.
- [225] J. Tuck, L. Ceze, and J. Torrellas, "Scalable Cache Miss Handling for High Memory-Level Parallelism," in *MICRO*, 2006.
- [226] S. Phadke and S. Narayanasamy, "MLP Aware Heterogeneous Memory System," in *DATE*, 2011.
- [227] K. Van Craeynest, S. Eyerman, and L. Eeckhout, "MLP-Aware Runahead Threads in a Simultaneous Multithreading Processor," in *HiPEAC*, 2009.
- [228] S. Everman and L. Eeckhout, "A Memory-Level Parallelism Aware Fetch Policy for SMT Processors," in *HPCA*, 2007.
- [229] G. Patsilaras, N. K. Choudhary, and J. Tuck, "Efficiently Exploiting Memory Level Parallelism on Asymmetric Coupled Cores in the Dark Silicon Era," *TACO*, 2012.
- [230] Y. Eckert, N. Jayasena, and G. H. Loh, "Thermal Feasibility of Die-Stacked Processing in Memory," in *WoNDP*, 2014.
- [231] K. Nilakant, V. Dalibard, A. Roy, and E. Yoneki, "PrefEdge: SSD Prefetcher for Large-Scale Graph Traversal," in *SYSTOR*, 2014.
- [232] A. M. Kaushik, G. Pekhimenko, and H. Patel, "Gretch: A Hardware Prefetcher for Graph Analytics," *TACO*, 2021.
- [233] S. Ainsworth and T. M. Jones, "Graph Prefetching Using Data Structure Knowledge," in *ICS*, 2016.
- [234] A. Basak, S. Li, X. Hu, S. M. Oh, X. Xie, L. Zhao, X. Jiang, and Y. Xie, "Analysis and Optimization of the Memory Hierarchy for Graph Processing Workloads," in *HPCA*, 2019.
- [235] M. Yan, X. Hu, S. Li, A. Basak, H. Li, X. Ma, I. Akgun, Y. Feng, P. Gu, L. Deng et al., "Alleviating Irregularity in Graph Analytics Acceleration: A Hardware/Software Co-Design Approach," in *MICRO*, 2019.
- [236] S. T. Srinivasan, R. Rajwar, H. Akkary, A. Gandhi, and M. Upton, "Continual Flow Pipelines," in *ASPLOS*, 2004.
- [237] M. Annamalai, J. M. Patel, and E. S. Davidson, "Data Prefetching by Dependence Graph Precomputation," in *ISCA*, 2001.
- [238] J. D. Collins, D. M. Tullsen, H. Wang, and J. P. Shen, "Dynamic Speculative Precomputation," in *ICS*, 2001.
- [239] J. Dundas and T. Mudge, "Improving Data Cache Performance by Pre-Executing Instructions Under a Cache Miss," in *ICS*, 1997.
- [240] O. Mutlu, J. Stark, C. Wilkerson, and Y. N. Patt, "Runahead Execution: An Effective Alternative to Large Instruction Windows," *IEEE Micro*, 2003.
- [241] T. Ramirez, A. Pajuelo, O. J. Santana, and M. Valero, "Runahead Threads to Improve SMT Performance," in *HPCA*, 2008.
- [242] T. Ramirez, A. Pajuelo, O. J. Santana, O. Mutlu, and M. Valero, "Efficient Runahead Threads," in *PACT*, 2010.
- [243] W. Zhang, D. M. Tullsen, and B. Calder, "Accelerating and Adapting Precomputation Threads for Efficient Prefetching," in *HPCA*, 2007.
- [244] R. Cooksey, S. Jourdan, and D. Grunwald, "A Stateless, Content-Directed Data Prefetching Mechanism," in *ASPLOS*, 2002.
- [245] A. Roth and G. S. Sohi, "Effective Jump-Pointer Prefetching for Linked Data Structures," in *ISCA*, 1999.
- [246] P. C. Santos, G. F. Oliveira, J. P. Lima, M. A. Alves, L. Carro, and A. C. Beck, "Processing in 3D Memories to Speed Up Operations on Complex Data Structures," in *DATE*, 2018.
- [247] S. Ghose, H. Lee, and J. F. Martinez, "Improving Memory Scheduling via Processor-Side Load Criticality Information," in *ISCA*, 2013.
- [248] O. Mutlu, "Lecture Notes for 18-447 Computer Architecture – Lecture 17: Memory Hierarchy and Caches," <https://course.ece.cmu.edu/~ece447/s15/lib/exe/fetch.php?media=onur-447-spring15-lecture17-memoryhierarchyandcaches-afterlecture.pdf>, 2015.
- [249] T. L. Johnson, D. A. Connors, M. C. Merten, and W.-m. W. Hwu, "Run-Time Cache Bypassing," *TC*, 1999.
- [250] A. Sembrant, E. Hagersten, and D. Black-Schaffer, "The Direct-to-Data (D2D) Cache: Navigating the Cache Hierarchy with a Single Lookup," in *ISCA*, 2014.
- [251] P.-A. Tsai, N. Beckmann, and D. Sanchez, "Jenga: Software-Defined Cache Hierarchies," in *ISCA*, 2017.
- [252] V. Seshadri, A. Bhowmick, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "The Dirty-Block Index," in *ISCA*, 2014.
- [253] 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," *TACO*, 2015.
- [254] T. L. Johnson and W.-M. W. Hwu, "Run-Time Adaptive Cache Hierarchy Management via Reference Analysis," in *ISCA*, 1997.
- [255] G. Tyson, M. Farrens, J. Matthews, and A. R. Pleszkun, "A Modified Approach to Data Cache Management," in *MICRO*, 1995.
- [256] G. Memik, G. Reinman, and W. H. Mangione-Smith, "Just Say No: Benefits of Early Cache Miss Determination," in *HPCA*, 2003.
- [257] M. Kharbutli and Y. Solihin, "Counter-Based Cache Replacement and Bypassing Algorithms," *TC*, 2008.
- [258] S. Gupta, H. Gao, and H. Zhou, "Adaptive Cache Bypassing for Inclusive Last Level Caches," in *IPDPS*, 2013.
- [259] L. Li, D. Tong, Z. Xie, J. Lu, and X. Cheng, "Optimal Bypass Monitor for High Performance Last-Level Caches," in *PACT*, 2012.
- [260] A. Sridharan and A. Seznec, "Discrete Cache Insertion Policies for Shared Last Level Cache Management on Large Multicores," in *IPDPS*, 2016.
- [261] 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.
- [262] 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.
- [263] H. Hassan, M. Patel, J. S. Kim, A. G. Yaglikci, N. Vijaykumar, N. M. Ghiasi, S. Ghose, and O. Mutlu, "CROW: A Low-Cost Substrate for Improving DRAM Performance, Energy Efficiency, and Reliability," in *ISCA*, 2019.
- [264] Y. Wang, A. Tavakkol, L. Orosa, S. Ghose, N. M. Ghiasi, M. Patel, J. S. Kim, H. Hassan, M. Sadrosadati, and O. Mutlu, "Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration," in *MICRO*, 2018.
- [265] J. Kim, M. Patel, H. Hassan, and O. Mutlu, "Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines," in *ICCD*, 2018.
- [266] A. Das, H. Hassan, and O. Mutlu, "VRL-DRAM: Improving DRAM Performance via Variable Refresh Latency," in *DAC*, 2018.
- [267] K. K. Chang, "Understanding and Improving the Latency of DRAM-Based Memory Systems," Ph.D. dissertation, Carnegie Mellon University, 2017.
- [268] H. Hassan, N. Vijaykumar, S. Khan, S. Ghose, 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.
- [269] D. Lee, "Reducing DRAM Latency at Low Cost by Exploiting Heterogeneity," Ph.D. dissertation, Carnegie Mellon University, 2016.
- [270] H. Luo, T. Shahroodi, H. Hassan, M. Patel, A. G. Yaglikci, L. Orosa, J. Park, and O. Mutlu, "CLR-DRAM: A Low-Cost DRAM Architecture Enabling Dynamic Capacity-Latency Trade-Off," in *ISCA*, 2020.
- [271] J. Choi, W. Shin, J. Jang, J. Suh, Y. Kwon, Y. Moon, and L.-S. Kim, "Multiple Clone Row DRAM: A Low Latency and Area Optimized DRAM," in *ISCA*, 2015.
- [272] Y. H. Son, O. Seongil, Y. Ro, J. W. Lee, and J. H. Ahn, "Reducing Memory Access Latency with Asymmetric DRAM Bank Organizations," in *ISCA*, 2013.
- [273] Micron, "RDRAM 2 and 3 Specifications," <https://www.micron.com/products/dram/rldram-memory>.
- [274] Y. Sato, T. Suzuki, T. Aikawa, S. Fujioka, W. Fujieda, H. Kobayashi, H. Ikeda, T. Nagasawa, A. Funyu, Y. Fuji et al., "Fast Cycle RAM (FCRAM): A 20-nm Random Row Access, Pipe-Lined Operating DRAM," in *VLSIC*, 1998.
- [275] L. Orosa, Y. Wang, M. Sadrosadati, J. Kim, M. Patel, I. Puddu, H. Luo, K. Razavi, J. Gómez-Luna, H. Hassan, N. M. Ghiasi, S. Ghose, and O. Mutlu, "CODIC: A Low-Cost Substrate for Enabling Custom In-DRAM Functionalities and Optimizations," in *ISCA*, 2021.
- [276] O. Seongil, Y. H. Son, N. S. Kim, and J. H. Ahn, "Row-Buffer Decoupling: A Case for Low-Latency DRAM Microarchitecture," in *ISCA*, 2014.
- [277] R. Ausavarungrun, K. K.-W. Chang, L. Subramanian, G. H. Loh, and O. Mutlu, "Staged Memory Scheduling: Achieving High Performance and Scalability in

- Heterogeneous Systems,” in *ISCA*, 2012.
- [278] 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.
- [279] 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.
- [280] 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.
- [281] S. Rixner, “Memory Controller Optimizations for Web Servers,” in *MICRO*, 2004.
- [282] S. Rixner, W. J. Dally, U. J. Kapasi, P. Mattson, and J. D. Owens, “Memory Access Scheduling,” in *ISCA*, 2000.
- [283] 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,” U.S. Patent 5 630 096, 1997.
- [284] T. Moscibroda and O. Mutlu, “Distributed Order Scheduling and Its Application to Multi-Core DRAM Controllers,” in *PODC*, 2008.
- [285] 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.
- [286] E. Ipek, O. Mutlu, J. F. Martinez, and R. Caruana, “Self-Optimizing Memory Controllers: A Reinforcement Learning Approach,” in *ISCA*, 2008.
- [287] I. Hur and C. Lin, “Adaptive History-Based Memory Schedulers,” in *MICRO*, 2004.
- [288] M. Xie, D. Tong, K. Huang, and X. Cheng, “Improving System Throughput and Fairness Simultaneously in Shared Memory CMP Systems via Dynamic Bank Partitioning,” in *HPCA*, 2014.
- [289] G. L. Yuan, A. Bakhoda, and T. M. Aamodt, “Complexity Effective Memory Access Scheduling for Many-Core Accelerator Architectures,” in *MICRO*, 2009.
- [290] H. Wang, R. Singh, M. J. Schulte, and N. S. Kim, “Memory Scheduling Towards High-Throughput Cooperative Heterogeneous Computing,” in *PACT*, 2014.
- [291] C. Kim, D. Burger, and S. W. Keckler, “An Adaptive, Non-Uniform Cache Structure for Wire-Delay Dominated On-Chip Caches,” in *ASPLOS*, 2002.
- [292] T. Moscibroda and O. Mutlu, “A Case for Bufferless Routing in On-Chip Networks,” in *ISCA*, 2009.
- [293] R. Das, O. Mutlu, T. Moscibroda, and C. R. Das, “Application-Aware Prioritization Mechanisms for On-Chip Networks,” in *MICRO*, 2009.
- [294] R. Das, O. Mutlu, T. Moscibroda, and C. R. Das, “Aergia: Exploiting Packet Latency Slack in On-Chip Networks,” in *ISCA*, 2010.
- [295] G. P. Nychis, C. Fallin, T. Moscibroda, O. Mutlu, and S. Seshan, “On-Chip Networks from a Networking Perspective: Congestion and Scalability in Many-Core Interconnects,” in *SIGCOMM*, 2012.
- [296] M. Besta, S. M. Hassan, S. Yalamanchili, R. Ausavarungnirun, O. Mutlu, and T. Hoefler, “Slim NoC: A Low-Diameter On-Chip Network Topology for High Energy Efficiency and Scalability,” in *ASPLOS*, 2018.
- [297] C. Fallin, C. Craik, and O. Mutlu, “CHIPPER: A Low-Complexity Bufferless Deflection Router,” in *HPCA*, 2011.
- [298] B. Grot, J. Hestness, S. W. Keckler, and O. Mutlu, “Kilo-NOC: A Heterogeneous Network-on-Chip Architecture for Scalability and Service Guarantees,” in *ISCA*, 2011.
- [299] L. Benini and G. De Micheli, “Networks on Chip: A New Paradigm for Systems on Chip Design,” in *DATE*, 2002.
- [300] D. Zhang, “ZSim++,” <https://github.com/dzhang50/zsim-plusplus>.
- [301] D. Zhang, X. Ma, M. Thomson, and D. Chioi, “Minnow: Lightweight Offload Engines for Worklist Management and Worklist-Directed Prefetching,” *ASPLOS*, 2018.
- [302] O. Mutlu, “Lecture Notes for Digital Design and Computer Architecture – Lecture 15a: Out-of-Order Execution,” <https://safari.ethz.ch/digitaltechnik/spring2020/lib/exe/fetch.php?media=onur-digitaldesign-2020-lecture15a-out-of-order-execution-beforelecture.pdf>, 2020.
- [303] R. Hadidi, L. Nai, H. Kim, and H. Kim, “CAIRO: A Compiler-Assisted Technique for Enabling Instruction-Level Offloading of Processing-in-Memory,” *TACO*, 2017.
- [304] H. Asghari-Moghaddam, A. Farmahini-Farahani, K. Morrow *et al.*, “Near-DRAM Acceleration with Single-ISA Heterogeneous Processing in Standard Memory Modules,” *IEEE Micro*, 2016.
- [305] M. A. Alves, M. Diener, P. C. Santos, and L. Carro, “Large Vector Extensions Inside the HMC,” in *DATE*, 2016.
- [306] E. Azarkhish, D. Rossi, I. Loi, and L. Benini, “Design and Evaluation of a Processing-in-Memory Architecture for the Smart Memory Cube,” in *ARC*, 2016.
- [307] H. Asghari-Moghaddam, Y. H. Son, J. H. Ahn, and N. S. Kim, “Chameleon: Versatile and Practical Near-DRAM Acceleration Architecture for Large Memory Systems,” in *MICRO*, 2016.
- [308] J. Picorel, D. Jevdjic, and B. Falsafi, “Near-Memory Address Translation,” in *PACT*, 2017.
- [309] Z. Liu, I. Calciu, M. Herlihy, and O. Mutlu, “Concurrent Data Structures for Near-Memory Computing,” in *SPAA*, 2017.
- [310] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning*, 2nd ed. Springer-Verlag, 2008.
- [311] E. Azarkhish, D. Rossi, I. Loi, and L. Benini, “High Performance AXI-4.0 Based Interconnect for Extensible Smart Memory Cubes,” in *DATE*, 2015.
- [312] R. Hadidi, B. Asgari, J. Young, B. A. Mudassar, K. Garg, T. Krishna, and H. Kim, “Performance Implications of NoCs on 3D-Stacked Memories: Insights from the Hybrid Memory Cube,” in *ISPASS*, 2018.
- [313] C. Min, J. Mao, H. Li, and Y. Chen, “NeuralHMC: An Efficient HMC-Based Accelerator for Deep Neural Networks,” in *ASP-DAC*, 2019.
- [314] G. Dai, T. Huang, Y. Chi, J. Zhao, G. Sun, Y. Liu, Y. Wang, Y. Xie, and H. Yang, “GraphH: A Processing-in-Memory Architecture for Large-Scale Graph Processing,” *TCAD*, 2018.
- [315] Y. S. Shao, B. Reagen, G.-Y. Wei, and D. Brooks, “Aladdin: A Pre-RTL, Power-Performance Accelerator Simulator Enabling Large Design Space Exploration of Customized Architectures,” in *ISCA*, 2014.
- [316] J. Jang, J. Heo, Y. Lee, J. Won, S. Kim, S. J. Jung, H. Jang, T. J. Ham, and J. W. Lee, “Charon: Specialized Near-Memory Processing Architecture for Clearing Dead Objects in Memory,” in *MICRO*, 2019.
- [317] C. Xie, S. L. Song, J. Wang, W. Zhang, and X. Fu, “Processing-in-Memory Enabled Graphics Processors for 3D Rendering,” in *HPCA*, 2017.
- [318] M. Lenjani, P. Gonzalez, E. Sadredini, S. Li, Y. Xie, A. Akel, S. Eilert, M. R. Stan, and K. Skadron, “Fulcrum: A Simplified Control and Access Mechanism Toward Flexible and Practical In-Situ Accelerators,” in *HPCA*, 2020.
- [319] C. D. Kersey, H. Kim, and S. Yalamanchili, “Lightweight SIMT Core Designs for Intelligent 3D Stacked DRAM,” in *MEMSYS*, 2017.
- [320] J. Huang, R. R. Puli, P. Majumder, S. Kim, R. Boyapati, K. H. Yum, and E. J. Kim, “Active-Routing: Compute on the Way for Near-Data Processing,” in *HPCA*, 2019.
- [321] S. M. Hassan, S. Yalamanchili, and S. Mukhopadhyay, “Near Data Processing: Impact and Optimization of 3D Memory System Architecture on the Uncore,” in *MEMSYS*, 2015.
- [322] Q. Guo, N. Alachiotis, B. Akin, F. Sadi, G. Xu, T. M. Low, L. Pileggi, J. C. Hoe, and F. Franchetti, “3D-Stacked Memory-Side Acceleration: Accelerator and System Design,” in *WoNDP*, 2014.
- [323] P. Liu, A. Hemani, K. Paul, C. Weis, M. Jung, and N. Wehn, “3D-Stacked Many-Core Architecture for Biological Sequence Analysis Problems,” *IJPP*, 2017.
- [324] Q. Zhu, B. Akin, H. E. Sumbul *et al.*, “A 3D-Stacked Logic-in-Memory Accelerator for Application-Specific Data Intensive Computing,” in *3DIC*, 2013.
- [325] S. Padmanabha, A. Lukefahr, R. Das, and S. Mahlke, “Mirage Cores: The Illusion of Many Out-of-Order Cores Using In-Order Hardware,” in *MICRO*, 2017.
- [326] C. Fallin, C. Wilkerson, and O. Mutlu, “The Heterogeneous Block Architecture,” in *ICCD*, 2014.
- [327] C. Villavieja, J. A. Joao, R. Miftakhutdinov, and Y. N. Patt, “Yoga: A Hybrid Dynamic VLIW/OoO Processor,” Univ. of Texas at Austin, High Performance Systems Group, Tech. Rep. TR-HPS-2014-001, 2014.
- [328] M. A. Suleman, M. Hashemi, C. Wilkerson, Y. N. Patt *et al.*, “Morphcore: An Energy-Efficient Microarchitecture for High Performance ILP and High Throughput TLP,” in *MICRO*, 2012.
- [329] E. Ipek, M. Kirman, N. Kirman, and J. F. Martinez, “Core Fusion: Accommodating Software Diversity in Chip Multiprocessors,” in *ISCA*, 2007.
- [330] M. A. Suleman, O. Mutlu, J. A. Joao, and Y. N. Patt, “Data Marshaling for Multi-Core Architectures,” in *ISCA*, 2010.
- [331] P. Petrica, A. M. Izraelevitz, D. H. Albonese, and C. A. Shoemaker, “Flicker: A Dynamically Adaptive Architecture for Power Limited Multicore Systems,” in *ISCA*, 2013.
- [332] C. Kim, S. Sethumadhavan, M. S. Govindan, N. Ranganathan, D. Gulati, D. Burger, and S. W. Keckler, “Composable Lightweight Processors,” in *MICRO*, 2007.
- [333] A. Lukefahr, S. Padmanabha, R. Das, F. M. Sleiman, R. Dreslinski, T. F. Wenisch, and S. Mahlke, “Composite Cores: Pushing Heterogeneity Into a Core,” in *MICRO*, 2012.
- [334] M. A. Suleman, O. Mutlu, M. K. Qureshi, and Y. N. Patt, “Accelerating Critical Section Execution with Asymmetric Multi-Core Architectures,” in *ASPLOS*, 2009.
- [335] J. A. Joao, M. A. Suleman, O. Mutlu, and Y. N. Patt, “Bottleneck Identification and Scheduling in Multithreaded Applications,” in *ASPLOS*, 2012.
- [336] S. Chaudhry, R. Cypher, M. Ekman, M. Karlsson, A. Landin, S. Yip, H. Zeffer, and M. Tremblay, “Simultaneous Speculative Threading: A Novel Pipeline Architecture Implemented in Sun’s ROCK Processor,” in *ISCA*, 2009.
- [337] J. A. Joao, M. A. Suleman, O. Mutlu, and Y. N. Patt, “Utility-Based Acceleration of Multithreaded Applications on Asymmetric CMPs,” in *ISCA*, 2013.
- [338] D. Tarjan, M. Boyer, and K. Skadron, “Federation: Repurposing Scalar Cores for Out-of-Order Instruction Issue,” in *DAC*, 2008.
- [339] M. Alipour, S. Kaxiras, D. Black-Schaffer, and R. Kumar, “Delay and Bypass: Ready and Criticality Aware Instruction Scheduling in Out-of-Order Processors,” in *HPCA*, 2020.
- [340] R. Kumar, M. Alipour, and D. Black-Schaffer, “Freeway: Maximizing MLP for Slice-Out-of-Order Execution,” in *HPCA*, 2019.

- [341] M. Alipour, R. Kumar, S. Kaxiras, and D. Black-Schaffer, "FIFOOrder Micro-Architecture: Ready-Aware Instruction Scheduling for OoO Processors," in *DATE*, 2019.
- [342] R. Kumar, K. I. Farkas, N. P. Jouppi, P. Ranganathan, and D. M. Tullsen, "Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction," in *MICRO*, 2003.
- [343] R. Kumar, D. M. Tullsen, P. Ranganathan, N. P. Jouppi, and K. I. Farkas, "Single-ISA Heterogeneous Multi-Core Architectures for Multithreaded Workload Performance," in *ISCA*, 2004.
- [344] H. Ahmed, P. C. Santos, J. P. Lima, R. F. Moura, M. A. Alves, A. C. Beck, and L. Carro, "A Compiler for Automatic Selection of Suitable Processing-in-Memory Instructions," in *DATE*, 2019.
- [345] S. Baskaran and J. Sampson, "Decentralized Offload-Based Execution on Memory-Centric Compute Cores," in *MEMSYS*, 2020.
- [346] J. Li, X. Wang, A. Tumeo, B. Williams, J. D. Leidel, and Y. Chen, "PIMS: A Lightweight Processing-in-Memory Accelerator for Stencil Computations," in *MEMSYS*, 2019.
- [347] G. Ayers, H. Litz, C. Kozyrakis, and P. Ranganathan, "Classifying Memory Access Patterns for Prefetching," in *ASPLOS*, 2020.
- [348] J. D. Collins, H. Wang, D. M. Tullsen, C. Hughes, Y.-F. Lee, D. Lavery, and J. P. Shen, "Speculative Precomputation: Long-Range Prefetching of Delinquent Loads," in *ISCA*, 2001.
- [349] J. González and A. González, "Speculative Execution via Address Prediction and Data Prefetching," in *ICS*, 1997.
- [350] T.-F. Chen and J.-L. Baer, "Effective Hardware-Based Data Prefetching for High-Performance Processors," *TC*, 1995.
- [351] R. Bera, A. V. Nori, O. Mutlu, and S. Subramoney, "DSPatch: Dual Spatial Pattern Prefetcher," in *MICRO*, 2019.
- [352] M. Bakhshalipour, P. Lotfi-Kamran, and H. Sarbazi-Azad, "Domino Temporal Data Prefetcher," in *HPCA*, 2018.
- [353] J. W. Fu, J. H. Patel, and B. L. Janssens, "Stride Directed Prefetching in Scalar Processors," in *MICRO*, 1992.
- [354] Y. Ishii, M. Inaba, and K. Hiraki, "Access Map Pattern Matching for Data Cache Prefetch," in *ISC*, 2009.
- [355] M. Hashemi, K. Swersky, J. A. Smith, G. Ayers, H. Litz, J. Chang, C. Kozyrakis, and P. Ranganathan, "Learning Memory Access Patterns," in *ICML*, 2018.
- [356] L. Orosa, R. Azevedo, and O. Mutlu, "AVPP: Address-First Value-Next Predictor with Value Prefetching for Improving the Efficiency of Load Value Prediction," *TACO*, 2018.
- [357] A. Jog, O. Kayiran, A. K. Mishra, M. T. Kandemir, O. Mutlu, R. Iyer, and C. R. Das, "Orchestrated Scheduling and Prefetching for GPGPUs," in *ISCA*, 2013.
- [358] C. J. Lee, O. Mutlu, V. Narasiman, and Y. N. Patt, "Prefetch-Aware Memory Controllers," *TC*, 2011.
- [359] T. M. Austin and G. S. Sohi, "Zero-Cycle Loads: Microarchitecture Support for Reducing Load Latency," in *MICRO*, 1995.
- [360] L. Ceze, K. Strauss, J. Tuck, J. Torrellas, and J. Renau, "CAVA: Using Checkpoint-Assisted Value Prediction to Hide L2 Misses," *TACO*, 2006.
- [361] D. Kado, J. Kim, P. Sharma, R. Panda, P. Gratz, and D. Jimenez, "B-Fetch: Branch Prediction Directed Prefetching for Chip-Multiprocessors," in *MICRO*, 2014.
- [362] N. Kirman, M. Kirman, M. Chaudhuri, and J. F. Martinez, "Checkpointed Early Load Retirement," in *HPCA*, 2005.
- [363] D. Joseph and D. Grunwald, "Prefetching Using Markov Predictors," in *ISCA*, 1997.
- [364] E. Ebrahimi, C. J. Lee, O. Mutlu, and Y. N. Patt, "Prefetch-Aware Shared Resource Management for Multi-Core Systems," in *ISCA*, 2011.
- [365] E. Ebrahimi, O. Mutlu, C. J. Lee, and Y. N. Patt, "Coordinated Control of Multiple Prefetchers in Multi-Core Systems," in *MICRO*, 2009.
- [366] C. J. Lee, V. Narasiman, O. Mutlu, and Y. N. Patt, "Improving Memory Bank-Level Parallelism in the Presence of Prefetching," in *MICRO*, 2009.
- [367] C. J. Lee, O. Mutlu, V. Narasiman, and Y. N. Patt, "Prefetch-Aware DRAM Controllers," in *MICRO*, 2008.
- [368] M. J. Charney and A. P. Reeves, "Generalized Correlation-Based Hardware Prefetching," Cornell Univ. Tech. Rep., 1995.
- [369] M. Charney, "Correlation-Based Hardware Prefetching," Ph.D. dissertation, Cornell University, 1995.
- [370] O. Mutlu, H. Kim, D. N. Armstrong, and Y. N. Patt, "Using the First-Level Caches as Filters to Reduce the Pollution Caused by Speculative Memory References," *IJPP*, 2005.
- [371] A. Yazdanbakhsh, G. Pekhimenko, B. Thwaites, H. Esmailzadeh, O. Mutlu, and T. C. Mowry, "RFVP: Rollback-Free Value Prediction with Safe-to-Approximate Loads," *TACO*, 2016.
- [372] R. J. Eickemeyer and S. Vassiliadis, "A Load-Instruction Unit for Pipelined Processors," *IBM JRD*, 1993.
- [373] F. A. Endo, A. Perais, and A. Seznez, "On the Interactions Between Value Prediction and Compiler Optimizations in the Context of EOLE," *TACO*, 2017.
- [374] M. H. Lipasti and J. P. Shen, "Exceeding the Dataflow Limit via Value Prediction," in *MICRO*, 1996.
- [375] M. H. Lipasti, C. B. Wilkerson, and J. P. Shen, "Value Locality and Load Value Prediction," in *ASPLOS*, 1996.
- [376] B. Calder, G. Reinman, and D. M. Tullsen, "Selective Value Prediction," in *ISCA*, 1999.
- [377] K. Wang and M. Franklin, "Highly Accurate Data Value Prediction Using Hybrid Predictors," in *MICRO*, 1997.
- [378] F. Gabbay and A. Mendelson, "Speculative Execution Based on Value Prediction," Technion - Israel Institute of Technology, Tech. Rep. 1080, 1996.
- [379] M. Burtcher and B. G. Zorn, "Exploring Last N Value Prediction," in *PACT*, 1999.
- [380] C.-Y. Fu, M. D. Jennings, S. Y. Larin, and T. M. Conte, "Value Speculation Scheduling for High Performance Processors," in *ASPLOS*, 1998.
- [381] B. Goeman, H. Vandierendonck, and K. De Bosschere, "Differential FCM: Increasing Value Prediction Accuracy by Improving Table Usage Efficiency," in *HPCA*, 2001.
- [382] T. Nakra, R. Gupta, and M. L. Soffa, "Global Context-Based Value Prediction," in *HPCA*, 1999.
- [383] T. Sato and I. Arita, "Low-Cost Value Predictors Using Frequent Value Locality," in *HPCA*, 2002.
- [384] Y. Sazeides and J. E. Smith, "The Predictability of Data Values," in *MICRO*, 1997.
- [385] N. Tuck and D. M. Tullsen, "Multithreaded Value Prediction," in *HPCA*, 2005.
- [386] D. M. Tullsen and J. S. Seng, "Storageless Value Prediction Using Prior Register Values," in *ISCA*, 1999.
- [387] E. S. Tune, D. M. Tullsen, and B. Calder, "Quantifying Instruction Criticality," in *PACT*, 2002.
- [388] G. Pekhimenko, V. Seshadri, O. Mutlu, M. A. Kozuch, P. B. Gibbons, and T. C. Mowry, "Base-Delta-Immediate Compression: Practical Data Compression for On-Chip Caches," in *PACT*, 2012.
- [389] J. Dussert, T. Piquet, and A. Seznez, "Zero-Content Augmented Caches," in *ICS*, 2009.
- [390] J. Yang, Y. Zhang, and R. Gupta, "Frequent Value Compression in Data Caches," in *MICRO*, 2000.
- [391] A. R. Alameldeen and D. A. Wood, "Adaptive Cache Compression for High-Performance Processors," in *ISCA*, 2004.
- [392] Y. Zhang, J. Yang, and R. Gupta, "Frequent Value Locality and Value-Centric Data Cache Design," in *ASPLOS*, 2000.
- [393] 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.
- [394] G. Pekhimenko, E. Bolotin, M. O'Connor, O. Mutlu, T. C. Mowry, and S. W. Keckler, "Toggle-Aware Compression for GPUs," *CAL*, 2015.
- [395] G. Pekhimenko, E. Bolotin, M. O'Connor, O. Mutlu, T. C. Mowry, and S. W. Keckler, "Energy-Efficient Data Compression for GPU Memory Systems," in *ASPLOS*, 2015.
- [396] 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.
- [397] G. Pekhimenko, V. Seshadri, Y. Kim, H. Xin, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "Linearly Compressed Pages: A Low-Complexity, Low-Latency Main Memory Compression Framework," in *MICRO*, 2013.
- [398] X. Chen, L. Yang, R. P. Dick, L. Shang, and H. Lekatsas, "C-Pack: A High-Performance Microprocessor Cache Compression Algorithm," *TVLSI*, 2009.
- [399] E. G. Hallnor and S. K. Reinhardt, "A Unified Compressed Memory Hierarchy," in *HPCA*, 2005.
- [400] D. W. Hammerstrom and E. S. Davidson, "Information Content of CPU Memory Referencing Behavior," in *ISCA*, 1977.
- [401] M. M. Islam and P. Stenstrom, "Zero-Value Caches: Cancelling Loads That Return Zero," in *PACT*, 2009.
- [402] A. Arelakis and P. Stenstrom, "SC2: A Statistical Compression Cache Scheme," in *ISCA*, 2014.
- [403] M. Ekman and P. Stenstrom, "A Robust Main-Memory Compression Scheme," in *ISCA*, 2005.
- [404] N. Vijaykumar, G. Pekhimenko, A. Jog, A. Bhowmick, R. Ausavarungrun, 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.
- [405] J. Gaur, A. R. Alameldeen, and S. Subramoney, "Base-Victim Compression: An Opportunistic Cache Compression Architecture," in *ISCA*, 2016.
- [406] J. S. Miguel, J. Albericio, A. Moshovos, and N. E. Jerger, "Doppelgänger: A Cache for Approximate Computing," in *ISCA*, 2015.
- [407] G. F. Oliveira, L. R. Gonçalves, M. Brandalero, A. C. S. Beck, and L. Carro, "Employing Classification-Based Algorithms for General-Purpose Approximate Computing," in *DAC*, 2018.
- [408] N. Vijaykumar, A. Jain, D. Majumdar, K. Hsieh, G. Pekhimenko, E. Ebrahimi, N. Hajinazar, P. B. Gibbons, and O. Mutlu, "A Case for Richer Cross-Layer Abstractions: Bridging the Semantic Gap with Expressive Memory," in *ISCA*, 2018.
- [409] A. Yazdanbakhsh, B. Thwaites, H. Esmailzadeh, G. Pekhimenko, O. Mutlu, and T. C. Mowry, "Mitigating the Memory Bottleneck with Approximate Load Value

- Prediction," *IEEE Design & Test*, 2016.
- [410] P.-A. Tsai, Y. L. Gan, and D. Sanchez, "Rethinking the Memory Hierarchy for Modern Languages," in *MICRO*, 2018.
- [411] M. A. Z. Alves, P. C. Santos, M. Diener, and L. Carro, "Opportunities and Challenges of Performing Vector Operations Inside the DRAM," in *MEMSYS*, 2015.
- [412] Z. Sura, A. Jacob, T. Chen, B. Rosenburg, O. Sallenave, C. Bertolli, S. Antao, J. Brunheroto, Y. Park, K. O'Brien *et al.*, "Data Access Optimization in a Processing-in-Memory System," in *CF*, 2015.
- [413] D. P. Zhang, N. Jayasena, A. Lyashevsky *et al.*, "A New Perspective on Processing-in-Memory Architecture Design," in *MSPC*, 2013.
- [414] N. S. Kim, D. Chen, J. Xiong, and W. H. Wen-mei, "Heterogeneous Computing Meets Near-Memory Acceleration and High-Level Synthesis in the Post-Moore Era," *IEEE Micro*, 2017.
- [415] P.-A. Tsai and D. Sanchez, "Compress Objects, Not Cache Lines: An Object-Based Compressed Memory Hierarchy," in *ASPLOS*, 2019.
- [416] N. Vijaykumar, E. Ebrahimi, K. Hsieh, P. B. Gibbons, and O. Mutlu, "The Locality Descriptor: A Holistic Cross-Layer Abstraction to Express Data Locality in GPUs," in *ISCA*, 2018.
- [417] O. Mutlu, "Intelligent Architectures for Intelligent Computing Systems," in *DATE*, 2020.
- [418] G. Narancic, P. Judd, D. Wu, I. Atta, M. Elnacouzi, J. Zebchuk, J. Albericio, N. E. Jerger, A. Moshovos, K. Kutulakos *et al.*, "Evaluating the Memory System Behavior of Smartphone Workloads," in *SAMOS*, 2014.
- [419] K. Yan, X. Zhang, and X. Fu, "Characterizing, Modeling, and Improving the QoE of Mobile Devices with Low Battery Level," in *MICRO*, 2015.
- [420] D. Shingari, A. Arunkumar, and C.-J. Wu, "Characterization and Throttling-Based Mitigation of Memory Interference for Heterogeneous Smartphones," in *IISWC*, 2015.
- [421] C. Kozyrakis, A. Kansal, S. Sankar, and K. Vaid, "Server Engineering Insights for Large-Scale Online Services," *IEEE Micro*, 2010.
- [422] A. Yasin, Y. Ben-Asher, and A. Mendelson, "Deep-Dive Analysis of the Data Analytics Workload in CloudSuite," in *IISWC*, 2014.
- [423] U. Gupta, X. Wang, M. Naumov, C.-J. Wu, B. Reagen, D. Brooks, B. Cottel, K. Hazelwood, B. Jia, H.-H. S. Lee *et al.*, "The Architectural Implications of Facebook's DNN-Based Personalized Recommendation," in *HPCA*, 2020.
- [424] Y. Gan and C. Delimitrou, "The Architectural Implications of Cloud Microservices," *CAL*, 2018.
- [425] G. Ayers, J. H. Ahn, C. Kozyrakis, and P. Ranganathan, "Memory Hierarchy for Web Search," in *HPCA*, 2018.
- [426] R. C. Murphy, P. M. Kogge, and A. Rodrigues, "The Characterization of Data Intensive Memory Workloads on Distributed PIM Systems," in *Intelligent Memory Systems*. Springer-Verlag, 2001.
- [427] A. Limaye and T. Adegbiya, "A Workload Characterization of the SPEC CPU2017 Benchmark Suite," in *ISPASS*, 2018.
- [428] S. Corda, G. Singh, A. J. Awan, R. Jordans, and H. Corporaal, "Memory and Parallelism Analysis Using a Platform-Independent Approach," in *SCOPES*, 2019.
- [429] D. Patterson, T. Anderson, N. Cardwell *et al.*, "A Case for Intelligent RAM," *IEEE Micro*, 1997.
- [430] H. Li and R. Durbin, "Fast and Accurate Short Read Alignment with Burrows-Wheeler Transform," *Bioinformatics*, 2009.
- [431] J. Beranek, "Hardware Effects," <https://github.com/Kobzol/hardware-effects>.
- [432] L.-N. Pouchet, "PolyBench: The Polyhedral Benchmark Suite," <https://www.cs.colostate.edu/~pouchet/software/polybench/>.
- [433] M. Chen, S. Mao, and Y. Liu, "Big Data: A Survey," *Mobile Networks and Applications*, 2014.
- [434] H. Mo, "Mesoscale Modeling and Direct Simulation of Explosively Dispersed Granular Materials," Ph.D. dissertation, University of Waterloo, 2019.
- [435] S. Thomas, C. Gohkale, E. Tanuwidjaja, T. Chong, D. Lau, S. Garcia, and M. B. Taylor, "CortexSuite: A Synthetic Brain Benchmark Suite," in *IISWC*, 2014.
- [436] M. J. Chaisson and G. Tesler, "Mapping Single Molecule Sequencing Reads Using Basic Local Alignment with Successive Refinement (BLASR): Application and Theory," *Bioinformatics*, 2012.
- [437] H. Li, "Aligning Sequence Reads, Clone Sequences and Assembly Contigs with BWA-MEM," arXiv:1303.3997 [q-bio.GN], 2013.
- [438] WebM, "VP8/VP9 Codec SDK," <https://github.com/ittiamvpx/libvpx>.
- [439] C. Balkesen, G. Alonso, J. Teubner, and M. T. Özsu, "Multi-Core, Main-Memory Joins: Sort vs. Hash Revisited," *Vldb*, 2013.
- [440] Standard Performance Evaluation Corp., "SPEC CPU2006 Benchmarks," <http://www.spec.org/cpu2006/>.
- [441] S. Ghosh, O. Vinyals, B. Strope, S. Roy, T. Dean, and L. Heck, "Contextual LSTM (CLSTM) Models for Large Scale NLP Tasks," arXiv:1602.06291 [cs.CL], 2016.
- [442] A. Buluç and J. R. Gilbert, "The Combinatorial BLAS: Design, Implementation, and Applications," *IJHPCA*, 2011.
- [443] M. Karasikov, H. Mustafa, D. Danciu, M. Zimmermann, C. Barber, G. Ratsch, and A. Kahles, "Metagraph: Indexing and Analysing Nucleotide Archives at Petabase-Scale," bioRxiv 2020.10.01.322164, 2020.
- [444] Intel Corp., "Accelerate Fast Math with Intel oneAPI Math Kernel Library," <https://software.intel.com/content/www/us/en/develop/tools/oneapi/components/onemkl.html>.
- [445] Standard Performance Evaluation Corp., "SPEC CPU2017 Benchmarks," <http://www.spec.org/cpu2017/>.
- [446] R. O. Nascimento and P. R. Maciel, "DBT-5: An Open-Source TPC-E Implementation for Global Performance Measurement of Computer Systems," *Computing and Informatics*, 2010.
- [447] X. Yu, G. Bezerra, A. Pavlo, S. Devadas, and M. Stonebraker, "Staring Into the Abyss: An Evaluation of Concurrency Control with One Thousand Cores," *Vldb*, 2014.
- [448] M. Naumov, D. Mudigere, H. M. Shi, J. Huang, N. Sundaraman, J. Park, X. Wang, U. Gupta, C. Wu, A. G. Azzolini, D. Dzulgakov, A. Mallevech, I. Cherniavskii, Y. Lu, R. Krishnamoorthi, A. Yu, V. Kondratenko, S. Pereira, X. Chen, W. Chen, V. Rao, B. Jia, L. Xiong, and M. Smelyanskiy, "Deep Learning Recommendation Model for Personalization and Recommendation Systems," arXiv:1906.00091 [cs.IR], 2019.
- [449] R. Palomar, J. Gómez-Luna, F. A. Cheikh, J. Olivares-Bueno, and O. J. Elle, "High-Performance Computation of Bézier Surfaces on Parallel and Heterogeneous Platforms," *IJPP*, 2018.
- [450] H. Kasture and D. Sanchez, "Tailbench: A Benchmark Suite and Evaluation Methodology for Latency-Critical Applications," in *IISWC*, 2016.
- [451] S. Nalli, S. Haria, M. D. Hill, M. M. Swift, H. Volos, and K. Keeton, "An Analysis of Persistent Memory Use with WHISPER," in *ASPLOS*, 2017.
- [452] K. Kara, D. Alistarh, G. Alonso, O. Mutlu, and C. Zhang, "FPGA-Accelerated Dense Linear Machine Learning: A Precision-Convergence Trade-Off," in *FCCM*, 2017.

APPENDIX

A Application Functions in the DAMOV Benchmark Suite

We present the list of application functions in each one of the six classes of data movement bottlenecks we identify using our new methodology.

Our benchmark suite is composed of 144 different application functions, collected from 74 different applications. These applications belong to a different set of previously published and widely used benchmark suites. In total, we collect applications from 16 benchmark suites, including: BWA [430], Chai [199], Darknet [215], GASE [208], Hardware Effects [431], Hashjoin [209], HPCC [206], HPCG [207], Ligra [212], PARSEC [202], Parboil [201], PolyBench [432], Phoenix [213], Rodinia [203], SPLASH-2 [205], STREAM [120]. The 144 application functions that are part of DAMOV are listed across six tables, each designating one of the six classes we identify in Section 3.3:

- Table 2 lists application functions in Class 1a, i.e., that are DRAM bandwidth-bound (characterized in Section 3.3.1);
- Table 3 lists application functions in Class 1b, i.e., that are DRAM latency-bound (characterized in Section 3.3.2);
- Table 4 lists application functions in Class 1c, i.e., that are bottlenecked by the available L1/L2 cache capacity (characterized in Section 3.3.3);
- Table 5 lists application functions in Class 2a, i.e., that are bottlenecked by L3 cache contention (characterized in Section 3.3.4);
- Table 6 lists application functions in Class 2b, i.e., that are bottlenecked by L1 cache size (characterized in Section 3.3.5);
- Table 7 lists application functions in Class 2c, i.e., that are compute-bound (characterized in Section 3.3.6).

In each table we list the benchmark suite, the application name, and the function name. We also list the input size/problem size we use to evaluate each application function.

Table 2: List of application functions in Class 1a.

Class	Suite	Benchmark	Function	Input Set/ Problem Size	Representative Function?
1a	Chai [199]	Transpose	cpu	-m 1024 -n 524288	No
1a	Chai [199]	Vector Pack	run_cpu_threads	-m 268435456 -n 16777216	No
1a	Chai [199]	Vector Unpack	run_cpu_threads	-m 268435456 -n 16777216	No
1a	Darknet [215]	Yolo	gemm	ref	Yes
1a	Hardware Effects [431]	Bandwidth Saturation - Non Temporal	main	ref	No
1a	Hardware Effects [431]	Bandwidth Saturation - Temporal	main	ref	No
1a	Hashjoin [209]	NPO	knuth	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	NPO	ProbeHashTable	-r 12800000 -s 12000000 -x 12345 -y 54321	Yes
1a	Hashjoin [209]	PRH	knuth	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	PRH	lock	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	PRHO	knuth	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	PRHO	radix	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	PRO	knuth	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	PRO	parallel	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	PRO	radix	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Hashjoin [209]	RJ	knuth	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1a	Ligra [212]	Betweenness Centrality	edgeMapSparse	rMat	No
1a	Ligra [212]	Breadth-First Search	edgeMapSparse	rMat	No
1a	Ligra [212]	Connected Components	compute	rMat	No
1a	Ligra [212]	Connected Components	compute	USA	No
1a	Ligra [212]	Connected Components	edgeMapDense	USA	No
1a	Ligra [212]	Connected Components	edgeMapSparse	USA	Yes
1a	Ligra [212]	K-Core Decomposition	compute	rMat	No
1a	Ligra [212]	K-Core Decomposition	compute	USA	No
1a	Ligra [212]	K-Core Decomposition	edgeMapDense	USA	No
1a	Ligra [212]	K-Core Decomposition	edgeMapSparse	rMat	No
1a	Ligra [212]	Maximal Independent Set	compute	rMat	No
1a	Ligra [212]	Maximal Independent Set	compute	USA	No
1a	Ligra [212]	Maximal Independent Set	edgeMapDense	USA	No
1a	Ligra [212]	Maximal Independent Set	edgeMapSparse	rMat	No
1a	Ligra [212]	Maximal Independent Set	edgeMapSparse	USA	No
1a	Ligra [212]	PageRank	compute	rMat	No
1a	Ligra [212]	PageRank	compute	USA	No
1a	Ligra [212]	PageRank	edgeMapDense	USA	Yes
1a	Ligra [212]	Radii	compute	rMat	No
1a	Ligra [212]	Radii	compute	USA	No
1a	Ligra [212]	Radii	edgeMapSparse	USA	No
1a	Ligra [212]	Triangle Count	edgeMapDense	rMat	Yes
1a	SPLASH-2 [205]	Oceancp	relax	simlarge	No
1a	SPLASH-2 [205]	Oceanncp	relax	simlarge	No
1a	STREAM [120]	Add	Add	50000000	Yes
1a	STREAM [120]	Copy	Copy	50000000	Yes
1a	STREAM [120]	Scale	Scale	50000000	Yes
1a	STREAM [120]	Triad	Triad	50000000	Yes

Table 3: List of application functions in Class 1b.

Class	Suite	Benchmark	Function	Input Set/ Problem Size	Representative Function?
1b	Chai [199]	Canny Edge Detection	gaussian	ref	No
1b	Chai [199]	Canny Edge Detection	supression	ref	No
1b	Chai [199]	Histogram - input partition	run_cpu_threads	ref	Yes
1b	Chai [199]	Select	run_cpu_threads	-n 67108864	No
1b	GASE [208]	FastMap	2occ4	Wg2	No
1b	GASE [208]	FastMap	occ4	Wg2	No
1b	Hashjoin [209]	PRH	HistogramJoin	-r 12800000 -s 12000000 -x 12345 -y 54321	Yes
1b	Phoenix [213]	Linear Regression	linear_regression_map	key_file_500MB	No
1b	Phoenix [213]	PCA	main	ref	No
1b	Phoenix [213]	String Match	string_match_map	key_file_500MB	Yes
1b	PolyBench [432]	linear-algebra	lu	LARGE_DATASET	Yes
1b	Rodinia [203]	Kmeans	euclidDist	819200.txt	No
1b	Rodinia [203]	Kmeans	find	819200.txt	No
1b	Rodinia [203]	Kmeans	main	819200.txt	No
1b	Rodinia [203]	Streamcluster	pengain	ref	No
1b	SPLASH-2 [205]	Oceancp	slave2	simlarge	Yes

Table 4: List of application functions in Class 1c.

Class	Suite	Benchmark	Function	Input Set/ Problem Size	Representative Function?
1c	BWA [430]	Align	bwa_aln_core	Wg1	No
1c	Chai [199]	Breadth-First Search	comp	USA-road-d	No
1c	Chai [199]	Breadth-First Search	fetch	USA-road-d	No
1c	Chai [199]	Breadth-First Search	load	USA-road-d	No
1c	Chai [199]	Breadth-First Search	run_cpu_threads	USA-road-d	No
1c	Chai [199]	Canny Edge Detection	hystresis	ref	No
1c	Chai [199]	Canny Edge Detection	sobel	ref	No
1c	Chai [199]	Histogram - output partition	run_cpu_threads	ref	No
1c	Chai [199]	Padding	run_cpu_threads	-m 10000 -n 9999	Yes
1c	Chai [199]	Select	fetch	-n 67108864	No
1c	Chai [199]	Stream Compaction	run_cpu_threads	ref	No
1c	Darknet [215]	Resnet	gemm	ref	Yes
1c	Hashjoin [209]	NPO	lock	-r 12800000 -s 12000000 -x 12345 -y 54321	No
1c	Ligra [212]	BFS-Connected Components	edgeMapSparse	rMat	No
1c	Ligra [212]	Triangle Count	compute	rMat	No
1c	Ligra [212]	Triangle Count	compute	USA	No
1c	Ligra [212]	Triangle Count	edgeMapDense	USA	No
1c	PARSEC [202]	Blackscholes	BlkSchlsEqEuroNoDiv	simlarge	No
1c	PARSEC [202]	Fluidaminate	ProcessCollision2MT	simlarge	Yes
1c	PARSEC [202]	Streamcluster	DistL2Float	simlarge	No
1c	Rodinia [203]	Myocyte	find	1000000	No
1c	Rodinia [203]	Myocyte	master	1000000	No

Table 5: List of application functions in Class 2a.

Class	Suite	Benchmark	Function	Input Set/ Problem Size	Representative Function?
2a	HPCC [206]	RandomAccess	main	ref	No
2a	HPCC [206]	RandomAccess	update	ref	No
2a	Ligra [212]	Betweenness Centrality	Compute	rMat	No
2a	Ligra [212]	Betweenness Centrality	Compute	USA	No
2a	Ligra [212]	Betweenness Centrality	edgeMapDense	rMat	No
2a	Ligra [212]	Betweenness Centrality	*edgeMapSparse	USA	Yes
2a	Ligra [212]	BFS-Connected Components	Compute	rMat	No
2a	Ligra [212]	BFS-Connected Components	Compute	USA	No
2a	Ligra [212]	BFS-Connected Components	edgeMapSparse	USA	Yes
2a	Ligra [212]	Breadth-First Search	compute	rMat	No
2a	Ligra [212]	Breadth-First Search	compute	USA	No
2a	Ligra [212]	Breadth-First Search	edgeMapDense	rMat	No
2a	Ligra [212]	Breadth-First Search	edgeMapSparse	USA	Yes
2a	Ligra [212]	Connected Components	edgeMapDense	rMat	No
2a	Ligra [212]	Maximal Independent Set	edgeMapDense	rMat	No
2a	Ligra [212]	PageRank	edgeMapDense(Rmat)	rMat	No
2a	Phoenix [213]	WordCount	main	word_100MB	No
2a	PolyBench [432]	linear-algebra	gramschmidt	LARGE_DATASET	Yes
2a	Rodinia [203]	CFD Solver	main	fvcorr.domn.193K	No
2a	SPLASH-2 [205]	FFT2	Reverse	simlarge	Yes
2a	SPLASH-2 [205]	FFT2	Transpose	simlarge	Yes
2a	SPLASH-2 [205]	Oceancp	jacobcalc	simlarge	No
2a	SPLASH-2 [205]	Oceancp	laplaccalc	simlarge	No
2a	SPLASH-2 [205]	Oceanncp	jacobcalc	simlarge	Yes
2a	SPLASH-2 [205]	Oceanncp	laplaccalc	simlarge	Yes
2a	SPLASH-2 [205]	Oceanncp	slave2	simlarge	No

Table 6: List of application functions in Class 2b.

Class	Suite	Benchmark	Function	Input Set/ Problem Size	Representative Function?
2b	Chai [199]	Bezier Surface	main_thread	ref	Yes
2b	Hardware Effects [431]	False Sharing - Isolated	main	ref	No
2b	PolyBench [432]	convolution	convolution-2d	LARGE_DATASET	No
2b	PolyBench [432]	linear-algebra	gemver	LARGE_DATASET	Yes
2b	SPLASH-2 [205]	Lucb	Bmod	simlarge	Yes
2b	SPLASH-2 [205]	Radix	slave2	simlarge	Yes

Table 7: List of application functions in Class 2c.

Class	Suite	Benchmark	Function	Input Set/ Problem Size	Representative Function?
2c	BWA [430]	Align	bwa_aln_core	Wg2	No
2c	Chai [199]	Transpose	run_cpu_threads	-m 1024 -n 524288	No
2c	Darknet [215]	Alexnet	gemm	ref	No
2c	Darknet [215]	vgg16	gemm	ref	No
2c	Hardware Effects [431]	False Sharing - Shared	main	ref	No
2c	HPCG [207]	HPCG	ComputePrologation	ref	Yes
2c	HPCG [207]	HPCG	ComputeRestriction	ref	Yes
2c	HPCG [207]	HPCG	ComputeSPMV	ref	Yes
2c	HPCG [207]	HPCG	ComputeSYMGS	ref	Yes
2c	Ligra [212]	K-Core Decomposition	edgeMapDense	rMat	No
2c	Ligra [212]	Radii	edgeMapSparse	rMat	No
2c	Parboil [201]	Breadth-First Search	BFS_CPU	ref	No
2c	Parboil [201]	MRI-Gridding	CPU_kernels	ref	No
2c	Parboil [201]	Stencil	cpu_stencil	ref	No
2c	Parboil [201]	Two Point Angular Correlation Function	doCompute	ref	No
2c	PARSEC [202]	Bodytrack	FilterRow	ref	No
2c	PARSEC [202]	Ferret	DistL2Float	ref	Yes
2c	Phoenix [213]	Kmeans	main	ref	No
2c	PolyBench [432]	linear-algebra	3mm	LARGE_DATASET	Yes
2c	PolyBench [432]	linear-algebra	doitgen	LARGE_DATASET	Yes
2c	PolyBench [432]	linear-algebra	gemm	LARGE_DATASET	Yes
2c	PolyBench [432]	linear-algebra	symm	LARGE_DATASET	Yes
2c	PolyBench [432]	stencil	fdtd-apml	LARGE_DATASET	Yes
2c	Rodinia [203]	Back Propagation	adjustweights	134217728	No
2c	Rodinia [203]	Back Propagation	layerforward	134217728	No
2c	Rodinia [203]	Breadth-First Search	main	graph1M_6	Yes
2c	Rodinia [203]	Needleman-Wunsch	main	32768	Yes
2c	Rodinia [203]a	Srad	FIN	ref	No
2c	SPLASH-2 [205]	Barnes	computeForces	simlarge	No
2c	SPLASH-2 [205]	Barnes	gravsub	simlarge	No

B Representative Application Functions

Table 8: 44 representative application functions studied in detail in this work.*

Suite	Benchmark	Function	Short Name	Class	%
Chai [199]	Bezier Surface	Bezier	CHABsBez	2b	100
	Histogram	Histogram	CHAHsti	1b	100
	Padding	Padding	CHAOpad	1c	75.1
Darknet [215]	Resnet 152	gemm_nn	DRKRes	1c	95.2
	Yolo	gemm_nn	DRKYolo	1a	97.1
Hashjoin [209]	NPO	ProbeHashTable	HSJNPO	1a	47.8
	PRH	HistogramJoin	HSJPRH	1b	53.1
HPCG [207]	HPCG	ComputeProlongation	HPGProl	2c	34.3
	HPCG	ComputeRestriction	HPGRes	2c	42.1
	HPCG	ComputeSPMV	HPGSpm	2c	30.5
	HPCG	ComputeSYMGS	HPGSyms	2c	63.6
Ligra [212]	Betweenness Centrality	EdgeMapSparse (USA [217])	LIGBcEms	2a	78.9
	Breadth-First Search	EdgeMapSparse (USA)	LIGBfsEms	2a	67.0
	BFS-Connected Components	EdgeMapSparse (USA)	LIGBfscEms	2a	68.3
	Connected Components	EdgeMapSparse (USA)	LIGCompEms	1a	25.6
	PageRank	EdgeMapDense (USA [217])	LIGPrkEmd	1a	57.2
	Radii	EdgeMapSparse (USA)	LIGRadiEms	1a	67.0
	Triangle	EdgeMapDense (Rmat)	LIGTriEmd	1a	26.7
PARSEC [202]	Ferret	DistL2Float	PRSFerr	2c	18.6
	Fluidaminate	ProcessCollision2MT	PRSFlu	1c	23.9
Phoenix [213]	Linear Regression	linear_regression_map	PHELinReg	1b	76.2
	String Matching	string_match_map	PHEStrMat	1b	38.3
PolyBench [432]	Linear Algebra	3 Matrix Multiplications	PLY3mm	2c	100.0
	Linear Algebra	Multi-resolution analysis kernel	PLYDoitgen	2c	98.3
	Linear Algebra	Matrix-multiply $C = \alpha \cdot A \cdot B + \beta \cdot C$	PLYgemm	2c	99.7
	Linear Algebra	Vector Mult. and Matrix Addition	PLYgemver	2b	44.4
	Linear Algebra	Gram-Schmidt decomposition	PLYGramSch	2a	100.0
	Linear Algebra	LU decomposition	PLYalu	1b	100.0
	Linear Algebra	Symmetric matrix-multiply	PLYSymm	2c	99.9
	Stencil	2D Convolution	PLYcon2d	2b	100.0
	Stencil	2-D Finite Different Time Domain	PLYdtd	2c	39.8
Rodinia [433]	BFS	BFSGraph	RODBfs	2c	100.0
	Needleman-Wunsch	runTest	RODNw	2c	84.9
SPLASH-2 [205]	FFT	Reverse	SPLFftRev	2a	12.7
	FFT	Transpose	SPLFftTra	2a	8.0
	Lucb	Bmod	SPLLucb	2b	77.6
	Oceanncp	jacobcalc	SPLOcnpJac	2a	30.7
	Oceanncp	laplaccalc	SPLOcnpLap	2a	23.4
	Oceancp	slave2	SPLOcpSlave	1b	24.4
	Radix	slave_sort	SPLRad	2b	41.1
STREAM [120]	Add	Add	STRAdd	1a	98.4
	Copy	Copy	STRCpy	1a	98.3
	Scale	Scale	STRSca	1a	97.5
	Triad	Triad	STRTriad	1a	99.1

* Short names are encoded as XXXYyyZzz, where XXX is the source application suite, Yyy is the application name, and Zzz is the function (if more than one per benchmark). For graph processing applications from Ligra, we test two different input graphs, so we append the graph name to the short benchmark name as well. The % column indicates the percentage of clock cycles that the function consumes as a fraction of the execution time of the entire benchmark.

C Complete List of Evaluated Applications

Table 9: List of the evaluated 345 applications.

Benchmark Suite	Application	Benchmark Suite	Application	Benchmark Suite	Application
ArtraCFD [434]	ArtraCFD	HPCG [207]	Global Dot Product	SD-VBS - Vision [435]	disparity
blastr [436]	Long read aligner		Multigrid preconditione		localization
BWA [437]	aln		Sparse Matrix Vector Multiplication (SpMV)		mser
	fastmap		Symmetric Gauss-Seidel smoother (SymGS)		multi_ncut
Chai [199]	BFS	IMPICA Workloads [55]	Vector Update	sort-merge-joins [439]	pca
	BS		btree		sift
	CEDD		hashtable		stitch
	HSTI		llubenchmark		svm
	HSTO	libvpx [438]	VP8		texture_synthesis
	OOPPAD		VP9		tracking
	OOPTRNS		BC		m-pass
	SC		BellmanFord		m-way
	SELECT	Ligra [212]	BFS	SPEC CPU2006 [440]	400.perlbench
	TRNS		BFS-Bitvector		401.bzip2
clstm [441]	VPACK		BFS-CC		403.gcc
	VUPACK		CF		410.bwaves
CombBLAS [442]	clstm		Components		416.gamess
	BetwCent		KCore		429.mcf
	BipartiteMatchings		MIS		433.milc
	CC		PageRank		434.zeusmp
	DirOptBFS		PageRankDelta		435.gromacs
	FilteredBFS		Radii		436.cactusADM
	FilteredMIS		Triangle		437.leslie3d
	MCL3D	Metagraph [443]	annotate		444.namd
	Ordering/RCM		classify		445.gobmk
	TopDownBFS		ASUM		447.deall
CORAL [200]	AMG2013	MKL [444]	AXPY		450.soplex
	CAM-SE		DOT		453.povray
	Graph500		GEMM		454.calculix
	HACC		GEMV		456.hammer
	Hash	Parboil [201]	mri-q		458.sjeng
	homme1_3_6		BFS		459.GemsFDTD
	Integer Sort		cutcp		462.libquantum
	KMI		histo		464.h264ref
	LSMS		lbm		465.tonto
	LULESH		mri-gridding		470.lbm
	MCB		sad		471.omnetpp
	miniFE		sgemm		473.astar
	Nekbone		spmv		481.wrf
	QBOX	PARSEC [202]	stencil		482.sphinx3
Darknet [215]	SNAP		tpacf		483.xalancbmk
	SPECint2006"peak"		blacksholes		500.perlbench_r
	UMT2013		bodytrack		502.gcc_r
	AlexNet		canneal		503.bwaves_r
	Darknet19		dedup		505.mcf_r
	Darknet53		facesim		507.cactuBSSN_r
	Densenet 201		ferret		508.namd_r
	Extraction		fluidanimate		510.parest_r
	Resnet 101		fraqmine		511.povray_r
	Resnet 152	Phoenix [213]	raytrace		519.lbm_r
DBx1000 [447]	Resnet 18		streamcluster		520.omnetpp_r
	Resnet 34		swaptions		521.wrf_r
	Resnet 50		vips		523.xalancbmk_r
	ResNeXt 101		x264		525.x264_r
	ResNext 152		histogram		526.blender_r
	ResNeXt50		kmeans		527.cam4_r
	VGG-16		linear-regression		531.deepsjeng_r
	Yolo		matrix multiply		538.imagick_r
	TPC-E	PolyBench [432]	pca		541.leela_r
	TPCC DL_DETECT		string_match		544.nab_r
DLRM [448]	TPCC HEKATON		word_count		548.exchange2_r
	TPCC NO_WAIT		2mm		549.fotonik3d_r
	TPCC SILO		3mm		554.roms_r
	TPCC TICTOC		atax		557.xz_r
	YCSB DL_DETECT		bicg		600.perlbench_s
	YCSB HEKATON		cholesky		602.gcc_s
	YCSB NO_WAIT		convolution-2d		603.bwaves_s
	YCSB SILO		convolution-3d		605.mcf_s
	YCSB TICTOC		correlation		607.cactuBSSN_s
	RM1-large [133]		covariance		619.lbm_s
GASE [208]	RM1-small [133]	GraphMat [211]	doitgen	SPEC CPU2017 [445]	620.omnetpp_s
	RM2-large [133]		durbin		621.wrf_s
GraphMat [211]	RM2-small [133]		fdtd-apm		623.xalancbmk_s
	FastMap		gemm		625.x264_s
	gale_aln		gemver		627.cam4_s
	BFS		gramschmidt		628.pop2_s
	DeltaStepping		gramschmidt		631.deepsjeng_s
	Incremental PageRank		lu		638.imagick_s
	LDA		lu		641.leela_s
	PageRank		mvt		644.nab_s
	SDG		symm		648.exchange2_s
	SSSP		syr2k		649.fotonik3d_s
	Topological Sort		syrk		654.roms_s
	Triangle Counting		trmm		657.xz_s

Benchmark Suite	Application	Benchmark Suite	Application	Benchmark Suite	Application
Hardware Effects [431]	4k aliasing	resectionvolume [449]	resectionvolume	SPLASH-2 [205]	barnes
	bandwidth saturation non-temporal		b+tree		cholesky
	bandwidth saturation temporal		backprop		fft
	branch misprediction sort		bfs		fmm
	branch misprediction unsort		cfld		lu_cb
	branch target misprediction		heartwall		lu_ncb
	cache conflicts		hotspot		ocean_cp
	cache/memory hierarchy bandwidth		hotspot3D		ocean_ncp
	data dependencies		kmeans		radiosity
	denormal floating point numbers		lavaMD		radix
	denormal floating point numbers flush		leukocyte		raytrace
	DRAM refresh interval		lud		volrend
	false sharing		mummergepu		water_nsquared
	hardware prefetching		myocyte		water_spatial
	hardware prefetching shuffle		nn	Tailbench [450]	img-dnn
	hardware store elimination		nw		masstree
	memory-bound program		particlefilter		moses
	misaligned accesses		pathfinder		shore
	non-temporal stores		sradi		silo
	software prefetching		streamcluster		specjbb
	store buffer capacity		lda		sphinx
	write combining		libl		xapian
	NPO	SD-VBS- Cortex [435]	me		ctree
	PRH		pca	WHISPER [451]	echo
Hashjoin [209]	PRHO		rbm		exim
	PRO		sphinx		hashmap
	RJ		srr		memcached
	RandomAccesses		svd		nfs
HPCC [206]					redis
					sql
					tpcc
					vacation
					ycsb
					SGD
					STREAM



GERALDO F. OLIVEIRA received a B.S. degree in computer science from the Federal University of Viçosa, Viçosa, Brazil, in 2015, and an M.S. degree in computer science from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil, in 2017. Since 2018, he has been working toward a Ph.D. degree with Onur Mutlu at ETH Zürich, Zürich, Switzerland. His current research interests include

system support for processing-in-memory and processing-using-memory architectures, data-centric accelerators for emerging applications, approximate computing, and emerging memory systems for consumer devices. He has several publications on these topics.



JUAN GÓMEZ-LUNA is a senior researcher and lecturer at SAFARI Research Group @ ETH Zürich. He received the BS and MS degrees in Telecommunication Engineering from the University of Sevilla, Spain, in 2001, and the PhD degree in Computer Science from the University of Córdoba, Spain, in 2012. Between 2005 and 2017, he was a

faculty member of the University of Córdoba. His research interests focus on processing-in-memory, memory systems, heterogeneous computing, and hardware and software acceleration of medical imaging and bioinformatics. He is the lead author of PRIM (<https://github.com/CMU-SAFARI/prim-benchmarks>), the first publicly-available benchmark suite for a real-world processing-in-memory architecture, and Chai (<https://github.com/chai-benchmarks/chai>), a benchmark suite for heterogeneous systems with CPU/GPU/FPGA.



LOIS OROSA is a senior researcher in the SAFARI research group at ETH Zurich. His current research interest are in computer architecture, hardware security, memory systems, and machine learning accelerators. He obtained his PhD from the University of Santiago de Compostela, and he was a PostDoc in the Institute of Computing at University of Campinas. He was a visiting scholar at University of Illinois at Urbana-Champaign and Universidade NOVA de Lisboa, and he acquired industrial experience at several companies.



SAUGATA GHOSE is an assistant professor in the Department of Computer Science at the University of Illinois Urbana-Champaign. He holds M.S. and Ph.D. degrees in electrical and computer engineering from Cornell University, and dual B.S. degrees in computer science and in computer engineering from Binghamton University, State University of New York. Prior to joining Illinois, he was a

postdoc and later a systems scientist at Carnegie Mellon University. He received the best paper award from DFRWS-EU in 2017 for work on solid-state drive forensics, and was a 2019 Wimmer Faculty Fellow at CMU. His current research interests include data-oriented computer architectures and systems, new interfaces between systems software and architectures, low-power memory and storage systems, and architectures for emerging platforms and domains. For more information, please visit his website at <https://ghose.web.illinois.edu/>.



NANDITA VIJAYKUMAR is an assistant professor in the Computer Science Department at the University of Toronto and the Department of Computer and Mathematical Sciences at the University of Toronto Scarborough. She is also affiliated with the Vector Institute for Artificial Intelligence. Before joining the University of Toronto, she was a research scientist in the Memory Architecture and Accelerator Lab at Intel Labs. She

received her Ph.D. and M.S. in 2019 from Carnegie Mellon University where she was advised by Prof. Onur Mutlu and Prof. Phil Gibbons. She also worked with the Systems Group in the Computer Science Department at ETH Zurich as a visiting student. In the past, She have also worked for AMD, Intel, Microsoft, and Nvidia. Her research interests lie in the general area of computer architecture, compilers, and systems with a focus on the interaction between programming models, systems, and architectures. Her current interests are in the system-level and programming challenges of robotics and large-scale machine learning. For more information, please visit his website at <http://www.cs.toronto.edu/~nandita/>.



IVAN FERNANDEZ received his B.S. degree in computer engineering and his M.S. degree in mechatronics engineering from University of Malaga in 2017 and 2018, respectively. He is currently working toward the Ph.D. degree at the University of Malaga. His current research interests include processing in memory, near-data processing, stacked memory architectures, high-performance computing, transprecision computing, and time series analysis.



MOHAMMAD SADROSADATI received the B.Sc., M.Sc., and Ph.D. degrees in Computer Engineering from Sharif University of Technology, Tehran, Iran, in 2012, 2014, and 2019, respectively. He spent one year from April 2017 to April 2018 as an academic guest at ETH Zurich hosted by Prof. Onur Mutlu during his Ph.D. program. He is currently a postdoctoral researcher at ETH Zurich work-

ing under the supervision of Prof. Onur Mutlu. His research interests are in the areas of heterogeneous computing, processing-in-memory, memory systems, and interconnection networks. Due to his achievements and impact on improving the energy efficiency of GPUs, he won Khwarizmi Youth Award, one of the most prestigious awards, as the first laureate in 2020, to honor and embolden him to keep taking even bigger steps in his research career.



ONUR MUTLU is a Professor of Computer Science at ETH Zurich. He is also a faculty member at Carnegie Mellon University, where he previously held the Strecker Early Career Professorship. His current broader research interests are in computer architecture, systems, hardware security, and bioinformatics. A variety of techniques he, along with his group and collaborators, has invented over the years have influenced industry and

have been employed in commercial microprocessors and memory/storage systems. He obtained his PhD and MS in ECE from the University of Texas at Austin and BS degrees in Computer Engineering and Psychology from the University of Michigan, Ann Arbor. He started the Computer Architecture Group at Microsoft Research (2006-2009), and held various product and research positions at Intel Corporation, Advanced Micro Devices, VMware, and Google. He received the IEEE High Performance Computer Architecture Test of Time Award, the IEEE Computer Society Edward J. McCluskey Technical Achievement Award, ACM SIGARCH Maurice Wilkes Award, the inaugural IEEE Computer Society Young Computer Architect Award, the inaugural Intel Early Career Faculty Award, US National Science Foundation CAREER Award, Carnegie Mellon University Ladd Research Award, faculty partnership awards from various companies, and a healthy number of best paper or "Top Pick" paper recognitions at various computer systems, architecture, and security venues. He is an ACM Fellow, IEEE Fellow for, and an elected member of the Academy of Europe (Academia Europaea). His computer architecture and digital logic design course lectures and materials are freely available on YouTube (<https://www.youtube.com/OnurMutluLectures>), and his research group makes a wide variety of software and hardware artifacts freely available online (<https://safari.ethz.ch/>). For more information, please see his webpage at <https://people.inf.ethz.ch/omutlu/>.