EXMA: A Genomics Accelerator for Exact-Matching

Lei Jiang Indiana University Bloomington jianq60@iu.edu

Abstract—Genomics is the foundation of precision medicine, global food security and virus surveillance. Exact-match is one of the most essential operations widely used in almost every step of genomics such as alignment, assembly, annotation, and compression. Modern genomics adopts Ferragina-Manzini Index (FM-Index) augmenting space-efficient Burrows-Wheeler transform (BWT) with additional data structures to permit ultra-fast exactmatch operations. However, FM-Index is notorious for its poor spatial locality and random memory access pattern. Prior works create GPU-, FPGA-, ASIC- and even process-in-memory (PIM)based accelerators to boost FM-Index search throughput. Though they achieve the state-of-the-art FM-Index search throughput, the same as all prior conventional accelerators, FM-Index PIMs process only one DNA symbol after each DRAM row activation, thereby suffering from poor memory bandwidth utilization.

In this paper, we propose a hardware accelerator, EXMA, to enhance FM-Index search throughput. We first create a novel EXMA table with a multi-task-learning (MTL)-based index to process multiple DNA symbols with each DRAM row activation. We then build an accelerator to search over an EXMA table. We propose 2-stage scheduling to increase the cache hit rate of our accelerator. We introduce dynamic page policy to improve the row buffer hit rate of DRAM main memory. We also present CHAIN compression to reduce the data structure size of EXMA tables. Compared to state-of-the-art FM-Index PIMs, EXMA improves search throughput by $4.9\times$, and enhances search throughput per Watt by $4.8\times$.

Index Terms—Domain-Specific Hardware Accelerator, Genomics, Exact-Matching

I. INTRODUCTION

Because of the huge advancement of sequencing technologies such as Illumina [1], PacBio SMRT [2], and Oxford Nanopore [3], sequencing a entire human genome requires only < 1 day. The big genomic data has been a cornerstone to enabling personalized healthcare [4], and ensuring global food security [5]. Recently, genome sequencing becomes a powerful tool to fight virus outbreaks, e.g., Ebola [6], Zika [7] and COVID-19 [8].

However, it is challenging to process and analyze huge volumes of genomic data generated by high throughput sequencers that scale faster than Moore's Law [9]. For instance, thousands of USB-drive-size Oxford Nanopore Minion sequencers are deployed to monitor virus outbreaks [6]–[8] in the wild by generating several terabytes data per day. Analyzing a single genome may take hundreds of CPU hours [10], [11] even on high-end servers. To overcome the looming crisis of big genomic data, the application-specific hardware acceleration has become essential for genomics.

This work was partially supported by the National Science Foundation (NSF) through awards CCF-1908992 and CCF-1909509.

Farzaneh Zokaee Indiana University Bloomington fzokaee@iu.edu



Fig. 1. Execution time breakdown of human genome analysis (DynPro means dynamic programming).

A genome sequencing pipeline [4] sequences organic genomes, archives genomic data, analyzes genome sequences, and generates genetic variants that can used for patient treatment. Therefore, the *latency* of genome sequencing is a matter of life and death. Read alignment [12], which aligns reads, i.e., small DNA fragments, against a long genome reference, is identified as one of the most time-consuming steps [10], [11], [13]–[15] in genome analysis. Read alignment adopts the seed-and-extend paradigm [12], [16], and thus includes two major stages, i.e., seeding and seed extension. During seeding, parts of each read are mapped to their exactly matched positions, i.e., seeds, of the long reference by hash tables [10], [11], [17] or Ferragina-Manzini Index (FM-Index) [12], [16]. Seed extension pieces together a larger sequence with seeds and edit distance errors, i.e., insertions, deletions (indels) and substitutions, by dynamic programming [18]-[20].

State-of-the-art read alignment applications such as BWA-MEM [12], MA [16] and SOAP [21] use FM-Index to build super-maximal exact matches (SMEMs) during seeding, since it augments the space-efficient Burrows-Wheeler transform (BWT) [22] with accessory data structures that permit ultra-fast exact-match operations. SMEMs generated by FM-Index guarantee each seed does not overlap other seeds and has the maximal length that cannot be further extended. Compared to hash tables, FM-Index reduces not only the number of errors in output genome mappings but also the durations of seed extension substantially [23].

Besides read alignment, FM-Index is widely used for exactmatch operations in other time-consuming steps of genome analysis such as genome assembly [24], annotation [25] and compression [26]. Figure 1 shows the execution time breakdown of various genome analysis applications on a human genome¹. On average, FM-Index searches cost 31% $\sim 81\%$ of the execution time of these genome analysis applications. Since Illumina, Nanopore and PacBio genome sequencers generate reads having different lengths and error rates, aligning and assembling these reads require different

¹The experimental methodology is elaborated in §V.

amounts of time for FM-Index searches. The reads produced by Illumina machines have lower error rates, so an Illumina dataset invokes FM-Index searches more frequently.

However, FM-Index is notorious for its poor spatial locality and random memory access pattern [27]. The kernel of conventional FM-Index search is pointer chasing. After activating one DRAM row, FM-Index processes only one DNA symbol, thereby greatly decreasing DRAM bandwidth utilization. Although a recent algorithmic work, LISA [28], uses a learned index [29] to search multiple DNA symbols after each row activation, the learned index accuracy is low. LISA has to search many unnecessary entries, and thus achieves only moderate search throughput improvement. Beyond CPUs [12] and GPUs [21], prior work creates FPGA [13], [30]-, ASIC [31]-, and even processing-in-memory (PIM) [14], [15]-based designs to accelerate conventional FM-Index searches processing only one DNA symbol after each row activation. Therefore, these accelerators are fundamentally limited by the poor off-chip memory bandwidth utilization. Instead of searching multiple DNA symbols after each row activation, a recent DRAM PIM, MEDAL [15], achieves the state-of-the-art search throughput by enabling DRAM chip-level parallelism, where each chip can independently activate a partial row to process a DNA symbol. However, we observe that there are a lot of conflicts on the DDR4 address bus shared by all chips in a rank. The shared address bus seriously limits search throughput of MEDAL.

In this paper, we propose an algorithm and hardware codesigned accelerator, **EXMA**, to process **EXact-MA**tch operations during genome analysis. Our contributions are summarized as follows:

- An EXMA table with a MTL-based index We propose a novel data structure, EXMA table, that can process k DNA symbols, i.e., a k-mer, in a DRAM row in each FM-Index search iteration. We further present a multi-tasklearning (MTL)-based index to accelerate searches over an EXMA table. The MTL-based index trained with multiple kmers uses less neural network parameters, but obtains higher accuracy over learning to search each k-mer independently.
- A hardware accelerator We build an accelerator to search an EXMA table with a MTL-based index. We present a 2stage scheduling to increase the hit rate of on-chip caches of our accelerator for the table and its index. We also propose dynamic page policy to improve the row buffer hit rate of DRAM main memory. At last, we introduce CHAIN compression to greatly reduce the data structure size of an EXMA table.
- Search throughput and throughput per Watt We evaluated and compare EXMA to prior CPU-, GPU-, FPGA-, ASIC-, and PIM-based FM-Index accelerators. Compared to the state-of-the-art DRAM PIM MEDAL, EXMA improves search throughput by 4.9×, and enhances search throughput per Watt by 4.8×.



A. Read Alignment

Seed-and-Extend: As one of the bottlenecks in genome analysis, read alignment may consume hundreds of CPU hours [10], [11], [13]. During read alignment, DNA reads generated by various sequencing machines, e.g., Illumina, PacBio SMRT, and Oxford Nanopore, are mapped to a preexisting genome reference, as shown in Figure 2(a). Read alignment is complicated by the fact that there are genetic variations in the human population, and sequencing machines also introduce sequencing errors [32]. The overall variation of human population has been estimated as 0.1% [9], while the sequencing error rate of various sequencing machines is $0.2\% \sim 30\%$ [32]. To reduce sequencing errors, a sequencing machine produces $30 \sim 50$ reads to cover every position in the genome. As Figure 2(b) shows, read alignment adopts the seed-and-extend paradigm [12], [16], [24], [33]-[35] to accommodate sequencing errors and genetic variations. During seeding, a read is divided into multiple smaller parts that are aligned against the reference. If a part is exactly matched, it becomes a seed. The computationally expensive Smith-Waterman algorithm [10], [11], [13] is invoked only around seeds to handle sequencing errors and genetic variations.

Exact-Match Operation: The alphabet \sum of DNA includes A, C, G and T. Given a genome reference $\mathcal{G} \in \sum^*$ of length $|\mathcal{G}|$ and a query $Q \in \sum^*$ of length |Q|, the seeding, aka exact-match problem, is to find all occurrences of Q in \mathcal{G} . A naïve algorithm of exhausting all possible positions for Q will take $\mathcal{O}(|\mathcal{G}||Q|)$ comparisons, which is infeasible for large genome. It is possible to use a hash table [10], [11], [17] to support exact-match operations with hundreds of gigabytes DRAM, but the hash-table-based seeding not only degrades genome mapping quality but also prolongs seed extension durations [23]. State-of-the-art alignment algorithms [12], [16], [21] use FM-Index for seeding. To search a query Q over the reference genome \mathcal{G} , FM-Index occupies $\mathcal{O}(|\mathcal{G}|log(|\mathcal{G}|))$ DRAM space and does $\mathcal{O}(|Q|)$ comparisons during a search.

B. FM-Index

1) Burrows-Wheeler Transform

FM-index is built upon BWT [22]. To compute the BWT of a genome reference \mathcal{G} , we can list all its circularly shifted sequences. For instance, Figure 3(a) shows $\mathcal{G} = CATAGA$, where \$ indicates the end of the sequence and it is the lexicographically smallest symbol. Circularly shifted sequences of $\mathcal{G} = CATAGA$ \$ can be listed as CATAGA, A\$CATAGA, ..., and CATAGA\$, which can be sorted in the *lexicographical* order to form a Burrows-Wheeler (BW)-matrix. The last column of the BW-matrix is the BWT of \mathcal{G} , i.e.,



Fig. 3. FM-Index overview: (a) the BWT of a genome reference $\mathcal{G} = CATAGA$; (b) an Occ(s, i) table; (c) a Count(s) table; (d) backward search; (e) a search example; and (f) a bucket-based data structure

 $BWT(\mathcal{G}) = AGTC$ \$AA. The sub-sequence ending with \$ in each row of the BW-matrix is a suffix of \mathcal{G} , which can be denoted by an integer (SA in Figure 3(a)) recording its starting position in the reference. For example, ATAGA\$ is SA[3] = 1, which means it starts from the position 1 of \mathcal{G} .

2) FM-index

The data structure and search algorithm of FM-Index can be summarized as:

- Occ and Count. FM-index searches are implemented with two functions Occ(s, i) and Count(s) over the BWT of G. As Figure 3(b) exhibits, Occ(s, i) returns the number of symbol s in the BWT from the position 0 to the position i 1, e.g., Occ(C, 5) = 1, which means that there is only 1 C from the position 0 to the position 4 of BWT(G) = AGTC\$AA. Count(s) shown in Figure 3(c) computes the number of symbols in the BWT that are lexicographically smaller than the symbol s, e.g., Count(T) = 6, which indicates that there are 6 symbols in BWT(G) = AGTC\$AA lexicographically smaller than T.
- Backward Search. An exact-match operation is implemented by backward search, whose algorithm can be viewed in Figure 3(d). The interval (low, high) covers a range of indices in the BW-matrix where the suffixes have the same prefix. The pointer low locates the index in the BW-matrix where the pattern is first found as a prefix, while the pointer high provides the index after the one where the pattern is last found. At first, low and high are initialized to the minimum and maximum indexes of the Occ table respectively. And then, they iterate from the last symbol in a query Q to the first. The pointer pos is updated by

$$Count(Q[i]) + Occ(Q[i], pos)$$
(1)

where Q[i] indicates the i_{th} symbol in the query Q. The pointer pos can be low or high, as shown from the lines 2 to 3 in Figure 3(d). The computations of low and high are **pointer chasing** and thus suffer from poor spatial locality [14], [15], [30]. Finally, the interval (low, high) gives the range of indexes in the BW-matrix where the suffixes have the target query as a prefix. These indexes are converted to reference genome positions using SA. Figure 3(e) illustrates an example of searching a query TAG in the reference $\mathcal{G} = CATAGA$. Before a search happens,

(low, high) is initialized to (0, 7). In the iteration 0, the last symbol G is processed, and then (low, high) is updated to (5, 6). After three iterations, (low, high) eventually equals (6, 7). By looking up SA[6] = 2 in Figure 3(a), we find that the query TAG in reference $\mathcal{G} = CATAGA$ starts from the position 2.

• Bucket-based Storage. Both Count(s) and Occ(s, i) can be pre-calculated and stored. However, the storage overhead of Occ(s, i) is proportional to the genome reference length $|\mathcal{G}|$, and thus significant. To keep the storage overhead in check, the Occ(s, i) values are sampled into buckets of width d shown in Figure 3(f) (d = 4). The Occ(s, i) values are stored each d positions as markers to reduce the storage overhead by a factor of d. The omitted Occ(s, i) values can be reconstructed by summing the previous marker and the number of symbol s from the remaining positions in the BWT bucket. To simplify searches, Count(s) values of each symbol are added to corresponding markers. Markers and BWT buckets are interleaved to build a FM-Index.

3) Multi-step FM-Index

i	AA	AC… TG T1	
0	0	0 … 0 0	1
1	0	1 … 0 0	
2	1	1 … 0 0	
3	2	1 … 0 0	
4	2	1 … 1 0	
5	2	1…11	
6	3	1 … 1 1	
÷	÷	: : :	
١G	f ₀	$f_1 \cdots f_{14} f_{15}$	

During an iteration of the FM-Index backward search, two memory accesses for *low* and *high* are issued for each symbol in a query Q. Totally, 2|Q| memory accesses are required for an exact match operation of a query. The FM-Index backward search performance is seriously limited by random memory ac-

Fig. 4. A 2-step table. cesses [14], [15], [30], since each access opens a DRAM row but fetches only 64B. k-step FM-Index [36] is proposed to reduce the number of memory accesses to $\frac{2|Q|}{k}$ by updating a k-mer, i.e., k DNA symbols, in each search iteration. The idea of k-step FM-Index is to enlarge the alphabet size from Σ to Σ^k . For instance, if k = 2, instead of single DNA symbols, as Figure 4 shows, the enlarged alphabet includes 16 2-mers: AA, AC, \ldots, TT . We can construct a BWT with the enlarged alphabet and its corresponding FM-Index to perform k-step backward search in the same way. The trade-off for k-step FM-Index is the increase in its size, which is calculated as

$$F = \frac{\lceil log_2(|\mathcal{G}|)\rceil \cdot |\mathcal{G}| \cdot |\Sigma|^k}{8d} + \frac{|\mathcal{G}| \cdot \lceil log_2(|\Sigma|^k + 1)\rceil}{8} \quad (2)$$





where k is the number of DNA symbols updated in each search iteration. The size of multi-step FM-Index exponentially increases with an enlarging k.

4) Learned Indexes for Sequence Analysis

To support multi-step searches with smaller DRAM overhead, a recent work proposes Learned Indexes for Sequence Analysis (LISA) [28] consisting of an Index-Paired BWT (IP-BWT) array and a learned index.

- **IP-BWT**. In Figure 5(a), each entry of the IP-BWT is a pair of [k-mer, N], where k-mer is the first k symbols of the corresponding BW-matrix row, and N is the row number of the sequence with the first k and the last $|\mathcal{G}| k$ symbols swapped in the BW-matrix. For example, if k = 2, the k-mer of the row 0 of the IP-BWT can be derived from the row 0 of the BW-matrix, CATAGA, shown in Figure 3(a) using only the first 2 symbols of CATAGA, we can have ATAGA\$C, which is the row 3 of the BW-matrix. So the row 0 of the IP-BWT is [\$C, 3].
- Backward Search. The backward search of LISA finds the lower bound position of a [k-mer, N] pair in the IP-BWT. Since the IP-BWT is sorted, LISA adopts binary search for backward searches. As Figure 5(b) shows, to search the query TAG, we first break it into TA and G, since each iteration can process a 2-mer. In the first iteration, we start with G. The padding algorithm [28] of LISA converts G to G for low and GT for high. low and high are initialized to 0 and 6 respectively. ① To search [G\$, 0], a binary search is performed over the IP-BWT. ^(D) During the binary search, G is first compared against AT, i.e., the row 3 of the IP-BWT. **③** Because G\$ > AT, the binary search goes to the row 5, i.e., GA of the IP-BWT. **4** Finally, it ends with the row 4 of the IP-BWT, i.e., CA. **6** Since G > CA, the new low is calculated as 6+1=7. high can be computed in the same way. Each search iteration requires $log_2(|\mathcal{G}|)$ comparisons due to binary search.
- Learned Index. To reduce the number of comparisons during binary searches, LISA adopts a learned index [29], i.e., a model hierarchy consisting of multiple neural network models, as shown in Figure 5(c), where m_i is the neural network model *i*. The learned index enables LISA to do only one comparison during each iteration in the best case. To search [G\$, 0] by the learned index, we can traverse down lower-level neural network models based on the output of

the higher-level neural network models. Finally, a leaf neural network model predicts the position of [G\$, 0] in the IP-BWT. However, if the predicted position does not contain [G\$, 0], a linear search over the IP-BWT starts from the predicted position to find its actual position.

III. RELATED WORK AND DESIGN MOTIVATION

It is challenging to achieve high-throughput and powerefficient FM-Index searches by state-of-the-art FM-Index algorithms and accelerators.

A. Algorithm Inefficiency

Intractable Size of k-step FM-Index. We recorded the row IDs of 200 consecutive 1-step FM-Index search iterations in Figure 6(a), where 197 different rows are accessed. Because 1step FM-Index processes only 1 DNA symbol in each iteration, in most cases, searching a DNA symbol by 1-step FM-Index requires one row activation. Prior accelerators [14], [15] for 1-step FM-Index expect no row buffer hit and thus adopt close-page policy. Though k-step FM-Index can search k DNA symbols by activating a row, as Figure 6(b) shows, its data structure size exponentially increases with the step number k. For instance, 5-step FM-Index costs 105GB, while 6-step FM-Index occupies 374GB. As a result, 5-step FM-Index (FM-5) improves search throughput by only $1.21 \times$ over 1step FM-Index, as shown in Figure 6(d). Further enlarging the step number of k-step FM-Index (FM-6) decreases its search throughput improvement, due to more TLB misses.

Weakness of LISA Learned Index. LISA can search k DNA symbols after each row activation by its IP-BWT. Moreover, as Figure 6(b) shows, the size of LISA linearly increases with the step number k. However, the learned index of LISA suffers from low accuracy and low cache hit rate. First, the LISA learned index has to index $|\mathcal{G}|$ IP-BWT entries, where $|\mathcal{G}|$ is the length of reference genome. For a human genome, $|\mathcal{G}| = 3G$. When the learned index of LISA predicts a wrong position, LISA has to linearly search up to $|\mathcal{G}|$ IP-BWT entries. As Figure 6(c) describes, on average, LISA has to search $\sim 3K$ extra IP-BWT entries during each iteration, due to the low accuracy of its learned index. Consequently, compared to 1-step FM-Index, 21-step LISA (LISA-21) improves search throughput by only $2.15 \times$ in Figure 6(d). But if LISA has a perfect learned index (LISA-21P) which always predicts the right position, compared to 1-step FM-Index, LISA-21P improves search throughput by $5.1 \times$. Further increasing the step number of LISA (LISA-32) also introduces more TLB misses, thereby achieving smaller search throughput improvement. Second, traversing the learned index's hierarchical models is also pointer chasing. The LISA learned index consumes $\sim 1.5GB$. If we assume a perfect cache (100% hit) for LISA-21P (LISA-21PC), LISA-21PC improves search throughput by $8.53 \times$ over 1-step FM-Index.

Algorithmic Takeaways. (1) Implementing k-step search with moderately enlarged data structure is important for FM-Index. However, a larger step number, i.e., k, may not necessarily lead to better search throughput. (2) A more



Fig. 6. The inefficiency of prior FM-Index algorithms (A human genome dataset is used). In Figure 6(c), we present the maximal and minimal errors, the mean of errors, and the 25_{th} , 50_{th} , 75_{th} percentiles of errors. In Figure 6(d), FM-X means X-step conventional FM-Index; LISA-X means X-step LISA; LISA-XP means X-step LISA with a 100% accurate index; and LISA-XPC denotes X-step LISA with a 100% accurate index.

CK	0 1	2 3	4 5	[6 Z	89	10 11	12 13	14 15	16 17 1	8 19
RAS	RAS ₀	RAS ₁	RAS ₂	RAS ₃	dela	ay!!!···>	RAS ₃]		
CAS			•••••	CAS ₀	CAS ₁	CAS ₂]			
ADDR	R-A ₀	R-A ₁	R-A ₂	C-A ₀	C-A1	C-A ₂	R-A ₃]		
Chip ₀				:!				dat	a ₀	
Chip ₁									data	L (
Chip ₂								[data ₂	2 (
Chin										



accurate learned index reduces linear search overhead to improve search throughput. (3) A higher cache hit rate of learned index also improves search throughput by reducing redundant memory accesses to learned index.

B. Hardware Incapacity

FPGAs and ASICs. Most prior FM-Index hardware artifacts such as CPUs [12], GPUs [21], FPGAs [13], and ASICs [37] can accelerate only 1-step FM-Index searches. A recent FPGA design [30] supports 2-step FM-Index searches, while k-step LISA is built merely on CPUs. Existing FM-Index application-specific accelerators including FPGAs [13], [30], and ASICs [37] can search only 1 or 2 DNA symbols after each DRAM row activation. Therefore, they cannot fully exploit the maximal DRAM bandwidth.

Processing-In-Memories. Though recent works [14], [15], [38], [39] propose PIM accelerators to process FM-Index searches, they cannot fully utilize the available DRAM bandwidth either. Most NVM-based PIMs [14], [38], [39] focus on processing simple arithmetic computations of FM-Index by NVM arrays, but ignore optimizing external memory accesses. For instance, a ReRAM-based FM-Index PIM, FindeR [14], has only 2.6GB memory arrays, so it still suffers from low DRAM bandwidth utilization when fetching FM-Index buckets that cannot fit into its internal arrays from external DRAMs. Only a DRAM PIM, MEDAL [15], modifies its DRAM main memory for higher FM-Index search throughput. Instead of processing multiple DNA symbols during a DRAM row activation, MEDAL enables chip-level parallelism where each chip in a rank can independently process a DNA symbol by opening its partial row. In ideal case, 16 chips in a rank can enlarge FM-Index search throughput by $16\times$, and each chip has only 1/16 row size. However, we find the DDR4 address bus shared by all chips in a rank seriously limits search throughput of MEDAL. The DDR4 address bus is only 17bit wide [40]. During each access, the row activation and the column access serially pass their addresses via the same 17-



Fig. 8. A 2-step EXMA table (MAX means the end of increments of a k-mer; and f_i indicates the number of increments of the i_{th} k-mer).

bit bus. As Figure 7 shows, MEDAL can sequentially activate a 1/16 partial row in $chip_{0\sim2}$. But when activating a 1/16 partial row in $chip_3$, its row address $(R-A_3)$ and the column address of $chip_0$ $(C-A_0)$ compete for the same address bus. The activation in $chip_3$ has to be delayed to CK_{12} . And idle bubbles appear in the DDR4 data bus. Because of address bus conflicts, although MEDAL claims a $68 \times$ search throughput improvement over a multi-core CPU baseline, we observe it improves search throughput by only $11 \times$.

IV. EXMA

We first create a row-buffer-friendly alternative to FM-Index, EXMA table, to process multiple DNA symbols in each search iteration. And then, we present a Multi-task-Learning (MTL)-based index to accelerate searches over an EXMA table. Compared to LISA learned index, the MTL-based index is more accurate, although it has less parameters.

A. EXMA Table

Data Structure. The major reason why the learned index of LISA is not accurate is that it has to index $|\mathcal{G}|$ IP-BWT entries. To reduce the problem size for a learned index, we propose a novel data structure, EXMA table. In each row of the Occ table, only one k-mer can increase its value by one. For instance, in the 2-step Occ table shown in Figure 4, only "AC" increases its value from 0 to 1 in the row 1. Based on this observation, as Figure 8 shows, for each k-mer, an EXMA table stores only the row numbers of the Occ table, where its value increases. For example, for "AA", its value increases in the row 2, 3, 6, and others. We store these row number as the *increments* of "AA". Totally, we have f_0 "AA"s, f_1 "AC"s, ..., and f_{15} "TT"s in the *Occ* table, where f_0, \ldots, f_{15} are non-negative integers and $\sum_{i=0}^{15} f_i = |\mathcal{G}|$. So in the EXMA table, each 2-mer has f_i increments, where $0 \le i \le 15$. We add a symbol MAX to indicate the end of the increments of a k-mer, where $MAX = |\mathcal{G}| + 1$. The increments of an



EXMA table has the space complexity of $\mathcal{O}(|\mathcal{G}|log(|\mathcal{G}|))$. For a 3G-base human genome, the increments occupy 12GB. We can consecutively store the increments of all k-mers to take advantage of the row buffer locality. Each k-mer needs a *base* to point to its first increment. For instance, the base of "AC" is $f_0 + 1$ indicating its first increment is in the position $f_0 + 1$. Totally, a k-step EXMA table stores 4^k bases, each of which is related to a k-mer. Even if a k-mer has no increment, e.g., "TC", its base is set to $MAX = |\mathcal{G}| + 1$. The bases of a k-step EXMA table have the space complexity of $\mathcal{O}(4^k log(|\mathcal{G}|))$.

Backward Search. Each iteration in a backward search computes Equation 1. The entire *Count* table costs only several bytes, so the bottleneck is the *Occ* table lookups. To compute Occ(k-mer, pos) in each search iteration, we compare *pos* against all increments of the *k*-mer and find the first increment larger than *pos*, where *pos* can be *low* or *high*. For instance, to compute Occ(AA, 4), we first read the base of "AA", which is 0. And then, we initialize a counter and start a linearly search from the position 0 of "AA" to *MAX*. If an increment is smaller than 4, we increase the counter by 1. When 6 is found, we stop, since 6 > 4. At last, the counter value is 2, i.e., Occ(AA, 4) = 2.

Naïve Adoption of Learned Index. Since all increments of each k-mer are sorted, similar to LISA, we can adopt learned index [29] to perform only one comparison to compute Occ(k-mer, pos) in the best case. We build a learned index model hierarchy for each k-mer that has > 256 increments. As Figure 9(a) shows, similar to LISA [28], to build learned indexes, we use a fixed ratio between the number of parameters of a learned index model and the number of increments that need to be indexed. As a result, if a k-mer has more increments, its learned index model has more parameters. For instance, the model of "TT" (m_{TT}) owns more weights and biases than that of "AA" (m_{AA}), since "TT" has more increments. To compute Occ(AA, pos) in a search iteration, we first read all the parameters of m_{AA} , and input only pos to the model. If the prediction is not correct, we start a linear search from the predicted position to find the correct position. However, since an EXMA table have to index totally $|\mathcal{G}|$ increments, the learned index of EXMA does not have higher accuracy than that of LISA.

Step # of an EXMA Table. We tuned the step number k of an EXMA table to balance the DRAM overhead and search throughput. For a genome reference \mathcal{G} , the size of increments in an EXMA table is constant, while the size of bases in the table is proportional to 4^k . Although a small k has few bases, the search throughput is low. For instance, for a 3G-base human genome, if k = 2, the bases require only 32-byte. But each time, only a 2-mer can be processed by a search



(a) 15 "A"s (b) 7 "AC"s and "A" (c) 7 "AT"s and "G" Fig. 11. Increment distributions of 15-mers.

iteration. In contrast, a large k improves the search throughput but significantly increases the number of bases and thus the size of an EXMA table. For a human genome, as Figure 10(a) shows, a 15-step EXMA table (15) costs 29.5GB, while a 16step EXMA (16) occupies 41.5GB. Increasing k from 15 to 16 increases 12GB DRAM overhead. As Figure 10(b) describes, EXMA (EXMA-15) achieves its best search throughput with 15-step. Compared to LISA-21, EXMA-15 degrades search throughput by 7.3%, since it has a smaller step number.

B. Multi-task-Learning Index for EXMA

MTL-based Index. To more accurately approximate the cumulative distribution function of increments of each kmer in an EXMA table with less parameters, we propose a Multi-Task-Learning [41]-[43] (MTL)-based index for the increments of an EXMA table. As Figure 11 shows, the increments of various k-mers in 15-EXMA exhibit similar random distributions. Based on Stein's paradox [44], it is more accurate to approximate many independent random distributions using samples from all of them rather than approximating them separately. The MTL-based index is trained with the increments of multiple k-mers to obtain superior accuracy over learning the increments of each k-mer independently. We adopt the hard-parameter-sharing MTL [43] that shares a subset of parameters between the learned index models of k-mers having different numbers of increments. For instance, as Figure 9(b) shows, "AA" uses the smallest model (m_{AA}) to index its increments, since it has the fewest increments among all 2-mers. "TT" has more increments, and thus a larger model (m_{TT}) which contains the entire m_{AA} . Besides m_{AA} , m_{TT} comprises more levels of nodes to index its increments more accurately. Compared to the naïve learned index, we add more neurons in the hidden layers of each non-leaf node of a MTLbased index to accommodate two inputs, i.e., pos and k-mer. But the size of a MTL-based index is smaller that of the naïve learned index, since the k-mers share most parameters of a MTL-based index.

Inference. The MTL-based index of an EXMA table effectively approximates Occ(k-mer, pos) as

$$p = F(k\text{-}mer, pos) * f_{k\text{-}mer}$$
(3)



Fig. 12. Profiling EXMA-15 : (a) increment #; (b) search time breakdown. where p is the predicted position of pos in the increments of the k-mer; F(.,.) is the neural network model hierarchy to estimate the probability to observe an increment $\leq pos$; f_{k-mer} is the number of increments of the k-mer, and can be stored along the k-mer base. After each inference, we read both p and p+1 to check whether the prediction is correct. If not, we start a linear search to find the correct position. Therefore, the accuracy of a MTL-based index decides search throughput of FM-Index, but has no impact on the quality of final DNA mapping. A MTL-based index model hierarchy is a tree structure. To keep the index size in check, we deploy simple linear regression models [45] as leaf nodes in the model hierarchy of the MTL-based index. A linear regression model contains only one weight and one bias. Each non-leaf node is a neural network having a fully-connected layer, each of which contains 10 neurons with sigmoid activation.

Training. The MTL-based index is built for the increments of p k-mers $\{\mathcal{T}_i\}_{i=1}^p$. For a k-mer $\{\mathcal{T}_i\}$, its training dataset includes f_i increments $\{\mathbf{x}_{i,j}\}_{j=1}^{f_i}$ and their positions $\{y_{i,j}\}_{j=1}^{f_i}$. The learning function for the k-mer \mathcal{T}_i is defined as $\mathbf{w}_i^T \mathbf{x} + b_i$. Based on [42], [43], [46], we formulate the loss function to learn the relations between k-mers as

$$\min_{\mathbf{W},\mathbf{b}} \sum_{i=1}^{p} \frac{\beta_i}{f_i} \sum_{j=1}^{J_i} l(\mathbf{w}_i^T \mathbf{x}_{i,j} + b_i, y_{i,j})$$
(4)

where $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_p)$; $\mathbf{b} = (b_1, \dots, b_p)^T$; $l(\cdot, \cdot)$ means the cross-entropy loss function; β_i is the importance of kmer $\{\mathcal{T}_i\}$. We trained the MTL-based index to minimize this equation by an Adam optimizer. Similar to LISA [28], the training and testing of a MTL-based index use the same dataset, an EXMA table of a reference. Training a MTL-based index for a reference typically takes $1 \sim 2$ days. Once a MTLbased index is trained, billions of exact match operations can happen on its reference.



MTL-based Index Performance. We profiled the performance of EXMA-15 equipped with a naïve learned index in Figure 12. As Figure 12(a) shows, 2.5E-5% and 4E-6% of 15-mers have $64K\sim256K$ and >1M increments, respectively. However, searching the 15-mers having $64K\sim256K$ and >1M increments consumes 36% and 20.5% of the search time respectively, as shown in Figure 12(b). This is because the naïve learned index predicts a lot wrong



Fig. 14. The architecture of EXMA accelerator.

positions, and the linear search overhead is large. The prediction errors of the naïve learned index for the 15-mers having 64K~256K (learn-256K) and >1M (learn-1M) increments are shown in Figure 13. On average, we have to search 917 and 2133 more increments to find the correct one for the 15-mers having 64K~256K and >1M increments respectively. The MTL-based index greatly improves index prediction accuracy by simultaneously learning from multiple 15-mers having similar amounts of increments. The MTLbased index (MTL) further reduces the mean of prediction errors to 45 and 182 for the 15-mers having 64K~256K and >1M increments respectively. As a result, the MTL-based index (EXMA-15M) improves search throughput by 75% over LISA-21 with only a half number of parameters, as shown in Figure 10(b).

C. EXMA Accelerator

We propose a hardware accelerator to process search operations over an EXMA table using a MTL-based index. We integrate two caches for the bases and the MTL-based index of an EXMA table to reduce unnecessary DRAM accesses. And then, we present EXMA scheduling to improve cache hit rate. We also introduce dynamic page policy to improve DRAM row buffer hit rate during searching the increments of an EXMA table. At last, we create CHAIN compression to reduce the EXMA table size.

1) Accelerator Architecture

The architecture of our EXMA accelerator is shown in Figure 14. The kernel of the EXMA accelerator is an inference engine computing predictions of a MTL-based index. We adopt the state-of-the-art Tangram neural network accelerator [47] as the inference engine. The Tangram accelerator consists of a number of processing elements (PEs) organized in a 2D array. Each PE includes a simple ALU for multiplyaccumulate (MAC) operations and a small register file of 32B. A larger SRAM buffer is shared by all PEs. We add two small caches for the bases and MTL-based index of an EXMA table. Data fetched and stored by the accelerator goes through a de/compression unit (§IV-C4) that de/compresses the bases and increments of an EXMA table. We integrate a scheduling queue into the EXMA accelerator to implement 2stage EXMA scheduling (§IV-C2). At last, the dynamic page management (§IV-C3) switching between open and close page policies is implemented in the CPU memory controller.

2) EXMA Scheduling

Poor Locality of MTL-based Index. The conventional first-ready first-come-first-serve (FR-FCFS) policy is adopted by almost all accelerators [13]–[15], [30], [37], [38] to schedule FM-Index searches. However, FR-FCFS significantly degrades the hit rates of our base cache and index cache.

	4-mer base	4-mer, pos		4-mer pos	base cache	index cache
	AAAA 0x0000 0000	m ₀	Ro	TTTT 998	miss	miss
	AAAC 0x000C 0011	$\overline{m_1}$ $\overline{m_2}$	R₁	AAAA 29	TTTG TTTT miss	$[m_0,m_2,m_{18}]$ miss
		$m_1 m_4 \dots m_{10}$	R_2	TTTG 1	AAAA AAAC miss	$[m_0, m_1, m_3]$ hit
	TTTT Ox1FFF 00D1	$715 \cdots 103121$	R ₃	AAAC 99	TTTG TTTT miss	m_{0}, m_{1}, m_{3} miss
	(a) base array	(b) MTL index			(c) FR-FCFS	
		Fig. 15. The low c	ache	e hit rate of MTI	-based index.	
	base cache	index cache		*201100	a having small differ	man between their nee values
R_1	miss	R ₂ miss		the in	lev cache can expect	more hits. This is because the
R₃	AAAA AAAC hit	R_1 m_0, m_1, m_3 hit		smalle	r difference the pos	values of two requests exhibit.
R_2	AAAA AAAC miss	R ₃ m ₀ ,m ₂ ,m ₁₈ miss		the mo	ore likely these two re	quests use the same MTL index
R ₀	TTTG TTTT hit	R ₀ m ₀ ,m ₂ ,m ₁₈ hit		nodes	during searches. As Fig	gure 16(b) shows, the index cache
-	(a) the 1 st stage	(b) the 2 nd stage		has tw	o misses and two h	ts. Totally, the 2-stage EXMA
	Fig. 16. The 2-stage E	XMA scheduling.		schedu	ling has 4 misses and	4 hits.

We show an example of FR-FCFS in Figure 15, where there are 4 FM-Index requests in the form of [k-mer, pos], i.e., $R_0 = [TTTT, 998], R_1 = [AAAA, 29], R_2 = [TTTG, 1]$ and $R_3 = [AAAC, 99]$. As Figure 15(a) shows, all bases are stored consecutively in the lexicographical order of kmers. Each base occupies 4-byte. So the bases of AAAA and AAAC (TTTC and TTTT) are stored in the same 64byte. The MTL-based index used by four requests is shown in Figure 15(b). To predict the increments of pos 1 and 29 (99 and 998), the MTL index nodes of m_0 , m_1 and m_3 $(m_0, m_2 \text{ and } m_{18})$ are required. Assume the base cache can contain only one 64-byte line, while the index cache can store three index nodes. Four requests are scheduled by FR-FCFS in Figure 15(c). When R_0 arrives, both caches are empty and thus have a miss. And then, the bases of "TTTG" and "TTTT" are stored in the base cache, while the index nodes of m_0, m_2 and m_{18} used by R_0 are stored in the index cache. For R_1 , both caches also suffer from a miss. The bases of "AAAA" and "AAAC" are fetched to the base cache, while m_0, m_1 and m_3 used by R_1 are installed into the index cache. For R_2 , the base cache has a miss, but the index cache has a hit. At last, both cache have a miss for R_3 . Totally, four misses happen in the base cache, while three misses occur in the index cache.

2-Stage Scheduling. We propose a 2-stage EXMA scheduling technique to enhance the hit rates of the base and index caches. Unlike FR-FCFS scheduling requests based on their addresses and order, EXMA re-orders the requests according to their data including k-mers and positions (pos). In the first stage, EXMA sorts the FM-Index requests based on their k-mers. By consecutively issuing FM-Index requests in the lexicographic order of their k-mers, the hit rate of the base cache increases, since the bases of the k-mers sorted in the lexicographic order tend to have stronger spatial locality. As Figure 16(a) shows, the first stage of EXMA scheduling issues four requests in the order of R_1 , R_3 , R_2 , and R_0 . The base cache has two hits and two misses, but the index cache has all four misses. This is why EXMA needs to do the second stage of scheduling. During the second stage of EXMA scheduling, four requests are ranked according to their pos values. Through consecutively computing inferences of MTL indexes of the

Implementation. Our 2-stage EXMA scheduling is implemented with the scheduling queue that is a content-addressable memory (CAM). A CAM can perform sorting operations [48]. The k-mer and pos of a request can be stored in one row of the CAM. Each DNA symbol in a k-mer is denoted by 3 bits, since we need to encode \$, A, C, T and G. Requests can be sorted in the CAM based on their k-mers or pos values.

3) Dynamic Page Policy

Dynamic Page Policy. Prior FM-Index accelerators [13]-[15], [30], [37], [38] adopt close-page policy in their DRAM main memories. They always pre-charge a DRAM row after each access, since conventional 1-step FM-Index searches have little spatial locality, as shown in Figure 6(a). On the contrary, our EXMA table stores the increments of a k-mer consecutively in DRAM rows. As the algorithm of FM-Index backward search (line $2 \sim 3$ in Figure 3(d)) indicates, each iteration issues two requests searching the same k-mer but with different position values. In a search iteration, we compute Occ(k-mer, low) and Occ(k-mer, high). Both search the increments of the same k-mer that are very likely stored in the same row. So our accelerator asks the CPU memory controller (MC) to keep the DRAM row open after the first request in a search iteration is processed, since we expect the second request can hit in the row buffer. However, the row will be pre-charged, after the second is processed.

Implementation. The dynamic page policy can be implemented in the CPU MC. When searching a request, if there is another request pending in the scheduling queue of the accelerator searches the same k-mer, the CPU MC keeps the DRAM row open after the ongoing request completes. Otherwise, it pre-charges that DRAM row. The CPU MC maintains a register to indicate whether all rows are closed, and another register to record which row is open for each bank. 4) CHAIN compression

 $B\Delta I$. The state-of-the-art cache compression technique, $B\Delta I$ [49], breaks a 64-byte cache line into eight 8-byte data sections. As Figure 17(a) shows, to compress the cache line, $B\Delta I$ records the first data section (data₀) and stores only the difference (Δ_i) between each data section $(data_i)$ and $data_0$. To decompress a data section, B ΔI simply calculates



 $data_i = data_0 + \Delta_i$. B Δ I typically reduces data size of SPEC-CPU2006 applications by ~ 50%. In these applications, data sections in a cache line are not sorted.

CHAIN. Since both increments and bases of each k-mer are sorted and stored consecutively in a DRAM row, we believe they are more compressible than the data of SPEC-CPU2006 applications. Therefore, we propose a novel compression technique, CHAIN, to more aggressively compress an EXMA table. As Figure 17(b) describes, to compress the increments in a 64B memory line, CHAIN first stores the first increment $(incr_0)$. And then, it stores the value difference (Δ_i) between two consecutive increments. So we have $\Delta_i = incr_i - incr_{i-1}$. To decompress an increment $incr_i$ in a 64B memory line, CHAIN simply computes $incr_i = incr_0 + \sum_{j=1}^{i} \Delta_j$. Bases of an EXMA table can be (de)compressed in the same way.

Implementation. The CHAIN compression and decompression require only 64-bit adders. Multiple adders concurrently operate for the CHAIN compression, while the CHAIN decompression requires only one adder for accumulations. The CHAIN decompression slightly prolongs FM-Index search latency but greatly increases FM-Index search throughput.

5) Putting all together

As Figure 14 describes, **0** after receiving FM-Index requests from the CPU, the accelerator stores them in its scheduling queue and performs the first stage scheduling. 2 Based on their k-mers, the accelerator checks whether the bases of the requests stored in the queue are in the base cache or not. 3 If misses occur, DRAM accesses are issued to fetch the bases. Otherwise, the accelerator does the second stage scheduling. **4** The accelerator checks whether the MTL index nodes required by the requests in the queue are in the index cache or not, according to both k-mer and pos values. **6** If misses happen, DRAM accesses are issued to fetch MTL index nodes. Otherwise, the inference engine computes with MTL index nodes. 6 Until the inference of a leaf node is finished, the accelerator issues a DRAM access to read the increment in the predicted position. If the increment is correct, the computation of Occ(k-mer, pos) is completed. Otherwise, it linearly searches DRAM rows to find the correct increment. All DRAM accesses from the EXMA accelerator are managed by its DMA controller communicating with the CPU MC, which decides whether to pre-charge opened rows based on the requests in the scheduling queue of the accelerator.

TABLE I THE HARDWARE CONFIGURATION OF EXMA.

Component	Description	Area (mm^2)	Energy/Op (pJ)					
Infer. engine	4.8×8 PE arrays	0.512	0.25					
Sch. queue	SRAM, 128-bit×256	0.023	1.9					
Index cache	SRAM, 32KB, 16-way	0.084	2.62					
Base cache	eDRAM, 1MB, 8-way	0.667	17.2					
De/compress	32 64-bit adders	0.091	0.21					
Sch. & row	2-stage sch. & dyn. page	0.035	1.02					
DMA ctrl	adopted from [52]	0.21	3.42					
EXMA accelerator: area $1.62mm^2$, and leakage $223.8mW$								
CPU 2.5GHz, 16-core, 40MB LLC, 64 LLC MSHRs								
DRAM	AM DDR4-2400, 384GB, 4 channels, 3 DIMMs / channel,							
main	4 ranks / DIMM, 2 bank groups / rank, 16 chips / rank,							
memory	2 banks / bank group, FR-FCFS, close-page, row size 2KB							
system	system 2 chips / data buffer, t_{RCD} - t_{CAS} - t_{RP} : 16-16-16							

6) System Integration

Our EXMA is connected to a CPU processor as a looselycoupled *non-coherent* accelerator [50], [51] by a Network-on-Chip (NoC). EXMA accesses DRAM via two DMA-dedicated planes of the NoC, bypassing the cache hierarchy of the CPU. The EXMA data region must be flushed from the CPU cache hierarchy before FM-Index searches start. We chose the noncoherent model [50] for better performance, since the memory footprint of FM-Index searches is always larger than the CPU LLC capacity.

D. Design Overhead

The training overhead of a MTL-based index is shown in the Section of Training in §IV-B. The hardware overhead of the EXMA accelerator is summarized in Table I. From [47], we adopted the inference engine, which runs at 800MHz and is synthesized with Synopsys 28nm generic library. We quantized the model hierarchy of MTL index with 8-bit without accuracy degradation. A PE has an 8-bit MAC ALU and a 32B register file. The inference engine contains 4.8×8 PE arrays, each of which has a 16KB shared SRAM buffer and an activation unit. The EXMA accelerator also includes a scheduling queue (SRAM CAM) with 512 128-bit entries, a 32KB 16-way SRAM index cache, and a 1MB 8-way eDRAM base cache. We modeled the area and power of memory components including registers, buffers and caches by CACTI [53]. We used the same DDR4-2400 DRAM main memory configuration as the recent FM-Index PIM MEDAL [15]. The EXMA accelerator connects to four DRAM channels, with a total 384GB capacity. We adopted DRAMPower [54] to model the power consumption of our DDR4 main memory.

V. EXPERIMENTAL METHODOLOGY

Simulation. We used gem5-aladdin [55] to model our CPU baseline and our EXMA accelerator. The configuration of our CPU baseline is shown in Table I. We used McPAT [56] to calculate the power consumption of the CPU processor. We implemented the main memory system using DRAMsim2 [57].

Accelerator Baselines. Besides CPU, we also ran the FM-Index search kernel on an Nvidia Tesla P100 GPU [58], a Stratix-V FPGA [30], and a 28*nm* ASIC design [37]. We compared EXMA against recent FM-Index PIM designs,



MEDAL [15] and FindeR [14]. We used gem5-gpu [59] to simulate the GPU, and gem5-aladdin to model the computing units of FPGA, ASIC and PIMs. All their DRAM main memories are implemented by DRAMsim2. The power data of the GPU, FPGA, ASIC and PIMs are adopted from [14], [15], [30], [37], [58]. The power of DRAM main memories is also modeled by DRAMPower.

Workloads. To evaluate EXMA, we adopted FM-Indexbased genome analysis applications: BWA-MEM [12] for short read alignment, MA [16] for long read alignment, SGA [24] for read assembly, ExactWordMatch [25] for annotation and a reference-based compression algorithm [26]. SGA for long reads uses the FM-Index-based error correcting scheme [33] to reduce errors.

Datasets. For alignment, annotation and compression, we used human (human, 3G-bp), picea glauca (picea, 20G-bp), and pinus lambertiana (pinus, 31G-bp) genomes as reference genomes. To study short reads, we adopted DWGSim [60] to generate 101-bp short reads with $50 \times$ coverage. To evaluate long reads, we created long reads (with length of 1K-bp) by PBSIM [61]. The error profiles of reads is summarized in the format of (name, mismatch%, insertion%, deletion%, total error%), i.e., (Illumina, 0.18%, 0.01%, 0.01%, 0.2%) [62], (PacBio, 1.50%, 9.02%, 4.49%, 15.01%), and (ONT_2D, 16.50%, 5.10%, 8.40%, 30.0%) [10].

Schemes. The schemes we studied can be summarized as:

- CPU: We ran LISA-21 for FM-Index searches in genome applications on our CPU baseline. We also applied $B\Delta I$ [49] compression on LISA data for three datasets.
- EXMA-15: EXMA-15 with the MTL-based index and CHAIN compression is used to replace LISA-21 in CPU.
- EX-acc: We ran EXMA-15 on the EXMA accelerator.
- EX-2stage: 2-stage scheduling is added to EX-acc.
- EXMA: Dynamic page policy is enabled on EX-2stage.

VI. RESULTS AND ANALYSIS

Throughput Comparison against CPU. As Figure 18 shows, we compare FM-Index search throughput of EXMA and CPU by running the seeding of short read alignment, since FM-Index searches consume 99% of the seeding time in short read alignment. Compared to CPU, on average, EXMA-15 improves search throughput by 80%. Our MTL-based index achieves high accuracy on picea, since the increment distributions of its different *k*-mers are more similar to each other. EX-acc improves search throughput by $7.25 \times$ over CPU. Our EXMA accelerator can support more concurrent search operations, while CPU has only a limited number of LLC MSHRs. Compared to CPU, EX-2stage increases search throughput by $15 \times$. Pinus with EX-2stage has the smallest throughput improvement. Because the size of the pinus MTL-



Fig. 19. The speedup of EXMA in genome analysis (norm. to CPU). based index is the largest among 3 datasets, and thus its index cache has the lowest hit rate. On average, EXMA increases search throughput by $23.6 \times$ over CPU.

Performance Comparison against CPU. We report and compare the speedup achieved by EXMA in various genome applications in Figure 19, where we list 3 sets of "alignment and assembly" for reads generated by Illumina, Nanopore, and PacBio respectively. For each application, although EXMA obtains smaller FM-Index search throughput improvement on larger datasets (Figure 18), e.g., pinus, EXMA improves the application performance more significantly on larger datasets. This is because CPU consumes a larger portion of the execution time of a genome analysis application to perform FM-Index searches when processing larger datasets that introduce more TLB and data cache misses. EXMA achieves larger performance improvement on alignment and assembly for Illumina, annotation, and compression, since FM-Index searches dominate the execution of these applications. On average, EXMA improves the performance of genome applications by $2.5 \times \sim 3.2 \times$, when processing various datasets.



Fig. 20. The energy reduction of EXMA in genome analysis (norm. to CPU. DRAM-chip/IO indicates the energy of DRAM chips/DDR4 interface. EXMA-leak/dyn is the static/dynamic energy of the EXMA accelerator).

Energy Comparison against CPU. We compare the energy reduction obtained by EXMA in various genome applications in Figure 20, where we list 3 sets of "align(ment) and asse(mbly)" for reads generated by Illumina, Nanopore, and PacBio respectively. On average, EXMA reduces total energy consumption of genome analysis by $61\% \sim 70\%$ when processing different datasets. The major part of the energy reduction comes from voiding using the CPU processor during FM-Index searches. The more time FM-Index searches consume in a genome analysis application, the more energy reduction EXMA can achieve in that application. On average, the EXMA accelerator consumes only < 3% of the total energy consumption of various genome applications. The vast majority of energy consumption is consumed by the DRAM main memory and the CPU handling non-FM-Index-search parts in genome analysis applications.

Comparison against Accelerators. We evaluated EXMA and compare it against various hardware accelerators including a GPU [58], a FPGA [30], an ASIC [37], and two PIMs [14], [15] when processing pinus in Table II. Not all accelerators can run the whole genome applications, so we use "million



6.9

THE COMPARISON OF ACCELERATORS WHEN PROCESSING PINUS.									
	GPU	FPGA [30]	ASIC [37]	MEDAL [15]	FindeR [14]	EXMA			
lgorithm	LISA-21	FM-2	FM-1	FM-1	FM-1	EXMA-15			
Aem (GB)	384	384	384	384	384	384			
Acc Power (W)	182	11	9.4	0.011	0.28	0.89			
Aem Power (W)	72	72	72	72	72	72			
/lbase/s	157	96	34	102	93	504			

0.42

1.42

A

Mbase/s/Watt

0.61

base per second" (Mbase/s) as the performance metric to evaluate only FM-Index search throughput. Different accelerators adopt different search algorithms. We implemented LISA-21 on the Tesla P100 GPU. The FPGA design [30] supports FM-2 (conventional 2-step) searches. EXMA performs EXMA-15 searches. The other accelerators can conduct only FM-1 searches. Since the capacity of internal memories of the GPU (16GB HBM) and the PIM FindeR (2.6GB ReRAM) is smaller than the pinus data size. We provide the same DDR4 main memory configuration shown in Table I to all accelerators. The power values of both the accelerator (Acc Power) and the DDR4 main memory (Mem Power) are listed in Table II. The search performance is decided by two factors, i.e., the memory bandwidth utilization, and the number of DNA symbols searched during each iteration. For the bandwidth utilization, only EXMA supports dynamic page policy, while the other use only close page policy. MEDAL can support chip-level parallelism, but its search throughput is limited by the address bus. So EXMA achieves the highest bandwidth utilization. For the number of DNA symbols searched during each iteration, only the GPU and EXMA can process > 2 DNA symbols in each iteration. But the low accuracy of learned index makes the GPU to search many unnecessary IP-BWT entries to find the correct one. Therefore, EXMA obtains the best search throughput. The throughput per Watt of the GPU and FPGA designs is low, since the power consumption of their computing parts is not trivial. For the PIMs and EXMA, the power consumption is dominated by the DRAM main memory. Compared to the PIM MEDAL, EXMA improves search throughput per Watt by $4.8 \times$.

Bandwidth utilization. Figure 21 shows the comparison of bandwidth utilization, which is defined as the ratio between the data capacity fetched from DRAM and total DRAM bandwidth. ASIC using FM-1 has only 26% of the total DRAM bandwidth, since it uses close-page policy and fetches only 64B after activating a 2KB row. GPU implementing LISA-21 fetches entire rows from host memory, so it achieves higher bandwidth utilization. MEDAL increases bandwidth utilization by activating each individual chips. However, due to conflicts on the address bus, MEDAL uses only 67% of the DRAM bandwidth. In contrast, EXMA obtains 91% bandwidth utiliza-

tion by switching between close-page and open-page policies.

DIMM number. We studied the sensitivity of EXMA and MEDAL to the DIMM number in Figure 22. With 2 DIMMs, EXMA improves search throughput by 29% over MEDAL. By having 3 DIMMs, EXMA linearly scales its throughput up by 40%, since a single EXMA accelerator can maintain all the DIMMs. However, MEDAL increases its throughput by only 14.5% with 3 DIMMs. Each MEDAL PIM accelerator sits on a DIMM. More DIMMs bring more ranks. MEDAL suffers from non-trivial inter-DIMM communication overhead. When the number of DIMM increases to 4, the search throughput of neither EXMA nor MEDAL increases significantly. The data bus (bandwidth utilization) of EXMA is saturated, while the address bus of MEDAL is saturated.

PE Array number. We show search throughput of EXMA with a varying number of PE arrays in Figure 22. Two PE arrays of EXMA already achieve 89% of the search throughput of the configuration with four PE arrays. This is because the computations of MTL-based indexes are not intensive. Further increasing the PE array number to 8 increases search throughput by only 3% over the configuration with four PE arrays. So we selected 4 PE arrays in our baseline configuration.

CAM & Cache. We explored search throughput of EXMA with varying CAM and base cache sizes in Figure 22. We use a CAM consisting of 256, 512, and 1024 entries to serve as the scheduling queue of EXMA. A 256-entry CAM cannot fully satisfy 2-stage scheduling and scheduling for dynamic page policy, and thus achieves only 77% of search throughput of the configuration with a 512-entry CAM. Further increasing the CAM entry number to 1K improves search throughput by only 2% over the configuration with a 512-entry CAM. Compared to the index cache, the search throughput is more sensitive to the capacity of the base cache, since the global buffer and register file of PE arrays can temporarily store MTL-based index nodes. We tried 512KB, 1MB and 2MB for the base cache. The search throughput stops increasing significantly, when the base cache capacity reaches 1MB. So we selected a 512-entry CAM and a 1MB base cache in our baseline configuration.

CHAIN. We show the compression result of CHAIN on pinus in Figure 23. Since the size of the IP-BWT table of LISA-21 is proportional to its step number, the total data size of LISA-21 (original) is $2.2 \times$ larger than that of EXMA-15 (original). After $B\Delta I$ compresses the data size of LISA-21 to 50%, i.e., 152GB. On the contrary, CHAIN compresses the data size of EXMA-15 to only 25%, i.e., 40GB. We observed similar compression rates of $B\Delta I$ and CHAIN on the other genome datasets.

VII. CONCLUSION

Though state-of-the-art genome analysis adopts FM-Index to process exact-matches, FM-Index is notorious of random memory access patterns. In this paper, we first present a row-buffer-friendly and highly-compressible EXMA table with a MTL-based index to process multiple DNA symbols by activating a DRAM row during each search iteration. And then, we build a hardware accelerator to process FM-Index searches on a EXMA table. Compared to the state-of-the-art FM-Index PIM MEDAL, EXMA improves search throughput by $4.9\times$, and enhances search throughput per Watt by $4.8\times$.

REFERENCES

- M. Schirmer, U. Z. Ijaz, R. D'Amore, N. Hall, W. T. Sloan, and C. Quince, "Insight into biases and sequencing errors for amplicon sequencing with the illumina miseq platform," *Nucleic acids research*, vol. 43, no. 6, pp. e37–e37, 2015.
- [2] J. J. Mosher, B. Bowman, E. L. Bernberg, O. Shevchenko, J. Kan, J. Korlach, and L. A. Kaplan, "Improved performance of the pacbio smrt technology for 16s rdna sequencing," *Journal of microbiological methods*, vol. 104, pp. 59–60, 2014.
- [3] M. Eisenstein, "Oxford nanopore announcement sets sequencing sector abuzz," *Nature Biotechnology*, vol. 30, no. 4, pp. 295–297, 2012.
- [4] J. D. Merker, A. M. Wenger, T. Sneddon, M. Grove, Z. Zappala, L. Fresard, D. Waggott, S. Utiramerur, Y. Hou, K. S. Smith *et al.*, "Long-read genome sequencing identifies causal structural variation in a mendelian disease," *Genetics in Medicine*, vol. 20, no. 1, p. 159, 2018.
- [5] X. Ma, M. Mau, and T. F. Sharbel, "Genome editing for global food security," *Trends in biotechnology*, 2017.
- [6] T. Hoenen, A. Groseth, K. Rosenke, R. J. Fischer, A. Hoenen, S. D. Judson, C. Martellaro, D. Falzarano, A. Marzi, R. B. Squires *et al.*, "Nanopore sequencing as a rapidly deployable ebola outbreak tool," *Emerging infectious diseases*, vol. 22, no. 2, p. 331, 2016.
- [7] J. Quick, N. D. Grubaugh, S. T. Pullan, I. M. Claro, A. D. Smith, K. Gangavarapu, G. Oliveira, R. Robles-Sikisaka, T. F. Rogers, N. A. Beutler et al., "Multiplex pcr method for minion and illumina sequencing of zika and other virus genomes directly from clinical samples," *nature protocols*, vol. 12, no. 6, p. 1261, 2017.
- [8] N. Zhu, D. Zhang, W. Wang, X. Li, B. Yang, J. Song, X. Zhao, B. Huang, W. Shi, R. Lu *et al.*, "A novel coronavirus from patients with pneumonia in china, 2019," *New England Journal of Medicine*, 2020.
- [9] S. Canzar and S. L. Salzberg, "Short read mapping: An algorithmic tour," *Proceedings of the IEEE*, vol. 105, no. 3, pp. 436–458, 2017.
- [10] Y. Turakhia, G. Bejerano, and W. J. Dally, "Darwin: A genomics coprocessor provides up to 15,000x acceleration on long read assembly," in ACM Architectural Support for Programming Languages and Operating Systems, 2018.
- [11] D. Fuijiki, A. Subramaniyan, T. Zhang, Y. Zheng, R. Das, D. Blaauw, and S. Narayanasamy, "Genax: A genome sequencing accelerator," in *IEEE/ACM International Symposium on Computer Architecture*, 2018.
- [12] H. Li, "Aligning sequence reads, clone sequences and assembly contigs with bwa-mem," arXiv preprint arXiv:1303.3997, 2013.
- [13] M. C. F. Chang, Y. T. Chen, J. Cong, P. T. Huang, C. L. Kuo, and C. H. Yu, "The smem seeding acceleration for dna sequence alignment," in *IEEE International Symposium on Field-Programmable Custom Computing Machines FCCM*, 2016, pp. 32–39.
- [14] F. Zokaee, M. Zhang, and L. Jiang, "Finder: Accelerating fm-indexbased exact pattern matching in genomic sequences through reram technology," in *International Conference on Parallel Architectures and Compilation Techniques*, 2019, pp. 284–295.
- [15] W. Huangfu, X. Li, S. Li, X. Hu, P. Gu, and Y. Xie, "Medal: Scalable dimm based near data processing accelerator for dna seeding algorithm," in *IEEE/ACM International Symposium on Microarchitecture*, 2019, p. 587–599.
- [16] M. Schmidt, K. Heese, and A. Kutzner, "Accurate high throughput alignment via line sweep-based seed processing," *Nature communications*, vol. 10, no. 1, pp. 1–10, 2019.
- [17] A. Nag, C. Ramachandra, R. Balasubramonian, R. Stutsman, E. Giacomin, H. Kambalasubramanyam, and P.-E. Gaillardon, "Gencache: Leveraging in-cache operators for efficient sequence alignment," in

IEEE/ACM International Symposium on Microarchitecture, 2019, pp. 334–346.

- [18] R. Kaplan, L. Yavits, R. Ginosar, and U. Weiser, "A resistive cam processing-in-storage architecture for dna sequence alignment," *IEEE Micro*, 2017.
- [19] A. Madhavan, T. Sherwood, and D. Strukov, "Race logic: A hardware acceleration for dynamic programming algorithms," in *IEEE/ACM International Symposium on Computer Architecture*, 2014.
- [20] E. Rucci, C. Garcia, G. Botella, A. De Giusti, M. Naiouf, and M. Prieto-Matias, "Accelerating smith-waterman alignment of long dna sequences with opencl on fpga," in *International Conference on Bioinformatics and Biomedical Engineering*. Springer, 2017, pp. 500–511.
- [21] R. Luo, T. Wong, J. Zhu, C.-M. Liu, X. Zhu, E. Wu, L.-K. Lee, H. Lin, W. Zhu, D. W. Cheung *et al.*, "Soap3-dp: fast, accurate and sensitive gpu-based short read aligner," *PloS one*, vol. 8, no. 5, p. e65632, 2013.
- [22] M. Burrows and D. J. Wheeler, "A block-sorting lossless data compression algorithm," 1994.
- [23] N. Ahmed, K. Bertels, and Z. Al-Ars, "A comparison of seed-andextend techniques in modern dna read alignment algorithms," in *IEEE International Conference on Bioinformatics and Biomedicine*, 2016, pp. 1421–1428.
- [24] J. T. Simpson and R. Durbin, "Efficient de novo assembly of large genomes using compressed data structures," *Genome research*, vol. 22, no. 3, 2012.
- [25] J. Healy, E. E. Thomas, J. T. Schwartz, and M. Wigler, "Annotating large genomes with exact word matches," *Genome research*, vol. 13, no. 10, pp. 2306–2315, 2003.
- [26] P. Prochazka and J. Holub, "Compressing similar biological sequences using fm-index," in *Data Compression Conference*, 2014, pp. 312–321.
- [27] A. Chacon, S. Marco-Sola, A. Espinosa, P. Ribeca, and J. C. Moure, "Boosting the fm-index on the gpu: Effective techniques to mitigate random memory access," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 12, no. 5, pp. 1048–1059, 2015.
- [28] D. Ho, J. Ding, S. Misra, N. Tatbul, V. Nathan, V. Md, and T. Kraska, "Lisa: Towards learned dna sequence search," in *Workshop on Systems for ML at NeurIPS 2019*, 2019.
- [29] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis, "The case for learned index structures," in ACM International Conference on Management of Data, 2018, p. 489–504.
- [30] J. Arram, T. Kaplan, W. Luk, and P. Jiang, "Leveraging fpgas for accelerating short read alignment," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2017.
- [31] Y. C. Wu, C. H. Chang, J. H. Hung, and C. H. Yang, "A 135-mw fully integrated data processor for next-generation sequencing," *IEEE Transactions on Biomedical Circuits and Systems*, 2017.
- [32] M. A. Quail, M. Smith, P. Coupland, T. D. Otto, S. R. Harris, T. R. Connor, A. Bertoni, H. P. Swerdlow, and Y. Gu, "A tale of three next generation sequencing platforms: comparison of ion torrent, pacific biosciences and illumina miseq sequencers," *BMC genomics*, vol. 13, no. 1, p. 341, 2012.
- [33] J. R. Wang, J. Holt, L. McMillan, and C. D. Jones, "Fmlrc: Hybrid long read error correction using an fm-index," *BMC bioinformatics*, vol. 19, no. 1, p. 50, 2018.
- [34] Y.-T. Huang and Y.-W. Huang, "An efficient error correction algorithm using fm-index," *BMC bioinformatics*, vol. 18, no. 1, p. 524, 2017.
- [35] H. Li, "Exploring single-sample snp and indel calling with wholegenome de novo assembly," *Bioinformatics*, vol. 28, no. 14, pp. 1838– 1844, 2012.
- [36] A. Chacón, J. C. Moure, A. Espinosa, and P. Hernández, "n-step fmindex for faster pattern matching," *Procedia Computer Science*, vol. 18, pp. 70–79, 2013.
- [37] Y. Wang, X. Li, D. Zang, G. Tan, and N. Sun, "Accelerating fm-index search for genomic data processing," in *International Conference on Parallel Processing*, 2018.
- [38] S. Angizi, J. Sun, W. Zhang, and D. Fan, "Aligns: A processing-inmemory accelerator for dna short read alignment leveraging sot-mram," in ACM/IEEE Design Automation Conference, 2019, pp. 1–6.
- [39] S. Angizi, J. Sun, W. Zhang, and D. Fan, "Pim-aligner: A processing-inmram platform for biological sequence alignment," in *Design, Automation Test in Europe Conference & Exhibition*, 2020, pp. 1265–1270.
- [40] JEDEC, "Ddr4 sdram standard jesd79-4c," 2020. [Online]. Available: https://www.jedec.org/standards-documents/docs/jesd79-4a
- [41] S. Ruder, "An overview of multi-task learning in deep neural networks," CoRR, vol. abs/1706.05098, 2017.

- [42] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *IEEE conference on computer vision and pattern recognition*, 2018, pp. 7482– 7491.
- [43] H. Bilen and A. Vedaldi, "Integrated perception with recurrent multi-task neural networks," in *Advances in neural information processing systems*, 2016, pp. 235–243.
- [44] C. Stein, "Inadmissibility of the usual estimator for the mean of a multivariate normal distribution," Stanford University, Tech. Rep., 1956.
- [45] M. H. Kutner, C. J. Nachtsheim, J. Neter, W. Li et al., Applied linear statistical models. McGraw-Hill Irwin New York, 2005, vol. 5.
- [46] A. Argyriou, T. Evgeniou, and M. Pontil, "Multi-task feature learning," in Advances in neural information processing systems, 2007, pp. 41–48.
 [47] M. Gao, X. Yang, J. Pu, M. Horowitz, and C. Kozyrakis, "Tangram:
- [47] M. Gao, X. Yang, J. Pu, M. Horowitz, and C. Kozyrakis, "Tangram: Optimized coarse-grained dataflow for scalable nn accelerators," in ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2019, pp. 807–820.
- [48] I. Okabayashi, H. Kotani, and H. Kadota, "A proposed structure of a 4 mbit content-addressable and sorting memory," in *Symposium on VLSI Circuits Digest of Technical Papers*, 1990, pp. 109–110.
- [49] G. Pekhimenko, V. Seshadri, O. Mutlu, M. A. Kozuch, P. B. Gibbons, and T. C. Mowry, "Base-delta-immediate compression: Practical data compression for on-chip caches," in *IEEE International Conference on Parallel Architectures and Compilation Techniques*, 2012.
- [50] K. Bhardwaj, M. Havasi, Y. Yao, D. M. Brooks, J. M. H. Lobato, and G. Wei, "Determining optimal coherency interface for many-accelerator socs using bayesian optimization," *IEEE Computer Architecture Letters*, vol. 18, no. 2, pp. 119–123, 2019.
- [51] D. Giri, P. Mantovani, and L. P. Carloni, "Accelerators and coherence: An soc perspective," *IEEE Micro*, vol. 38, no. 6, pp. 36–45, 2018.
- [52] G. Ma and H. He, "Design and implementation of an advanced dma controller on amba-based soc," in *IEEE International Conference on ASIC*, 2009, pp. 419–422.
- [53] N. P. Jouppi, A. B. Kahng, N. Muralimanohar, and V. Srinivas, "Cactiio: Cacti with off-chip power-area-timing models," *IEEE Transactions* on Very Large Scale Integration Systems, vol. 23, no. 7, pp. 1254–1267, 2014.
- [54] K. Chandrasekar, C. Weis, Y. Li, B. Akesson, N. Wehn, and K. Goossens, "Drampower: Open-source dram power & energy estimation tool," 2012. [Online]. Available: https://github.com/tukl-msd/DRAMPower
- [55] Y. S. Shao, S. L. Xi, V. Srinivasan, G.-Y. Wei, and D. Brooks, "Co-designing accelerators and soc interfaces using gem5-aladdin," in *IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2016, pp. 1–12.
- [56] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "Mcpat: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *IEEE/ACM International Symposium on Microarchitecture*, 2009, pp. 469–480.
- [57] P. Rosenfeld, E. Cooper-Balis, and B. Jacob, "Dramsim2: A cycle accurate memory system simulator," *IEEE computer architecture letters*, vol. 10, no. 1, pp. 16–19, 2011.
- [58] NVIDIA, "Nvidia tesla p100," 2016. [Online]. Available: https://images.nvidia.com/content/pdf/tesla/whitepaper/pascalarchitecture-whitepaper.pdf
- [59] J. Power, J. Hestness, M. Orr, M. Hill, and D. Wood, "gem5-gpu: A heterogeneous cpu-gpu simulator," *Computer Architecture Letters*, vol. 13, no. 1, Jan 2014.
- [60] H. Li, B. Handsaker, A. Wysoker, T. Fennell, J. Ruan, N. Homer, G. Marth, G. Abecasis, and R. Durbin, "The sequence alignment/map format and samtools," *Bioinformatics*, vol. 25, no. 16, pp. 2078–2079, 2009.
- [61] Y. Ono, K. Asai, and M. Hamada, "Pbsim: Pacbio reads simulator—toward accurate genome assembly," *Bioinformatics*, vol. 29, no. 1, pp. 119–121, 2012.
- [62] M. Schirmer, R. D'Amore, U. Z. Ijaz, N. Hall, and C. Quince, "Illumina error profiles: resolving fine-scale variation in metagenomic sequencing data," *BMC bioinformatics*, vol. 17, no. 1, p. 125, 2016.