

# FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

Gagandeep Singh<sup>◇</sup> Mohammed Alser<sup>◇</sup> Damla Senol Cali<sup>✕</sup>

Dionysios Diamantopoulos<sup>▽</sup> Juan Gómez-Luna<sup>◇</sup>

Henk Corporaal<sup>\*</sup> Onur Mutlu<sup>◇✕</sup>

<sup>◇</sup>ETH Zürich <sup>✕</sup>Carnegie Mellon University

<sup>\*</sup>Eindhoven University of Technology <sup>▽</sup>IBM Research Europe


**Abstract**—Modern data-intensive applications demand high computation capabilities with strict power constraints. Unfortunately, such applications suffer from a significant waste of both execution cycles and energy in current computing systems due to the costly data movement between the computation units and the memory units. Genome analysis and weather prediction are two examples of such applications. Recent FPGAs couple a reconfigurable fabric with high-bandwidth memory (HBM) to enable more efficient data movement and improve overall performance and energy efficiency. This trend is an example of a paradigm shift to *near-memory computing*. We leverage such an FPGA with high-bandwidth memory (HBM) for improving the pre-alignment filtering step of genome analysis and representative kernels from a weather prediction model. Our evaluation demonstrates large speedups and energy savings over a high-end IBM POWER9 system and a conventional FPGA board with DDR4 memory. We conclude that FPGA-based near-memory computing has the potential to alleviate the data movement bottleneck for modern data-intensive applications.

■ **MODERN COMPUTING SYSTEMS** suffer from a large gap between the performance and energy efficiency of computation units and memory units. These systems follow a *processor-centric* approach where data has to move back and forth from the memory units using a relatively slow and power-hungry off-chip bus to the computation units for processing. As a result, data-intensive workloads (e.g., genome analysis [1–15] and weather modeling [16–19]) require continuous memory-CPU-memory data movement, which imposes an extremely large overhead in terms of execution time and energy efficiency [20].

We provide in Figure 1 the roofline model [21] on an IBM POWER9 CPU (IC922) [22] for the state-of-the-art pre-alignment filtering algorithm

for genome analysis [23] and two compound stencil kernels from the widely-used COSMO (Consortium for Small-Scale Modeling) weather prediction model [16]. A key observation is that both applications have low arithmetic intensity with complex memory access behavior. The pre-alignment filtering algorithm, *SneakySnake*, builds a special matrix (called a chip maze in Section “Case Study 1: Pre-Alignment Filtering in Genome Analysis”) used to calculate an optimal solution for the pre-alignment filtering problem. *SneakySnake* calculates only portions of this chip maze to maintain speed. This involves irregular visits to different entries of the chip maze, leading to a strong mismatch between the nature of data access patterns and the layout

of data in memory for SneakySnake. Such a mismatch leads to limited spatial locality and cache effectiveness, causing frequent data movement between the memory subsystem and the processing units. The weather kernels (`vadv` and `hdiff`) consist of compound stencils that perform a series of element-wise computations on a three-dimensional grid [24]. Such compound kernels are dominated by DRAM-latency-bound operations due to complex memory access patterns. As a result, the performance of these applications is significantly lower than the peak CPU performance. This is a common trend in various data-intensive workloads [25–53].



**Figure 1: Roofline for POWER9 (1-socket) showing pre-alignment filtering algorithm (SneakySnake), and vertical advection (`vadv`) and horizontal diffusion (`hdiff`) kernels from the COSMO weather prediction model for single-thread and 64-thread implementations. The plot shows also the rooflines of the FPGAs used in our work with peak DRAM and on-chip BRAM bandwidth.**

In this work, our **goal** is to overcome the memory bottleneck of two key real-world data-intensive applications, genome analysis and weather modeling, by exploiting near-memory computation capability on modern FPGA accelerators with high-bandwidth memory (HBM) [54] that are attached to a host CPU. The use of FPGAs can yield significant performance improvements, especially for parallel algorithms. Modern FPGAs provide four key trends:

- 1) The integration of high-bandwidth memory (HBM) on the same package with an FPGA allows us to implement our accelerator logic much closer to the memory with an order of magnitude more bandwidth than traditional DDR4-based FPGA boards. Thus,

these modern FPGAs adopt a more *data-centric* approach to computing.

- 2) FPGA manufacturers have introduced UltraRAM (URAM) [55] along with the Block RAM (BRAM) that offers large on-chip memory next to the logic.
- 3) Recent FPGA boards with new cache-coherent interconnects (e.g., IBM Coherent Accelerator Processor Interface (CAPI) [56], Cache Coherent Interconnect for Accelerators (CCIX) [57], and Compute Express Link (CXL) [58]) employ a shared memory space that allows tight integration of FPGAs with CPUs at high bidirectional bandwidth (on the order of tens of GB/s). This integration allows the FPGA to coherently access the host system’s memory using a pointer, rather than requiring multiple copies of the data.
- 4) Newer FPGAs are manufactured with an advanced technology node of 7-14nm FinFET [59,60] that offers higher performance.

These four trends suggest that modern FPGA architectures deliver unprecedented levels of integration and compute capability due to new advances and features, which provides an opportunity to largely alleviate the *memory bottleneck* of real-world data-intensive applications.


**To this end**, we demonstrate the capability of near-HBM FPGA-based accelerators for two key real-world data-intensive applications: (1) pre-alignment filtering in genome analysis, (2) representative kernels from a widely-used weather prediction application, COSMO. Pre-alignment filtering is one of the fundamental steps in most genome analysis tasks, where up to 98% of input genomic data is filtered out. Thus, accelerating this step would benefit the overall end-to-end execution time of genome analysis [1,5,7,11,23,61,62]. The weather kernels we evaluate are an essential part of climate and weather modeling and prediction [17], which is critical for a sustainable life ecosystem [63].

Our accelerator designs make use of a heterogeneous memory hierarchy (consisting of URAM, BRAM, and HBM). We evaluate the performance and energy efficiency of our accelerators, perform a scalability analysis, and compare them to a traditional DDR4-based FPGA board and a state-of-the-art multi-core IBM POWER9 system. Based

on our analysis, we show that our full-blown HBM-based designs of *SneakySnake*, *vadv*, and *hdiff* provide (1) 27.4 $\times$ , 5.3 $\times$ , and 12.7 $\times$  higher speedup, and (2) 133 $\times$ , 12 $\times$ , and 35 $\times$  higher energy efficiency, respectively, compared to a 16-core IBM POWER9 system.

## Near-memory Computation on FPGAs

We provide in Figure 2 a high-level schematic of our integrated system with an FPGA-based near-memory accelerator. The FPGA is connected to two HBM stacks, each of which has 16 *pseudo memory channels* [64]. A channel is exposed to the FPGA as a 256-bit wide interface, and the FPGA has 32 such channels in total. The HBM IP provides 8 memory controllers (per stack) to handle the data transfer to/from the HBM memory channels. This configuration enables high-bandwidth and low-latency memory accesses for near-memory computing. The FPGA is also connected to a host CPU, an IBM POWER9 processor, using OCAPI (OpenCAPI) [65].



**Figure 2: Heterogeneous platform with an IBM POWER9 system connected to an HBM-based FPGA board via OCAPI. We also show components of an FPGA: flip-flop (FF), lookup table (LUT), UltraRAM (URAM), and Block RAM (BRAM).**

The FPGA device implements an *accelerator functional unit* (AFU) that interacts with the host system through the TLx (Transaction Layer) and the DLx (Data Link Layer), which are the CAPI endpoints on the FPGA. An AFU comprises multiple *processing elements* (PEs) that accelerate a portion of an application.

## Modern Data-Intensive Applications

### Case Study 1: Pre-Alignment Filtering in Genome Analysis

One of the most fundamental computational steps in most genome analysis tasks is sequence

alignment [7, 8]. This step is formulated as an *approximate string matching* (ASM) problem [11, 66] and it calculates: (1) *edit distance* (the minimum number of edits needed to convert one sequence into the other) between two given sequences [66, 67], (2) type of each edit (i.e., insertion, deletion, or substitution), (3) location of each edit in one of the two given sequences, and (4) *alignment score* that is the sum of the scores (calculated using a user-defined scoring function) of all edits and matches between the two sequences.

Sequence alignment is a computationally-expensive step as it usually uses dynamic programming (DP)-based algorithms [6, 11, 68–71], which have quadratic time and space complexity (i.e.,  $O(m^2)$  for a sequence length of  $m$ ). In genome analysis, an overwhelming majority (>98%) of the sequence pairs examined during sequence alignment are highly dissimilar and their alignment results are *simply* discarded as such dissimilar sequence pairs are usually not useful for genomic studies [61, 72, 73]. To avoid examining dissimilar sequences using computationally-expensive sequence alignment algorithms, genome analysis pipelines typically use filtering heuristics that are called *pre-alignment filters* [7, 23, 73–76]. The key idea of pre-alignment filtering is to quickly estimate the number of edits between two given sequences and use this estimation to decide whether or not the computationally-expensive DP-based alignment calculation is needed — if not, a significant amount of time is saved by avoiding DP-based alignment. If two genomic sequences differ by more than an edit distance threshold,  $E$ , then the two sequences are identified as dissimilar sequences and hence DP calculation is not needed.

*SneakySnake* [23] is a recent highly-parallel and highly-accurate pre-alignment filter that works on *modern* high-performance computing architectures such as CPUs, GPUs, and FPGAs. The key idea of *SneakySnake* is to reduce the ASM problem to the *single net routing* (SNR) problem [77] in VLSI. The goal of the SNR problem is to find the shortest routing path that interconnects two terminals on the boundaries of VLSI chip layout while passing through the minimum number of obstacles. Solving the SNR problem is faster than solving the ASM

problem, as calculating the routing path after facing an obstacle is independent of the calculated path before this obstacle (checkpoints in Figure 3). This provides two key benefits: 1) It obviates the need for using computationally-costly DP algorithms to keep track of the subpath that provides the optimal solution (i.e., the one with the least possible routing cost). 2) The independence of the subpaths allows for solving many SNR subproblems in parallel by judiciously leveraging the parallelism-friendly architecture of modern FPGAs and GPUs to greatly speed up the SneakySnake algorithm. The number of obstacles faced throughout the found routing path represents a *lower bound* on the edit distance between two sequences, and hence this number is used to make accurate filtering decisions by SneakySnake.




Figure 3: An example of the SneakySnake chip maze for a reference sequence  $R = \text{'GGTGCAGAGCTC'}$ , a query sequence  $Q = \text{'GGTGAGAGTTGT'}$ , and an edit distance threshold ( $E$ ) of 3. Our SneakySnake algorithm quickly finds an optimal signal net (highlighted in orange) with 3 obstacles, each of which is located at the end of each arrow (subpath), and hence SneakySnake decides that sequence alignment for  $R$  and  $Q$  is needed, as the number of obstacles  $\leq E$ .

The SneakySnake algorithm includes three main steps [23]: 1) **Building the chip maze**. The chip maze,  $Z$ , is a matrix where each of its entries represents the pairwise comparison result of a character of one sequence with another character of the other sequence, as we show in Figure 3. Given two genomic sequences of length  $m$ , a reference sequence  $R[1 \dots m]$  and a query sequence  $Q[1 \dots m]$ , and an edit distance threshold  $E$ , SneakySnake calculates the entry  $Z[i, j]$  of the chip maze as follows:

$$\begin{cases} 0, & \text{if } i = E + 1, Q[j] = R[j], \\ 0, & \text{if } 1 \leq i \leq E, Q[j - i] = R[j], \\ 0, & \text{if } i > E + 1, Q[j + i - E - 1] = R[j], \\ 1, & \text{otherwise} \end{cases} \quad (1)$$

where an entry of value '1' represents an obstacle, an entry of value '0' represents an available path,  $1 \leq i \leq (2E + 1)$ , and  $1 \leq j \leq m$ .

2) **Finding the longest available path in each row of the chip maze**. SneakySnake finds the longest available path, which represents the longest common subsequence between two given sequences. It counts the consecutive entries of value 0 in each row starting from the previous checkpoint until it faces an obstacle, and then it examines the next rows in the same way. Once it examines all rows, SneakySnake compares the lengths of the found segments of consecutive zeros and considers the longest segment as the chosen path (arrows highlighted in orange in Figure 3), and places a checkpoint right after the obstacle that follows the chosen path. SneakySnake then starts the count from this checkpoint for each row of the chip maze. Thus, the CPU-based implementation consists of irregular memory access patterns.

3) **Finding the estimated number of edits**. SneakySnake estimates the number of edits to be equal to the number of obstacles found along the shortest routing path (3 obstacles for our example in Figure 3, each of which is located at the end of each chosen subpath). Thus, SneakySnake repeats the second step until either there are no more available entries to be examined, or the total number of obstacles passed through exceeds the allowed edit distance threshold (i.e.,  $E$ ).

## Case Study 2: Weather Modeling and Prediction

Accurate and fast weather prediction using detailed weather models is essential to make weather-dependent decisions in a timely manner. The Consortium for Small-Scale Modeling (COSMO) [16] built one such weather model to meet the high-resolution forecasting requirements of weather services. The COSMO model is a non-hydrostatic atmospheric prediction model that is widely used for meteorological purposes and research applications [24, 63].

The central part of the COSMO model (called *dynamical core* or *dycore*) solves the Euler equations on a curvilinear grid [78] and applies 1) implicit discretization (i.e., parameters are dependent on each other at the same time instance [79])



in the vertical dimension and 2) explicit discretization (i.e., a solution depends on the previous system state [79]) in the two horizontal dimensions. The use of different discretizations leads to three computational patterns [80]: 1) horizontal stencils, 2) tridiagonal solvers in the vertical dimension, and 3) point-wise computation. These computational kernels are compound stencil kernels that operate on a three-dimensional grid [24, 81]. A stencil operation updates values in a structured multidimensional grid (*row*, *column*, *depth*) based on the values of a fixed local neighborhood of grid points. *Vertical advection* (*vadv*) and *horizontal diffusion* (*hdiff*) are two such compound stencil kernels found in the *dycore* of the COSMO weather prediction model. They are similar to the kernels used in other weather and climate models [82–84].

These kernels are representative of the data access patterns and algorithmic complexity of the entire COSMO model, and are similar to the kernels used in other weather and climate models [18]. As shown in Figure 1, their performance is bounded by access to main memory due to their irregular memory access patterns and low arithmetic intensity that altogether result in an order-of-magnitude lower performance than the peak CPU performance.

The horizontal diffusion kernel iterates over a 3D grid performing *Laplacian* and *flux* stencils to calculate different grid points as shown in Figure 4. A single *Laplacian* stencil accesses the input grid at five memory offsets, the result of which is used to calculate the *flux* stencil. *hdiff* has purely horizontal access patterns and does not have dependencies in the vertical dimension. Thus, it can be fully parallelized in the vertical dimension. Vertical advection has a higher degree of complexity since it uses the Thomas algorithm [85] to solve a tridiagonal matrix of weather data (called *fields*, such as, air pressure, wind velocity, and temperature) along the vertical axis. *vadv* consists of a forward sweep that is followed by a backward sweep along the vertical dimension. *vadv* requires access to the weather data, which are stored as array structures while performing forward and sweep computations. Unlike the conventional stencil kernels, vertical advection has dependencies in the vertical direction, which leads to limited available parallelism and




Figure 4: Horizontal diffusion kernel composition using Laplacian and flux stencils in a two dimensional plane [19].

irregular memory access patterns. For example, when the input grid is stored by *row*, accessing data elements in the *depth* dimension typically results in many cache misses [86].

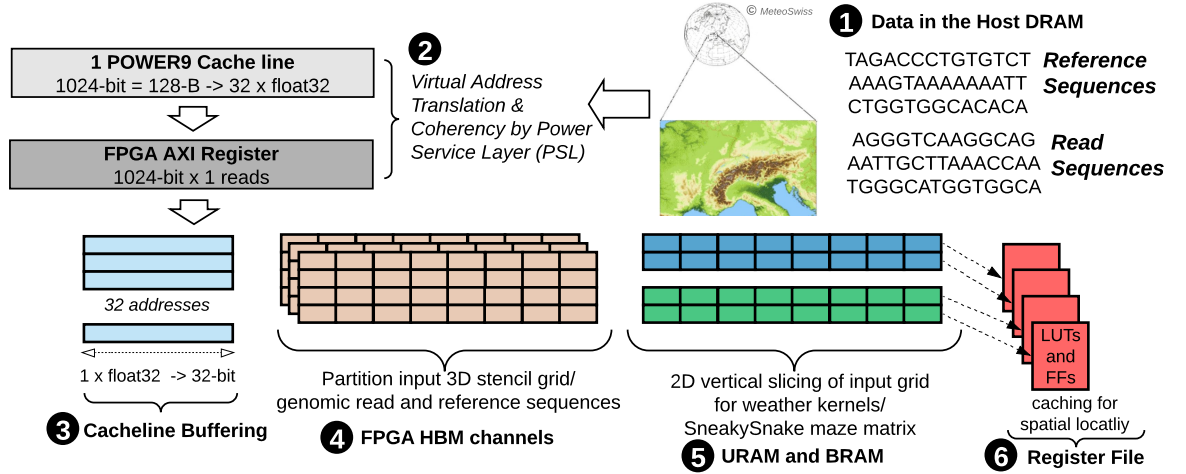
## Accelerator Implementation

We design and implement an accelerator on our HBM-based FPGA-board (Figure 2) for each of the three kernels (*SneakySnake*, *vadv*, and *hdiff*) in our two case studies. We use a High-Level Synthesis (HLS) [87] design flow to implement and map our accelerator design. We describe the design and the execution flow for our HBM-based accelerators.

Figure 5 shows the end-to-end data transfer from the host DRAM to the processing element on an FPGA. We make use of streams (`hls::streams`<sup>1</sup>) to connect different dataflow tasks that allow consumer functions to operate before the producer functions have been completed. Streaming simplifies address management as the data samples are sent in sequential order between two modules. Before feeding data to a processing element, we use the on-chip heterogeneous memory hierarchy to unpack the stream data in a way that matches the data access pattern of an application. Therefore, we implement our accelerator design following a dataflow approach in five steps.

First, the input data stored in the DRAM of the host system (❶ in Figure 5) is transferred to the FPGA over a 1024-bit wide OCAPI interface (❷) by a *data-fetch engine*. A single cache line stream of `float32` datawidth would have

<sup>1</sup>We use Vivado HLS C++ template class `hls::stream` to implement FIFO-based streaming interfaces.



**Figure 5: Data transfer flow from the host DRAM to the on-board FPGA memory via POWER9 cache lines. Heterogeneous memory partitioning of on-chip memory blocks enable low read/write latencies across the FPGA memory hierarchy.**

32 data elements. The data-fetch engine reads 1024-bit wide POWER9 cache line data over the OCAPI interface and pushes the data into a 1024-bit buffer before converting it to 256-bit HBM pseudo-channel bitwidth. For weather prediction, the input data is the atmospheric data collected from weather simulations based on the atmospheric model resolution grid. For genome analysis, the input data is the reference and read sequences for the pre-alignment filtering step of the genome analysis pipeline. Second, following the initial buffering (3), the *HBM-write engine* maps the data onto the HBM memory (4). We partition the data among HBM channels (5) to exploit data-level parallelism and to scale our design. Our evaluated workloads have limited locality, so to exploit locality, we cache certain parts of data into a register file made of LUTs and FFs (6). Third, we assign a dedicated HBM memory channel to a specific processing element (PE); therefore, we enable as many HBM channels as the number of PEs. This allows us to use the high HBM bandwidth effectively because each PE fetches from an independent 256-bit channel, which provides low-latency and high-bandwidth data access to each PE. An *HBM-read engine* reads data from a dedicated HBM memory channel and assigns data to a specialized PE. The HBM channel provides 256-bit data, which is a quarter of the OCAPI bitwidth (1024-bit). Therefore, to match the OCAPI bitwidth, we

introduce a stream converter logic that converts a 256-bit HBM stream to a 1024-bit wide stream, which is equal to the maximum OCAPI bitwidth.

Fourth, each PE performs computation (*SneakySnake* for genome analysis, and *vadvc* or *hdiff* for weather prediction) on the transferred data. In *SneakySnake*, we equally divide the number of read and reference sequences among the PEs. In *vadvc* and *hdiff*, each PE operates on a block of the input grid. For *SneakySnake*<sup>2</sup>, each row in the chip maze is stored as a register array of length equal to the read length. The registers are accessed simultaneously throughout the execution. In every iteration, we count consecutive zeros in each row until we find an obstacle (i.e., until we come across a 1). Following this, we shift all the bits by the maximum number of zeros in the chip maze. This shifting allows us to overcome the irregular array accesses while finding the longest possible path in the chip maze. Fifth, once the calculated results are available, the *HBM-write engine* writes calculated results to its assigned HBM memory channel, after which the *write-back engine* transfers the data back to the host system for further processing.

We create a specialized memory hierarchy from the pool of heterogeneous FPGA memories

<sup>2</sup>We open-source our HBM+OCAPI-based *SneakySnake* accelerator implementations (both single-channel-single PE and multi-channel-single PE): <https://github.com/CMU-SAFARI/SneakySnake/tree/master/SneakySnake-HLS-HBM>

(i.e., on-chip BRAM and URAM, and in-package HBM). By using a greedy algorithm, we determine the best-suited hierarchy for each kernel. Heterogeneous partitioning of on-chip memory blocks reduces read and write latencies across the FPGA memory hierarchy. To optimize a PE, we apply three optimization strategies. First, we exploit the inherent parallelism in a given algorithm using hardware pipelining. Second, we partition data arrays onto multiple physical on-chip memories (BRAM/URAM) instead of a single large memory to avoid stalling of our pipelined design, since the on-chip BRAM/URAM have only two read/write ports. On-chip memory reshaping is an effective technique for improving the bandwidth of BRAMs/URAMs. Third, we partition the input data between PEs, therefore, all PEs exploit data-level parallelism. We apply all three optimizations using source-code annotations via *Vivado HLS #pragma* directives [88].

## Evaluation

We evaluate our accelerator designs for *SneakySnake*, *vadv*, and *hdiff* in terms of performance, energy consumption, and FPGA resource utilization on two different FPGAs, and two different external data communication interfaces between the CPU and the FPGA board. We implement our accelerator designs on both 1) an Alpha-Data ADM-PCIE-9H7 card [89] featuring the Xilinx Virtex Ultrascale+ XCVU37P-FSVH2892-2-e [60] with 8GiB HBM2 [54] and 2) an Alpha-Data ADM-PCIE-9V3 card [90] featuring the Xilinx Virtex Ultrascale+ XCVU3P-FFVC1517-2-i with 8GiB DDR4 [60], connected to an IBM POWER9 host system. For the external data communication interface, we use both CAPI2 [56] and the state-of-the-art OCAPI (OpenCAPI) [65] interface. We compare these implementations to execution on a POWER9 CPU with 16 cores (using all 64 hardware threads). We run *SneakySnake* using the first 30,000 real genomic sequence pairs (text and query pairs) of 100bp\_2 dataset [23], which is widely used as in prior works [9, 62, 73, 75, 76]. The length of each sequence is 100 bp (base-pair) long. For weather prediction, we run our experiments using a  $256 \times 256 \times 64$ -point domain similar to the grid domain used by the COSMO model.

## Performance Analysis

We provide the execution time of *SneakySnake*, *vadv*, and *hdiff* on the POWER9 CPU with 64 threads and the FPGA accelerators (both DDR4-based and HBM-based) in Figure 6 (a), (b), and (c), respectively. For both FPGA designs, we scale the number of PEs from 1 to the maximum number that we can accommodate on the available FPGA resources. On the DDR4-based design, the maximum number of PEs is 4, 4, and 8 for *SneakySnake*, *vadv*, and *hdiff*, respectively. On the HBM-based design, we can fit up to 12, 14, and 16 PEs for *SneakySnake*, *vadv*, and *hdiff*, respectively. Based on our analysis, we make **four key observations**.

First, the full-blown HBM+OCAPI-based implementations (with the maximum number of PEs) of *SneakySnake*, *vadv*, and *hdiff* outperform the 64-thread IBM POWER9 CPU version by 27.4 $\times$ , 5.3 $\times$ , and 12.7 $\times$ , respectively. We achieve 28%, 37%, and 44% higher performance for *SneakySnake*, *vadv*, and *hdiff*, respectively, with OCAPI-based HBM design than CAPI2-based HBM design due to the following two reasons: 1) OCAPI provides double the bitwidth (1024-bit) of the CAPI2 interface (512-bit), which provides a higher bandwidth to the host CPU, i.e., 22.1/22.0 GB/s R/W versus 13.9/14.0 GB/s; and 2) with OCAPI, memory coherency logic is moved onto the IBM POWER CPU, which provides more FPGA area and allows us to run our accelerator logic at a higher clock frequency (250MHz for OCAPI versus 200MHz for CAPI2).

Second, for a single PE, the DDR4-CAPI2-based FPGA accelerator design is faster than the HBM-CAPI2-based design for all three kernels. This is because the HBM-based design uses one HBM channel per PE, as the bus width of the DDR4 channel (512 bits) is larger than that of an HBM channel (256 bits). Therefore, the HBM channel has a lower transfer rate of 0.8-2.1 GT/s (Gigatransfers per second) than for a DDR4 channel (2.1-4.3 GT/s) with a theoretical bandwidth of 12.8 GB/s and 25.6 GB/s per channel, respectively. One way to match the DDR4 bus width is to have a single PE fetch data from multiple HBM channels in parallel. As shown in Figure 6, in our multi-channel setting (HBM\_multi+OCAPI), we




Figure 6: Performance for (a) *SneakySnake*, (b) *vadvc*, and (c) *hdiff* as a function of accelerator PE count on the HBM- and DDR4-based FPGA boards. Energy efficiency for (d) *SneakySnake*, (e) *vadvc*, and (f) *hdiff* and on HBM- and DDR4-based FPGA boards. We also show the single socket (64 threads) performance and energy efficiency of an IBM POWER9 host system for *SneakySnake*, *vadvc*, and *hdiff*. For HBM-based design, we implement our accelerator with both the CAPI2 interface and the state-of-the-art OpenCAPI (OCAPI) interface (with both single channel and multiple channels per PE).

use 4 HBM pseudo channels per PE to meet the bitwidth of the OCAPI interface. We observe that by fetching more data from multiple channels, compared to our single-channel-single PE design, we achieve  $1.4\times$ ,  $1.2\times$ , and  $1.8\times$  performance improvement for *SneakySnake*, *vadvc*, and *hdiff*, respectively.

Third, as we increase the number of PEs, we divide the workload evenly across PEs. As a result, we observe linear scaling in the performance of HBM-based designs, where each PE reads and writes through a dedicated HBM channel. For multi-channel designs, we are able to accommodate only 3 PEs for the three evaluated kernels (i.e., 12 HBM channels) because adding more HBM channels leads to timing constraint violations. We observe that the best-performing multi-channel-single PE design (i.e., using 3 PEs with 12 HBM channels for all three workloads) has  $1.1\times$ ,  $4.7\times$ , and  $3.1\times$  lower performance than the best-performing single-channel-single PE design (i.e., 12 PEs for *SneakySnake*, 14 PEs for *vadvc*, and 16 PEs for *hdiff*, respectively).

This observation shows that there is a tradeoff between (1) enabling more HBM pseudo channels to provide each PE with more bandwidth, and (2) implementing more PEs in the available area. For *SneakySnake*, in our dataflow design, data transfer time dominates the computation time; therefore, adding more PEs does not lead to a linear reduction in performance. For *vadvc* and *hdiff*, both data transfer and computation take a comparable amount of time. Therefore, in such workloads, we are able to achieve a linear execution time reduction with the number of PEs.

Fourth, the performance of the DDR4-based designs scales non-linearly for *vadvc* and *hdiff* with the number of PEs, as all PEs access memory through the same channel. Multiple PEs compete for a single memory channel, which causes frequent memory stalls due to contention in the memory channel. For *SneakySnake*, which is the most memory-bound of the three kernels, we observe a constant execution time with the increase in PEs. With a single PE,



memory access time hides all computation time. Increasing the number of PEs reduces the time devoted to computation, but not the memory access time because there is a single channel. Therefore, memory bandwidth saturates with a single PE, and the total execution time does *not* reduce with the number of PEs.

We conclude that FPGA-based near-memory computing provides significant speedup (between  $5.3\times$ - $27.4\times$ ) to key data-intensive applications over a state-of-the-art CPU (POWER9).

### Energy Efficiency Analysis

We provide energy efficiency results for *SneakySnake*, *vadv*, and *hdiff* on the two FPGA designs and the POWER9 CPU in Figure 6 (d), (e), and (f), respectively. We express energy efficiency in Mseq/s/Watt (i.e., millions of read sequences per second per Watt) for *SneakySnake*, and in terms of GFLOPS/Watt for *hdiff* and *vadv*. For power measurement on the POWER9 system with an FPGA board, we use the AMESTER tool<sup>3</sup> to monitor built-in power sensors. We measure the active power consumption, i.e., the difference between the total power of a complete socket (including processor, memory, fans, and I/O) when running an application and when idle. Based on our analysis, we make **five key observations**.

First, our full-blown HBM+OCAPI-based accelerator designs (with 12 PEs for *SneakySnake*, 14 PEs for *vadv*, and 16 PEs for *hdiff*) improve energy efficiency by  $133\times$ ,  $12\times$ , and  $35\times$  compared to the POWER9 system for *SneakySnake*, *vadv*, and *hdiff*, respectively.

Second, the DDR4-CAPI2-based designs are slightly more energy efficient ( $1.1\times$  to  $1.5\times$ ) than the HBM-CAPI2-based designs when the number of PEs is small. This observation is in line with our discussion about performance with small PE counts in the previous section. However, as we increase the number of PEs, the HBM-based designs provide higher energy efficiency since they make use of multiple HBM channels. Third, compared to our single-channel-single PE design, our multi-channel-single PE design provides only  $1.1\times$ ,  $1\times$ , and  $1.5\times$  higher

energy efficiency for *SneakySnake*, *vadv*, and *hdiff*, respectively. This is because using more channels leads to higher power consumption ( $\sim 1$  Watt per channel) even though we get higher bandwidth per PE by using multiple channels.

Fourth, the energy efficiency of the HBM-based design for *hdiff* increases with the number of PEs until a saturation point (8 PEs). However, the energy efficiency of *SneakySnake* and *vadv* (HBM+CAPI2) designs decreases after using more than 4 PEs, and that of the *vadv* (HBM+OCAPI) design decreases after using more than 8 PEs. This is because every additional HBM channel increases power consumption by  $\sim 1$  Watt (for the HBM AXI3 interface operating at 250MHz with a logic toggle rate of  $\sim 12.5\%$ ). In case of *vadv*, there is a large amount of control flow that leads to large and inefficient resource consumption when increasing the PE count. This causes a high increase in power consumption with 14 PEs.

Fifth, as we increase the number of PEs, performance of *SneakySnake* increases, whereas energy efficiency may not follow the same trend. This is because *SneakySnake* spends significant amount of execution time in fetching data from memory, as mentioned in the discussion on performance analysis. Thus, this may lead to a reduction in energy efficiency, depending upon the number of HBM channels used.

We conclude that increasing the number of PEs and enabled HBM channels does *not* always increase energy efficiency. However, data-parallel kernels like *hdiff* can achieve much higher performance in an energy-efficient manner with more PEs and HBM channels.

### FPGA Resource Utilization

We list the resource utilization of *SneakySnake*, *vadv*, and *hdiff* on the FPGA board with HBM memory in Table 1. We make three observations. First, BRAM utilization is significantly higher than utilization of other resources. The reason is that we use `hls::streams` to implement input and output to different functions. Streams are FIFOs, which are implemented with BRAMs. Second, *SneakySnake* performs all computations using flip-flops (FF) and lookup table registers (LUT). It does not require digital signal processing units

<sup>3</sup><https://github.com/open-power/amester>

(DSP) since it does not execute floating-point operations. Third, resource consumption of `vadvc` is much higher than that of `hdiff` because it has higher computation complexity and requires a larger number of input parameters for the compound stencil computation. We observe that there are enough resources to accommodate more than 16 PEs for `hdiff`, but in this work, we use only a single HBM stack due to timing constraint violations. Therefore, in this work, the maximum number of PEs is 16 to match 16 memory channels offered by a single HBM stack.

**Table 1: FPGA resource utilization in the full-blown HBM+OCAPI-based designs for *SneakySnake* (12 PEs), *vadvc* (14 PEs), and *hdiff* (16 PEs).**

Algorithm	BRAM	DSP	FF	LUT	URAM
<i>SneakySnake</i>	58%	0%	18%	70%	1%
<i>vadvc</i>	90%	39%	37%	55%	53%
<i>hdiff</i>	96%	4%	10%	15%	8%

## Discussion

This paper presents our recent efforts to leverage near-memory computing capable FPGA-based accelerators to accelerate three major kernels taken from two data-intensive applications: (1) pre-alignment filtering in genome analysis, and (2) horizontal diffusion and vertical advection stencils from weather prediction. We identify key challenges for such acceleration and provide solutions to them. We summarize the most important insights and takeaways as follows.

First, our evaluation shows that High-Bandwidth Memory-based near-memory FPGA accelerator designs can improve performance by  $5.3\times$ - $27.4\times$  and energy efficiency by  $12\times$ - $133\times$  over a high-end 16-core IBM POWER9 CPU.

Second, our HBM-based FPGA accelerator designs employ a dedicated HBM channel per PE. This avoids memory access congestion, which is typical in DDR4-based FPGA designs and ensures memory bandwidth scaling with the number of PEs. As a result, in most of the data-parallel applications, performance scales linearly with the number of PEs.

Third, the maximum performance of our HBM-based design is reached using the maximum PE count that we can fit in the reconfigurable fabric, with each PE having a dedicated

HBM channel. However, adding more PEs could lead to timing constraint violations for HBM-based designs. HBM-based FPGAs consist of multiple super-logic regions (SLRs) [91], where an SLR represents a single FPGA die. All HBM channels are connected only to SLR0, while other SLRs have indirect connections to the HBM channels. Therefore, for a large design, if a PE is implemented in a non-SLR0 region, it might make timing closure difficult.

Fourth, the energy efficiency of our HBM-based designs tends to saturate (or even reduces) as we increase the number of PEs beyond some point. The highest energy efficiency is achieved with a PE count that is smaller than the highest-performing PE count. The major reason for a decrease in the energy efficiency is the increase in power consumption with every additional HBM channel.

We hope that the near-memory acceleration efforts for genome analysis and weather prediction we described and the challenges we identified and solved provide a foundation for accelerating modern and future data-intensive applications using powerful near-memory reconfigurable accelerators.

## ACKNOWLEDGMENTS

We thank the anonymous reviewers of IEEE Micro for their feedback. We thank the SAFARI Research Group members for valuable feedback and the stimulating intellectual environment they provide. This work is supported by funding from ASML, Google, Huawei, Intel, Microsoft, VMware, and the Semiconductor Research Corporation to Onur Mutlu. Special thanks to IBM Research Europe, Zurich for providing access to IBM systems. This work was partially funded by Eindhoven University of Technology and ETH Zürich.

## REFERENCES

1. Anirban Nag, et al. GenCache: Leveraging In-Cache Operators for Efficient Sequence Alignment. In *ISCA*, 2019.
2. Can Firtina, et al. Apollo: A Sequencing-Technology-Independent, Scalable and Accurate Assembly Polishing Algorithm. *Bioinformatics*, 2020.
3. Heng Li and Richard Durbin. Fast and Accurate Short

- Read Alignment with Burrows–Wheeler Transform. In *Bioinformatics*, 2009.
4. Shanshan Ren, et al. Efficient Acceleration of the Pair-HMMs Forward Algorithm for GATK HaplotypeCaller on Graphics Processing Units. In *Bioinformatics*, 2018.
  5. Yatish Turakhia, et al. Darwin: A Genomics Co-Processor Provides Up to 15,000X Acceleration on Long Read Assembly. In *ASPLOS*, 2018.
  6. Heng Li. minimap2: Pairwise Alignment for Nucleotide Sequences. In *Bioinformatics*, 2018.
  7. Mohammed Alser, et al. Accelerating Genome Analysis: A Primer on an Ongoing Journey. In *IEEE Micro*, 2020.
  8. Mohammed Alser, et al. Technology Dictates Algorithms: Recent Developments in Read Alignment. In *arXiv*, 2020.
  9. Jeremie S Kim, et al. GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-In-Memory Technologies. In *BMC Genomics*, 2018.
  10. Damla Senol Cali, et al. Nanopore Sequencing Technology and Tools for Genome Assembly: Computational Analysis of the Current State, Bottlenecks and Future Directions. In *Brief. Bioinform.*, 2019.
  11. Damla Senol Cali, et al. GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis. In *MICRO*, 2020.
  12. <https://www.illumina.com/products/by-type/informatics-products/dragen-bio-it-platform.html>.
  13. Nathan LaPierre, et al. Metalign: Efficient Alignment-based Metagenomic Profiling via Containment Min Hash. In *BioRxiv*. Cold Spring Harbor Laboratory, 2020.
  14. Jeremie S Kim, et al. AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes. *bioRxiv*, 2021.
  15. Nathan LaPierre, et al. MiCoP: Microbial Community Profiling Method for Detecting Viral and Fungal Organisms in Metagenomic Samples. *BMC Genomics*, 2019.
  16. G Doms and U Schättler. The Nonhydrostatic Limited-Area Model LM (Lokal-model) of the DWD. Part I: Scientific Documentation. In *DWD, GB Forschung und Entwicklung*, 1999.
  17. Felix Thaler, et al. Porting the COSMO Weather Model to Manycore CPUs. In *SC*, 2019.
  18. Gagandeep Singh, et al. NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling. In *FPL*, 2020.
  19. Gagandeep Singh, et al. NARMADA: Near-Memory Horizontal Diffusion Accelerator for Scalable Stencil Computations. In *FPL*, 2019.
  20. Onur Mutlu, et al. Processing Data Where It Makes Sense: Enabling In-Memory Computation. In *MICPRO*, 2019.
  21. Samuel Williams, et al. Roofline: An Insightful Visual Performance Model for Multicore architectures. In *CACM*, 2009.
  22. Satish Kumar Sadasivam, et al. IBM POWER9 Processor Architecture. In *IEEE Micro*, 2017.
  23. Mohammed Alser, et al. SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs. In *Bioinformatics*, 2020.
  24. Tobias Gysi, et al. MODESTO: Data-centric Analytic Optimization of Complex Stencil Programs on Heterogeneous Architectures. In *SC*, 2015.
  25. Geraldo Francisco Oliveira, et al. DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks. In *arXiv*, 2021.
  26. R. Nair, et al. Active Memory Cube: A Processing-in-Memory Architecture for Exascale Systems. In *IBM JRD*, 2015.
  27. Onur Mutlu, et al. A Modern Primer on Processing in Memory. In *Emerging Computing: From Devices to Systems-Looking Beyond Moore and Von Neumann*. Springer, 2021.
  28. Lifeng Nai, et al. GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks. In *HPCA*, 2017.
  29. Amin Farmahini-Farahani, et al. NDA: Near-DRAM Acceleration Architecture Leveraging Commodity DRAM Devices and Standard Memory Modules. In *HPCA*, 2015.
  30. Vincent T Lee, et al. Application Codesign of Near-Data Processing for Similarity Search. In *IPDPS*, 2018.
  31. Yangwook Kang, et al. Enabling Cost-Effective Data Processing with Smart SSD. In *MSST*, 2013.
  32. Milad Hashemi, et al. Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads. In *MICRO*, 2016.
  33. Berkin Akin, et al. Data Reorganization in Memory Using 3D-Stacked DRAM. In *ISCA*, 2015.
  34. Oreoluwatomiwa Babarinsa and Stratos Idreos. JAFAR: Near-Data Processing for Databases. In *SIGMOD*, 2015.
  35. Joo Hwan Lee, et al. BSSync: Processing Near Memory for Machine Learning Workloads with Bounded Staleness Consistency Models. In *PACT*, 2015.
  36. Ping Chi, et al. PRIME: A Novel Processing-In-memory Architecture for Neural Network Computation in ReRAM-based Main Memory. In *ISCA*, 2016.
  37. Duckhwan Kim, et al. Neurocube: A Programmable Digital Neuromorphic Architecture with High-Density 3D Memory. In *ISCA*, 2016.
  38. Hadi Asghari-Moghaddam, et al. Chameleon: Versatile

- and Practical Near-DRAM Acceleration Architecture for Large Memory Systems. In *MICRO*, 2016.
39. Amirali Boroumand, et al. LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory. In *CAL*, 2016.
  40. Vivek Seshadri, et al. Gather-Scatter DRAM: In-DRAM Address Translation to Improve the Spatial Locality of Non-unit Strided Accesses. In *MICRO*, 2015.
  41. Zhiyu Liu, et al. Concurrent Data Structures for Near-Memory Computing. In *SPAA*, 2017.
  42. Mingyu Gao, et al. Practical Near-Data Processing for In-Memory Analytics Frameworks. In *PACT*, 2015.
  43. Amir Morad, et al. GP-SIMD Processing-in-Memory. In *ACM TACO*, 2015.
  44. Amirali Boroumand, et al. Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks. In *ASPLOS*, 2018.
  45. Junwhan Ahn, et al. A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing. In *ISCA*, 2015.
  46. Junwhan Ahn, et al. PIM-Enabled Instructions: A Low-Overhead, Locality-Aware Processing-in-Memory Architecture. In *ISCA*, 2015.
  47. Kevin Hsieh, et al. Accelerating Pointer Chasing in 3D-Stacked Memory: Challenges, Mechanisms, Evaluation. In *ICCD*, 2016.
  48. Milad Hashemi, et al. Accelerating Dependent Cache Misses with an Enhanced Memory Controller. In *ISCA*, 2016.
  49. Gagandeep Singh, et al. Near-Memory Computing: Past, Present, and Future. In *MICPRO*, 2019.
  50. Ivan Fernandez, et al. NATSA: A Near-Data Processing Accelerator for Time Series Analysis. In *ICCD*, 2020.
  51. Nastaran Hajinazar, et al. SIMDRAM: A Framework for Bit-Serial SIMD Processing Using DRAM. In *ASPLOS*, 2021.
  52. Gagandeep Singh, et al. NAPEL: Near-Memory Computing Application Performance Prediction via Ensemble Learning. In *DAC*, 2019.
  53. Juan Gómez-Luna, et al. Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture. In *arXiv*, 2021.
  54. High Bandwidth Memory (HBM) DRAM (JESD235), [https://www.jedec.org/document\\_search?search\\_api\\_views\\_fulltext=jesd235](https://www.jedec.org/document_search?search_api_views_fulltext=jesd235).
  55. UltraScale Architecture Memory Resources, [https://www.xilinx.com/support/documentation/user\\_guides/ug573-ultrascale-memory-resources.pdf](https://www.xilinx.com/support/documentation/user_guides/ug573-ultrascale-memory-resources.pdf).
  56. Jeffrey Stuecheli, et al. CAPI: A Coherent Accelerator Processor Interface. In *IBM JRD*, 2015.
  57. Brad Benton. CCIX, Gen-Z, OpenCAPI: Overview and Comparison. In *OFA*, 2017.
  58. Debendra Das Sharma. Compute Express Link. In *CXL Consortium White Paper*, 2019.
  59. Brian Gaide, et al. Xilinx Adaptive Compute Acceleration Platform: VersalTM Architecture. In *FPGA*, 2019.
  60. Virtex UltraScale+, <https://www.xilinx.com/products/silicon-devices/fpga/virtex-ultrascale-plus.html>.
  61. Hongyi Xin, et al. Accelerating Read Mapping With FastHASH. In *BMC Genomics*, 2013.
  62. Evangelos Georganas, et al. merAligner: A Fully Parallel Sequence Aligner. In *IPDPS*, 2015.
  63. Christoph Schär, et al. Kilometer-Scale Climate Models: Prospects and Challenges. In *BAMS*, 2020.
  64. AXI High Bandwidth Memory Controller v1.0, [https://www.xilinx.com/support/documentation/ip\\_documentation/hbm/v1\\_0/pg276-axi-hbm.pdf](https://www.xilinx.com/support/documentation/ip_documentation/hbm/v1_0/pg276-axi-hbm.pdf).
  65. Jeffrey Stuecheli et al. IBM POWER9 Opens Up a New Era of Acceleration Enablement: OpenCAPI. In *IBM JRD*, 2018.
  66. Gonzalo Navarro. A Guided Tour to Approximate String Matching. In *CSUR*, 2001.
  67. Vladimir I Levenshtein. Binary Codes Capable of Correcting Deletions, Insertions, and Reversals. In *Soviet Dokl. Phys.*, 1966.
  68. Temple F Smith and Michael S Waterman. Identification of Common Molecular Subsequences. In *J. Mol. Biol*, 1981.
  69. Saul B Needleman and Christian D Wunsch. A General Method Applicable to the Search for Similarities in the Amino Acid Sequence of Two Proteins. In *J. Mol. Biol*, 1970.
  70. Martin Šošić and Mile Šikić. Edlib: A C/C++ Library for Fast, Exact Sequence Alignment Using Edit Distance. In *Bioinformatics*, 2017.
  71. Jeff Daily. Parasail: SIMD C Library for Global, Semi-Global, and Local Pairwise Sequence Alignments. In *BMC Bioinformatics*, 2016.
  72. Hongyi Xin, et al. Shifted Hamming Distance: A Fast and Accurate SIMD-Friendly Filter to Accelerate Alignment Verification in Read Mapping. In *Bioinformatics*, 2015.
  73. Mohammed Alser, et al. GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping. In *Bioinformatics*, 2017.
  74. Mohammed Alser, et al. MAGNET: Understanding and Improving the Accuracy of Genome Pre-Alignment Filtering. In *TIR*, 2017.
  75. Mohammed Alser, et al. Shouji: A Fast and Efficient



- Pre-Alignment Filter for Sequence Alignment. In *Bioinformatics*, 2019.
76. Zülal Bingöl, et al. GateKeeper-GPU: Fast and Accurate Pre-Alignment Filtering in Short Read Mapping. In *arXiv*, 2021.
  77. Jerry H. Lee, et al. Use of Steiner's Problem in Suboptimal Routing in Rectilinear Metric. In *TCAS*, 1976.
  78. Louis J Wicker and William C Skamarock. Time-Splitting Methods for Elastic Models Using Forward Time Schemes. In *Monthly Weather Review* 130, 2002.
  79. Luca Bonaventura. A Semi-implicit Semi-Lagrangian Scheme using the Height Coordinate for a Nonhydrostatic and Fully Elastic Model of Atmospheric Flows. In *JCP*, 2000.
  80. Felix Thaler, et al. Porting the COSMO Weather Model to Manycore CPUs. In *PASC*, 2019.
  81. Gagandeep Singh, et al. Low Precision Processing for High Order Stencil Computations. In *Springer LNCS*, 2019.
  82. Scott Kehler, et al. High Resolution Deterministic Prediction System (HRDPS) Simulations of Manitoba Lake Breezes. In *Atmosphere-Ocean*, 2016.
  83. Richard B Neale, et al. Description of the NCAR Community Atmosphere Model (CAM 5.0). In *NCAR Tech. Note*, 2010.
  84. Jun A. Zhang, et al. Evaluating the Impact of Improvement in the Horizontal Diffusion Parameterization on Hurricane Prediction in the Operational Hurricane Weather Research and Forecast (HWRF) Model. In *Weather and Forecasting*, 2018.
  85. Llewellyn Thomas. Elliptic Problems in Linear Differential Equations over a Network. In *Watson Sci. Comput. Lab. Rept.*, Columbia University, 1949.
  86. Jingheng Xu, et al. Performance Tuning and Analysis for Stencil-Based Applications on POWER8 Processor. In *ACM TACO*, 2018.
  87. Vivado High-Level Synthesis, <https://www.xilinx.com/products/design-tools/vivado/integration/esl-design.html>.
  88. HLS Pragmas, [https://www.xilinx.com/html\\_docs/xilinx2020\\_2/vitis\\_doc/hls\\_pragmas.html](https://www.xilinx.com/html_docs/xilinx2020_2/vitis_doc/hls_pragmas.html).
  89. ADM-PCIE-9H7-High-Speed Communications Hub, <https://www.alpha-data.com/dcp/products.php?product=adm-pcie-9h7>.
  90. ADM-PCIE-9V3-High-Performance Network Accelerator, <https://www.alpha-data.com/dcp/products.php?product=adm-pcie-9v3>.
  91. Virtex UltraScale+ HBM FPGA: A Revolutionary Increase in Memory Performance, [https://www.xilinx.com/support/documentation/white\\_papers/wp485-hbm.pdf](https://www.xilinx.com/support/documentation/white_papers/wp485-hbm.pdf).

**Gagandeep Singh** is with ETH Zürich, Zürich, Switzerland. Contact him at [gagan.gagandeepsingh@safari.ethz.ch](mailto:gagan.gagandeepsingh@safari.ethz.ch)

**Mohammed Alser** is with ETH Zürich, Zürich, Switzerland. Contact him at [alserm@ethz.ch](mailto:alserm@ethz.ch)

**Damla Senol Cali** is with Carnegie Mellon University, Pittsburgh, PA, USA. Contact her at [dsenol@andrew.cmu.edu](mailto:dsenol@andrew.cmu.edu)

**Dionysios Diamantopoulos** is with IBM Research Europe, Zürich Lab, Rüschlikon, Switzerland. Contact him at [did@zurich.ibm.com](mailto:did@zurich.ibm.com)

**Juan Gómez-Luna** is with ETH Zürich, Zürich, Switzerland. Contact him at [juan.gomez@safari.ethz.ch](mailto:juan.gomez@safari.ethz.ch)

**Henk Corporaal** is with Eindhoven University of Technology, Eindhoven, The Netherlands. Contact him at [h.corporaal@tue.nl](mailto:h.corporaal@tue.nl)

**Onur Mutlu** is with ETH Zürich, Zürich, Switzerland and Carnegie Mellon University, Pittsburgh, PA, USA. Contact him at [omutlu@ethz.ch](mailto:omutlu@ethz.ch)