

# Accelerating Genome Analysis: A Primer on an Ongoing Journey

Mohammed Alser<sup>◇</sup> Zülal Bingöl<sup>▽</sup> Damla Senol Cali<sup>✕</sup> Jeremie Kim<sup>◇✕</sup>

Saugata Ghose<sup>\*✕</sup> Can Alkan<sup>▽</sup> Onur Mutlu<sup>◇✕▽</sup>

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

<sup>\*</sup>University of Illinois at Urbana–Champaign

**Abstract**—Genome analysis fundamentally starts with a process known as *read mapping*, where sequenced fragments of an organism’s genome are compared against a reference genome. Read mapping is currently a major bottleneck in the entire genome analysis pipeline, because state-of-the-art genome sequencing technologies are able to sequence a genome much faster than the computational techniques employed to analyze the genome. We describe the ongoing journey in significantly improving the performance of read mapping. We explain state-of-the-art algorithmic methods and hardware-based acceleration approaches. Algorithmic approaches exploit the structure of the genome as well as the structure of the underlying hardware. Hardware-based acceleration approaches exploit specialized microarchitectures or various execution paradigms (e.g., processing inside or near memory). We conclude with the challenges of adopting these hardware-accelerated read mappers.

■ **GENOME ANALYSIS** is the foundation of many scientific and medical discoveries, and serves as a key enabler of personalized medicine. This analysis is currently limited by the inability of modern genome sequencing technologies to read an organism’s complete genome. Instead, sequencing machines extract smaller random fragments of an organism’s DNA sequence, known as *reads*. While the human genome contains over three billion *bases* (i.e., A, C, G, T in DNA), the length of a read is orders of magnitude smaller, ranging from a few hundred bases (for *short* reads) to a few million bases (for *long* reads). Computers are used to perform genome *assembly*, which reassembles read fragments back into an entire genome sequence. Genome assembly is

currently the bottleneck to quickly and accurately determining an individual’s entire genome, due to the complex algorithms and large datasets used for assembly.

A widely-used approach for genome assembly is to perform *sequence alignment*, which compares read fragments against a known *reference genome* (i.e., a complete representative DNA sequence for a particular species). A process known as *read mapping* matches each read generated from sequencing to one or more possible locations within the reference genome, based on the similarity between the read and the reference sequence segment at that location. Unfortunately, the bases in a read may not be identical to the bases in the reference genome at the location that

the read actually comes from. These differences may be due to (1) sequencing errors (up to 0.1% in short reads [1] and up to 20% in long reads [2], [3]) during extraction, and (2) genetic mutations that are specific to the individual organism's DNA and may not exist in the reference genome. Due to these potential differences, the similarity between a read and a reference sequence segment must be identified using an *approximate string matching* (ASM) algorithm. The possible genetic differences between the reference genome and the sequenced genome are then identified using *genomic variant calling* algorithms [4].

The ASM performed during read mapping typically uses a computationally-expensive dynamic programming (DP) algorithm. This time-consuming algorithm has long been a major bottleneck in the *entire* genome analysis pipeline, accounting for over 70% of the execution time of read mapping [5]. The vast majority of read mappers, such as the widely-used minimap2 [6], are implemented as software running on CPUs. We refer readers to a comprehensive survey [7] for a discussion of state-of-the-art CPU-based read mappers. Accelerating ASM can help bridge the wide performance gap between sequencing machines and CPU-based read mapping algorithms, but faces four key challenges:

- 1) Due to the large datasets that a read mapper operates on, it generates a large amount of *data movement* between the CPU and main memory. The CPU accesses off-chip main memory through a pin-limited bus known as the *memory channel*, and a high amount of data movement across the memory channel is extremely costly in terms of both execution time and energy [8], [9].
- 2) Modern sequencing machines generate read fragments at an exponentially higher rate than prior sequencing technologies, with their growth far outpacing the growth in computational power in recent years [10]. For example, the Illumina NovaSeq 6000 system can sequence about 48 human whole genomes at  $30\times$  genome coverage (the average number of times a genomic base is sequenced) in about two days. However, analyzing (performing mapping and variant calling) the sequencing data of a *single* hu-

man genome requires over 32 CPU hours on a 48-core Intel Xeon processor, 23 of which are spent on read mapping [11].

- 3) The first two challenges worsen when a *metagenomic* sample is profiled, where the sample donor is unknown. This requires matching the extracted reads to *thousands* of reference genomes [12].
- 4) There is also an urgent need for rapidly incorporating clinical DNA sequencing and analysis into clinical practice for rapid surveillance of disease outbreaks (e.g., COVID-19 [13]) and early diagnosis of genetic disorders in critically ill infants [14].

Increasing the number of CPUs used for genome analysis decreases the overall analysis time, but significantly increases energy consumption and hardware costs. Cloud computing platforms are a potential alternative to distribute the workload at a reasonable cost, but are disallowed due to data protection guidelines in many countries [15].

As a result, there is a dire need for new computational techniques that can quickly process and analyze a tremendous number of extracted reads in order to drive cutting-edge advances in the genetic applications space [16]. Many works boost the performance of existing and new read mappers using new algorithms, hardware/software co-design, and hardware accelerators. Our *goal* in this work is to survey a prominent set of these three types of acceleration efforts for guiding the design of new highly-efficient read mappers. *To this end*, we (1) discuss various state-of-the-art mechanisms and techniques that improve the execution time of read mapping using different modern high-performance computing architectures, and (2) highlight the challenges, in the last section, that system architects and programmers must address to enable the widespread adoption of hardware-accelerated read mappers.

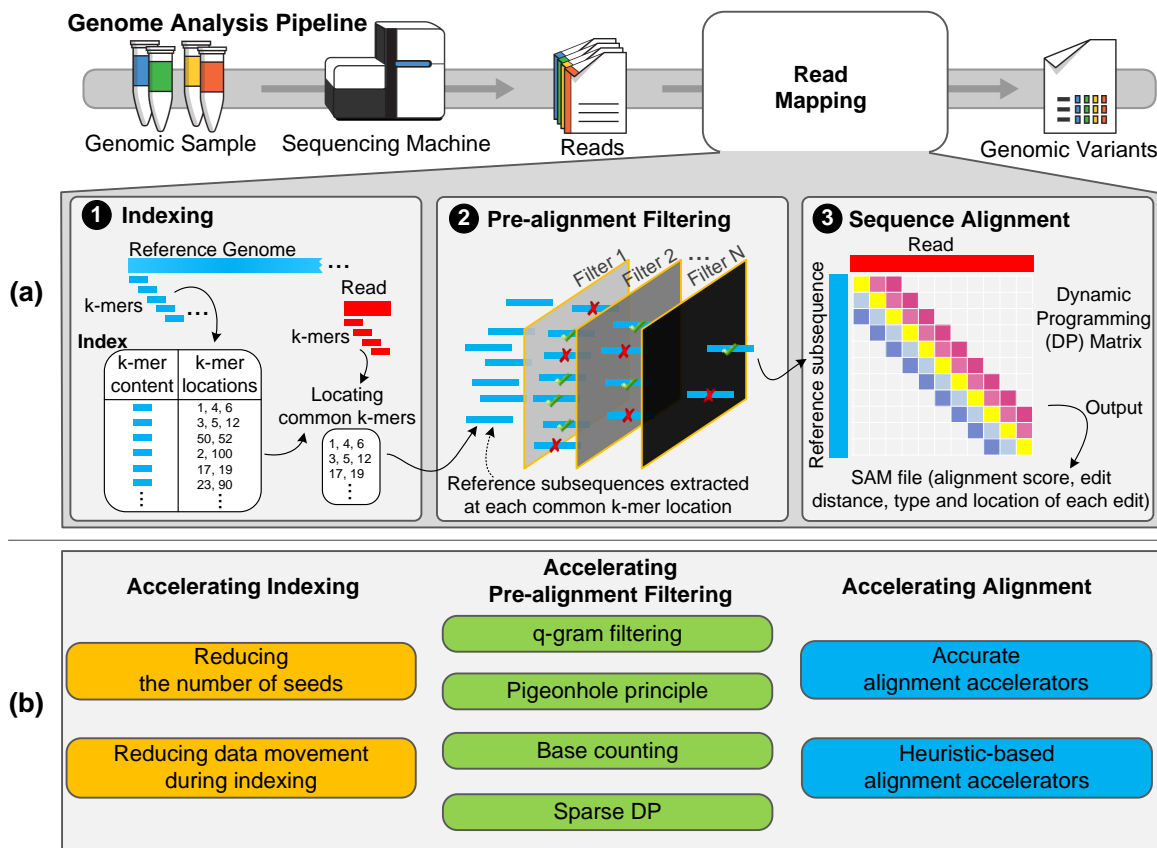
## Read Mapping

The main goal of read mapping is to locate possible subsequences of the reference genome sequence that are similar to the read sequence while allowing at most  $E$  edits, where  $E$  is the *edit distance threshold*. Commonly allowed edits include deletion, insertion, and substitution of characters in one or both sequences.

Edits can be as short as a single base pair (bp) alteration [17] or a much longer alteration (e.g., an insertion of about 600,000-base long region [18]). Mapping billions of reads to the reference genome is computationally expensive [16], [5], [19]. Therefore, most read mapping algorithms apply two key heuristic steps, *indexing* and *filtering*, to reduce the number of reference genome segments that need to be compared with each read. The three steps of read mapping are shown in Figure 1a. First, a read mapper indexes the reference genome by using substrings (called *seeds*) from each read to quickly identify all potential mapping locations of each read in the reference genome. Second, the mapper uses filtering heuristics to examine the similarity for every sequence pair (a read sequence and one potential matching segment in the reference genome identified during indexing). These filtering heuristics aim to eliminate most of the dissimilar sequence pairs. Third, the mapper performs sequence align-

ment (using ASM) to check whether or not the remaining sequence pairs that are identified by filtering to be similar are actually similar. The alignment step examines all possible *prefixes* of two sequences and tracks the prefixes that provide the highest possible *alignment score* (known as *optimal alignment*). The alignment score is a quantitative representation of the quality of an alignment for a given user-defined scoring function (computed based on the number of edits and/or matches).

Alignment algorithms typically use DP-based approaches to avoid re-examining the same prefixes many times. These DP-based algorithms provide the most accurate alignment results compared to other non-DP algorithms (such as the algorithm used in HISAT2 [20]), but they have quadratic time and space complexity (i.e.,  $O(m^2)$  for a sequence length of  $m$ ). Sequence alignment calculates information about the alignment such as the alignment score, *edit distance*, and the



**Figure 1.** (a) The three steps of read mapping in genome analysis: (1) indexing, (2) pre-alignment filtering, and (3) sequence alignment. (b) Overview of the existing approaches to accelerating each step of read mapping.

type of each edit. Edit distance is defined as the minimum number of changes needed to convert a sequence into the other sequence [21]. Such information is typically output by read mapping into a *sequence alignment/map* (SAM) file. Given the time spent on read mapping, all three steps have been targeted for acceleration. Figure 1b summarizes the different acceleration approaches, and we discuss a set of such works in the following sections.

### Accelerating Indexing

The indexing operation generates a table that is indexed by the contents of a seed, and identifies all locations where the seed exists in the reference genome. Indexing needs to be done only once for a reference genome, and eliminates the need to perform ASM across the entire genome. During read mapping, a seed from a read is looked up in the table, and only the corresponding locations are used for ASM (as only they can match the entire read). The major challenge with indexing is choosing the appropriate length and number of to-be-indexed seeds, as they can significantly impact the memory footprint and overall performance of read mapping [6]. Querying short seeds potentially leads to a large number of mapping locations that need to be checked for a string match. The use of long reads requires extracting from each read a large number of seeds, as the sequencing error rate is much higher in long reads. This affects (1) the number of times we query the index structure and (2) the number of retrieved mapping locations. Thus, there are two key approaches used for accelerating the indexing step (Figure 1b).

#### Reducing the Number of Seeds

Read mapping algorithms (e.g., minimap2 [6]) typically reduce the number of seeds that are stored in the index structure by finding the minimum representative set of seeds (called *minimizers*) from a group of adjacent seeds within a genomic region. The representative set can be calculated by imposing an ordering (e.g. lexicographically or by hash value) on a group of adjacent seeds and storing only the seed with the smallest order. Read mappers also apply heuristics to avoid examining the mapping locations of a seed that occur more times than a user-defined

threshold value [6]. Various data structures have been proposed and implemented to both reduce the storage cost of the indexing data structure and improve the algorithmic runtime of identifying the mapping locations within the indexing data structure. One example of such data structures is FM-index (implemented by Langarita et al. [22]), which provides a compressed representation of the full-text index, while allowing for querying the index without the need for decompression. This approach has two main advantages. 1) We can query seeds of arbitrary lengths, which helps to reduce the number of queried seeds. 2) It typically has less (by 1.5 – 2×) memory footprint compared to that of the indexing step of minimap2 [6]. However, one major bottleneck of FM-indexes is that locating the exact matches by querying the FM-index is significantly slower than that of classical indexes [23], [22]. BWA-MEM2 [23] proposes an uncompressed version of the FM-index that is at least 10× larger than the compressed FM-index to speed up the querying step by 2×.

#### Reducing Data Movement During Indexing

RADAR [24] observes that the indexing step is memory intensive, because the large number of random memory accesses dominates computation. The authors propose a processing-in-memory (PIM) architecture that stores the entire index inside the memory and enables querying the same index concurrently using a large number of ASIC compute units. The amount of data movement is reduced from tens of gigabytes to a few bytes for a single query task, allowing RADAR to balance memory accesses with computation, and thus provide speedups and energy savings.

### Accelerating Pre-Alignment Filtering

After finding one or more potential mapping locations of the read in the reference genome, the read mapper checks the similarity between each read and each segment extracted at these mapping locations in the reference genome. These segments can be *similar* or *dissimilar* to the read, though they share common seeds. To avoid examining dissimilar sequences using computationally-expensive sequence alignment algorithms, read mappers typically use filtering heuristics that are called *pre-alignment filters*.

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 the edit distance threshold, then the two sequences are identified as dissimilar sequences and hence DP calculation is not needed. In practice, only genomic sequence pairs with an edit distance less than or equal to a user-defined threshold (i.e.,  $E$ ) provide useful data for most genomic studies [25], [5], [7]. Pre-alignment filters use one of four major approaches to quickly filter out the dissimilar sequence pairs: (1) the pigeonhole principle, (2) base counting, (3)  $q$ -gram filtering, or (4) sparse DP. Long read mappers typically use  $q$ -gram filtering or sparse DP, as their performance scales linearly with read length and independently of the edit distance.

### Pigeonhole Principle

The pigeonhole principle states that if  $E$  items are put into  $E+1$  boxes, then one or more boxes would be empty. This principle can be applied to detect dissimilar sequences and discard them from the candidate sequence pairs used for ASM. If two sequences differ by  $E$  edits, then they should share at least a single subsequence (free of edits) among  $E+1$  non-overlapping subsequences [5], where  $E$  is the edit distance threshold. For a read of length  $m$ , if there are no more than  $E$  edits between the read and the reference segment, then the read and reference segment are considered similar if they share at most  $E+1$  non-overlapping subsequences, with a total length of at least  $m - E$ .

The problem of identifying these  $E+1$  non-overlapping subsequences is highly parallelizable, as these subsequences are independent of each other. Shouji [5] exploits the pigeonhole principle to reduce the search space and provide a scalable architecture that can be implemented for any values of  $m$  and  $E$ , by examining common subsequences independently and rapidly with high parallelism. Shouji accelerates sequence alignment by 4.2-18.8 $\times$  without affecting the alignment accuracy. We refer the reader to the sidebar for a brief discussion of several other related works.

### Sidebar: Related Works on Pre-Alignment Filtering Using the Pigeonhole Principle

Pigeonhole-filtering-based pre-alignment filtering can accelerate read mappers even without specialized hardware. For example, the Adjacency Filter [1] accelerates sequence alignment by up to 19 $\times$ . The accuracy and speed of pre-alignment filtering with the pigeonhole principle have been rapidly improved over the last seven years. Shifted Hamming Distance (SHD) [2] uses SIMD-capable CPUs to provide high filtering speed, but supports a sequence length up to only 128 base pairs due to the SIMD register widths. GateKeeper [3] utilizes the large amounts of parallelism offered by FPGA architectures to accelerate SHD and overcome such sequence length limitations. MAGNET [4] provides a comprehensive analysis of all sources of filtering inaccuracy of GateKeeper and SHD. Shouji [5] leverages this analysis to improve the accuracy of pre-alignment filtering by up to two orders of magnitude compared to both GateKeeper and SHD, using a new algorithm and a new FPGA architecture. SneakySnake [6] achieves up to four orders of magnitude higher filtering accuracy compared to GateKeeper and SHD by mapping the pre-alignment filtering problem to the single net routing (SNR) problem in VLSI chip layout. SNR finds the shortest routing path that interconnects two terminals on the boundaries of a VLSI chip layout in the presence of obstacles. SneakySnake is the only pre-alignment filter that works on CPUs, GPUs, and FPGAs. GenCache [7] proposes to perform highly-parallel pre-alignment filtering inside the CPU cache to reduce data movement and improve energy efficiency, with about 20% cache area overhead. GenCache shows that using different existing pre-alignment filters together (similar approach to [26]), each of which operates only for a given edit distance threshold (e.g., using SHD only when  $E$  is between 1 and 5), provides a 2.5 $\times$  speedup over GenCache with a single pre-alignment filter.

### REFERENCES

1. Hongyi Xin et al. Accelerating Read Mapping with FastHASH. *BMC Genomics*, 2013.
2. Hongyi Xin et al. Shifted Hamming Distance: A Fast and Accurate SIMD-Friendly Filter to Accelerate Alignment Verification in Read Mapping. *Bioinformatics*, 2015.
3. Mohammed Alser et al. GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping. *Bioinformatics*, 2017.
4. Mohammed Alser et al. MAGNET: Understanding and Improving the Accuracy of Genome Pre-Alignment Filtering. *Transactions on Internet Research*, 2017.
5. Mohammed Alser et al. Shouji: A Fast and Efficient Pre-Alignment Filter for Sequence Alignment. *Bioinformatics*, 2019.
6. Mohammed Alser et al. SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs. arXiv:1910.09020 [q-bio.GN], 2019.
7. Anirban Nag et al. GenCache: Leveraging In-Cache Operators for Efficient Sequence Alignment. In *MICRO*, 2019.



### Base Counting

The base counting filter compares the numbers of bases (A, C, G, T) in the read with the corresponding base counts in the reference segment. If one sequence has, for example, three more Ts than another sequence, then their alignment has at least three edits. If the difference in count is greater than  $E$ , then the two sequences are dissimilar and the reference segment is discarded. The base counting filter is used in mrsFAST-Ultra [27] and GASSST [26]. Such a simple filtering approach rejects a significant fraction of dissimilar sequences (e.g., 49.8%–80.4% of sequences, as shown in GASSST [26]) and thus avoids a large fraction of expensive verification computations required by sequence alignment algorithms.

### $q$ -gram Filtering Approach

The  $q$ -gram filtering approach considers all of the sequence's possible overlapping substrings of length  $q$  (known as  $q$ -grams). Given a sequence of length  $m$ , there are  $m - q + 1$  overlapping  $q$ -grams that are obtained by sliding a window of length  $q$  over the sequence. A single difference in one of the sequences can affect at most  $q$  overlapping  $q$ -grams. Thus,  $E$  differences can affect no more than  $q \cdot E$   $q$ -grams, where  $E$  is the edit distance threshold. The *minimum number* of shared  $q$ -grams between two similar sequences is therefore  $(m - q + 1) - (q \cdot E)$ . This filtering approach requires very simple operations (e.g., sums and comparisons), which makes it attractive for hardware acceleration, such as in GRIM-Filter [25]. GRIM-Filter exploits the high memory bandwidth and computation capability in the logic layer of 3D-stacked memory to accelerate  $q$ -gram filtering in the DRAM chip itself, using a new representation of reference genome that is friendly to in-memory processing.  $q$ -gram filtering is generally robust in handling only a small number of edits, as the presence of edits in any  $q$ -gram is significantly underestimated (e.g., counted as a single edit) [28].

### Sparse Dynamic Programming

Sparse DP algorithms exploit the exact matches (seeds) shared between a read and a reference segment to reduce execution time. These algorithms exclude the corresponding locations

of these seeds from estimating the number of edits between the two sequences, as they were already detected as exact matches during indexing. Sparse DP filtering techniques apply DP-based alignment algorithms only between every two non-overlapping seeds to quickly estimate the total number of edits. This approach is also known as *chaining*, and is used in minimap2 [6] and rHAT [29]. The recent work in [30] presents GPU and FPGA accelerators that achieve  $7\times$  and  $28\times$  acceleration, respectively, compared to the sequential implementation (executed with 14 CPU threads) of the chaining algorithm used in minimap2.

### Accelerating Sequence Alignment

After filtering out most of the mapping locations that lead to dissimilar sequence pairs, read mapping calculates the sequence alignment information for every read and reference segment extracted at each mapping location. Sequence alignment calculation is typically accelerated using one of two approaches: (1) accelerating the DP-based algorithms using hardware accelerators without altering algorithmic behavior, and (2) developing heuristics that sacrifice the optimality of the alignment score solution in order to reduce alignment time.

Despite more than three decades of attempts to accelerate sequence alignment, the fastest known edit distance algorithm [31] has a nearly quadratic running time,  $O(m^2/\log_2 m)$  for a sequence length of  $m$ , which is proven to be a tight bound [32]. With maintaining the algorithmic behavior as in the first approach, it is challenging to rapidly calculate sequence alignment of long reads with high parallelism. As long reads have high sequencing error rates (up to 20% of the read length), the edit distance threshold for long reads is typically higher than that for short reads, which results in calculating more entries in the DP matrix compared to that of short reads. The use of heuristics (i.e., the second approach) helps to reduce the number of calculated entries in the DP matrix and hence allows both the execution time and memory footprint to grow only linearly with read length (as opposed to quadratically with classical DP). Next, we describe the two approaches in detail.

## Accurate Alignment Accelerators

From a hardware perspective, sequence alignment acceleration has five directions: (1) using SIMD-capable CPUs, (2) using multicore CPUs and GPUs, (3) using FPGAs, (4) using ASICs, and (5) using processing-in-memory architectures. Traditional DP-based algorithms are typically accelerated by computing only the necessary regions (i.e., diagonal vectors) of the DP matrix rather than the entire matrix, as proposed in Ukkonen’s banded algorithm [33]. This reduces the search space of the DP-based algorithm and reduces computation time. The number of diagonal bands required for computing the DP matrix is  $2E+1$ , where  $E$  is the edit distance threshold. For example, the number of entries in the banded DP matrix for a 2 Mb long read can be 1.2 trillion. Parasail [34] and KSW2 (used in minimap2 [6]) exploit both Ukkonen’s banded algorithm and SIMD-capable CPUs to compute banded alignment for a sequence pair with a configurable scoring function. SIMD instructions offer significant parallelism to the matrix computation by executing the same vector operation on multiple operands at once. KSW2 is nearly as fast as Parasail when KSW2 does not use heuristics (explained in the next subsection).

The multicore architecture of CPUs and GPUs provides the ability to compute alignments of many independent sequence pairs concurrently. GASAL2 [35] exploits the multicore architecture of both CPUs and GPUs for highly-parallel computation of sequence alignment with a user-defined scoring function. Unlike other GPU-accelerated tools, GASAL2 transfers the bases to the GPU, without encoding them into binary format, and hides the data transfer time by overlapping GPU and CPU execution. GASAL2 is up to  $20\times$  faster than Parasail (when executed with 56 CPU threads). BWA-MEM2 [23] accelerates the banded sequence alignment of its predecessor (BWA-MEM [36]) by up to  $11.6\times$ , by leveraging multicore and SIMD parallelism. However, to achieve such levels of acceleration, BWA-MEM2 builds an index structure that is  $6\times$  larger than that of minimap2.

Other designs, such as FPGASW [37], exploit the very large number of hardware execution units in FPGAs to form a linear systolic array [38].

Each execution unit in the systolic array is responsible for computing the value of a single entry of the DP matrix. The systolic array computes a single vector of the matrix at a time. The data dependency between the entries restricts the systolic array to computing the vectors sequentially (e.g., top-to-bottom, left-to-right, or in an anti-diagonal manner). FPGASW has a similar execution time as its GPU implementation, but is  $4\times$  more power efficient.

Specialized hardware accelerators (i.e., ASIC designs) provide application-specific, power- and area-efficient solutions to accelerate sequence alignment. For example, GenAx [39] is composed of SillaX, a sequence alignment accelerator, and a second accelerator for finding seeds. SillaX supports both a configurable scoring function and traceback operations. SillaX is more efficient for short reads than for long reads, as it consists of an automata processor whose performance scales quadratically with the edit distance. GenAx is  $31.7\times$  faster than the predecessor of BWA-MEM2 (i.e., BWA-MEM [36]) for short reads.

Recent processing-in-memory architectures such as RAPID [40] exploit the ability to perform computation inside or near the memory chip to enable efficient sequence alignment. RAPID modifies the DP-based alignment algorithm to make it friendly to in-memory parallel computation by calculating two DP matrices: one for calculating substitutions and exact matches and another for calculating insertions and deletions. RAPID claims that this approach efficiently enables higher levels of parallelism compared to traditional DP algorithms. The main two benefits of RAPID and such PIM-based architectures are higher performance and higher energy efficiency [8], [9], as they alleviate the need to transfer data between the main memory and the CPU cores through slow and energy hungry buses, while providing high degree of parallelism with the help of PIM. RAPID is on average  $11.8\times$  faster and  $212.7\times$  more power efficient than 384-GPU cluster of GPU implementation of sequence alignment, known as CUDAlign [41].

## Heuristic-Based Alignment Accelerators

The second direction is to *limit* the functionality of the alignment algorithm or *sacrifice* the optimality of the alignment solution in order

to reduce execution time. The use of restrictive functionality and heuristics limits the possible applications of the algorithms that utilize this direction. Examples of limiting functionality include limiting the scoring function, or only taking into account accelerating the computation of the DP matrix without performing the backtracking step [42]. There are several existing algorithms and corresponding hardware accelerators that limit scoring function flexibility. Levenshtein distance [21] and Myers’s bit-vector algorithm [43] are examples of algorithms whose scoring functions are fixed, such that they penalize all types of edits equally when calculating the total alignment score. Restrictive scoring functions reduce the total execution time of the alignment algorithm and reduce the bit-width requirement of the register that accommodates the value of each entry in the DP matrix. ASAP [44] accelerates Levenshtein distance calculation by up to  $63.3\times$  using FPGAs compared to its CPU implementation. The use of a fixed scoring function as in Edlib [45], which is the state-of-the-art implementation of Myers’s bit-vector algorithm, helps to outperform Parasail (which uses a flexible scoring function) by  $12\text{--}1000\times$ . One downside of fixed function scoring is that it may lead to the selection of a suboptimal sequence alignment.

There are other algorithms and hardware architectures that provide low alignment time by trading off accuracy. Darwin [16] builds a customized hardware architecture to speed up the alignment process, by dividing the DP matrix into overlapping submatrices and processing each submatrix independently using systolic arrays. Darwin provides three orders of magnitude speedup compared to Edlib [45]. Dividing the DP matrix (known as the *Four-Russians Method* [46]) enables significant parallelism during DP matrix computation, but it leads to suboptimal alignment calculation [26]. Darwin claims that choosing a large submatrix size ( $\geq 320 \times 320$ ) and ensuring sufficient overlap ( $\geq 128$  entries) between adjacent submatrices may provide optimal alignment calculation for some datasets.

There are other proposals that limit the number of calculated entries of the DP matrix based on one of two approaches: (1) using sparse DP or (2) using a greedy approach to maintain a high alignment score. Both approaches

suffer from providing suboptimal alignment calculation [47], [48]. The first approach uses the same sparse DP algorithm used for pre-alignment filtering but as an alignment step, as done in the `exonerate` tool [47]. The second approach is employed in *X-drop* [48], which (1) avoids calculating entries (and their neighbors) whose alignment scores are more than  $X$  below the highest score seen so far (where  $X$  is a user-specified parameter), and (2) stops early when a high alignment score is not possible. The *X-drop* algorithm is guaranteed to find the optimal alignment between relatively-similar sequences for *only some* scoring functions [48]. A similar algorithm (known as *Z-drop*) makes KSW2 at least  $2.6\times$  faster than Parasail. A recent GPU implementation [49] of the *X-Drop* algorithm is  $3.1\text{--}120.4\times$  faster than KSW2.

There are also a large number of edit distance approximation algorithms that provide a reduction in time complexity (e.g.,  $O(m^{1.647})$  instead of  $O(m^2)$ ), but they suffer from providing overestimated edit distance [50], [51], [52], [53].

## Discussion and Future Opportunities

Despite more than two decades of attempts, bridging the performance gap between sequencing machines and read mapping is still challenging. We summarize four main challenges below.

First, we need to accelerate the entire read mapping process rather than its individual steps. Accelerating only a single step of read mapping limits the overall achieved speedup according to Amdahl’s Law. Illumina and NVIDIA have recently started following a more holistic approach, and they claim to accelerate genome analysis by more than  $48\times$ , mainly by using specialization and hardware/software co-design. Illumina has built an FPGA-based platform, called DRAGEN (<https://www.illumina.com/products/by-type/informatics-products/dragen-bio-it-platform.html>), that accelerates all steps of genome analysis, including read mapping and variant calling. DRAGEN reduces the overall analysis time from 32 CPU hours to only 37 minutes [11]. NVIDIA has built Parabricks, a software suite accelerated using the company’s latest GPUs. Parabricks (<https://developer.nvidia.com/clara-parabricks>) can analyze whole human genomes at  $30\times$  coverage in about 45 minutes.



Second, we need to reduce the high amount of data movement that takes place during genome analysis. Moving data (1) between compute units and main memory, (2) between multiple hardware accelerators, and (3) between the sequencing machine and the computer performing the analysis incurs high costs in terms of execution time and energy. These costs are a significant barrier to enabling efficient analysis that can keep up with sequencing technologies, and some recent works try to tackle this problem [25], [8], [9]. GenASM [19] is a framework that uses bitvector-based ASM to accelerate multiple steps of the genome analysis pipeline, and is designed to be implemented inside 3D-stacked memory. Through a combination of hardware–software co-design to unlock parallelism, and processing-in-memory to reduce data movement, GenASM can perform (1) pre-alignment filtering for short reads, (2) sequence alignment for both short and long reads, and (3) whole genome alignment, among other use cases. For short/long read alignment, GenASM achieves  $111\times/116\times$  speedup over state-of-the-art software read mappers while reducing power consumption by  $33\times/37\times$ . DRAGEN reduces data movement between the sequencing machine and the computer performing analysis by adding specialized hardware support inside the sequencing machine for data compression. However, this still requires movement of compressed data. Performing read mapping inside the sequencing machine itself can significantly improve efficiency by eliminating sequencer-to-computer movement, and embedding a single specialized chip for read mapping within a portable sequencing device can potentially enable new applications of genome sequencing (e.g., rapid surveillance of diseases such as Ebola [54] and COVID-19 [13], near-patient testing, bringing precision medicine to remote locations). Unfortunately, efforts in this direction remain very limited.

Third, we need to develop flexible hardware architectures that do not conservatively limit the range of supported parameter values at design time. Commonly-used read mappers (e.g., `minimap2`) have different input parameters, each of which has a wide range of input values. For example, the edit distance threshold is typically

user defined and can be very high (15-20% of the read length) for recent long reads. A configurable scoring function is another example, as it determines the number of bits needed to store each entry of the DP matrix (e.g., DRAGEN imposes a restriction on the maximum frequency of seed occurrence). Due to rapid changes in sequencing technologies (e.g., high sequencing error rate and longer read lengths) [55], [56], these design restrictions can quickly make specialized hardware obsolete. Thus, read mappers need to adapt their algorithms and their hardware architectures to be modular and scalable so that they can be implemented for any sequence length and edit distance threshold based on the sequencing technology.

Fourth, we need to adapt existing genomic data formats for hardware accelerators or develop more efficient file formats. Most sequencing data is stored in the FASTQ/FASTA format, where each base takes a single byte (8 bits) of memory. This encoding is inefficient, as only 2 bits (3 bits when the ambiguous base, N, is included) are needed to encode each DNA base. The sequencing machine converts sequenced bases into FASTQ/FASTA format, and hardware accelerators convert the file contents into unique (for each accelerator) compact binary representations for efficient processing. This process that requires multiple format conversions wastes time. For example, only 43% of the sequence alignment time in BWA-MEM2 [23] is spent on calculating the DP matrix, while 33% of the sequence alignment time is spent on pre-processing the input sequences for loading into SIMD registers, as provided in [23]. To address this inefficiency, we need to widely adopt efficient hardware-friendly formats, such as UCSC’s 2bit format (<https://genome.ucsc.edu/goldenPath/help/twoBit>), to maximize the benefits of hardware accelerators and reduce resource utilization. We are not aware of any recent read mapper that uses such formats.

The acceleration efforts we highlight in this work represent state-of-the-art efforts to reduce current bottlenecks in the genome analysis pipeline. We hope that these efforts and the challenges we discuss provide a foundation for future work in accelerating read mappers and developing other genome sequence analysis tools.

## Acknowledgments

The work of Onur Mutlu's SAFARI Research Group was supported by funding from Intel, the Semiconductor Research Corporation, VMware, and the National Institutes of Health (NIH).

## ■ REFERENCES

1. Travis C Glenn. Field Guide to Next-Generation DNA Sequencers. *Molecular Ecology Resources*, 2011.
2. Simon Ardui, Adam Ameer, Joris R Vermeesch, and Matthew S Hestand. Single Molecule Real-Time (SMRT) Sequencing Comes of Age: Applications and Utilities for Medical Diagnostics. *Nucleic Acids Research*, 2018.
3. Miten Jain, Sergey Koren, Karen H Miga, Josh Quick, Arthur C Rand, Thomas A Sasani, John R Tyson, Andrew D Beggs, Alexander T Dilthey, Ian T Fiddes, et al. Nanopore sequencing and assembly of a human genome with ultra-long reads. *Nature biotechnology*, 2018.
4. Steve S Ho, Alexander E Urban, and Ryan E Mills. Structural variation in the sequencing era. *Nature Reviews Genetics*, pages 1–19, 2019.
5. Mohammed Alser, Hasan Hassan, Akash Kumar, Onur Mutlu, and Can Alkan. Shouji: A Fast and Efficient Pre-Alignment Filter for Sequence Alignment. *Bioinformatics*, 2019.
6. Heng Li. minimap2: Pairwise Alignment for Nucleotide Sequences. *Bioinformatics*, 2018.
7. Mohammed Alser, Jeremy Rotman, Kodi Taraszka, Huwenbo Shi, Pelin Icer Baykal, Harry Taegyung Yang, Victor Xue, Sergey Knyazev, Benjamin D. Singer, Brunilda Balliu, David Koslicki, Pavel Skums, Alex Zelikovsky, Can Alkan, Onur Mutlu, and Serghei Mangul. Technology Dictates Algorithms: Recent Developments in Read Alignment. arXiv:2003.00110 [q-bio.GN], 2020.
8. Onur Mutlu, Saugata Ghose, Juan Gómez-Luna, and Rachata Ausavarungnirun. Processing data where it makes sense: Enabling in-memory computation. *Microprocessors and Microsystems*, 67:28–41, 2019.
9. Saugata Ghose, Amiralı Boroumand, Jeremie S Kim, Juan Gómez-Luna, and Onur Mutlu. Processing-in-memory: A workload-driven perspective. *IBM Journal of Research and Development*, 63(6):3–1, 2019.
10. Zachary D. Stephens, Skylar Y. Lee, Faraz Faghri, Roy H. Campbell, Chengxiang Zhai, Miles J. Efron, Ravishankar Iyer, Michael C. Schatz, Saurabh Sinha, and Gene E. Robinson. Big Data: Astronomical or Genomical? *PLoS Biology*, 2015.
11. Amit Goyal, Hyuk Jung Kwon, Kichan Lee, Reena Garg, Seon Young Yun, Yoon Hee Kim, Sunghoon Lee, and Min Seob Lee. Ultra-Fast Next Generation Human Genome Sequencing Data Processing Using DRAGEN® Bio-IT Processor for Precision Medicine. *Open Journal of Genetics*, 2017.
12. Nathan LaPierre, Mohammed Alser, Eleazar Eskin, David Koslicki, and Serghei Mangul. Metalign: Efficient alignment-based metagenomic profiling via containment min hash. *BioRxiv*, 2020.
13. Joshua S Bloom, Eric M Jones, Molly Gasperini, Nathan B Lubock, Laila Sathe, Chetan Munugala, A Sina Boeshaghi, Oliver F Brandenburg, Longhua Guo, James Boockock, et al. Swab-seq: A high-throughput platform for massively scaled up sars-cov-2 testing. *medRxiv*, 2020.
14. Jan M Friedman, Yvonne Bombard, Martina C Cornel, Conrad V Fernandez, Anne K Junker, Sharon E Plon, Zornitza Stark, and Bartha Maria Knoppers. Genome-wide sequencing in acutely ill infants: genomic medicine's critical application? *Genetics in Medicine*, 21(2):498–504, 2019.
15. Ben Langmead and Abhinav Nellore. Cloud Computing for Genomic Data Analysis and Collaboration. *Nature Reviews Genetics*, 2018.
16. Yatish Turakhia, Gill Bejerano, and William J. Dally. Darwin: A Genomics Co-Processor Provides Up to 15,000X Acceleration on Long Read Assembly. In *ASPLOS*, 2018.
17. Rasmus Nielsen, Joshua S Paul, Anders Albrechtsen, and Yun S Song. Genotype and SNP calling from next-generation sequencing data. *Nature Reviews Genetics*, 12(6):443–451, 2011.
18. Sébastien Jacquemont, Alexandre Reymond, Flore Zuferey, Louise Harewood, Robin G Walters, Zoltán Kutalik, Danielle Martinet, Yiping Shen, Armand Valsesia, Noam D Beckmann, et al. Mirror extreme BMI phenotypes associated with gene dosage at the chromosome 16p11.2 locus. *Nature*, 478(7367):97–102, 2011.
19. Damla Senol Cali, Gurpreet S. Kalsi, Zülal Bingöl, Lavanya Subramanian, Can Firtina, Jeremie S. Kim, Rachata Ausavarungnirun, Mohammed Alser, Anant Nori, Juan Gómez Luna, Amiralı Boroumand, Allison Scibisz, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu. GenASM: A Low-Power, Memory-Efficient Approximate String Matching Acceleration Framework for Genome Sequence Analysis. In *Proc. 53rd Int. Symp. Microarchitecture (MICRO)*, 2020.
20. Daehwan Kim, Joseph M Paggi, Chanhee Park, Christopher Bennett, and Steven L Salzberg. Graph-based genome alignment and genotyping with HISAT2 and HISAT-genotype. *Nature biotechnology*, 37(8):907–915, 2019.

21. Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet Physics-Doklady*, volume 10, pages 707–710, 1966.
22. Ruben Langarita, Adria Armejach, Javier Setoain, Pablo Enrique Ibanez Marin, Jesús Alastruey-Benedé, and Miquel Moreto Planas. Compressed Sparse FM-Index: Fast Sequence Alignment Using Large K-Steps. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2020.
23. Md Vasimuddin, Sanchit Misra, Heng Li, and Srinivas Aluru. Efficient Architecture-Aware Acceleration of BWA-MEM for Multicore Systems. In *IPDPS*, 2019.
24. Wenqin Huangfu, Shuangchen Li, Xing Hu, and Yuan Xie. RADAR: A 3D-ReRAM Based DNA Alignment Accelerator Architecture. In *DAC*, 2018.
25. Jeremie S. Kim, Damla Senol Cali, Hongyi Xin, Donghyuk Lee, Saugata Ghose, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, and Onur Mutlu. GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies. *BMC Genomics*, 2018.
26. Guillaume Rizk and Dominique Lavenier. GASSST: Global Alignment Short Sequence Search Tool. *Bioinformatics*, 2010.
27. Faraz Hach, Iman Sarrafi, Farhad Hormozdiari, Can Alkan, Evan E Eichler, and S Cenk Sahinalp. mrsFAST-Ultra: a compact, SNP-aware mapper for high performance sequencing applications. *Nucleic acids research*, 42(W1):W494–W500, 2014.
28. David Weese, Manuel Holtgrewe, and Knut Reinert. RazerS 3: faster, fully sensitive read mapping. *Bioinformatics*, 28(20):2592–2599, 2012.
29. Bo Liu, Dengfeng Guan, Mingxiang Teng, and Yadong Wang. rHAT: fast alignment of noisy long reads with regional hashing. *Bioinformatics*, 32(11):1625–1631, 2016.
30. Licheng Guo, Jason Lau, Zhenyuan Ruan, Peng Wei, and Jason Cong. Hardware acceleration of long read pairwise overlapping in genome sequencing: A race between FPGA and GPU. In *2019 IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 127–135. IEEE, 2019.
31. William J Masek and Michael S Paterson. A faster algorithm computing string edit distances. *Journal of Computer and System sciences*, 20(1):18–31, 1980.
32. Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM Journal on Computing*, 47(3):1087–1097, 2018.
33. Esko Ukkonen. Algorithms for approximate string matching. *Information and control*, 64(1-3):100–118, 1985.
34. Jeff Daily. Parasail: SIMD C Library for Global, Semi-Global, and Local Pairwise Sequence Alignments. *BMC Bioinformatics*, 2016.
35. Nauman Ahmed, Jonathan Lévy, Shanshan Ren, Hamid Mushtaq, Koen Bertels, and Zaid Al-Ars. GASAL2: A GPU Accelerated Sequence Alignment Library for High-Throughput NGS Data. *BMC Bioinformatics*, 2019.
36. Heng Li. Aligning sequence reads, clone sequences and assembly contigs with bwa-mem. *arXiv preprint arXiv:1303.3997*, 2013.
37. Xia Fei, Zou Dan, Lu Lina, Man Xin, and Zhang Chunlei. FPGASW: Accelerating Large-Scale Smith–Waterman Sequence Alignment Application with Backtracking on FPGA Linear Systolic Array. *Interdisciplinary Sciences: Computational Life Sciences*, 2018.
38. HT Kung. Why Systolic Architectures? *IEEE Computer*, (1):37–46, 1982.
39. Daichi Fujiki, Arun Subramanian, Tianjun Zhang, Yu Zeng, Reetuparna Das, David Blaauw, and Satish Narayanasamy. GenAx: A Genome Sequencing Accelerator. In *ISCA*, 2018.
40. Saransh Gupta, Mohsen Imani, Behnam Khaleghi, Venkatesh Kumar, and Tajana Rosing. RAPID: A ReRAM Processing in-Memory Architecture for DNA Sequence Alignment. In *2019 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, pages 1–6. IEEE, 2019.
41. Edans Flavius de Oliveira Sandes, Guillermo Miranda, Xavier Martorell, Eduard Ayguade, George Teodoro, and Alba Cristina Magalhaes Melo. CUDAlign 4.0: Incremental speculative traceback for exact chromosome-wide alignment in GPU clusters. *IEEE Transactions on Parallel and Distributed Systems*, 27(10):2838–2850, 2016.
42. Peng Chen, Chao Wang, Xi Li, and Xuehai Zhou. Accelerating the Next Generation Long Read Mapping with the FPGA-Based System. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2014.
43. Gene Myers. A fast bit-vector algorithm for approximate string matching based on dynamic programming. *Journal of the ACM (JACM)*, 46(3):395–415, 1999.
44. Subho Sankar Banerjee, Mohamed El-Hadedy, Jong Bin Lim, Zbigniew T. Kalbarczyk, Deming Chen, Steven S. Lumetta, and Ravishankar K. Iyer. ASAP: Accelerated Short-Read Alignment on Programmable Hardware. *IEEE Transactions on Computers*, 2019.
45. Martin Šošić and Mile Šikić. Edlib: A C/C++ Library

- for Fast, Exact Sequence Alignment Using Edit Distance. *Bioinformatics*, 2017.
46. Vladimir L'vovich Arlazarov, Yefim A Dinitz, MA Kronrod, and Igor Aleksandrovich Faradzhev. On economical construction of the transitive closure of an oriented graph. In *Doklady Akademii Nauk*, volume 194, pages 487–488. Russian Academy of Sciences, 1970.
  47. Guy St. C. Slater and Ewan Birney. Automated Generation of Heuristics for Biological Sequence Comparison. *BMC Bioinformatics*, 2005.
  48. Zheng Zhang, Scott Schwartz, Lukas Wagner, and Webb Miller. A Greedy Algorithm for Aligning DNA Sequences. *Journal of Computational Biology*, 2000.
  49. Alberto Zeni, Giulia Guidi, Marquita Ellis, Nan Ding, Marco D. Santambrogio, Steven Hofmeyr, Aydın Buluç, Leonid Olikier, and Katherine Yelick. LOGAN: High-Performance GPU-Based X-Drop Long-Read Alignment. arXiv:2002.05200 [q-bio.GN], 2020.
  50. Tuğkan Batu, Funda Ergun, and Genk Sahinalp. Oblivious string embeddings and edit distance approximations. In *Proceedings of the seventeenth annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 792–801. Society for Industrial and Applied Mathematics, 2006.
  51. Alexandr Andoni and Krzysztof Onak. Approximating edit distance in near-linear time. *SIAM Journal on Computing*, 41(6):1635–1648, 2012.
  52. Diptarka Chakraborty, Debarati Das, Elazar Goldenberg, Michal Koucky, and Michael Saks. Approximating edit distance within constant factor in truly sub-quadratic time. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 979–990. IEEE, 2018.
  53. Moses Charikar, Ofir Geri, Michael P Kim, and William Kuszmaul. On Estimating Edit Distance: Alignment, Dimension Reduction, and Embeddings. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
  54. Joshua Quick, Nicholas J Loman, Sophie Duraffour, Jared T Simpson, Ettore Severi, Lauren Cowley, Joseph Akoi Bore, Raymond Koundouno, Gytis Dudas, Amy Mikhail, et al. Real-time, portable genome sequencing for Ebola surveillance. *Nature*, 530(7589):228–232, 2016.
  55. Damla Senol Cali, Jeremie S Kim, Saugata Ghose, Can Alkan, and Onur Mutlu. Nanopore sequencing technology and tools for genome assembly: Computational analysis of the current state, bottlenecks and future directions. *Briefings in Bioinformatics*, 20(4):1542–1559, 2019.
  56. Can Firtina, Jeremie S Kim, Mohammed Alser, Damla Senol Cali, A Ercument Cicek, Can Alkan, and Onur Mutlu. Apollo: a sequencing-technology-independent, scalable and accurate assembly polishing algorithm. *Bioinformatics*, 36(12):3669–3679, 2020.
- Mohammed Alser** is with ETH Zürich, Switzerland. Contact him at [alserm@inf.ethz.ch](mailto:alserm@inf.ethz.ch)
- Zülal Bingöl** is with Bilkent University, Turkey. Contact her at [zulal.bingol@bilkent.edu.tr](mailto:zulal.bingol@bilkent.edu.tr)
- Damla Senol Cali** is with Carnegie Mellon University, USA. Contact her at [dsenol@andrew.cmu.edu](mailto:dsenol@andrew.cmu.edu)
- Jeremie Kim** is with ETH Zürich, Switzerland, and Carnegie Mellon University, USA. Contact him at [jeremie.kim@inf.ethz.ch](mailto:jeremie.kim@inf.ethz.ch)
- Saugata Ghose** is with the University of Illinois at Urbana–Champaign and Carnegie Mellon University, USA. Contact him at [ghose@illinois.edu](mailto:ghose@illinois.edu)
- Can Alkan** is with Bilkent University, Turkey. Contact him at [calkan@cs.bilkent.edu.tr](mailto:calkan@cs.bilkent.edu.tr)
- Onur Mutlu** is with ETH Zürich, Switzerland, Carnegie Mellon University, USA, and Bilkent University, Turkey. Contact him at [omutlu@gmail.com](mailto:omutlu@gmail.com)