

RESEARCH ARTICLE

Casper: Accelerating Stencil Computations Using Near-Cache Processing

ALAIN DENZLER¹, GERALDO F. OLIVEIRA¹, NASTARAN HAJINAZAR, RAHUL BERA, GAGANDEEP SINGH¹, JUAN GÓMEZ-LUNA¹, (Member, IEEE), AND ONUR MUTLU, (Fellow, IEEE)

Department of Information Technology and Electrical Engineering (D-ITET), ETH Zürich, 8092 Zürich, Switzerland

Corresponding author: Juan Gómez-Luna (juang@ethz.ch)

ABSTRACT Stencil computations are commonly used in a wide variety of scientific applications, ranging from large-scale weather prediction to solving partial differential equations. Stencil computations are characterized by three properties: 1) low arithmetic intensity, 2) limited temporal data reuse, and 3) regular and predictable data access pattern. As a result, stencil computations are typically bandwidth-bound workloads, which experience only limited benefits from the deep cache hierarchy of modern CPUs. In this work, we propose Casper, a near-cache accelerator consisting of specialized stencil computation units connected to the last-level cache (LLC) of a traditional CPU. Casper is based on two key ideas: 1) avoiding the cost of moving rarely reused data throughout the cache hierarchy, and 2) exploiting the regularity of the data accesses and the inherent parallelism of stencil computations to increase overall performance. With small changes in LLC address decoding logic and data placement, Casper performs stencil computations at the peak LLC bandwidth. We show that by tightly coupling lightweight stencil computation units near LLC, Casper improves performance of stencil kernels by $1.65\times$ on average (up to $4.16\times$) compared to a commercial high-performance multi-core processor, while reducing system energy consumption by 35% on average (up to 65%). Casper provides $37\times$ (up to $190\times$) improvement in performance-per-area compared to a state-of-the-art GPU.

INDEX TERMS Stencil computation, near-cache processing, processing-in-memory, near-data processing, memory systems, caches.

I. INTRODUCTION

A stencil operation [1] defines a computation pattern where elements in a multidimensional grid are updated based on the values of a fixed pattern of neighboring points. Computations using stencil operations (called *stencil computations*) are a key building block of important high-performance computing (HPC) applications [2] and are used in a wide range of workloads including climate modeling [3], seismic imaging [4], fluid dynamics [5], and electromagnetic simulations [6]. Stencil computations encompass a large percentage of the overall runtime of such applications [7], [8], [9], [10], and [11]. For example, stencil computations represent over

The associate editor coordinating the review of this manuscript and approving it for publication was Nitin Nitin¹.

90% and 60% of the overall runtime in a computational fluid dynamics solver [11] and the COSMO climate simulation model [10], [12], respectively. Consequently, a large body of research [1], [3], [9], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38] highlights the need for highly efficient stencil computations. However, the current compute-centric processing systems, such as multi-core CPUs and GPUs, fail to fully utilize their on-chip resources (e.g., deep cache hierarchy, high throughput floating-point engines) when computing stencil operations [14], [33]. This results in low performance and low energy efficiency for stencil computations in current systems. In this work, we show that a careful domain-specific hardware/software co-design can improve the performance and energy efficiency of

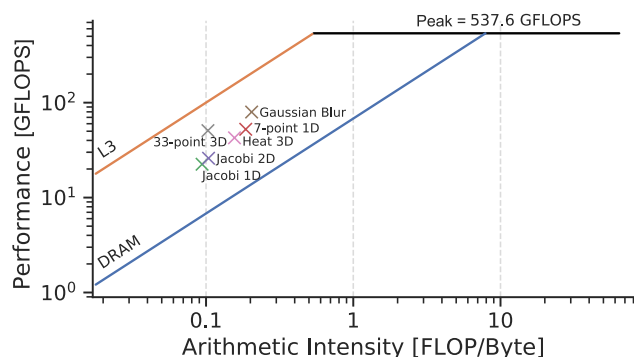


FIGURE 1. Roofline plot for a multi-core system [40], [41] running six stencil kernels.

stencil computations, at a much lower overhead than existing general-purpose solutions.

A. WHY A STENCIL ACCELERATOR?

Figure 1 shows the roofline plot [39] of important stencil kernels on a server-class CPU (a 16-core Intel Xeon CPU [40], [41]). The horizontal line is the peak floating-point performance of the system. The stencil kernels are multi-threaded and vectorized, and operate on double-precision floating-point values. The DRAM and L3 lines show the peak main memory and last-level cache (LLC) bandwidth, respectively. We make two observations. First, the stencil kernels are *not compute bound*, having low arithmetic intensity ranging from 0.09 FLOP/B to 0.2 FLOP/B. Second, all the kernels are *bounded by memory resources*, i.e., located on the left side of the inflection point and below the memory lines. More precisely, all the kernels are located below the L3 line and above the DRAM line, which shows that the stencils are bound by the LLC bandwidth rather than the main memory bandwidth. Given these observations, we conclude that the high number of LLC accesses is the main bottleneck for stencil computations, which causes such computations to experience only limited benefits from the deep cache hierarchies in modern CPU architectures.

Prior works [42], [43], [44], [45] also show that stencil kernels are bottlenecked by memory due to their low arithmetic intensity, leading to under-utilization of computational resources in compute-centric platforms (e.g., CPU and GPU). These prior works demonstrate that stencil kernels can leverage only 21.8% [42], 46% [46], 34% [46], and 46% [43] of the computational resources of a multi-core CPU, a 2-socket server-grade CPU, a 4-socket CPU, and a GPU, respectively, even in presence of code optimizations (e.g., temporal blocking). These percentages are inline with Figure 1, where all tested stencils achieve less than 20% of the peak performance (537.6 GFLOPS). Therefore, existing general-purpose processors cannot deliver high performance and high energy efficiency for stencil computations, thus opening up the space for custom stencil-based accelerators [1], [13], [14], [15], [16], [18], [20], [26], [27], [28], [29], [30], [31], [32], [32], [33], [34], [35], [36], [37], [38], [47].

Processing-in-Memory (PIM) is a promising paradigm for accelerating memory-bandwidth-bound workloads, which have low arithmetic intensity [34], [48], [49], [50], [51], [52], [53], [54], [55], [56], [57], [58]. The key idea of the PIM paradigm is to move computation close to or even into the memory devices where the data resides (i.e., caches [48], [59], [60], [61], [62], [63], [64], [65], DRAM [33], [34], [49], [50], [51], [52], [53], [54], [55], [56], [57], [58], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79], [80], [81], [82], [83], [84], [85], [86], [87], [88], [89], [90], [91], [92], [93], [94], [95], [96], [97], [98], [99], [100], [101], [102], [103], [104], [105], [106], [107], [108], storage [109], [110], [111], [112], [113], [114], [115], [116], [117]), eliminating the need to move the data to the processor and resulting in higher performance and lower energy consumption. Stencil computations are a prime candidate for acceleration using the PIM paradigm. In this work, we explore the opportunity to improve the performance and energy efficiency of stencil computations in traditional multi-core CPUs using computation near the LLC.

B. WHY NEAR LLC?

We exploit LLC as the prime location for computation, as opposed to offloading the computation to the off-chip main memory [30], [31], [33], [34], [37] or higher levels of caches (e.g., L2 [61]) for three main reasons. First, the per-thread datasets for stencil kernels in widely-deployed scientific applications are typically tiled to fit inside the LLC of traditional workstation-class CPUs [3], [118]. Hence, placing computation near LLC minimizes unnecessary data transfers between main memory and caches and cores. Second, the on-chip LLC bandwidth is multi-fold (e.g., about 10 \times in Intel Xeon [41]) higher than the traditional DDR-based DRAM main memory bandwidth. Third, though computing in higher-level caches (e.g., L2) can theoretically provide higher bandwidth, the additional data transfers required by cache management protocols (e.g., back invalidations, write backs, etc.) reduces the effective bandwidth significantly. Moreover, bringing data with low reuse (common in stencil computation data) to the higher-level caches results in major energy waste that can be alleviated by keeping such data in lower-level caches and performing the computation there.

Our goal in this paper is to design a near-LLC accelerator that improves the performance and energy efficiency of stencil computations by minimizing the unnecessary data movement between the memory and CPU, and within the cache hierarchy.

To this end, we propose Casper, a novel hardware/software codesign specifically targeted at stencil computations. We minimize data movement by placing a set of stencil processing units (SPUs) near the LLC of a traditional CPU architecture. Casper provides novel mechanisms to incorporate data mapping changes and support unaligned loads needed for high-performance stencil computations. Computation is mapped to SPUs in such a way that each stencil

processing unit (SPU) operates on the data that is located in the closest LLC slice. This reduces the overall data access latency and energy consumption while matching the compute performance to the peak bandwidth of the LLC.

Placing a stencil accelerator next to the LLC of a CPU introduces two key challenges. The first challenge is to maximize LLC bandwidth utilization. To address this challenge, Casper leverages the notion of *streams* [36], [119], [120], [121], [122] to expose the memory level parallelism that exists in the stencil computation to the SPUs. Each stream represents a set of consecutive memory accesses with a fixed stride. The notion of streams enables Casper to maximize memory bandwidth utilization without requiring complex structures (e.g., those that exist in high-performance cores to perform dynamic instruction scheduling [123], [124], [125]).

The second challenge is to minimize the data movement between different cache slices. In stencil computations, the neighboring grid points need to be accessed to compute the stencil operation for each grid point. However, current systems employ an address mapping scheme that distributes data over different LLC slices and provides load balance and fairness across CPU cores [126]. Such a mapping scheme can potentially map neighboring grid points to different LLC slices, introducing data transfers over the network-on-chip (NoC) and thereby eliminating the benefits of near-cache computing. To address this challenge, Casper uses a new hash function to align the data mapping to the grid structure of the stencil computation and place neighboring grid points in the same slice of the LLC.

We evaluate Casper using six widely used stencil kernels with up-to 3-dimensional grid domains. Casper outperforms a commercial multi-core CPU, on average, by $1.65\times$ (up to $4.16\times$) and reduces the energy consumption by 35% (up to 65%). Compared to a state-of-the-art GPU, Casper improves performance-per-area, on average, by $37\times$ (up to $190\times$).

We make the following key contributions:

- We propose Casper, the first near-cache accelerator for stencil computations. Casper addresses the memory bottleneck in stencil computations by (1) eliminating the need to move data to the processing core for computation, (2) minimizing the data movement within the cache hierarchy, and (3) maximizing the utilization of the LLC bandwidth. Casper achieves this with an area overhead of less than 1% for 16 SPUs in a Marvell ThunderX 2 [127], a server-class ARM CPU.
- We present a memory-centric execution model that maximizes the LLC bandwidth utilization by exposing the memory level parallelism that exists in the stencil computation to the near-cache accelerators.
- We provide hardware support to minimize data movement between different cache slices using a new mapping scheme that improves spatial locality for stencil data.
- We evaluate the effectiveness of Casper using six widely used stencil kernels and demonstrate that Casper

outperforms a commercial multi-core CPU, on average, by $1.65\times$ (up to $4.16\times$) and reduces the energy consumption by 35% (up to 65%). Compared to a state-of-the-art GPU, Casper improves performance-per-area, on average, by $37\times$ (up to $190\times$).

II. BACKGROUND

A. STENCIL COMPUTATIONS

Stencil computations [1] update the points in a data grid based on a fixed pattern of neighboring points. This fixed pattern, called the *stencil*, is applied on the complete grid iteratively until either a convergence criterion or a fixed number of steps are reached. Stencil computations are widespread in scientific computing, and are considered one of the thirteen dwarfs of scientific computing [2].

Stencil computations exhibit several common properties. We explain these properties using a *Jacobi* stencil [128] that is commonly used to solve discretized partial differential equations, as an example. Figure 2 shows the source code and data access pattern of a 2-dimensional Jacobi stencil. The computation performs arithmetic mean of each point in the grid and its immediate neighbors in all directions. This implementation uses three nested loops where the outermost loop iterates over time steps and the two inner loops sweep over the complete 2D grid. From this example, we can identify four common properties of stencil computations. First, the computation is embarrassingly parallel because the read and write data sets are disjoint. Second, the computation is regular and statically analyzable. The data dependencies and the structure of the computation can be analyzed ahead-of-time and do not depend on a dynamic input. Third, the arithmetic intensity of the stencil is low. Fourth, only a few types of operations are needed to compute the stencil. For example, in case of *Jacobi* stencil, only a floating-point multiply-accumulate (MAC) operation is performed for each input grid point (e.g., multiply $A[i][j]$ by 0.2 and accumulate) when computing an output grid point ($B[i][j]$). These properties make stencil computations a suitable candidate for high-performance acceleration (due to the highly-parallel and regular computation) while making them a natural fit for near-memory acceleration (due to the low arithmetic intensity and relying on only few types of operations).

```

for (int t = 0; t < T; t++) {
  for (int i = 1; i < N-1; i++) {
    for (int j = 1; j < N-1; j++) {
      B[i][j] = 0.2 * (A[i][j]
                     + A[i][j-1]
                     + A[i][j+1]
                     + A[i-1][j]
                     + A[i+1][j]);
    }
  }
  // swap pointers
  temp = A; A = B; B = temp;
}

```




FIGURE 2. 2D Jacobi stencil pseudo-code and data access pattern.

B. MEMORY-CENTRIC ARCHITECTURES: OVERVIEW AND LIMITATIONS

Processing-in-Memory (PIM) architectures (e.g., [33], [34], [48], [49], [50], [51], [52], [53], [54], [55], [56], [57], [58], [59], [60], [61], [62], [63], [64], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79], [80], [81], [82], [83], [84], [85], [86], [87], [88], [89], [90], [91], [92], [93], [94], [95], [96], [97], [98], [99], [100], [101], [103], [104], [109], [110], [111], [112], [113], [114], [115], [116], [117], [129], [130], [131], [132]) attempt to address the memory bottleneck issue by performing computation in proximity to the memory and thereby reducing the overheads of data movement in the system. There are two approaches to PIM [52]: (1) processing-using-memory (PuM), which performs computation inside the memory array itself, using the analog operational properties of the memory cells [33], [34], [48], [49], [51], [52], [53], [54], [55], [56], [57], [58], [59], [60], [61], [62], [63], [64], [68], [71], [74], [79], [81], [82], [83], [86], [87], [88], [89], [90], [91], [92], [93], [94], [97], [98], [99], [100], [101], [103], [104], [105], [106], [107], [109], [110], [111], [112], [113], [114], [115], [116], [117], [129], [130], [131], [133], [134], [135]; and (2) processing-near-memory (PnM), which employs compute elements close to the memory arrays, for example on the logic layer of a 3D-stacked DRAM device [50], [66], [67], [69], [70], [72], [73], [75], [76], [77], [78], [80], [84], [85], [95], [96], [102], [136], [137], [138], [139], [140], [141], [142], [143], [144], [145], [146], [147], [148], [149], [150], [151], [152], [153].

The PuM proposals, either in DRAM [49], [71], [74], [81], [98], [99], [100], [105], [106], [107], [130], [133], [134], [135] or SRAM [48], [59], [60], [149], perform bulk bitwise and arithmetic (e.g., addition, multiplication, reduction) operations. Even though these works benefit from the large internal bandwidth provided by the memory device (DRAM or SRAM), they require the application to follow a rigid data layout and data mapping scheme to align the operands correctly within the memory arrays. Such data layout and alignment management are not straightforward and remain an open research problem. In contrast, our work introduces hardware and software optimizations to efficiently orchestrate the data movement or handle unaligned operands.

PnM proposals span a wide range of applications, such as neural networks [64], [70], [101], [108], databases [102], [103], [148], mobile workloads [75], bioinformatics [67], [154], image processing [104], and graph processing [61], [66], [69], [153]. PIMS [34] proposes a PnM approach (in the logic layer of Hybrid Memory Cube (HMC) [155]) to accelerate stencil operations. The proposed solution leverages the computational capabilities that are already present in an HMC device¹ to execute the addition operation that is commonly used in stencil computations. Despite the performance improvement compared to a CPU-centric model, PIMS suffers from two important shortcomings. First, prior

¹The HMC specification [155] defines a series of 8 to 16-byte atomic operations that can be executed directly within the memory device.

works [156], [157] show that the atomic operations provided by the HMC device can only exploit a small fraction of the total internal memory bandwidth, creating a bottleneck for the bandwidth-hungry stencil computations. Second, due to the limited area and power budget available inside the logic layer of HMC, PIMS only supports the addition operation and continues to rely on the host processor to execute other more complicated operations such as multiplication. Therefore, PIMS still incurs a high memory traffic between the host and the memory device to compute a single stencil operation.

To our knowledge, Casper is the first work to tightly integrate specialized compute units into the LLC of a traditional CPU to accelerate stencil computations. In contrast to prior works, Casper accounts for the *fundamental* properties of stencil computations, such as data access pattern, aiming to provide a hardware/software co-design that can fully exploit the high memory bandwidth provided by the cache, while efficiently orchestrating the data movement required to execute the stencil operation.

III. Casper: OVERALL ARCHITECTURE

We propose Casper, a novel near-cache accelerator, to improve the performance and energy efficiency of stencil computations. This section introduces the high-level overview of our architecture (Section III-A), the execution model of Casper (Section III-B), and stencil processing unit (Section III-C).

A. OVERVIEW

Casper is a near-cache accelerator that leverages a memory-centric execution model and hardware support to accelerate stencil operations. A typical shared last-level cache (LLC) is partitioned into multiple cache slices, connected through the network on chip (NoC). Casper's hardware is composed of a set of stencil processing units (SPUs), placed in each cache slice of the LLC. In addition to accessing the local cache slice, each SPU can load data from remote LLC slices through the NoC. Similar to a regular LLC, SPUs can load data from other levels of the cache hierarchy into the LLC.

Computing at the LLC's peak throughput is only possible when each SPU in the system loads data from its local LLC slice, thus avoiding overheads related to NoC traffic. However, the existing mapping of memory addresses to the LLC slices does not guarantee mapping consecutive cache lines to the same slice. In fact, while mapping scheme information remains undisclosed, prior work shows that consecutive cache lines are usually mapped to different LLC slices in order to provide fairness and load balancing across CPU cores [158]. Due to the streaming nature of stencil computations and the dependency on a small group of neighboring grid points, scattering data through the cache slices would increase the number of remote data requests made by each SPU, penalizing performance and energy consumption. To address this issue, Casper maps blocks of consecutive memory addresses to the same LLC slice. To this end, we introduce hardware support for a *stencil segment* i.e., a

contiguous region of physical memory that contains the data set of the stencil kernel. Casper customizes the mapping of memory addresses to LLC slices for the stencil segment in order to optimize data locality when computing stencils. Data mapped outside the stencil segment follows the conventional address mapping scheme.

B. EXECUTION MODEL

At a high-level, each SPU performs the stencil operation on each grid point sequentially. The points are accessed in row-major order, following their layout in memory. For each grid point, the SPU loads the input data from LLC. Once the data arrives at the SPU, the SPU’s execution unit performs a single multiply operation with a predefined constant, whose output is added to the accumulator. The final result is stored in the results array after all computations for the grid point have been performed. Then, the SPU proceeds to calculate the next grid point.

To accelerate this process, Casper abstracts data movement through a series of *stream* operations, which enable feeding data to the executing units at a high throughput. A *stream* is a sequence of data elements of the same type located in consecutive physical memory addresses. Each stream is characterized by a start address, data type width, a number of elements, and a position pointer. When iterating over a stream, initially, a position pointer indicates the start of a stream. Upon receiving a control signal, this pointer is incremented to point to the stream’s next element. The stream can then be advanced until the final element in the stream is reached. The stream abstraction facilitates iterating over stencil data by simply incrementing the position pointer, as opposed to loading each data element using absolute memory addresses. Since a stencil computation moves through the entire grid at the same pace, maintaining the same relative offsets for the stencil pattern, it is naturally possible to use streams to abstract such an access pattern. Accordingly, all the streams can be advanced at the same pace, marshaled by the SPU moving through the grid.

C. STENCIL PROCESSING UNIT

Figure 3 shows the architecture of a SPU. The main building blocks of an SPU are ❶ the instruction buffer, ❷ the load queue, ❸ the stream buffer, ❹ the constant buffer, and ❺ the execution unit. The complete SPU is pipelined to maintain single-cycle instruction throughput. The SPU controls the complete computation, from instruction fetch until retire. Therefore, it does not require any interaction with the CPU.

1) INSTRUCTION BUFFER

The instruction buffer has a capacity to hold 64 *compressed* instructions to compute a stencil. Every instruction encodes the operands and the control signals necessary for one type of stencil operation. The same sequence of instructions is applied to every stencil grid point due to the regular nature of the stencil computation. We introduce the instruction set used in Casper in Section V.

2) LOAD QUEUE

The load queue is responsible for issuing data requests to the memory subsystem. While the SPU processes instructions in-order, the varying memory access latency of the memory subsystem can result in responses that arrive out-of-order. The load queue acts as a buffer to hold the data until all the previous longer-latency memory requests have been completed. This buffer space ensures that the correct instruction order is maintained. The load queue is sized to hide the latency of accessing the LLC’s local slice because as it is the preferred location for fetching the data.

3) STREAM AND CONSTANT BUFFERS

The stream buffer holds the current state of the streams. For every decoded instruction, this state is loaded from the stream buffer, and the effective address is calculated to issue the memory requests to the LLC. The constant buffer holds the constant factors needed for the MAC operation. The stream and constant buffers are initialized by API calls (Section V-B) before the SPU starts the stencil computation.

4) EXECUTION UNIT

As shown in Figure 3, the execution unit of a SPU is comprised of a 512-bit vector unit which operates on vectors of 8 double-precision floating-point elements.

For every input grid point, the SPU loads the data elements from the neighboring points, multiplies them with a constant factor, and accumulates the results. This final accumulated result is then stored in the output grid point. Therefore, each SPU execution unit only computes one kind of instruction: a double-precision floating-point MAC operation.

IV. Casper: MICRO-ARCHITECTURAL SUPPORT

The design and implementation of Casper require us to solve several key system integration challenges, as we discuss in this section. First, the SPUs access frequently data from two cache lines. We introduce two simple changes to the LLC row




FIGURE 3. Components of one SPU.

decoder for efficient access (Section IV-A). Second, the cache lines accessed by an SPU should preferably reside in the same LLC slice. We introduce lightweight remapping to ensure that contiguous blocks of stencil data map to the same LLC slice (Section IV-B). Third, Casper should be compatible with existing cache coherency mechanisms (Section IV-C). Fourth, Casper should be able to concurrently run with other CPU processes (Section IV-D).

A. SUPPORTING UNALIGNED LOADS

While the stream abstraction offers efficient memory accesses for SPUs, the LLC architecture only supports data accesses that are aligned to the 64 B cache line boundaries. As the relative offsets used in streams are not necessarily aligned to cache line boundary, the loaded data might need to be realigned before it can be used by MAC compute. For example, in Figure 4 each SPU computes eight 64-bit output grid points ($B[i]$) for a 3-point Jacobi-1D stencil. The access to the center point of the stencil $A[i]$ is correctly aligned to the 64 B boundary of the cache line (shaded in yellow) such that the data for this point can be used for computation as soon as it arrives at the SPU. However, to gather several input grid points at indices $+3$ ($A[i+3]$) and -3 ($A[i-3]$), the data coming from the cache needs to be correctly shifted (shaded in orange and red), and two cache lines need to be combined to assemble the operands for the computation. As a result, preparing the operand for the MAC unit involves two loads, one shift, and one combine operation on the data from the two cache lines. These additional operations to resolve unaligned data accesses lead to two main inefficiencies: (1) underutilization of MAC units, and (2) cache bandwidth overhead. The reason is that the number of load/store operations per MAC operation increases. For example, the sequential code shown at the top of Figure 4 needs 4 load/store operations per 3 MAC. However, the vectorized execution of this stencil (bottom part of Figure 4) would need 6 load/store operations per 3 MAC, i.e., two cache line loads for $A[5]$ to $A[12]$, one cache line load for $A[8]$ to $A[15]$, two cache line loads for $A[11]$ to $A[19]$, and one cache line store for the elements of B .

a vectorized execution of the stencil:

```
for (int i = 3; i < N-3; i++) {
  B[i] = 0.2 * A[i-3] + 0.2 * A[i] + 0.2 * A[i+3]
}
```

using the data layout:

$A[i] = [0 \quad 8 \quad 16]$

leads to memory accesses:

$B[8..16] = 0.2 * A[5 \quad 12]$
 $+ 0.2 * A[8 \quad 15]$
 $+ 0.2 * A[11 \quad 19]$

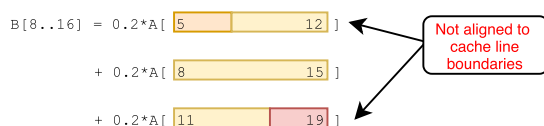


FIGURE 4. Unaligned loads occurring during the vectorized execution of a stencil.

To address the above challenges, we introduce two simple modifications to the LLC row decoding logic to (1) support

loading data aligned to any 8 B boundary, and (2) correctly align data for SPU execution. These two modifications allow us to load values from two adjacent cache lines (e.g., values $A[5]$ to $A[12]$ in Figure 4) in one access. Our modified LLC row decoding logic reduces SPU's complexity and area footprint by avoiding the need to add a large register file and/or extra logic for shifting and packing the partial cache lines inside the SPU.

1) IMPLEMENTATION CHALLENGES

Loading data aligned to 8 B boundaries can potentially result in loading data located on two cache lines. This introduces two key challenges. First, *two* tags need to be matched to locate the two cache lines in one access (i.e., using one address). Second, the correct data from each of the two cache lines must be loaded using only one provided address.

To address the first challenge, we enable the cache to match two tags in parallel by adding a second read port to the tag array. Since the two cache lines involved in an unaligned load are always at consecutive addresses (e.g., values $A[5]$ to $A[12]$ in Figure 4), they are guaranteed to be mapped to different cache sets, and thus, there are no conflicts during the tag matching process. If at least one cache line is not in a readable state, a regular cache miss is generated, and the request is stalled until the miss resolves.

To address the second challenge, we make changes into the LLC row decoding logic to load data either from the requested address or one of the cache lines located inside the 64B-vicinity of the requested address. More specifically, we add one 3-to-1 multiplexer for each LLC SRAM row. The inputs to the multiplexer are set to the row decoder output of the current row and both the adjacent rows. The multiplexer selects the appropriate row(s) based on the row selection signal.

We explain next the solutions to both challenges.

2) LOADING SHIFTED CACHE LINES

Figure 5 shows the execution sequence of an unaligned access to the LLC. In this example, we consider the request to the cache line holding elements 8 to 15, shifted to the right by three positions. This corresponds to the access to input grid points $A[i-3]$ in Figure 4 (i.e., values $A[5]$ to $A[12]$). First, we consider the case where the two cache lines involved in the access are mapped to the *same* cache way. One cache way consists of 4×32 kB data arrays, each having 2×16 kB subarrays. One 16 kB subarray holds 64 consecutive bits of a cache line for all the 2048 sets. Thus, each 64-bit segment of a cache line is stored in a different SRAM subarray. As shown in Figure 5, all the elements required to build the requested 64 B block of data are stored in different subarrays of a single cache way (orange and green elements on the left part of the figure). Therefore, by selecting the correct data at the subarray level (using shift direction *shdr* and amount *shamt*) it is possible to load all the elements using only one command, while maintaining the same throughput as a regularly aligned load (any extra latency is negligible, since there are no

conflicts during the matching of the two tags, and pipelined accesses hide it). To complete the unaligned access, we need to rotate the data such that the correct element is at the first position of the 64 B block of response data. To accomplish this, we use a rotate network that performs this operation before the final output. This approach provides the SPUs with the ability to load data that is aligned to arbitrary 8B boundaries from the LLC.

If the two cache lines involved in an unaligned load are mapped to *different* ways, the load sequence is similar. As the data and tag access happen in parallel inside the cache, the data load is initiated on all cache ways before the way hit has been confirmed. Thus for any access, all the ways present the data for the requested set on their output bus. Based on the way hit, the data from the correct way is selected for the output. We modify the way hit selection to select each subarray's output independently, depending on the shift amount and direction that is provided with the request.




FIGURE 5. Loading unaligned data from the LLC. The two cache lines involved in the access are mapped to the same cache way. One cache way consists of 4×32kB data arrays, each having 2×16kB subarrays.

3) ROW DECODING

Locally at each SRAM array, the shift amount and the direction are used to determine whether to load the data from the requested address or one of the two rows immediately adjacent to it. To this end, we include a local selection signal that reroutes the row selection signal to the correct row. Figure 6 shows the layout of the added logic for local selection signal which consists of one 3:1 multiplexer for each SRAM row. The inputs to each multiplexer are the row decoder output for the current row and the signals for both of the adjacent rows. The multiplexer then selects which input to forward. If the subarray needs to select elements from the cache lines specified in the request's address, the multiplexer forwards the middle signal (the row activation does not change). If the element from the adjacent cache line should be loaded, the output depends on the shift direction. If the shift direction is to the right (left), the row at index -1 ($+1$) should be activated. This shifts the row selection signal by one position in the requested direction, and loads data stored on the adjacent cache line. We include logic at the edge of each subarray to compute the select signal for the multiplexers based on the

shift direction, shift amount, and the subarray ID (i.e., the position of the bits stored in this subarray within a complete cache line). All the multiplexers for one SRAM subarray are controlled using the same select signal.




FIGURE 6. Additional logic added between the row decoder and the SRAM array. All the multiplexers receive the same select signal.

B. LLC DATA MAPPING

Sliced LLCs employ address mapping schemes that aim to provide load balance and fairness across CPU cores. These mapping schemes use hash-functions that are not disclosed by CPU vendors [158], but are shown to map consecutive cache lines to different LLC slices. This is a challenge for Casper, since the performance of the SPU can only be maximized if the SPU mainly accesses data stored on the local LLC slice. The computation for one data point depends on input data from its neighboring grid points which as consecutive cache lines can be mapped to different LLC slices. This leads to significant data transfers over the NoC, thus reducing the benefits of near-cache computing.

To address this challenge and enable maximum performance for each SPU, we introduce a mechanism to adapt the mapping of data to LLC slices to the needs of stencil computation. The new data mapping scheme maps blocks of consecutive memory addresses to the same slice, allowing neighboring grid points to be stored in the same slice of the LLC. Thus, remote slice accesses are reduced to only accesses to points that lie on the other side of the boundaries between the blocks.

We enable remapping of the data used for the stencil computation without affecting the mapping for other system data by introducing a *stencil segment*. Following the proposal by [159], a physically contiguous block of memory is used by the system to hold stencil data. Casper ensures that the physically contiguous blocks of data in the stencil segment are mapped to the same LLC slice. At each NoC injection point, the system checks whether or not the address is part of the stencil segment. In case the address belongs to a stencil segment a new hash function is applied to issue a stencil segment memory request to the LLC. Otherwise, the conventional hash function is applied. By deciding which mapping function to use based on a memory request's physical address, each address is mapped to exactly one cache slice, regardless of whether or not it contains stencil data.

Sizing blocks to map stencil data to the LLC comes with a trade-off. Smaller blocks introduce more remote slice

accesses for multidimensional grids, which hurts performance. The boundary elements are all located on separate cache lines, and cannot be loaded using the unaligned load mechanism as the affected cache lines are stored on separate LLC slices. On the other hand, smaller blocks allow partitioning physically contiguous arrays across these blocks and thus distributing the data evenly to cache slices. Therefore, data for the same grid points from different source arrays can map to the same slice. This mapping scheme improves locality for multi-source stencils at the expense of more remote slice accesses on the blocks' boundary. For example, in a 16 core system, data aligned to a 2MB boundary is mapped to the same slice, and the system can consume data from up to 16 arrays for one grid point from the local slice. In this work, we design the hash function to map stencil data as a linear hash that statically maps contiguous blocks of 128kB to LLC slices in round-robin fashion.²

C. COHERENCE SUPPORT

An SPU loads data from the LLC by directly injecting a load/store request in the LLC controller's request queues. Casper does not impact the current cache coherency mechanism, as requests from the SPU are injected into the same request queues as conventional requests from the CPU cores and the private caches. If a write request from an SPU targets a cache line that is not in writable state in the LLC, the coherency mechanism will trigger necessary state transitions and invalidations to allow the write to complete.

D. CONCURRENT EXECUTION WITH CPU AND CONTEXT SWITCHING

To enable concurrent execution with the CPU, we reserve one way of the LLC for other applications running in the CPU, similar to prior work [48]. Additionally, high priority is assigned to the SPU process to minimize the occurrence of context switches. If a context switch happens, the state of the SPU remains unchanged as it is not allowed to start a new stencil computation before the current computation finishes.

V. Casper: PROGRAMMING MODEL SUPPORT

A. CasperISA

Casper makes use of specialized instructions to compute each point involved in the stencil operation. Figure 7 shows the layout of a Casper instruction. Every Casper instruction is 15-bits wide and comprises of (1) 4b constant buffer index, (2) 4b stream buffer index, (3) 1b shift direction, (4) 3b shift amount, and (5) 3b control bits. 15-bits instructions lead to a compact stencil code in Casper.³ Note that Casper reuses the instructions for all the grid points, thus reducing the instruction count.

After decoding a new instruction, the SPU uses the 4-bits present in the constant field to index the constant buffer

²128kB blocks provide a good tradeoff across our evaluated stencils. We leave the design of a configurable hash function for future work.

³Common complex stencils have input sizes in the order of 30-40 points.




FIGURE 7. Instruction layout for stencil computation.

(a small SRAM buffer), loading the requested double-precision constant value that is used by the execution unit during the computation. The SPU uses the 4-bits in the stream field to index the stream buffer and find the memory address of the requested stream. The shift direction and the shift amount fields are used to assemble the correct memory address that point to the appropriate data. The final 3 bits are control bits *clear accumulator*, *enable output*, and *advance stream* that are used to reset the accumulators in the execution unit, enable the contents of the accumulator to be stored into the results array, and to advance the stream pointers, respectively.

We also provide a programming library that allows the user to easily generate Casper instructions from user-level code. The generated instructions to execute a given stencil computation are then stored in a contiguous array that holds all the stencil instructions. This array is then stored in the instruction buffer of all the SPUs. In this paper, we statically analyze stencil operations and generate the appropriate set of Casper instructions using our library. However, this step could be fully automated by a compiler due to stencil operations' regular nature.

B. CasperAPI

Table 1 shows Casper API functions. Calls to API functions are mapped directly to ISA instructions, which are broadcasted to all SPUs. These instructions are integrated into the existing ISA (e.g., x86 [160], ARM [161], RISC-V [162]) using spare instruction opcodes. We use the ARM ISA in our implementation (Section VII-A). Similar to other offload-based accelerators like GPUs, the CPU is not allowed to modify stencil data while the SPUs are running, to avoid corrupting the data. Enforcing this policy is the responsibility of the programmer.

VI. AN EXAMPLE: JACOBI-2D STENCIL

We illustrate how to program Casper using the code in Figure 8 as an example. The code implements the Jacobi-2D stencil presented in Section II-A. We explain the example using a system consisting of 4 LLC slices and SPUs.

First, a stencil segment covering 4 MB is allocated (line 4). Then, we define the start of the arrays A and B such that the same grid point of both arrays is mapped to the same LLC slice (lines 6-8). The programmer then initializes the arrays with the stencil data. Furthermore, the constant for the multiplication (line 12) and the stencil instructions (line 14) are sent to the SPU. The streams for all the SPU are configured inside the loop (lines 22-29). In this example,

TABLE 1. Casper Programmer API.

Function	Parameters	Description
<code>initStencilSegment</code>	<code>int size</code>	Requests a physically contiguous memory region of the specified size from the system to hold the stencil data
<code>initStencilcode</code>	<code>addr A, int length</code>	Takes a pointer to the microcode and the length of the code. After generating the code with helper functions provided by the programming library, the code is then broadcast to the SPU
<code>initConstant</code>	<code>double const, int index</code>	Initializes constant values that will be used during the multiplication step of a stencil operation. The function sets the specified constant value at the given index in the constant buffer
<code>initStream</code>	<code>addr A, int streamID, int accID</code>	Configures the streams used in the stencil code. The streams are configured per SPU to enable the programmer to tune the data layout to the grid's structure and minimize the number of off-slice cache accesses
<code>setNElements</code>	<code>int n, int accID</code>	To communicate to each SPU how many elements to compute. All the SPUs maintain a counter to keep track of their progress and stop when the desired number of elements has been computed
<code>startAccelerator</code>	-	The final function starts with the execution of all the SPU. After calling the start function, one of the SPUs acts as the leader and maintains a state to track the progress of all SPUs. Once all the SPUs finish their computation, the leader signals the completion to the CPU

```

1 void computeStencil(void* code) {
2
3     /* stencil segment holding 4MB */
4     uint64_t segment = initializeStencilSegment(4194304);
5
6     double* A = (double*) segment;
7     /* results start halfway through segment */
8     double* B = (double*) segment + 2097152;
9
10    /* initialize data ... */
11
12    initConstant(0.2, 0);
13
14    initStencilcode(code, 5);
15
16    /* 128kB block holds 16384 doubles */
17    int blockSize = 16384;
18
19    /* rows of 128 elements */
20    int rowLength = 128;
21
22    for (int i = 0; i < 4; i++) {
23        initStream(&A[i*blockSize-rowLength], 1, i);
24        initStream(&A[i*blockSize], 2, i);
25        initStream(&A[i*blockSize+rowLength], 3, i);
26        initStream(&B[i*blockSize], 0, i);
27        /* 512kB is 65536 doubles */
28        setNElements(65536, i);
29    }
30    startAccelerator();
31 }

```

FIGURE 8. Program code for Jacobi 2D stencil.

four streams are configured: three input streams to load the elements at $A[j-1][i]$, $A[j][i]$ and $A[j+1][i]$ (lines 23-25), and one output stream to store the result (line 26). As the elements $A[j][i-1]$, $A[j][i]$, and $A[j][i+1]$ are laid out in consecutive memory addresses, we can reuse the same stream and leverage the support for unaligned loads by shifting the access to the left (right) by one element to load the data. Finally, the number of elements to compute for each SPU (line 28) is configured, and the computation starts.

Figure 9 shows the stencil instructions executed on the SPU. As the stencil loads data from five input points, the instruction sequence consists of five instructions. All the input points are multiplied by the constant 0.2. This means all the instructions encode the same constant factor $c0$. The instructions load data from three different streams: the first stream points to the value $A[j-1][i]$, aligned to a cache line boundary. The values $A[j][i-1]$, $A[j][i]$, and

```

1     // 0.2 * A[j-1][i]
2     c0, s1, 0, 0, 1, 0, 1 //no shift, clear acc, inc stream
3     // 0.2 * A[j][i-1]
4     c0, s2, 1, 1, 0, 0, 0 //shift right by 1
5     // 0.2 * A[j][i]
6     c0, s2, 0, 0, 0, 0, 0 //no shift
7     // 0.2 * A[j][i+1]
8     c0, s2, 0, 1, 0, 0, 1 //shift left, inc stream
9     // 0.2 * A[j+1][i]
10    c0, s3, 0, 0, 0, 1, 1 //enable output, inc stream

```

FIGURE 9. Instruction sequence for the Jacobi-2D stencil.

$A[j][i+1]$ are stored in consecutive memory addresses. These three accesses all use the same stream, but using a shift amount and direction included in the request for the elements at indices $i+1$ and $i-1$. This loads the correctly aligned data using the unaligned load mechanism. Finally, the third stream is configured to load the value at $A[j+1][i]$.

As mentioned earlier (Section V), the final 3 bits of each Casper instruction hold control information for the SPU. The first bit, i.e., clear accumulator bit, must be set by the first instruction of a grid point (line 2). The next bit is the enable output bit which is set in the last instruction of the sequence (line 10) and generates a store request. The final bit or the advance stream bit signals the SPU to advance the stream pointer and has to be set in the last instruction consuming data from each stream (lines 2, 8, 10).

VII. METHODOLOGY

A. EXPERIMENTAL SETUP

We simulate the performance of Casper using the gem5 simulator (v20.0) [163], [164] in syscall emulation mode. Table 2 describes our system configuration in details. We use the ARM ISA and the Ruby memory model. The system is based on a generic modern server-class CPU, consisting of 16 out-of-order cores and three levels of cache, having a sliced LLC with 2 MB per slice. Our baseline configuration uses the same system configuration without the SPU and the LLC changes we propose. Also, we include stride prefetchers at all levels of the cache hierarchy. We evaluate the performance and energy benefit of Casper against the baseline CPU architecture, an NVIDIA Titan V GPU [165]. We use an energy model based on CACTI 7.0 [166] and

energy models proposed by prior works [167], [168]. We use the area model presented in [169] scaled down to 22nm to estimate the area of the SPU.⁴ For the GPU performance/area comparisons, we use the complete die size of 815mm² of the Titan V [171] because typical GPU-accelerated systems also need a host CPU to perform the computation. As a result, the total area for the end-user is either the complete GPU or the Casper hardware modifications, added to the area of the existing host CPU.

TABLE 2. Simulation Parameters for the baseline CPU and Casper.

Component	Configuration
<i>Casper</i>	16 SPUs, 1 SIMD unit/SPU (512-bits wide)
<i>CPU</i>	10 entry load queue, 0.016 nJ/instruction 16 out-of-order cores, 2 GHz, 8-wide issue, 72 entry load queue, 64 entry store queue, 1 SIMD unit/core (512-bits wide)
<i>L1/D Cache</i>	224 entry ROB, 0.08 nJ/instruction 32 kB, private, 8-way, 16 MSHRs, 4 cycle round-trip latency, 2 load ports, 1 store port
<i>L2 Cache</i>	15/33 pJ per hit/miss [167] 256 kB, private, 8-way, 16 MSHRs, 12 cycle round-trip latency, 1 load port, 1 store port
<i>L3 Cache</i>	46/93 pJ per hit/miss [167] 32 MB, shared, 16-way, 16 slices, 32 MSHRs/slice 36 cycle round-trip latency, 1 load/store port per slice
<i>Coherence Protocol</i>	MESI
<i>Replacement Policy</i>	LRU replacement
<i>Hardware Prefetchers</i>	Stride prefetchers at all levels of the cache
<i>On-Chip Network</i>	mesh, XY-routing, 64 B/cycle per direction
<i>Main Memory</i>	16 GB, DDR4, 4 channels, 160nJ per read/write [168]

B. BENCHMARKS

We evaluate Casper on six stencil benchmarks, including up to 3-dimensional stencils with varying data reuse characteristics, ranging from 3 input points (Jacobi 1D) up to 33 points for the 33-point 3D stencil. We use Jacobi 1D, -2D, 7-point 3D (heat diffusion) from Polybench [172], a 5×5 Gaussian blur filter [173], the 7-point 1D kernel from [174], and a 33-point 3D kernel to represent higher-order scientific simulations [43], [175]. All the benchmarks use Jacobi-style stencils with disjoint read and write data sets. They all operate on double-precision floating-point values. These benchmarks are elementary stencils that can be conjugated to form more complex stencils occurring in real-world applications. In addition to the standard data set sizes (which fit inside the LLC of the CPU), we also evaluate Casper on data sets that (1) exceed the size of the LLC (named *DRAM*), and (2) fit within the private L2 caches of the CPU (*L2*). Table 3 lists the domain sizes. In the appendix, Table 4 shows the dynamic instruction counts of all evaluated stencils and datasets for the baseline CPU and for Casper.

VIII. RESULTS

This section analyzes the performance, the energy consumption, and the hardware cost of Casper, and compares to the baseline CPU, the baseline GPU, and a PnM accelerator for

TABLE 3. Domain size used for evaluations.

Level	1D	2D	3D
<i>L2</i>	131,072	512×256	$64 \times 64 \times 32$
<i>L3</i>	1,048,576	1024×1024	$128 \times 128 \times 64$
<i>DRAM</i>	4,194,304	2048×2048	$256 \times 256 \times 64$

stencil operations. The appendix includes Table 5 and Table 6, which contain our detailed measurements.

A. PERFORMANCE

Figure 10 shows the speedup of Casper compared to a 16-core CPU baseline system. For the datasets that fit within the LLC, which represent typical data set sizes for stencil computations, we observe an average speedup of 1.65 \times . The 1- and 2-dimensional stencils achieve speedups between 1.66 \times and 3.0 \times . However, the 3-dimensional stencils cannot achieve the same performance and even experience a slowdown in the 33-point stencil case. The reason for this performance loss is twofold. First, 3-dimensional stencils need to load a significant part of their input data from remote LLC slices, which introduces longer access latencies and lowers the throughput of the SPU. Second, the 33-point 3D stencil has good L1 cache behavior in the baseline, achieving a hit-rate of 95%, making it a good fit for execution on a traditional CPU. We conclude that the performance benefits of Casper are larger on lower-dimensional stencils that load most of their input data from the local LLC slice.

The average performance improvement for the smaller data sets that fit in the L2 caches of the CPU (i.e., L2 in Figure 10) is 1.89 \times . Even though the data is stored closer to the core and does not need to travel through the entire cache hierarchy, the speedups are similar to the larger data sets that fit in the LLC. This is due to the fact that the access latency from CPU to the L2 cache is similar to the latency between SPU and the closest LLC slice (12 vs 8 cycles load-to-use).

Our results show that for large data sets that exceed the size of the LLC (i.e., *DRAM* in Figure 10), Casper improves performance by 1.4 \times , on average. The highest speedups are achieved by the 2-dimensional stencils, with Blur 2D reaching 4.16 \times . We explain this by the fact that the baseline CPU implementation of Blur 2D has a very low LLC hit-rate (only 2%), and thus the number of main memory accesses is 4 \times higher compared to Casper. The main reason for this low hit-rate of the baseline CPU implementation is that the prefetchers are interfering with demand accesses, evicting cache lines out of the LLC before they are used. The remaining stencils (i.e., Jacobi 1D, 33-point 3D) perform similar to the baseline or even experience slowdowns. The reason for this is the fact that the main memory bandwidth is the main bottleneck. The 33-point 3D stencil experiences a slowdown because it is well-suited for computation with smaller private caches for each core. We conclude that even though Casper cannot alleviate the main memory bandwidth bottleneck in the case of large data sets, it does not lead to significant slowdowns for most stencils.

⁴We conservatively scale down the model as analyzed in [170].




FIGURE 10. Speedup compared to the baseline multi-core system.




FIGURE 11. Normalized energy consumption compared to the 16-core baseline.

B. ENERGY CONSUMPTION

Figure 11 shows the normalized energy consumption of Casper compared to the baseline CPU. For the data sets that fit into the LLC (i.e., LLC in Figure 11), Casper reduces energy consumption by 55%, on average. Casper reduces energy consumption by 49% even for the 33-point 3D stencil, whose performance is slower in Casper than in the baseline. The reduction in energy consumption is larger for simpler stencils (Jacobi 1D/2D, 7-point 3D), reaching up to as 65% for 7-point 3D. This is due to the fact that complex stencils perform more LLC accesses. Furthermore, since the SPUs are situated close to the LLC slice, accessing LLC is not as energy-efficient as accessing the smaller L1 cache. This results in lower energy savings for the more complex stencils because of higher L1 reuse in baseline CPUs.

For both the smaller and larger data sets (i.e., L2 and DRAM in Figure 11), Casper reduces energy consumption by 26% and 23%, on average respectively. We make two observations. First, Casper increases the energy consumption of the 1-dimensional benchmarks (Jacobi 1D and 7-point 1D) when compared to the baseline, for both small and large data sets. For the large data sets, this is the case because the CPU cores can be idle for most of the runtime, waiting for memory. For the smaller data set, the baseline’s energy consumption is very low because there are very few LLC accesses. Since Casper loads the data from the shared LLC, it increases energy consumption in such cases. Second, for all the other benchmarks, Casper reduces energy consumption significantly. Casper is more energy-efficient even for the benchmarks that the CPU baseline outperforms since our SPU design is more energy-efficient than the CPU baseline. Thus, we conclude that Casper achieves significant reductions in energy consumption compared to a traditional out-of-order CPU.

C. COMPARISON WITH GPU

Figure 12 shows Casper’s performance and performance-per-area normalized to an NVIDIA Titan V GPU. GPU outperforms Casper by 3.71×, 2.89×, and 36.64× on average across all stencils that fit inside L2, LLC, and DRAM

respectively. However, in all stencil kernels, Casper provides higher performance-per-area (up to 190× compared to GPU). The reason for this large performance-per-area improvement is that 16 SPUs occupy 349× less area than the Titan V, i.e., $16 \times 0.146 \text{ mm}^2$ (see Section VIII-F) versus 815 mm^2). We observe that the L2- and LLC-sized data sets achieve performance-per-area improvements of 47× and 60×, respectively. For these data set sizes, Casper has the advantage of its tight integration into the large LLC. At the same time, the data does not fit into the GPU caches. For the large DRAM-sized data sets, however, the average improvement in performance/area is only 4.78×. In this case, GPU improves relative performance due to higher main memory bandwidth.




FIGURE 12. Performance/area compared to an NVIDIA Titan V GPU for three different domain sizes.

D. COMPARISON WITH PIMS

PIMS [34] proposes a PnM accelerator targeting stencil operations. PIMS accelerator leverages the atomic operations available in the HMC architecture to compute addition instructions in a stencil. Since PIMS represents the closest related work to Casper, we compare both accelerators using our stencil kernels. To evaluate the performance of PIMS, we conservatively consider only the latency of executing the atomic add operations, without accounting for the extra overhead of (1) loading the results back from the HMC device, and (2) executing the multiply operations required by each stencil on the host CPU. We use as a reference in our analysis the peak throughput of the HMC atomic operations reported by prior work [157].

Figure 13 shows the speedup of Casper in comparison to PIMS. We make the following observations. First, Casper provides an average speedup of 5.5× (up to 10×) compared against PIMS, for the data set sizes that fit inside the on-chip caches. This happens because of the low throughput of the atomic operations HMC provides, which becomes the main bottleneck. Second, the stencils that do not fit into the caches perform worse using Casper compared to PIMS. We attribute this to the fact that computing on the logic layer of the HMC has much higher memory bandwidth than the off-chip memory bus connected to the CPU for these bandwidth-bound stencils. We conclude that Casper performs significantly better than PIMS on typical stencil datasets, i.e., those that fit within the LLC.

E. EFFECT OF INDIVIDUAL OPTIMIZATIONS

The SPUs in Casper take advantage of two key optimizations: (1) a custom *data mapping* in LLC, which improves




FIGURE 13. Speedup compared to PIMS [34].

performance by increasing the locality of stencil data in the LLC and reducing the need for SPUs to access remote LLC slices, and (2) the *near-cache* (near-LLC) location of the SPUs, which minimizes data access latency and leverages the peak bandwidth of the LLC. In this section, we evaluate the contribution of each of these two optimizations to the overall performance of Casper. The baseline for this analysis is a system where the SPUs are located next to the private L1 caches of CPU cores. The baseline LLC data mapping conventionally places consecutive cache lines in consecutive LLC slices (similar to [158]). First, we apply only the data mapping optimization and compare the Casper against the baseline. Next, we apply both the data mapping optimization and the near-cache optimization. Then, we normalize to 100% the speedup that results from the two optimizations together, and obtain the percentage that comes from the data mapping and from the near-cache location. Figure 14 shows bars with blue and green parts. The blue part represents the percentage of the speedup due to the data mapping, while the green part represents the percentage of the speedup due to the near-cache location.

We make two observations from the results in Figure 14. First, computing near-cache (green portion of the bars) is the major contributor to the speedup. Second, the custom data mapping (blue portion of the bars) produces up to 30% of the speedup (Jacobi 1D with LLC data set), but its effect is negligible (or even negative) in several cases (1D and 2D benchmarks with L2/DRAM data sets, 7-point 3D with LLC data set). In these cases, the custom data mapping results in a number of accesses to remote LLC slices which is similar to the baseline data mapping.




FIGURE 14. Contribution of custom data mapping and near-cache SPU location to the speedup of Casper over a baseline system with SPUs near L1 for L2-/LLC-/DRAM-sized data sets.

F. HARDWARE COST

1) STENCIL PROCESSING UNIT

The area of one SPU scaled to 22 nm technology is 0.146 mm². The most significant contributors to this area

are the execution unit and the SRAM array used to buffer complete memory requests.

2) UNALIGNED LOADS

Our hardware mechanism to support unaligned loads consumes an additional 0.14 mm² compared to the baseline 2 MB SRAM cache slice. This amounts to a 5% increase in area per LLC slice. Almost the complete area overhead is attributed to the second read port of the tag array, which consumes 0.12 mm² of space. The remaining hardware overhead of one 3:1 multiplexer per SRAM row, the barrel shifter to rotate the final output, and the split multiplexers for way selection are minimal compared to the tag array overhead.

3) ADDRESS TO LLC SLICE MAPPING

Identifying the stencil segment requires two registers to store the start and the length of the segment. The address comparison needs one adder and one comparator. The new mapping is a simple bit-select from the physical address, and thus requires minimal additions. This hardware is introduced at every NoC injection point.

In summary, the hardware additions proposed in this paper require an additional 4.65 mm² of die area for a system using 16 SPUs. This amounts to a 0.77% area increase compared to the Marvell ThunderX 2 CPU [127], a server-class ARM CPU implemented in 16 nm hosting 32 MB of LLC.

IX. DISCUSSION

1) WHY A STENCIL ACCELERATOR?

Because of their large contribution to the overall runtime of HPC workloads, improving the performance and energy-efficiency of stencil computations is critical. A wide body of research [1], [3], [9], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38] highlights the need for highly efficient stencil computations. While more general-purpose solutions based on GPUs, FPGAs, and 3D-stacked memory attempt to trade off generality for performance and efficiency, we show that a careful domain-specific hardware/software co-design can improve the performance and energy efficiency even further, at a much lower overhead compared to the existing general-purpose solutions.

2) OTHER WORKLOADS FOR Casper

Apart from stencils, the near-LLC execution model is well-suited for applications with the following properties: (1) their memory access pattern shows temporal locality, (2) they operate on datasets that exceed the capacity of private caches, and (3) have streaming memory access pattern, and as a result do not benefit from deep cache hierarchies. Two examples of workloads that satisfy these properties are high performance computing (HPC) workloads operating on structured grids [176], [177], and dense linear algebra computations [178]. While we study our proposal specifically

for stencils, which are one of the most widely used kernels in HPC domain, supporting a wider set of applications is possible by redesigning the SPU pipeline to add support for data-dependent divisions that are present in some other HPC workloads. Together with the MAC operations (that Casper already supports), this extends Casper to a wider set of use-cases and applications.

X. RELATED WORK

To our knowledge, we present the first work that tightly integrates specialized computation units into the last-level cache of a CPU to perform stencil computations. In this section, we succinctly compare prior works against Casper.

Due to the high contribution of stencil computation to the overall execution time of HPC applications, a wide body of research has focused on studying and analyzing stencil computations [1], [3], [9], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38]. Prior works show the applicability of four different processing types in accelerating stencil computations: (1) Near-memory, (2) CPU, (3) GPU, (4) FPGA.

1) NEAR-MEMORY

PIMS [34] exploits the high-bandwidth provided by 3D-stacked memories (e.g., HMC [155], Hybrid Bandwidth Memory (HBM) [179], [180]) to accelerate stencils. Casper, being a near-LLC accelerator, can be integrated with any commodity processor without requiring costly interfacing using through-silicon vias.

2) CPU

Szustak et al. [181] accelerate the MPDATA stencil kernel on a multi-core CPU. Thaler et al. [118] use a many-core system to accelerate weather stencil kernels. Szustak and Bratek [46] propose parametric optimization techniques for the MPDATA application on shared-memory systems.

3) GPU

GPUs [3], [182], [183] have been shown to increase performance due to the high degree of parallelism present in the stencil computation. Pattnaik, et al. [184] develop an analytical performance model for choosing an optimal GPU-based execution strategy for various scientific stencil kernels. Gysi et al. [1] provide guidelines for optimizing stencil kernels for CPU–GPU systems.

4) FPGA

More recently, the use of FPGAs to accelerate stencils has been proposed [29], [30], [31], [32], [33], [47], [185], [186], [187], [188], [189], [190]. Augmenting general-purpose cores with specialized FPGA accelerators is a promising approach to enhance overall system performance. However, designing and deploying FPGA-based stencil accelerators have three inherent drawbacks compared to integrating Casper's SPUs near LLC and programming them. First, data

needs to be moved to these off-chip external FPGA-based accelerators. The fraction of the total execution time needed for this data movement is not negligible, and it may become larger if the entire stencil data does not fit in the limited FPGA memory (typically smaller than the host main memory). Second, taking full advantage of FPGAs for accelerating a workload is not a trivial task, as this requires sufficient FPGA programming skills to map the workload and optimize the design for the FPGA microarchitecture. Third, bitstream generation is a very time-consuming process, especially for high-end FPGAs. In contrast, Casper integrates compute units close to the LLC, which avoids unnecessary data movement to an external accelerator. Casper programming is easier and faster than FPGA programming, even if high-level synthesis tools (e.g., OpenCL [191], [192]) are used, because Casper only needs a small number of API functions (Table 1) and does not require time-consuming bitstream generation process.

XI. CONCLUSION

We present Casper, the first near-cache acceleration mechanism for stencil computations. Casper enables high performance and energy-efficient execution of stencil computations by (1) placing throughput-optimized stencil processing units near the last-level cache (LLC) and eliminating the need to move data to the processor, (2) orchestrating data accesses to minimize data movement within the cache hierarchy, and (3) maximizing the utilization of the LLC bandwidth. Casper achieves this with an area overhead of less than 1% for 16 SPUs in a Marvell ThunderX 2 [127], a server-class ARM CPU. We evaluate Casper using six widely used stencil kernels with up-to 3-dimensional grid domains. We show that Casper outperforms a commercial multi-core CPU, on average, by $1.65\times$ (up to $4.16\times$) and reduces the energy consumption by 35% (up to 65%). Compared to a state-of-the-art GPU, Casper improves performance-per-area, on average, by $37\times$ (up to $190\times$). We conclude that Casper is a promising near-cache-processing mechanism for accelerating stencil computations and addressing the memory bottleneck for such computations. We believe and hope that future work builds on Casper to further ease accelerating important high-performance computing applications that use stencil computations.

APPENDIX

This appendix presents some detailed measurements, which correspond to our analyses in Section VIII. Table 4 shows the dynamic instruction count for all evaluated stencils and datasets on the baseline CPU (16 cores) and Casper (16 SPUs). Table 5 shows the number of execution cycles for all evaluated stencils and datasets on the baseline CPU (16 cores), the baseline GPU, and Casper (16 SPUs). Table 6 shows the energy consumption for all evaluated stencils and datasets on the baseline CPU (16 cores) and Casper (16 SPUs).

TABLE 4. Dynamic Instruction Count for the Baseline CPU (16 cores) and Casper (16 SPUs).

	Jacobi 1D			7-point 1D			Jacobi 2D			Blur 2D			7-point 3D			33-point 3D		
	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM
CPU (16 cores)	165840	1312867	5245651	297277	2361924	9440116	537100	4311784	17255191	1804260	16552680	66329169	736767	6083864	24330380	2452622	20958248	83845023
Casper (16 SPUs)	3106	23038	3034882	26470	211402	3422962	5482	186718	12640918	38350	337858	4135498	20002	198730	21826798	261562	1050790	9321778

TABLE 5. Execution Cycles for the Baseline CPU (16 cores), the Baseline GPU, and Casper (16 SPUs).

	Jacobi 1D			7-point 1D			Jacobi 2D			Blur 2D			7-point 3D			33-point 3D		
	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM
CPU (16 cores)	13358	95251	3838447	14702	125138	5715526	26457	178032	8720011	95428	742734	22729495	39029	296436	7986968	115884	1009021	9060219
GPU	4030	36134	135360	4108	36594	139320	4646	37248	140160	6950	41318	153480	5184	36633	140856	6758	52491	278784
Casper (16 SPUs)	4569	33220	4370993	8449	66393	4514872	7658	58734	3931701	55764	446300	5454431	29572	286675	6784185	100243	1385955	13420984

TABLE 6. Energy Consumption (J) for the Baseline CPU (16 cores) and Casper (16 SPUs).

	Jacobi 1D			7-point 1D			Jacobi 2D			Blur 2D			7-point 3D			33-point 3D		
	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM	L2	LLC	DRAM
CPU (16 cores)	0.00012	0.00113	0.2631221	0.000144	0.00145	0.28253	0.000256	0.002	0.3483945	0.0009	0.0075	0.64639877	0.000386	0.003364	0.469465	0.0011542	0.010266	0.4424779
Casper (16 SPUs)	0.000468	0.00341	0.3114322	0.000629	0.00469	0.59888	0.00073	0.0055	0.8809648	0.0015	0.0118	1.19655244	0.001737	0.014002	1.4752518	0.0028739	0.027749	1.8090142

ACKNOWLEDGMENT

The authors would like to thank the SAFARI Research Group members for valuable feedback and the stimulating intellectual environment they provide. They also acknowledge support from the SAFARI Research Group's industrial partners, especially ASML, Google, Huawei, Intel, Microsoft, VMware, Xilinx, and Semiconductor Research Corporation. An earlier version of this work [193] was posted on arxiv.org (<https://arxiv.org/abs/2112.14216>) on December 28, 2021.

REFERENCES

- [1] T. Gysi, T. Grosser, and T. Hoefler, "MODESTO: Data-centric analytic optimization of complex stencil programs on heterogeneous architectures," in *Proc. SC*, Jun. 2015, pp. 177–186.
- [2] P. Colella. (2004). *Defining Software Requirements for Scientific Computing*. [Online]. Available: <https://www.krellinst.org/doecsgf/conf/2013/pres/pcollella.pdf>
- [3] O. Fuhrer, T. Chadha, T. Hoefler, G. Kwasniewski, X. Lapillonne, D. Leutwyler, D. Lüthi, C. Osuna, C. Schär, T. C. Schulthess, and H. Vogt, "Near-global climate simulation at 1 km resolution: Establishing a performance baseline on 4888 GPUs with COSMO 5.0," *Geosci. Model Develop.*, vol. 11, no. 4, pp. 1665–1681, May 2018.
- [4] G. A. McMechan, "Migration by extrapolation of time-dependent boundary values," *Geophys. Prospecting*, vol. 31, no. 3, pp. 413–420, Jun. 1983.
- [5] J. D. Anderson and J. Wendt, *Computational Fluid Dynamics an Introduction*, vol. 206. 1995.
- [6] A. Taflove, "Review of the formulation and applications of the finite-difference time-domain method for numerical modeling of electromagnetic wave interactions with arbitrary structures," *Wave Motion*, vol. 10, no. 6, pp. 547–582, Dec. 1988.
- [7] N. Maruyama and T. Aoki, "Optimizing stencil computations for NVIDIA Kepler GPUs," in *Proc. HiStencils*, 2014, pp. 89–95.
- [8] O. Fuhrer, C. Osuna, X. Lapillonne, T. Gysi, M. Bianco, and T. Schulthess, "Towards GPU-accelerated operational weather forecasting," in *Proc. GTC*, 2013, pp. 18–21.
- [9] C. Olschanowsky, M. M. Strout, S. Guzik, J. Loffeld, and J. Hittinger, "A study on balancing parallelism, data locality, and recomputation in existing PDE solvers," in *Proc. SC*, Nov. 2014, pp. 793–804.
- [10] *COSMO Physics Modules are Accelerated as a Part of HP2C Initiative Using OpenACC*, OpenACC-Toolkit Documentation, 2016. [Online]. Available: https://developer.download.nvidia.com/compute/OpenACC-Toolkit/docs/215493.7_OpenACC_COSMO_SS_v3.pdf
- [11] B. Mostafazadeh, F. Marti, F. Liu, and A. Chandramowlishwaran, "Roofline guided design and analysis of a multi-stencil CFD solver for multicore performance," in *Proc. IPDPS*, May 2018, pp. 753–762.
- [12] G. Doms and U. Schättler, "The nonhydrostatic limited-area model LM (lokal-model) of the DWD. Part I: Scientific documentation," DWD, GB Forschung und Entwicklung, 1999.
- [13] M. Christen, O. Schenk, and H. Burkhart, "PATUS: A code generation and autotuning framework for parallel iterative stencil computations on modern microarchitectures," in *Proc. IPDPS*, May 2011, pp. 676–687.
- [14] K. Datta, S. Kamil, S. Williams, L. Oliker, J. Shalf, and K. Yelick, "Optimization and performance modeling of stencil computations on modern microprocessors," *SIAM Rev.*, vol. 51, no. 1, pp. 129–159, Feb. 2009.
- [15] R. Strzodka, M. Shaheen, D. Pajak, and H.-P. Seidel, "Cache oblivious parallelograms in iterative stencil computations," in *Proc. SC*, Jun. 2010, pp. 49–59.
- [16] Y. Tang, R. A. Chowdhury, B. C. Kuzmaul, C.-K. Luk, and C. E. Leiserson, "The pochoir stencil compiler," in *Proc. SPAA*, Jun. 2011, pp. 117–128.
- [17] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, "Halide: A language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines," in *Proc. PLDI*, Jun. 2013, pp. 519–530.
- [18] J. Meng and K. Skadron, "A performance study for iterative stencil loops on GPUs with ghost zone optimizations," *Int. J. Parallel Program.*, vol. 39, no. 1, pp. 115–142, Feb. 2011.
- [19] A. Nguyen, N. Satish, J. Chhugani, C. Kim, and P. Dubey, "3.5-D blocking optimization for stencil computations on modern CPUs and GPUs," in *Proc. SC*, Nov. 2010, pp. 1–13.
- [20] T. Henretty, K. Stock, L.-N. Pouchet, F. Franchetti, J. Ramanujam, and P. Sadayappan, "Data layout transformation for stencil computations on short-vector SIMD architectures," in *Proc. CC*, 2011, pp. 225–245.
- [21] J. Jaeger and D. Barthou, "Automatic efficient data layout for multi-threaded stencil codes on CPU and GPUs," in *Proc. HiPC*, Dec. 2012, pp. 1–10.
- [22] M. Frigo and V. Strumpen, "The memory behavior of cache oblivious stencil computations," *J. Supercomput.*, vol. 39, no. 2, pp. 93–112, Mar. 2007.
- [23] S. Kamil, P. Husbands, L. Oliker, J. Shalf, and K. Yelick, "Impact of modern memory subsystems on cache optimizations for stencil computations," in *Proc. MSP*, 2005, pp. 36–43.
- [24] H. Stengel, J. Treibig, G. Hager, and G. Wellein, "Quantifying performance bottlenecks of stencil computations using the execution-cache-memory model," in *Proc. ICS*, Jun. 2015, pp. 207–216.
- [25] T. Brandvik and G. Pullan, "SBLOCK: A framework for efficient stencil-based PDE solvers on multi-core platforms," in *Proc. ICCIT*, Jun. 2010, pp. 1181–1188.
- [26] A. Armejach, H. Caminal, J. M. Cebrian, R. González-Alberquilla, C. Adeniyi-Jones, M. Valero, M. Casas, and M. Moretó, "Stencil codes on a vector length agnostic architecture," in *Proc. PACT*, Nov. 2018, pp. 1–12.
- [27] H. M. Waidyasooriya, Y. Takei, S. Tatsumi, and M. Hariyama, "OpenCL-based FPGA-platform for stencil computation and its optimization methodology," *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 5, pp. 1390–1402, May 2017.
- [28] J. de Fine Licht, M. Blott, and T. Hoefler, "Designing scalable FPGA architectures using high-level synthesis," in *Proc. PPOPP*, Feb. 2018, pp. 403–404.

- [29] K. Sano, Y. Hatsuda, and S. Yamamoto, "Multi-FPGA accelerator for scalable stencil computation with constant memory bandwidth," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 3, pp. 695–705, Mar. 2014.
- [30] J. V. Lunteren, R. Luijten, D. Diamantopoulos, F. Auernhammer, C. Hagleitner, L. Chelini, S. Corda, and G. Singh, "Coherently attached programmable near-memory acceleration platform and its application to stencil processing," in *Proc. DATE*, Mar. 2019, pp. 668–673.
- [31] G. Singh, D. Diamantopoulos, C. Hagleitner, S. Stuijk, and H. Corporaal, "NARMADA: Near-memory horizontal diffusion accelerator for scalable stencil computations," in *Proc. FPL*, Sep. 2019, pp. 263–269.
- [32] Y. Chi, J. Cong, P. Wei, and P. Zhou, "SODA: Stencil with optimized dataflow architecture," in *Proc. ICCAD*, Nov. 2018, pp. 1–8.
- [33] G. Singh, D. Diamantopoulos, C. Hagleitner, J. Gómez-Luna, S. Stuijk, O. Mutlu, and H. Corporaal, "NERO: A near high-bandwidth memory stencil accelerator for weather prediction modeling," in *Proc. FPL*, Aug. 2020, pp. 9–17.
- [34] 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 *Proc. ISMS*, 2019, pp. 41–52.
- [35] H. M. Waidyasooriya and M. Hariyama, "Multi-FPGA accelerator architecture for stencil computation exploiting spatial and temporal scalability," *IEEE Access*, vol. 7, pp. 53188–53201, 2019.
- [36] R. Cattaneo, G. Natale, C. Sicignano, D. Sciuto, and M. D. Santambrogio, "On how to accelerate iterative stencil loops: A scalable streaming-based approach," *ACM Trans. Archit. Code Optim.*, vol. 12, no. 4, pp. 1–26, Jan. 2016.
- [37] H. E. Yantur, A. M. Eltawil, and K. N. Salama, "Efficient acceleration of stencil applications through in-memory computing," *Micromachines*, vol. 11, no. 6, p. 622, Jun. 2020.
- [38] R. Wester and J. Kuper, "Deriving stencil hardware accelerators from a single higher-order function," in *Proc. CPA*, 2014, pp. 1–14.
- [39] S. Williams, A. Waterman, and D. Patterson, "Roofline: An insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, pp. 65–76, 2009.
- [40] Intel. (2016). *Intel Xeon Processor E7-4850 V4*. [Online]. Available: <https://ark.intel.com/content/www/us/en/ark/products/93806/intel-xeon-processor-e74850-v4-40m-cache-2-10-ghz.html>
- [41] Intel. *APP Metrics for Intel Microprocessors*. Accessed: Dec. 1, 2022. [Online]. Available: <https://www.intel.com/content/dam/support/us/en/documents/processors/APP-for-Intel-Xeon-Processors.pdf>
- [42] W. Augustin, V. Heuveline, and J.-P. Weiss, "Optimized stencil computation using in-place calculation on modern multicore systems," in *Proc. ECPP*, 2009, pp. 772–784.
- [43] K. Datta, M. Murphy, V. Volkov, S. Williams, J. Carter, L. Oliker, D. Patterson, J. Shalf, and K. Yelick, "Stencil computation optimization and auto-tuning on state-of-the-art multicore architectures," in *Proc. SC*, Nov. 2008, pp. 1–12.
- [44] E. H. Phillips and M. Fatica, "Implementing the Himeno benchmark with CUDA on GPU clusters," in *Proc. IPDPS*, 2010, pp. 1–10.
- [45] G. Singh, D. Diamantopoulos, S. Stuijk, C. Hagleitner, and H. Corporaal, "Low precision processing for high order stencil computations," in *Proc. 19th Int. Conf. Embedded Comput. Syst., Archit., Modeling, Simulation (SAMOS)*. Samos, Greece: Springer, Jul. 2019, pp. 403–415.
- [46] L. Szustak and P. Bratek, "Performance portable parallel programming of heterogeneous stencils across shared-memory platforms with modern Intel processors," *Int. J. High Perform. Comput. Appl.*, vol. 33, no. 3, pp. 534–553, May 2019.
- [47] G. Singh, D. Diamantopoulos, J. Gómez-Luna, C. Hagleitner, S. Stuijk, H. Corporaal, and O. Mutlu, "Accelerating weather prediction using near-memory reconfigurable fabric," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 15, no. 4, pp. 1–27, Dec. 2022.
- [48] C. Eckert, X. Wang, J. Wang, A. Subramaniyan, R. Iyer, D. Sylvester, D. Blaauw, and R. Das, "Neural cache: Bit-serial in-cache acceleration of deep neural networks," in *Proc. ISCA*, Jun. 2018, pp. 383–396.
- [49] 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 *Proc. MICRO*, Oct. 2017, pp. 273–287.
- [50] J. Ahn, S. Yoo, O. Mutlu, and K. Choi, "PIM-enabled instructions: A low-overhead, locality-aware processing-in-memory architecture," in *Proc. ISCA*, Jun. 2015, pp. 336–348.
- [51] M. Imani, S. Gupta, Y. Kim, and T. Rosing, "FloatPIM: In-memory acceleration of deep neural network training with high precision," in *Proc. ISCA*, Jun. 2019, pp. 802–815.
- [52] 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*. Singapore: Springer, 2022, pp. 171–243.
- [53] S. Ghose, A. Boroumand, J. S. Kim, J. Gómez-Luna, and O. Mutlu, "Processing-in-memory: A workload-driven perspective," *IBM J. Res. Develop.*, vol. 63, no. 6, pp. 3:1–3:19, Nov. 2019.
- [54] J. Gómez-Luna, I. El 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," 2021, *arXiv:2105.03814*.
- [55] J. Gomez-Luna, I. E. Hajj, I. Fernandez, C. Giannoula, G. F. Oliveira, and O. Mutlu, "Benchmarking a new paradigm: Experimental analysis and characterization of a real processing-in-memory system," *IEEE Access*, vol. 10, pp. 52565–52608, 2022.
- [56] G. F. Oliveira, J. Gómez-Luna, L. Orosa, S. Ghose, N. Vijaykumar, I. Fernandez, M. Sadrosadati, and O. Mutlu, "DAMOV: A new methodology and benchmark suite for evaluating data movement bottlenecks," *IEEE Access*, vol. 9, pp. 134457–134502, 2021.
- [57] S. Ghose, K. Hsieh, A. Boroumand, R. Ausavarungnirun, and O. Mutlu, "Enabling the adoption of processing-in-memory: Challenges, mechanisms, future research directions," 2018, *arXiv:1802.00320*.
- [58] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "Processing data where it makes sense: Enabling in-memory computation," *Microprocessors Microsyst.*, vol. 67, pp. 28–41, Jun. 2019.
- [59] D. Fujiki, S. Mahlke, and R. Das, "Duality cache for data parallel acceleration," in *Proc. ISCA*, Jun. 2019, pp. 397–410.
- [60] S. Aga, S. Jeloka, A. Subramaniyan, S. Narayanasamy, D. Blaauw, and R. Das, "Compute caches," in *Proc. HPCA*, Feb. 2017, pp. 481–492.
- [61] E. Lockerman, A. Feldmann, M. Bakhshalipour, A. Stanescu, S. Gupta, D. Sanchez, and N. Beckmann, "Livia: Data-centric computing throughout the memory hierarchy," in *Proc. ASPLOS*, Mar. 2020, pp. 417–433.
- [62] A. Subramaniyan, J. Wang, E. R. M. Balasubramanian, D. Blaauw, D. Sylvester, and R. Das, "Cache automaton," in *Proc. MICRO*, Oct. 2017, pp. 259–272.
- [63] A. Nag, C. N. Ramachandra, R. Balasubramonian, R. Stutsman, E. Giacomini, H. Kambalashubramanyam, and P.-E. Gaillardon, "Gen-Cache: Leveraging in-cache operators for efficient sequence alignment," in *Proc. MICRO*, Oct. 2019, pp. 334–346.
- [64] A. V. Nori, R. Bera, S. Balachandran, J. Rakshit, O. J. Omer, A. Abuhatzera, B. Kuttanna, and S. Subramoney, "REDUCT: Keep it close, keep it cool! Efficient scaling of DNN inference on multicore CPUs with near-cache compute," in *Proc. ISCA*, Jun. 2021, pp. 167–180.
- [65] A. Pattnaik, X. Tang, O. Kayiran, A. Jog, A. Mishra, M. T. Kandemir, A. Sivasubramanian, and C. R. Das, "Opportunistic computing in GPU architectures," in *Proc. ISCA*, Jun. 2019, pp. 210–223.
- [66] J. Ahn, S. Hong, S. Yoo, O. Mutlu, and K. Choi, "A scalable processing-in-memory accelerator for parallel graph processing," in *Proc. ISCA*, 2015, pp. 105–117.
- [67] 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*, vol. 19, no. S2, pp. 23–40, May 2018.
- [68] R. R. Puli, "Active routing: Compute on the way for near-data processing," Ph.D. dissertation, 2018.
- [69] L. Nai, R. Hadidi, J. Sim, H. Kim, P. Kumar, and H. Kim, "GraphPIM: Enabling instruction-level PIM offloading in graph computing frameworks," in *Proc. HPCA*, Feb. 2017, pp. 457–468.
- [70] M. Gao, J. Pu, X. Yang, M. Horowitz, and C. Kozyrakis, "TETRIS: Scalable and efficient neural network acceleration with 3D memory," in *Proc. ASPLOS*, 2017, pp. 751–764.
- [71] 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, "SIM-DRAM: A framework for bit-serial SIMD processing using DRAM," in *Proc. ASPLOS*, Apr. 2021, pp. 329–345.
- [72] A. Boroumand, S. Ghose, M. Patel, H. Hassan, B. Lucia, R. Ausavarungnirun, K. Hsieh, N. Hajinazar, K. T. Malladi, H. Zheng, and O. Mutlu, "CoNDA: Efficient cache coherence support for near-data accelerators," in *Proc. ISCA*, Jun. 2019, pp. 629–642.


- [73] A. Boroumand, S. Ghose, M. Patel, H. Hassan, B. Lucia, K. Hsieh, K. T. Malladi, H. Zheng, and O. Mutlu, "LazyPIM: An efficient cache coherence mechanism for processing-in-memory," *IEEE Comput. Archit. Lett.*, vol. 16, no. 1, pp. 46–50, Jan. 2017.
- [74] M. Besta, R. Kanakagiri, G. Kwasiński, R. Ausavarungnirun, J. Beránek, K. Kanellopoulos, K. Janda, Z. Vonarburg-Shmaria, L. Gianinazzi, I. Stefan, J. G. Luna, J. Golinowski, M. Copik, L. Kapp-Schwoerer, S. Di Girolamo, N. Blach, M. Konieczny, O. Mutlu, and T. Hoefler, "SISA: Set-centric instruction set architecture for graph mining on processing-in-memory systems," in *Proc. MICRO*, Oct. 2021, pp. 282–297.
- [75] A. Boroumand, S. Ghose, Y. Kim, R. Ausavarungnirun, E. Shiu, R. Thakur, D. Kim, A. Kuusela, A. Knies, P. Ranganathan, and O. Mutlu, "Google workloads for consumer devices: Mitigating data movement bottlenecks," in *Proc. ASPLOS*, 2018, pp. 316–331.
- [76] 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 *Proc. ICCD*, Oct. 2020, pp. 120–129.
- [77] D. S. Cali, G. S. Kalsi, Z. Bingol, C. Firtina, L. Subramanian, J. S. Kim, R. Ausavarungnirun, M. Alser, J. Gomez-Luna, A. Boroumand, A. Nori, A. Scibisz, S. Subramoney, C. Alkan, S. Ghose, and O. Mutlu, "GenASM: A high-performance, low-power approximate string matching acceleration framework for genome sequence analysis," in *Proc. MICRO*, Oct. 2020, pp. 951–966.
- [78] M. Hashemi, Khubaib, E. Ebrahimi, O. Mutlu, and Y. N. Patt, "Accelerating dependent cache misses with an enhanced memory controller," in *Proc. ISCA*, Jun. 2016, pp. 444–455.
- [79] G. Singh, M. Alser, D. S. Cali, D. Diamantopoulos, J. Gómez-Luna, H. Corporaal, and O. Mutlu, "FPGA-based near-memory acceleration of modern data-intensive applications," *IEEE Micro*, vol. 41, no. 4, pp. 39–48, Jul. 2021.
- [80] G. Singh, J. Gómez-Luna, G. Mariani, G. F. Oliveira, S. Corda, S. Stuijk, O. Mutlu, and H. Corporaal, "NAPEL: Near-memory computing application performance prediction via ensemble learning," in *Proc. DAC*, Jun. 2019, pp. 1–6.
- [81] V. Seshadri and O. Mutlu, "Simple operations in memory to reduce data movement," in *Advances in Computers*. Amsterdam, The Netherlands: Elsevier, 2017.
- [82] 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 *Proc. HPCA*, Feb. 2021, pp. 263–276.
- [83] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "Enabling practical processing in and near memory for data-intensive computing," in *Proc. DAC*, Jun. 2019, pp. 1–4.
- [84] 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, pp. 31–44.
- [85] K. Hsieh, E. Ebrahim, 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 *Proc. ISCA*, Jun. 2016, pp. 204–216.
- [86] H. S. Stone, "A logic-in-memory computer," *IEEE Trans. Comput.*, vol. C-19, no. 1, pp. 73–78, Jan. 1970.
- [87] M. Gokhale, B. Holmes, and K. Iobst, "Processing in memory: The terasys massively parallel PIM array," *Computer*, vol. 28, no. 4, pp. 23–31, Apr. 1995.
- [88] D. Patterson, T. Anderson, N. Cardwell, R. Fromm, K. Keeton, C. Kozyrakis, R. Thomas, and K. Yelick, "A case for intelligent RAM," *IEEE Micro*, vol. 17, no. 2, pp. 34–44, Mar./Apr. 1997.
- [89] R. Nair et al., "Active memory cube: A processing-in-memory architecture for exascale systems," *IBM J. Res. Develop.*, vol. 59, nos. 2–3, pp. 17:1–17:14, 2015.
- [90] M. Oskin, F. T. Chong, and T. Sherwood, "Active pages: A computation model for intelligent memory," in *Proc. ISCA*, 1998, pp. 192–203.
- [91] Y. Kang, W. Huang, S.-M. Yoo, D. Keen, Z. Ge, V. Lam, P. Pattnaik, and J. Torrellas, "FlexRAM: Toward an advanced intelligent memory system," in *Proc. ICCD*, Sep. 2012, pp. 5–14.
- [92] J. Draper, J. Chame, M. Hall, C. Steele, T. Barrett, J. LaCoss, J. Granacki, J. Shin, C. Chen, C. W. Kang, I. Kim, and G. Daglikoca, "The architecture of the DIVA processing-in-memory chip," in *Proc. SC*, Jun. 2002, pp. 14–25.
- [93] D. G. Elliott, M. Stumm, W. M. Snelgrove, C. Cojocar, and R. McKenzie, "Computational RAM: Implementing processors in memory," *IEEE Design Test Comput.*, vol. 16, no. 1, pp. 32–41, Jan./Mar. 1999.
- [94] J. B. Brockman, S. Thoziyoor, S. K. Kuntz, and P. M. Kogge, "A low cost, multithreaded processing-in-memory system," in *Proc. WMPPI*, 2004, pp. 16–22.
- [95] A. Farmahini-Farahani, J. H. Ahn, K. Morrow, and N. S. Kim, "NDA: Near-DRAM acceleration architecture leveraging commodity DRAM devices and standard memory modules," in *Proc. HPCA*, Feb. 2015, pp. 283–295.
- [96] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay, "Neurocube: A programmable digital neuromorphic architecture with high-density 3D memory," in *Proc. ISCA*, Jun. 2016, pp. 380–392.
- [97] D. Zhang, N. Jayasena, A. Lyashevsky, J. L. Greathouse, L. Xu, and M. Ignatowski, "TOP-PIM: Throughput-oriented programmable processing in memory," in *Proc. HPDC*, Jun. 2014, pp. 85–98.
- [98] S. Li, D. Niu, K. T. Malladi, H. Zheng, B. Brennan, and Y. Xie, "DRISA: A DRAM-based reconfigurable in-situ accelerator," in *Proc. MICRO*, Oct. 2017, pp. 288–301.
- [99] Q. Deng, L. Jiang, Y. Zhang, M. Zhang, and J. Yang, "DrAcc: A DRAM based accelerator for accurate CNN inference," in *Proc. DAC*, Jun. 2018, pp. 1–6.
- [100] S. Angizi and D. Fan, "GraphiDe: A graph processing accelerator leveraging In-DRAM-computing," in *Proc. GLSVLSI*, May 2019, pp. 45–50.
- [101] G. F. Oliveira, P. C. Santos, M. A. Alves, and L. Carro, "NIM: An HMC-based machine for neuron computation," in *Proc. ARC*, 2017, pp. 28–35.
- [102] M. Drumond, A. Daglis, N. Mirzadeh, D. Ustiugov, J. Picorel, B. Falsafi, B. Grot, and D. Pnevmatikatos, "The Mondrian data engine," in *Proc. ISCA*, Jun. 2017, pp. 639–651.
- [103] 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 *Proc. DATE*, Mar. 2017, pp. 710–715.
- [104] P. Gu, X. Xie, Y. Ding, G. Chen, W. Zhang, D. Niu, and Y. Xie, "IPIM: Programmable in-memory image processing accelerator using near-bank architecture," in *Proc. ISCA*, May 2020, pp. 804–817.
- [105] V. Seshadri and O. Mutlu, "In-DRAM bulk bitwise execution engine," 2019, *arXiv:1905.09822*.
- [106] 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," 2016, *arXiv:1611.09988*.
- [107] 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," *IEEE Comput. Archit. Lett.*, vol. 14, no. 2, pp. 127–131, Jul. 2015.
- [108] A. Boroumand, S. Ghose, B. Akin, R. Narayanaswami, G. F. Oliveira, X. Ma, E. Shiu, and O. Mutlu, "Google neural network models for edge devices: Analyzing and mitigating machine learning inference bottlenecks," in *Proc. PACT*, 2021, pp. 159–172.
- [109] L. Yavits, R. Kaplan, and R. Ginosar, "GIRAF: General purpose in-storage resistive associative framework," *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 2, pp. 276–287, Feb. 2022.
- [110] J. Lee, H. Kim, S. Yoo, K. Choi, H. P. Hofstee, G.-J. Nam, M. R. Nutter, and D. Jasek, "ExtraV: Boosting graph processing near storage with a coherent accelerator," *Proc. VLDB Endowment*, vol. 10, no. 12, pp. 1706–1717, Aug. 2017.
- [111] S. Kang, J. An, J. Kim, and S.-W. Jun, "MithriLog: Near-storage accelerator for high-performance log analytics," in *Proc. MICRO*, Oct. 2021, pp. 434–448.
- [112] N. M. Ghiasi, J. Park, H. Mustafa, J. Kim, A. Olgun, A. Gollwitzer, D. S. Cali, C. Firtina, H. Mao, N. A. Alserr, R. Ausavarungnirun, N. Vijaykumar, M. Alser, and O. Mutlu, "GenStore: A high-performance in-storage processing system for genome sequence analysis," in *Proc. ASPLOS*, Feb. 2022, pp. 635–654.
- [113] K. Keeton, D. A. Patterson, and J. M. Hellerstein, "A case for intelligent disks (IDISKS)," *ACM SIGMOD Rec.*, vol. 27, no. 3, pp. 42–52, Sep. 1998.
- [114] A. Acharya, M. Uysal, and J. Saltz, "Active disks: Programming model, algorithms and evaluation," in *Proc. ASPLOS*, 1998, pp. 81–91.
- [115] 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 *Proc. ISCA*, Jun. 2016, pp. 14–26.

- [116] 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 *Proc. DAC*, Jun. 2016, pp. 1–6.
- [117] D. Fujiki, S. Mahlke, and R. Das, "In-memory data parallel processor," in *Proc. ASPLOS*, Mar. 2018, pp. 1–14.
- [118] F. Thaler, S. Moosbrugger, C. Osuna, M. Bianco, H. Vogt, A. Afanasyev, L. Mosimann, O. Fuhrer, T. C. Schulthess, and T. Hoefler, "Porting the COSMO weather model to manycore CPUs," in *Proc. PASC*, Jun. 2019, pp. 1–11.
- [119] Z. Wang and T. Nowatzki, "Stream-based memory access specialization for general purpose processors," in *Proc. ISCA*, Jun. 2019, pp. 736–749.
- [120] B. Khailany, W. J. Dally, U. J. Kapasi, P. Mattson, J. Namkoong, J. D. Owens, B. Towles, A. Chang, and S. Rixner, "Imagine: Media processing with streams," *IEEE Micro*, vol. 21, no. 2, pp. 35–46, Mar./Apr. 2001.
- [121] S. Ciricescu, R. Essick, B. Lucas, P. May, K. Moat, J. Norris, M. Schuette, and A. Saidi, "The reconfigurable streaming vector processor (RSVP)," in *Proc. MICRO*, 2003, pp. 141–150.
- [122] T. Nowatzki, V. Gangadhar, N. Ardalani, and K. Sankaralingam, "Stream-dataflow acceleration," in *Proc. ISCA*, Jun. 2017, pp. 416–429.
- [123] 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 *Proc. 9th Int. Symp. High-Perform. Comput. Archit. (HPCA)*, 2003, pp. 129–140.
- [124] O. Mutlu, J. Stark, C. Wilkerson, and Y. N. Patt, "Runahead execution: An effective alternative to large instruction windows," *IEEE Micro*, vol. 23, no. 6, pp. 20–25, Nov. 2003.
- [125] O. Mutlu, H. Kim, and Y. N. Patt, "Techniques for efficient processing in runahead execution engines," in *Proc. ISCA*, 2005, pp. 370–381.
- [126] H. R. Zohouri, A. Podobas, and S. Matsuoka, "Combined spatial and temporal blocking for high-performance stencil computation on FPGAs using OpenCL," in *Proc. ACM/SIGDA Int. Symp. Field-Program. Gate Arrays*, Feb. 2018, pp. 153–162.
- [127] Marvell. *ThunderX2 CPU*. Accessed: Dec. 1, 2022. [Online]. Available: <https://en.wikichip.org/wiki/cavium/thunderx2>
- [128] J. W. Demmel, *Applied Numerical Linear Algebra*. Philadelphia, PA, USA: SIAM, 1997.
- [129] A. Olgun, J. Gómez Luna, K. Kanellopoulos, B. Salami, H. Hassan, O. Ergin, and O. Mutlu, "PiDRAM: A holistic end-to-end FPGA-based framework for processing-in-DRAM," 2021, *arXiv:2111.00082*.
- [130] 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 *Proc. HPCA*, Feb. 2019, pp. 582–595.
- [131] 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 DRAM chips," in *Proc. ISCA*, Jun. 2021, pp. 944–957.
- [132] O. Mutlu, "Intelligent architectures for intelligent computing systems," in *Proc. DATE*, Feb. 2021, pp. 318–323.
- [133] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "RowClone: Fast and energy-efficient in-DRAM bulk data copy and initialization," in *Proc. 46th Annu. IEEE/ACM Int. Symp. Microarchitecture*, Dec. 2013, pp. 185–197.
- [134] 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 *Proc. HPCA*, Mar. 2016, pp. 568–580.
- [135] V. Seshadri and O. Mutlu, "The processing using memory paradigm: In-DRAM bulk copy, initialization, bitwise AND and OR," 2016, *arXiv:1610.09603*.
- [136] M. Hashemi, O. Mutlu, and Y. N. Patt, "Continuous runahead: Transparent hardware acceleration for memory intensive workloads," in *Proc. MICRO*, Oct. 2016, pp. 1–12.
- [137] 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 *Proc. HPEC*, Sep. 2013, pp. 1–6.
- [138] S. H. Pugsley, J. Jestes, H. Zhang, R. Balasubramonian, V. Srinivasan, A. Buyuktosunoglu, A. Davis, and F. Li, "NDC: Analyzing the impact of 3D-stacked memory+logic devices on MapReduce workloads," in *Proc. ISPASS*, Mar. 2014, pp. 190–200.
- [139] B. Akin, F. Franchetti, and J. C. Hoe, "Data reorganization in memory using 3D-stacked DRAM," in *Proc. ISCA*, Jun. 2015, pp. 131–143.
- [140] J. H. Lee, J. Sim, and H. Kim, "BSSync: Processing near memory for machine learning workloads with bounded staleness consistency models," in *Proc. PACT*, Oct. 2015, pp. 241–252.
- [141] B. Akin, F. Franchetti, and J. C. Hoe, "Data reorganization in memory using 3D-stacked DRAM," in *Proc. 42nd Annu. Int. Symp. Comput. Archit. (ISCA)*. New York, NY, USA: Association for Computing Machinery, 2015, pp. 131–143, doi: [10.1145/2749469.2750397](https://doi.org/10.1145/2749469.2750397).
- [142] 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 *Proc. ISCA*, Jun. 2016, pp. 27–39.
- [143] 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 *Proc. MICRO*, Dec. 2015, pp. 267–280.
- [144] Z. Liu, I. Calciu, M. Herlihy, and O. Mutlu, "Concurrent data structures for near-memory computing," in *Proc. SPAA*, Jul. 2017, pp. 235–245.
- [145] M. Gao, G. Ayers, and C. Kozyrakis, "Practical near-data processing for in-memory analytics frameworks," in *Proc. PACT*, Oct. 2015, pp. 113–124.
- [146] A. Morad, L. Yavits, and R. Ginosar, "GP-SIMD processing-in-memory," *ACM Trans. Archit. Code Optim.*, vol. 11, no. 4, pp. 53:1–53:26, Jan. 2015, doi: [10.1145/2686875](https://doi.org/10.1145/2686875).
- [147] G. Kim, N. Chatterjee, M. O'Connor, and K. Hsieh, "Toward standardized near-data processing with unrestricted data placement for GPUs," in *Proc. SC*, Nov. 2017, pp. 1–12.
- [148] S. L. Xi, A. Augusta, M. Athanassoulis, and S. Idreos, "Beyond the wall: Near-data processing for databases," in *Proc. 11th Int. Workshop Data Manage. New Hardw.*, May 2015, pp. 1–10.
- [149] M. Kang, M.-S. Keel, N. R. Shanbhag, S. Eilert, and K. Curewitz, "An energy-efficient VLSI architecture for pattern recognition via deep embedding of computation in SRAM," in *Proc. ICASSP*, May 2014, pp. 8326–8330.
- [150] 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 *Proc. 49th Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*. Piscataway, NJ, USA: IEEE Press, Oct. 2016, pp. 1–13.
- [151] 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," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 38, no. 4, pp. 640–653, Apr. 2019.
- [152] Y. Huang, L. Zheng, P. Yao, J. Zhao, X. Liao, H. Jin, and J. Xue, "A heterogeneous PIM hardware-software co-design for energy-efficient graph processing," in *Proc. IPDPS*, May 2020, pp. 684–695.
- [153] Y. Zhuo, C. Wang, M. Zhang, R. Wang, D. Niu, Y. Wang, and X. Qian, "GraphQ: Scalable PIM-based graph processing," in *Proc. MICRO*, Oct. 2019, pp. 712–725.
- [154] M. Alser, Z. Bingol, D. S. Cali, S. Kim, J. Ghose, C. Alkan, and O. Mutlu, "Accelerating genome analysis: A primer on an ongoing journey," *IEEE Micro*, vol. 40, no. 5, pp. 65–75, Sep./Oct. 2020.
- [155] Hybrid Memory Cube Consortium. *Hybrid Memory Cube Specification Rev. 2.0*. Accessed: Dec. 1, 2022. [Online]. Available: <http://www.hybridmemorycube.org/>
- [156] P. Rosenfeld, "Performance exploration of the hybrid memory cube," Ph.D. dissertation, 2014.
- [157] G. F. Oliveira, P. C. Santos, M. A. Z. Alves, and L. Carro, "A generic processing in memory cycle accurate simulator under hybrid memory cube architecture," in *Proc. SAMOS*, Jul. 2017, pp. 54–61.
- [158] Y. Yarom, Q. Ge, F. Liu, R. B. Lee, and G. Heiser, "Mapping the Intel last-level cache," *Cryptol. ePrint Arch.*, 2015.
- [159] A. Basu, J. Gandhi, J. Chang, M. D. Hill, and M. M. Swift, "Efficient virtual memory for big memory servers," in *Proc. ISCA*, Jun. 2013, pp. 237–248.
- [160] A. Saini, "Design of the Intel Pentium processor," in *Proc. ICCD*, 1993, pp. 258–261.
- [161] D. Jaggur, "Arm architecture and systems," *IEEE Micro*, vol. 17, no. 4, pp. 9–11, Jul. 1997.
- [162] A. S. Waterman, "Design of the RISC-V instruction set architecture," Univ. California, Berkeley, Berkeley, CA, USA, 2016.
- [163] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoabi, N. Vaish, M. D. Hill, and D. A. Wood, "The gem5 simulator," *ACM SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, May 2011.

- [164] J. Lowe-Power et al., “The gem5 simulator: Version 20.0+,” 2020, *arXiv:2007.03152*.
- [165] NVIDIA Corporation. *NVIDIA TITAN V*. Accessed: Dec. 1, 2022. [Online]. Available: <https://www.nvidia.com/en-us/titan/titan-v/>
- [166] R. Balasubramanian, A. B. Kahng, N. Muralimanohar, A. Shafiee, and V. Srinivas, “CACTI 7: New tools for interconnect exploration in innovative off-chip memories,” *ACM Trans. Archit. Code Optim.*, vol. 14, no. 2, pp. 1–25, Jun. 2017.
- [167] P.-A. Tsai, C. Chen, and D. Sanchez, “Adaptive scheduling for systems with asymmetric memory hierarchies,” in *Proc. MICRO*, Oct. 2018, pp. 641–654.
- [168] P.-A. Tsai, N. Beckmann, and D. Sanchez, “Jenga: Software-defined cache hierarchies,” in *Proc. ISCA*, Jun. 2017, pp. 652–665.
- [169] 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 *Proc. ISCA*, Jun. 2014, pp. 97–108.
- [170] S. Salehi and R. F. DeMara, “Energy and area analysis of a floating-point unit in 15 nm CMOS process technology,” in *Proc. SoutheastCon*, Apr. 2015, pp. 1–5.
- [171] NVIDIA Corporation. (2017). *The NVIDIA Titan V Preview*. [Online]. Available: <https://www.anandtech.com/show/12170/nvidia-titan-v-preview-titanomachy>
- [172] L.-N. Pouchet. *PolyBench: The Polyhedral Benchmark Suite*. Accessed: Dec. 1, 2022. [Online]. Available: <https://www.cs.colostate.edu/~pouchet/software/polybench/>
- [173] J. Canny, “A computational approach to edge detection,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.
- [174] J. Holewinski, L.-N. Pouchet, and P. Sadayappan, “High-performance code generation for stencil computations on GPU architectures,” in *Proc. SC*, Jun. 2012, pp. 311–320.
- [175] K. Datta and K. A. Yelick, *Auto-Tuning Stencil Codes for Cache-Based Multicore Platforms*. Berkeley, CA, USA: Univ. of California, Berkeley, 2009.
- [176] M. Schönherr, K. Kucher, M. Geier, M. Stiebler, S. Freudiger, and M. Krafczyk, “Multi-thread implementations of the lattice Boltzmann method on non-uniform grids for CPUs and GPUs,” *Comput. Math. Appl.*, vol. 61, no. 12, pp. 3730–3743, Jun. 2011.
- [177] G. Allen, T. Goodale, G. Lanfermann, T. Radke, E. Seidel, W. Benger, H. C. Hege, A. Merzky, J. Masso, and J. Shalf, “Solving Einstein’s equations on supercomputers,” *Computer*, vol. 32, no. 12, pp. 52–58, 1999.
- [178] H. Ltaief, P. Luszczek, and J. Dongarra, “Profiling high performance dense linear algebra algorithms on multicore architectures for power and energy efficiency,” *Comput. Sci., Res. Develop.*, vol. 27, no. 4, pp. 277–287, Nov. 2012.
- [179] 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, J. H. Cho, K. H. Kwon, M. J. Kim, J. Lee, K. W. Park, B. Chung, and S. Hong, “A 1.2 V 8 Gb 8-channel 128 GB/s high-bandwidth memory (HBM) stacked DRAM with effective microbump I/O test methods using 29 nm process and TSV,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 432–433.
- [180] D. Lee, S. Ghose, G. Pekhimenko, S. Khan, and O. Mutlu, “Simultaneous multi-layer access: Improving 3D-stacked memory bandwidth at low cost,” *ACM Trans. Archit. Code Optim.*, vol. 12, no. 4, pp. 1–29, Jan. 2016.
- [181] L. Szustak, K. Rojek, and P. Gepner, “Using Intel Xeon Phi coprocessor to accelerate computations in MPDATA algorithm,” in *Proc. PPAM*, 2013, pp. 582–592.
- [182] O. Anjum, G. D. G. Simon, M. Hidayetoglu, and W.-M. Hwu, “An efficient GPU implementation technique for higher-order 3D stencils,” in *Proc. HPCC*, Aug. 2019, pp. 552–561.
- [183] Q. Sun, Y. Liu, H. Yang, Z. Jiang, X. Liu, M. Dun, Z. Luan, and D. Qian, “CsTuner: Scalable auto-tuning framework for complex stencil computation on GPUs,” in *Proc. CLUSTER*, Sep. 2021, pp. 192–203.
- [184] M. Wahib and N. Maruyama, “Scalable kernel fusion for memory-bound GPU applications,” in *Proc. SC*, Nov. 2014, pp. 191–202.
- [185] J. de Fine Licht, A. Kuster, T. De Matteis, T. Ben-Nun, D. Hofer, and T. Hoefler, “StencilFlow: Mapping large stencil programs to distributed spatial computing systems,” in *Proc. CGO*, Feb. 2021, pp. 315–326.
- [186] A. Sohrabizadeh, C. H. Yu, M. Gao, and J. Cong, “AutoDSE: Enabling software programmers to design efficient FPGA accelerators,” *ACM Trans. Design Autom. Electron. Syst.*, vol. 27, no. 4, pp. 1–27, Jul. 2022.
- [187] S. Wang and Y. Liang, “A comprehensive framework for synthesizing stencil algorithms on FPGAs using OpenCL model,” in *Proc. DAC*, Jun. 2017, pp. 1–6.
- [188] E. Reggiani, E. Del Sozzo, D. Conficconi, G. Natale, C. Moroni, and M. D. Santambrogio, “Enhancing the scalability of multi-FPGA stencil computations via highly optimized HDL components,” *ACM Trans. Reconfigurable Technol. Syst.*, vol. 14, no. 3, pp. 1–33, Sep. 2021.
- [189] M. Koraei, O. Fatemi, and M. Jahre, “DCMI: A scalable strategy for accelerating iterative stencil loops on FPGAs,” *ACM Trans. Archit. Code Optim.*, vol. 16, no. 4, pp. 1–24, Dec. 2019.
- [190] X. Tian, Z. Ye, A. Lu, L. Guo, Y. Chi, and Z. Fang, “SASA: A scalable and automatic stencil acceleration framework for optimized hybrid spatial and temporal parallelism on HBM-based FPGAs,” 2022, *arXiv:2208.10770*.
- [191] S. Huang, L.-W. Chang, I. El Hajj, S. Garcia De Gonzalo, J. Gómez-Luna, S. R. Chalamalasetti, M. El-Hadedy, D. Milojicic, O. Mutlu, D. Chen, and W.-M. Hwu, “Analysis and modeling of collaborative execution strategies for heterogeneous CPU-FPGA architectures,” in *Proc. ICPE*, Apr. 2019, pp. 79–90.
- [192] J. Jiang, Z. Wang, X. Liu, J. Gómez-Luna, N. Guan, Q. Deng, W. Zhang, and O. Mutlu, “Boyi: A systematic framework for automatically deciding the right execution model of OpenCL applications on FPGAs,” in *Proc. FPGA*, Feb. 2020, pp. 299–309.
- [193] A. Denzler, R. Bera, N. Hajinazar, G. Singh, G. F. Oliveira, J. Gómez-Luna, and O. Mutlu, “Casper: Accelerating stencil computation using near-cache processing,” 2021, *arXiv:2112.14216*.



ALAIN DENZLER received the B.S. and M.S. degrees in computer science from ETH Zürich, in 2017 and 2020, respectively. He currently works as a Software Engineer with NVIDIA Switzerland.



GERALDO F. OLIVEIRA received the B.S. degree in computer science from the Federal University of Viçosa, Viçosa, Brazil, in 2015, and the M.S. degree in computer science from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil, in 2017. He is currently pursuing the Ph.D. degree with ETH Zürich, Zürich, Switzerland, under the supervision of Prof. Onur Mutlu. 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.



NASTARAN HAJINAZAR received the M.S. degree in computer hardware engineering from the Sharif University of Technology, Tehran, Iran, in 2011, and the Ph.D. degree in computer science from Simon Fraser University, British Columbia, Canada, in 2020. She is currently a Senior Researcher with ETH Zürich. Her research interests include several aspects of computer architecture with a significant focus on designing efficient high-performance computing systems, memory architectures, and intelligent memory management techniques.



RAHUL BERA received the master’s degree in computer science from the Indian Institute of Technology Kanpur, in 2017. He is currently pursuing the Ph.D. degree with ETH Zürich, Switzerland. He has worked with AMD and Intel Labs, India. His research interests include the broad areas of memory hierarchy design and applied machine learning in computer architecture.



GAGANDEEP SINGH received the joint M.Sc. degree (Hons.) in integrated circuit design from Technische Universität München (TUM), Germany, and Nanyang Technological University (NTU), Singapore, in 2017, and the Ph.D. degree from Technische Universiteit Eindhoven, The Netherlands, under the supervision of Prof. Henk Corporaal and Prof. Onur Mutlu, in 2021. From June 2018 to January 2020, he was a Predoctoral Researcher with IBM Research

Zürich, Switzerland. He has worked with Oracle, India, and also performed research with imec, Belgium. He is currently a Senior Researcher with ETH Zürich. His research interests include computer architecture, FPGA acceleration, processing-in-memory, bioinformatics, and machine learning.



JUAN GÓMEZ-LUNA (Member, IEEE) received the B.S. and M.S. degrees in telecommunication engineering from the University of Sevilla, Spain, in 2001, and the Ph.D. degree in computer science from the University of Córdoba, Spain, in 2012. From 2005 to 2017, he was a Faculty Member with the University of Córdoba. He is currently a Senior Researcher and a Lecturer with the SAFARI Research Group, ETH Zürich. 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. His research interests include processing-in-memory, memory systems, heterogeneous computing, and hardware and software acceleration of medical imaging and bioinformatics.



ONUR MUTLU (Fellow, IEEE) received the B.S. degree in computer engineering and psychology from the University of Michigan, Ann Arbor, and the M.S. and Ph.D. degrees in electrical and computer engineering from The University of Texas at Austin. He started the Computer Architecture Group, Microsoft Research (2006–2009) and has held various product and research positions with Intel Corporation, Advanced Micro Devices, VMware, and Google. He is currently a Professor of computer science with ETH Zürich. He is also a Faculty Member with Carnegie Mellon University, where he previously held the Strecker Early Career Professorship. A variety of techniques that he, along with his group and collaborators, has invented over the years have influenced the industry and have been employed in commercial microprocessors and memory/storage systems. His research interests include computer architecture, systems, hardware security, and bioinformatics. He is an ACM Fellow and an elected member of the Academy of Europe (Academia Europaea). He received the IEEE High-Performance Computer Architecture Test of Time Award, the IEEE Computer Society Edward J. McCluskey Technical Achievement Award, the ACM SIGARCH Maurice Wilkes Award, the inaugural IEEE Computer Society Young Computer Architect Award, the inaugural Intel Early Career Faculty Award, the U.S. National Science Foundation CAREER Award, the Carnegie Mellon University Ladd Research Award, the 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. 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/>.

...