

RAMPS: A Reconfigurable Architecture for Minimal Perfect Sequencing

Chad Nelson, Kevin R. Townsend, *Student Member, IEEE*, Osama G. Attia, *Student Member, IEEE*, Phillip H. Jones, *Member, IEEE*, and Joseph Zambreno, *Senior Member, IEEE*

Abstract—The alignment of many short sequences of DNA, called reads, to a long reference genome is a common task in molecular biology. When the problem is expanded to handle typical workloads of billions of reads, execution time becomes critical. In this paper we present a novel reconfigurable architecture for minimal perfect sequencing (RAMPS). While existing solutions attempt to align a high percentage of the reads using a small memory footprint, RAMPS focuses on performing fast exact matching. Using the human genome as a reference, RAMPS aligns short reads hundreds of thousands of times faster than current software implementations such as SOAP2 or Bowtie, and about a thousand times faster than GPU implementations such as SOAP3. Whereas other aligners require hours to preprocess reference genomes, RAMPS can preprocess the reference human genome in a few minutes, opening the possibility of using new reference sources that are more genetically similar to the newly sequenced data.

Index Terms—Hardware algorithms, FPGAs, bioinformatics, short-read aligner, convey HC-2, reconfigurable hardware

1 INTRODUCTION

AFTER the first human genome was sequenced in 2003, the next generation of sequencing technologies were developed with higher throughput and lower costs. Sequencing machines operate by determining the sequence of nucleotides of many small fragments of DNA, called *reads*, in parallel. To improve the quality of the process, the amount of genetic data that is sequenced is often enough to cover the entire genome multiple times. The increasing speed and adoption of next generation sequencing technologies has resulted in a great need for timely alignment of large amounts of short read data.

For example, a single Illumina HiSeq machine sequences 120 billion base pairs in a 27 hours run [1], and the Beijing Genomics Institute has at least 167 sequencing machines [2] and efforts with names like the Million Human Genomes Project [3]. World sequencing capacity was estimated at 15 quadrillion nucleotides per year in 2013 [4], and there is no sign of slowing. Because of the large amount of data produced by a single sequencing machine, it also becomes important for the processing of the data to occur on site. Networks have already become a bottleneck to offsite computer clusters, resulting in some scientists sending their data on hard drives via FedEx to be processed [2]. In fact, the growth in the amount of sequencing data is currently growing at a rate faster than Moore's law and is projected to continue for some time [5]. Moore's law states that the number of transistors in an integrated circuit will double approximately every two years [6]. Thus, new innovations

in algorithms and designs that utilize the transistors more efficiently are needed if the alignment of the data is to keep pace with the sequencing machines.

Numerous software projects have been developed to align the short reads from the sequencing machines with a reference genome, some of which are discussed in Section 2. Because these attempts only changed the hardware used, not the algorithm, the results show only a marginal increase in performance due to the use of better hardware. To achieve significant performance gains, more must be done than tweaking the same algorithms to work on better hardware. Since the growth of DNA sequencing data is currently outpacing the growth of Moore's law [6], simply using better hardware won't solve the problem in the long run. RAMPS breaks the tradition of the previous short read aligners by choosing a simple and new approach that is fast and energy efficient.

In this paper we introduce RAMPS, a reconfigurable architecture for minimal perfect sequencing that extends our earlier work in [7]. RAMPS offers a new approach utilizing the Convey HC-2, a hybrid core computing system. Using the human genome as a reference, RAMPS aligns short reads hundreds of thousands of times faster than software such as SOAP2 or Bowtie, and about a thousand times faster than GPU implementations such as SOAP3. By decreasing the preprocessing time, we hope to fundamentally change the alignment problem by allowing a higher percentage of the read data to be aligned exactly to the reference. The speed gains are made possible by selection of a fast algorithm, the use of a highly pipelined hardware design, and the large amounts of memory bandwidth provided by Convey's hybrid core computing system.

The remaining of the paper is organized as follows. Section 2 surveys the current software and hardware short read aligners and discusses the algorithms that is currently being used in the field. In Section 3 we introduce the preliminaries of the minimal perfect hashing algorithm and DNA

• The authors are with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011. E-mail: cnelson711@gmail.com, {ktown, ogamal, phjones, zambreno}@iastate.edu.

Manuscript received 24 July 2015; revised 9 Dec. 2015; accepted 23 Dec. 2015. Date of publication 29 Dec. 2015; date of current version 14 Sept. 2016.

Recommended for acceptance by M. Smith.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPDS.2015.2513053

sequencing in general. Then, we present the main components of RAMPS implementation in Section 4; a series of hardware pipelines to preprocess a reference genome into a hash table, and a hardware aligner for performing fast look-ups. In Section 5, we show the experimental results of our FPGA implementation. Finally, we draw our conclusions and discuss future work in Section 6.

2 RELATED WORK

A plethora of short read aligners have been published in the research field of bioinformatics. They vary greatly in the algorithms used for matching and their target hardware platforms. Different aligners will attempt to optimize one parameter at the cost of others. Some seek high sensitivity, the ability to find matches given mutations in the DNA, at the cost of speed. Others compress data structures to limit the required memory footprint. Projects that attempt to assemble fragments without a reference genome, called *de novo*, are not listed as their runtime is not linear with respect to the number of reads.

Early read aligners like BLAST [8] and MAQ [9] were slow and operated by selecting the best indices from a list of candidate locations. BLAST, the Basic Local Alignment Search Tool, was published in 1990 and was one of the first tools available for sequence alignment. MAQ [9] was an implementation that provided quality data along with its alignment results. These tools used a few base pairs of the larger read, called seeds, to generate a long list of candidate locations in the reference genome. The algorithm would then perform a computationally expensive Smith-Waterman [10] comparison on all candidate locations to determine the best spot for the read. Each comparison using the Smith-Waterman algorithm is costly, and a long list of candidate locations limits the performance of these algorithms. Unfortunately, many of the early tools are now too slow to handle the influx of data from next generation sequencing machines.

A newer approach, used in SOAP2 [11], BWA [12], and Bowtie [13], is to index a reference genome using an FM Index [14], which compresses the genome using the Bowler-Wheeler transform [15]. This scheme allows the genome to be compressed in a suffix tree, reducing the memory footprint, and allowing the use of commodity hardware. These tools use an algorithm in which alignment is the result of a pointer-based tree transversal. In cases of a mismatch, time-consuming backtracking is used to find segments that may match with high probability. Applications like SOAP2, BWA, and Bowtie are generally similar in their approach, but differ slightly in the way they construct their index of the reference genome and optimize the algorithm. There are ways to speed up tree traversal-based aligners that would result in a small drop in sensitivity, such as combining nodes on different levels of the tree or by using a hash table on the starting segment of the tree transversal. Using a hash table for the first few levels of the tree, was demonstrated by Arbabi et al. [16] and resulted in reasonable speed improvements without greatly compromising sensitivity.

Our choice of algorithm differs from prior attempts in that we attempt to reduce memory bandwidth, whereas prior approaches seem focused on using commodity hardware, reducing memory footprints, and providing algorithms for

finding the matches for reads that contained mismatches or fuzzy data. For aligners based on the FM index, the number of memory operations in the best case is proportional to the length of the read. In the worst case, aligners like SOAP and Bowtie search multiple paths of the tree in order to allow for mismatches and indels, and this process results in significantly more memory operations and thus a much slower execution time. By preprocessing the reference genome into a hash table, RAMPS makes the alignment of arbitrary reads a simple hash table lookup requiring only two memory operations.

Many implementations using alternative hardware are ports of mainstream algorithms onto new hardware, such as GPUs, a cluster of PCs, or FPGAs. Aluru and Jammula presented a good survey of hardware acceleration in sequence analysis [17]. The authors in [18] ported the slower RMAP algorithm to GPUs. Even though the RMAP algorithm was a good fit for the GPU architecture and they improved the performance tenfold, the newer tree traversal-based algorithms running on commodity processors are significantly faster. Torres et al. took the BWA source and ported it to a GPU [19] allowing for only exact matches. SOAP3 [20] uses the same approach, but allows for mismatches and higher sensitivity. BFAST [21], is an implementation of a seeding-based algorithm that runs on a cluster of PCs and is capable of aligning billions of reads per day. Convey's bioinformatics personalities [22] [23] improved the performance of BWA implementation 15 \times over a server and have even implemented designs for performing Smith-Waterman on the Convey FPGA-based accelerators. Recently, [24], [25] presented a hardware accelerator based on Burrows-Wheeler Transform that showed speedup over the software BWA. Fernandez et al. presented an FPGA-based accelerator for the Bowtie tool using the Convey HC-1/HC-2ex systems [26]. The work in [27] presented a short-read mapping accelerator based on a hash-index mapping, that allows for mismatches, implemented on the convey HC-1ex system. The accelerator uses 64 processing elements utilizing only about 30 percent of the Convey system's peak memory bandwidth. These examples attempt to speed up the chosen algorithm by using the higher memory bandwidth of GPUs, clusters of PCs, and custom accelerators.

3 PRELIMINARIES

3.1 Background on DNA Sequencing

Deoxyribonucleic acid (DNA) is double helix composed of four nitrogen based nucleobases: Adenine, Thymine, Guanine, and Cytosine (abbreviated ATGC). The DNA molecule actually contains two copies of the genetic information; an important attribute that allows it to be easily replicated by splitting the two chains of the double helix apart. The two chains run anti-parallel to each other, with one end of a single chain labeled 3' (three prime) and the other labeled 5' (five prime), depending on the direction of the third and fifth carbon atom on the sugar molecule. The bases (ATGC) of each chain pair with one another using hydrogen bonds. Adenine always pairs with Thymine (AT); Cytosine always pairs with Guanine (CG). While the relative proportion of the bases in DNA were known to be approximately equal in the base pair groups, Watson and Crick were the first to

proposed the double helix architecture in which the bases actually bonded in 1953 [28]. Due to the antiparallel nature of the chains and base pairing, given one chain of DNA you can compose the opposite chain by reversing the order and substituting base pairs. The base pairs are chained together using a five carbon sugar (ribose or 2-deoxyribose) called a nucleoside; the nucleoside and nucleobase are together referred to as a nucleotide. In the human genome, there are over 3 billion base pairs. They are grouped into 23 chromosomes, each containing hundreds of millions of base pairs. Full genome sequencing is the process of trying to discover the exact sequence of base pairs for a particular individual.

In general, sequencing DNA today involves breaking the DNA into small segments, amplifying, splitting the chains apart, and rebuilding one of the chains one base pair at a time. Amplification can take a single segment and multiply it millions of times using a polymerase chain reaction (PCR) machine, based on a process of repeated heating, cooling, and duplication accredited to Mullis [29]. The amplification is necessary so that when fluorescently marked bases are added to a sample when rebuilding one of the chains a single nucleotide at a time, a laser can detect which base pair was added. Newer models of sequencing machines perform “paired end” reads; i.e., both chains of the double helix are sequenced in order to improve quality.

Next generation sequencing devices, such as the ones from Illumina [1], can produce a massive amount of data: 300 million to 3 billion short paired reads on a typical 1 to 11 day run of the sequencer. Typically, the base pairs in the genome are sequenced multiple times to ensure that every portion of the genome is sequenced, since it is hard to know exactly how the DNA is broken into small segments. The amount of duplicate data, called coverage, is typically in the range of $2\times$ (low coverage) to $20\times$ (deep coverage) [30]. Thus, it is normal to have tens of billions of base pairs worth of information to align.

To align the short sequencing data, a reference genome is used. The reference is like having a picture of the puzzle while attempting to put the puzzle together. Aligners take each short read from the sequencing machine and attempt to ascertain its position in the reference genome. At best, this amounts to a simple string matching. At worst, it involves allowing for mutations in the read data such as mutations (single nucleotides with a different base), indels (insertions and deletions of single nucleotides), or gapped alignment (allowing for large gaps in either the read or the reference).

3.2 File Formats

Sequencing data typically comes in a text (ASCII encoded) data format known as FASTQ. RAMPS utilizes programs from the MEMOCODE 2012 reference design to compress the reference genome into a binary encoded form. A single ASCII encoded character (such as an ‘A’) typically consumes a single byte of data (8 bits). Since there are only four bases, each base can be binary encoded into 2 bits; thus, the sequence data can be compressed by at least a factor of 4. The data is further compressed by removing axillary comments and quality data. Since we designed and implemented our own hardware architecture using the Convey system, we are better able to take advantage of this compression, as the hardware can manipulate 2 bits at




Fig. 1. An example of a minimal perfect hash table.

a time whereas a normal computer must manipulate data by the byte (8 bits) or word (32 bits). The quality score is also ASCII encoded. Sequencing machines map the probability that a given base is incorrect into a set range of numbers, and then store them in the FASTQ file appropriately. The higher the number of a quality character means the higher the quality.

There also exists a format called FASTA in addition to the FASTQ format for sequence data. FASTA is used for storing pre-aligned reference data, such as the human genome. The FASTA format similarly uses ASCII encoding for the reference sequence. In addition to the regular nucleic acids (ACGT), there are another few extra characters in the reference we used (human g1k v37 [31]). The most common, ‘N’, represents unknown base pairs. In the reference human genome we used, roughly 8 percent of the genome was unknown.

3.3 Minimal Perfect Hash Algorithms

A hash table is a type of associative array for storing data based on the hashed value of a key. A hash function is a series of operations performed on the key that maps that key to a specific index in the table. Minimal perfect hashing is a way of building a hash table for a fixed set of keys without collisions (perfect) and without wasted space (minimal). Collisions occur when two different keys hash to the same index. More formally, a minimal perfect hash (MPH) is a hash function h from a set of keys S to a range of numbers range $[n] = \{0, \dots, n-1\}$, where h is one-to-one on S . For a small number of keys, finding such a hash function is possible by randomly searching and checking. For larger sets of keys, generalized algorithms use an intermediate table of values to adjust the hash function subtly in order to eliminate collisions and create the minimal perfect mapping.

Fig. 1 illustrates an example of a minimal perfect hash on a set of known keys. Each key is initially hashed once to retrieve a few bits of information stored in an intermediate table. The key is then either rehashed with the new seed from the intermediate table, or the offset from the intermediate table is added to the initial index. The MPH is created by choosing values in the intermediate table so that a given set of keys will not collide (i.e., two keys will not be directed to the same index).

For any given key, the probability of it hashing into any given index should be equal. Hash functions are designed so that they provide the same index given the same key. However, many hash functions allow the use of a seed value. Changing the seed creates an entirely new mapping

from the key set to a set of indices. This becomes important when a collision occurs in the minimal perfect hashing scheme, as it allows the intermediate table to choose a new seed that will cause the colliding indices to not collide.

Linear time algorithms for the construction of general minimal perfect hashes became practical starting in the 1990s, but their construction was complex. Botelho's dissertation provides a history of these algorithms [32]. Schmidt and Siegel were the first to propose a linear time construction algorithm, though in practice it was impractical. The current dominant algorithm is Compress, Hash, Displace (CHD) [33], though there were other more complex algorithms prior. The majority of recent papers improve the main algorithm's storage complexity by using various compression schemes [34], [35], [36].

The basic approach used in the hash and displace algorithm of creating a minimal perfect hash is outlined in the Algorithm 1 below. Essentially, one chooses values for the intermediate table starting with those locations that had the most key collisions. After assigning new seed values to locations that had collisions, the remaining keys are displaced into open locations in the table.

Algorithm 1. Pseudocode of MPH Creation

```

1: procedure CREATEMPH(keys, table)
2:   Count the number of keys that fall into each slot of the
   table
3:   Sort keys into buckets in falling order of the count from
   previous step
4:   for (each bucket in order of size (where size > 1)) do
5:     repeat bucket.seed ++
6:     until (All keys in the bucket fall into empty spots in
   the hash table)
7:     Record the new seed value
8:     Mark the location of the keys (using the new seed)
9:   end for
10:  for (each bucket in order of original location in table
   (where bucket size == 1)) do
11:    Let i be the location of next available index in the
   hash table
12:    Let j be the location of single key in the hash table
13:    Record the new offset value (i - j)
14:  end for
15: end procedure

```

4 IMPLEMENTATION

In this section we discuss the various algorithms and hardware designs used to create RAMPS. It is important to remember that RAMPS is composed of two major components: the creation of a minimal perfect hash table and a simple aligner that uses the hash table to perform look ups. These two components were implemented as both a software solution and a hardware solution.

4.1 Approach

When surveying the plethora of available aligners, it becomes clear that the memory bandwidth is the main constraint on alignment speed. Thus, RAMPS is optimized for reducing memory bandwidth. We choose hash tables as the

main data structure because it minimize memory operations overhead, requiring only a few memory operations per read to find an index in the genome. The Convey HC-2 coprocessor was chosen because of its large available memory bandwidth of 80 GB/sec.

The main issue with regular hash tables is the size of the table when scaled to bioinformatics projects. The table size is proportional to the size of the reference genome since there would be one entry per unique 100 base pair sequence in the reference. We used a reference genome from the 1,000 genomes project [31] which contained about 2.8 billion unique 100 base pair segments. This means that every byte of data per entry increases the size of the hash table by 2.8 GB.

If we were to store the key (the short read) in the hash table, we would need about 32 bytes per entry, or a table size of over 80 GB. In addition, any collisions would necessitate a linked list type of structure for collision handling, resulting in more increases in the size of the table. There is no need to store the key (actual read data) in the hash table. We store the read's index to the reference genome in the hash table and use the reference genome itself to make sure the read belonged in the hash table. Additionally, minimal perfect hash functions are only for a fixed set of keys. Short read alignment uses a static reference. Thus, we eliminated the need for collision detection by using a minimal perfect hash function for the table. Algorithm 3 discusses the use of minimal perfect hash tables and their use in performing short read alignments.

The software was written first along with a testing framework. The MPH creation algorithm closely follows that of Belazzougui et al. [33], except we do not compress the intermediate values. A suitable hash function was chosen, Jenkin's Spooky hash [37], that would map easily into hardware. A pictorial explanation of generalized minimal perfect hash construction can be found at Hanov's website [38]. The general process for creating the hardware was a multi-step process: pick good algorithms, draw the designs, create a software model for the Convey software simulator, create the components using Xilinx's Core Generator or Verilog, test the user made components for correctness and timing with a test bench, combine the components together and test using Convey's hardware simulator, synthesize and test the bit files. The design drawings of the hardware are presented later in this section.

4.2 Overview of Convey Architecture

RAMPS's hardware pipeline was built for use with the Convey hybrid-core computing platform containing a minimum of 32GB of coprocessor memory. The Convey HC-1 and HC-2 are hybrid computers, containing a regular motherboard and a coprocessor board that contains a set of 14 FPGAs. Eight FPGAs are wired as memory controllers (MCs), two are used as an Application Engine Hub (AEH), and the remaining four are programmable and called Application Engines (AEs). The host ($\times 86$) processor can send the AEH custom instructions, which will then load a custom bitfile (branded as a "personality") onto the AEs and execute the custom instruction. The coprocessor contains its own memory, though the host processor and coprocessor can share all memory in a cache coherent manner. The entire Convey

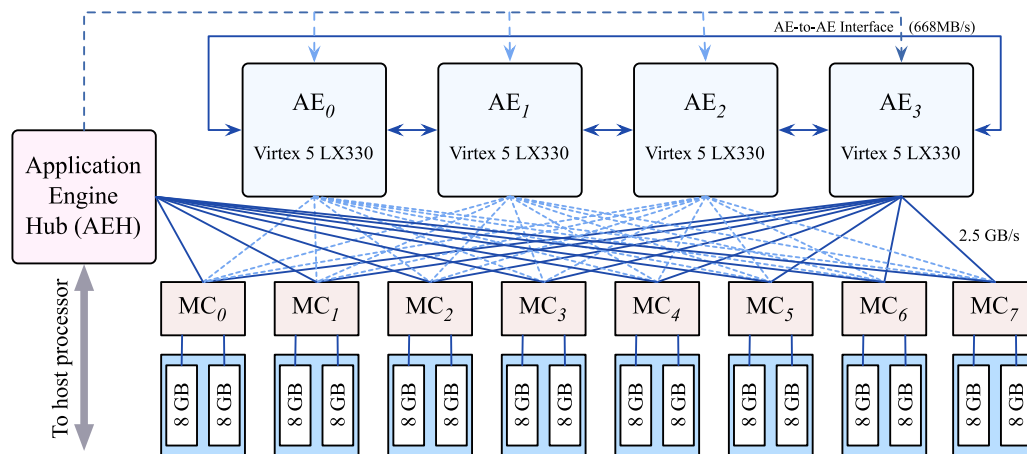


Fig. 2. The Convey HC-2 co-processor consists of four programmable Virtex 5 PFGAs, branded as Application Engines. The FPGAs are connected through a crossbar interface to a eight of memory controllers with a peak bandwidth of 80 GB/s.

HC-2 platform consumes approximately 600 watts with the coprocessor executing code [39].

The distinct competitive advantage that the Convey system provides is its 80 GB/s of memory bandwidth. Fig. 2 shows the 16 memory DIMMs available on the coprocessor board that allow it to access large amounts of memory quickly. In addition, the Convey system has *scatter-gather* DIMMs, allowing random access to memory locations with speed on par with sequential access to memory. In addition to raw hardware speed, Convey provides a rich set of hardware interfaces for accessing memory in its personality development kit (PDK). The PDK is built to allow developers to get their programs up and running quickly, without having to spend time reinventing the memory subsystems. Each AE is given access to 16 memory controller ports, which are multiplexed to the eight MCs. The AEs operate at a clock frequency of 150 MHz, allowing each memory controller port to make 150 million memory requests per second. Using these memory controllers, along with Convey's provided read order queue and crossbar switch, greatly simplified the hardware design by allowing each MC port to access any address and by creating an ordered data flow.

The development of a program on Convey starts with the software. First, a software emulator is written. This software emulator allows the developer to test the interface between host machine and the coprocessor's AEH. This step typically involves thinking about what memory should be moved or allocated on the coprocessor, what parameters to pass to the application engines, and how this data is transferred to and from the coprocessor through the AEH. It also allows the developer to create a model that can be used in the future to test hardware designs. The next step is to design the hardware and create a custom personality. This step involves diagramming hardware layouts, writing and testing individual modules, testing and fixing the larger top level module, and creating a bitfile by synthesizing the written HDL (hardware description language).

4.3 Minimal Perfect Hash Creation

RAMPS's minimal perfect hash is created using a fairly simple algorithm discussed in Algorithm 1. Before running the algorithm, the first step is to find the set of unique entries that need to be stored in the hash table. This can be done by

hashing each 100 base pair word in the reference genome in a number of rounds, throwing away duplicates and storing collisions for processing in the next round. After we have thrown away all the duplicates and processed the left over collisions, we are left with a set of unique keys. With a set of known keys, we can construct a minimal perfect hash using a generalized method called hash and displace [33].

After collecting a set of unique keys, all the keys are sorted using a two pass counting sort. During the first pass, a count of the size of each bucket of the hash table is taken; if a key collides with another key (they both hash to the same value), the bucket's size is incremented. Thus, after the first pass, the number of collisions for each entry in the hash table is known. During the second pass, keys are placed into a buckets array sorted by the size of the bucket. Since the hash function has an equal probability of choosing any given bucket, the size of a bucket is typically orders of magnitude smaller than the length of the hash table. For example, the largest bucket size we encountered was 13 when the algorithm was used with the 2.8 billion entries in the human genome. It is because of the small number of unique sizes that we are able to sort the keys in linear time with a two pass counting sort.

With the buckets now sorted, the algorithm begins to reseed buckets with a size of two or greater. Reseeding ensures that the hash function will be "perfect" and no longer contain collisions. Starting with the largest buckets, the keys in a bucket are reseeded such that they no longer collide and will occupy empty spots in the hash table. The reseed value is stored in an intermediate table and a bit field or bit array is updated to indicate that the locations in the hash table are now occupied. This process continues with the next largest bucket until all buckets containing two or more keys have been assigned a new seed.

Finally, once all buckets containing two or more keys have been reseeded, buckets containing a single key are placed into open spaces in the hash table. The offset from their original location is recorded in the intermediate table. This can be done without looking at the keys, using exclusively the information gained from the counting sort (the entries in the hash table with a single entry are marked) and the information gained from reseeding (the bit array containing a record of the empty and occupied slots in the hash

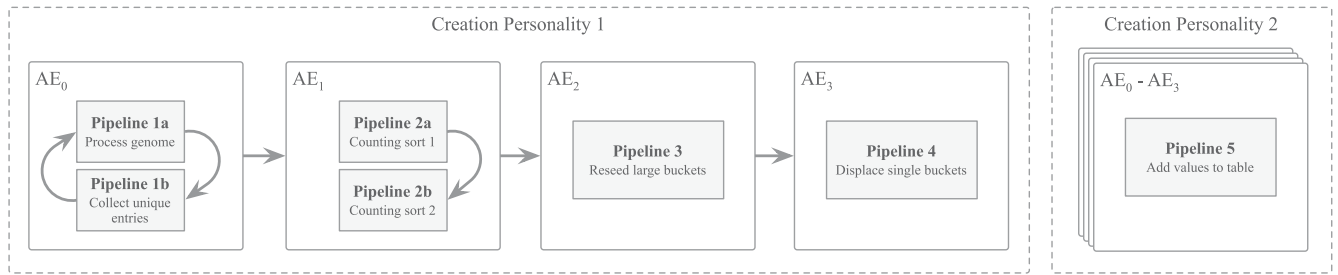


Fig. 3. RAMPS contains two personalities for creating the minimal perfect hash table. The first personality features four different bitfiles (total of six hardware pipelines). The data is processed by each hardware pipeline sequentially, starting with the first. Each major pipeline must process the data completely before moving on to the next phase. Here, AE_0 loops until all unique reads in the genome have been found. Next, AE_1 performs a two pass counting sort. AE_2 then reseeds large buckets, and AE_3 finally displaces the singular buckets. The second personality contains Pipeline 5 on all of its four AEs.

table). The intermediate table provides a minimal perfect hash for the given key set, and can now be used to add values to the table to complete the key-value association.

Since the algorithm for MPH creation naturally falls into five stages, the hardware for creating a minimal perfect hash table is broken into five major pipelines. Four of the major pipelines require proper memory ordering, atomic increment operations, or atomic test and set operations. These memory requirements can be solved by using a small cache, indicated by a lock on the memory controller in Figs. 7, 8, 9, and 10. Using our current design, it would be impractical to run these four major pipelines on all four application engines available on the Convey coprocessor board due to the complexity of maintaining cache coherence across all four chips. However, Botelho's dissertation [32] outlines ways of dividing the problem appropriately for a distributive version of the algorithm that scales nearly linearly with the number of nodes added.

The fifth major pipeline does not require special memory requirements, and thus it can be run on all application engines. All five pipelines for MPH construction occupy two personalities: the first personality contains Pipelines 1-4 and the second personality contains Pipeline 5. It should be noted that in Figs. 7, 8, 9, 10, 11, 12, and 13, the data flows from the left to the right of the diagram over the course of time and stages are marked on memory boundaries. Fig. 3 however, shows the data flow between application engines.

In RAMPS, the intermediate table and hash table are interleaved. With smaller table sizes (less than 20 million entries), it makes sense to have a separate intermediate table with compressed values so that the intermediate table can fit into the cache on a processor. In our case, a human sized genome requires a table size of nearly 3 billion entries, the extra complexity of compressing the intermediate table was skipped and it was instead placed along with the values of the hash table. In fact, this interleaving may improve performance marginally; offsets in the intermediate table are small and the result may be within a few entries in the table, thus if cache lines on the chip are larger than 8 bytes it may result in a cache hit or the memory controller can combine operations since the data may reside in the same row of a RAM module. The layout of a single entry can be seen in Fig. 4.

The software package of RAMPS, unlike the current version of the hardware, is configurable using a command line argument to allow creating hash tables of any number of base pairs that is a multiple of 4. The multiple of 4 was

chosen because the bit packed representation stuffs four base pairs per byte, thus keeping the data byte aligned. The default number of base pairs used for creating the hash table is 100. We use a command line argument to specify a different size in bytes. Additionally, a command line argument specifies if the reference genome contains paired or unpaired chains of the DNA and controls whether the hash function produces a normal hash or a paired hash. A paired hash means that both chains of the DNA at a given index will hash to the same value; e.g. $\text{hash}(\text{"AAAG"}) == \text{hash}(\text{"CTTT"})$. The software simply calculates a read's pair (reverse the sequence and swap base pairs) and hashes whichever input is greater. Paired hashing effectively doubles the number of reads you can align to a reference genome that only contains one of the DNA chains, and is enabled by default.

4.3.1 Hardware Hash Function

Most of the computation of the algorithm occurs in the hash function. Jenkin's Spooky Hash [37] was chosen because it is both fast in software and easy to implement in hardware due to its reliance on only shifts, adds, and XOR operations. RAMPS uses a slimmed down version Jenkins' Spooky Hash that has been stripped to work with only keys of size 25 bytes (100 base pairs), though adjustments could easily be made to handle keys of any given bytes size below a constant maximum. By stripping unnecessary branches and instructions, each hash requires 23 rotations, 23 XORs, 27 additions, and 1 mod operation. By grouping operations together, a 34 stage pipelined hash function was created in hardware.

The pipeline stages listed in Fig. 5 are registered, increasing the maximum allowed clock frequency for the unit.




Fig. 4. An entry in the Minimal Perfect Hash table. The table is an array with one 8 byte entry per unique read in the reference genome. The intermediate table is stored in 21 bits and is used to find a unique index into the table for each key. The first 43 bits contain the value associated with the key; i.e. the index of occurrence in the reference genome and a count of how many times it occurs.

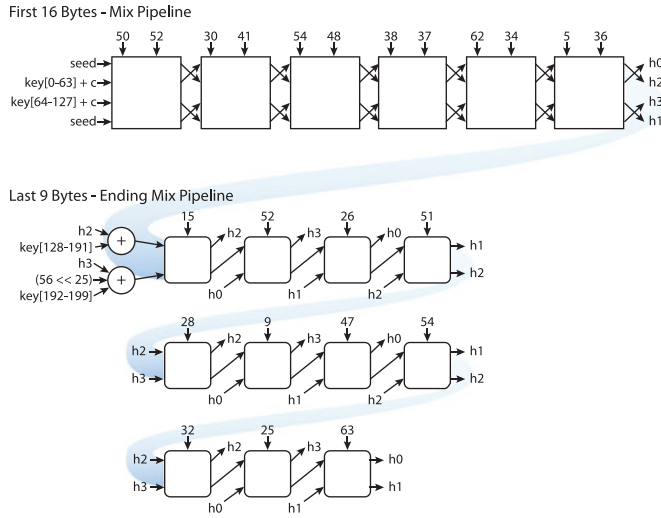


Fig. 5. The hardware hash pipeline is composed of several blocks which are further described in Fig. 6. The diagram here shows the mixing of the first 16 bytes followed by the last 9 bytes. The seemingly random constant values help the hash function create a waterfall, where a single bit change in the key being hashed causes all the bits to change with equal probability.

Further registers could be added within the mix blocks for additional speed improvements. The key is fixed at 25 bytes, or 100 base pairs. The hashing algorithm hashes the key 16 bytes at a time, and has specific ending mix blocks for the last 9 bytes.

Like the software version of the program, the hardware hash function supports paired hashing; i.e. hash("AAAG") is equal to hash("CTTT"). The hardware performs the necessary wire assignments and bitwise NOT operation as the first stage of the pipeline, and computes the hash of the larger of a read or the read's antiparallel pair. This can be disabled using the command line '-d' flag.

4.3.2 Pipeline 1a/b - Find Unique Entries

The first step is to find all the unique entries that need to be stored in the hash table. This can be done by hashing each 100 base pair word in the reference genome in a number of rounds, throwing away duplicates and storing collisions for processing in the next round. Algorithm 2 shows the process of removing duplicates.

Algorithm 2. Pseudocode of Pipeline 1a/b

```

1: procedure PIPELINE1A(list, genome, hashtable, collisions)
2:   for (i = 0; i < length(genome) - 99; i ++ ) do
3:     list[i] ← i;
4:   end for
5:   while list not empty do
6:     erase(hashtable);
7:     Pipeline1a(list, genome, hashtable, collisions);
8:     Pipeline1b(hashtable, unique);
9:     swap(list, collisions);
10:  end while
11: end procedure
    
```

During the first round, all valid indices in the genome are added to a list for processing. During the processing step (pipeline 1a), any portion of the genome that creates a collision in the hash table is set aside in a collisions list for

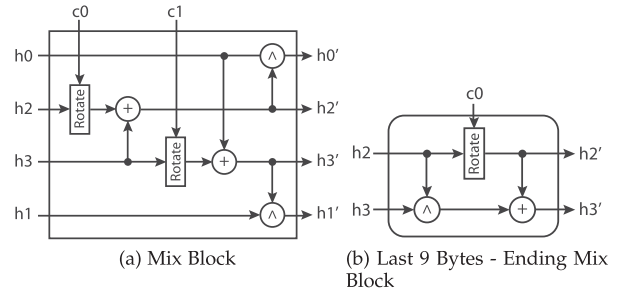


Fig. 6. The hash function pipeline in Fig. 5 is composed of several sub-components shown here. To the left (a) are the mix blocks used for hashing the first 16 bytes of data. On the right (b) is an ending mix block used to hash the last 9 bytes of the key.

processing in the next round. Due to the random nature of the hash function, roughly one-third of the indices collide and need to be processed in the next round, leading to a total runtime that is proportional to 1.5 times the length of the genome (sum of geometric series).

Fig. 7 illustrates this process in more detail. At the beginning of each round, the hash table is erased to remove left-over intermediate data from the previous round. First, an index into the genome is loaded from the list. Next, that index is used to load a 100 base pair read from the genome. While during the first round the indices are in order, the random nature of collisions will cause subsequent rounds to be unordered. In stage 3, the read is sent through the hashing pipeline in order to locate the correct index from the hash table. After finding the hash table entry, control flow breaks in one of two directions. Either the entry is empty and can be updated quickly, or there is a possible collision or duplication which must be checked by loading the original index stored in the hash table. Duplication happens when the genome contains multiple copies of the same 100 base pair sequence. Collisions occur when the hash of two different 100 base pair sequences are equal.

An important property of Pipeline 1a, 2a/b, and 3 is the ability to perform atomic read-write operations. This is indicated in the figures using a lock on the specified memory controller ports. The atomic operations are implemented by using an ordinary cache with the addition of a lock bit. Any read operations to a locked address are placed into a reply buffer and can receive the unlock signal from an incoming write operation with the same address.

After a list has been processed, Pipeline 1b (Fig. 8) quickly scans the hash table to collect non-empty entries into a tight array. Since the order of the unique reads, or keys, is unimportant, this step can utilize multiple processing units to achieve a significant speedup. This step is necessary; pipelines 2a, 2b, and 5 will each iterate over the list of keys. Removing the empty spaces now reduces the total memory requirement and eliminates memory operations in the future. After copying all the unique keys to a results array, the round ends. If there were any collisions, a new round begins by processing the list of indices added to the collisions list.

4.3.3 Pipeline 2a/b - Counting Sort

With a set of unique keys, construction of the minimal perfect hash can begin. As described in Algorithm 1, the first step is to use a counting sort to store the keys into buckets

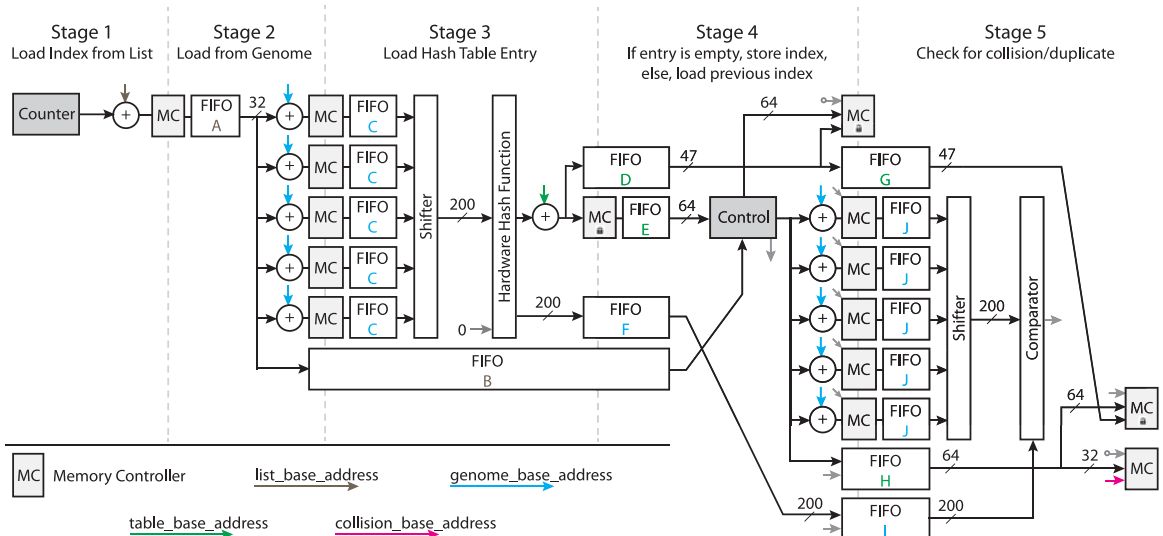


Fig. 7. Pipeline 1a is broken into five minor stages separated by memory operations, shown here flowing from left to right. The pipeline is responsible for processing a list of indices in the genome, removing duplicates and storing collisions for later processing.

sorted by the size of the bucket. This is accomplished in two passes, but the hardware pipeline is mostly the same for each. The exception is that the 2nd pass stores the unique index into the buckets array, while the first only calculates the size of each bucket.

Before beginning, the hash table is erased. During the first pass, each index is hashed and the corresponding entry in the hash table is incremented by one (indicating the number of indices that hash to a given bucket). In addition, registers on the FPGA hardware are used to keep a running total of the number of buckets of each size. This is accomplished by simply incrementing and decrementing running total counts as each key is hashed. Since the maximum expected bucket size is small [33], we can dedicate a small number of registers for this purpose (RAMPS uses 30). By keeping track of the number of buckets of each size, it allows the algorithm to use a counting sort that runs in time $O(n)$ with respect to the number of keys, n .

After calculating bucket size, the second pass sorts the keys. The initial stages of the pipeline are reused in order to load a unique read, or key, hash it, and load the corresponding value from the hash table. During the last stage of the pipeline, two possibilities exist every time a table entry is loaded: either this is the first key encountered for a given bucket, or it is not the first key. If it is the first key encountered for a given bucket, an appropriate index is calculated from the running totals calculated in the first pass. This index is then stored in the hash table, and a tally is incremented for the bucket. If there is a tally that is not 0, or an index stored in the hash table, then we know the key is not the first key encountered for the bucket. In this case, the tally is added to the bucket index stored in the hash table to calculate the appropriate index for storing the unique read, or key, in the sorted bucket's list. Fig. 9 shows this pipeline.

4.3.4 Pipeline 3 - Reseed Large Buckets

Starting with the largest buckets, the keys within each bucket are reseeded such that they no longer collide to the same spot in the hash table. The reseed value is stored in an intermediate section of the hash table. The hardware

pipeline was designed to handle buckets containing 2 to 5 keys. There are too few buckets with a size greater than 5 to necessitate designing another hardware pipeline. On a human genome scale, RAMPS software can reseed the buckets with a size greater than 5 in less than 1 second. Buckets containing a single key are displaced in the next hardware pipeline.

Fig. 10 illustrates the hardware pipeline for reseeding buckets. To reseed a single bucket, all the keys and segments of the genome are first loaded. The reads are rehashed with a new seed for the hash function, which results in the previously colliding keys having different indices. A bit array is checked to ensure that space is available in the hash table and to reserve the new slots for the

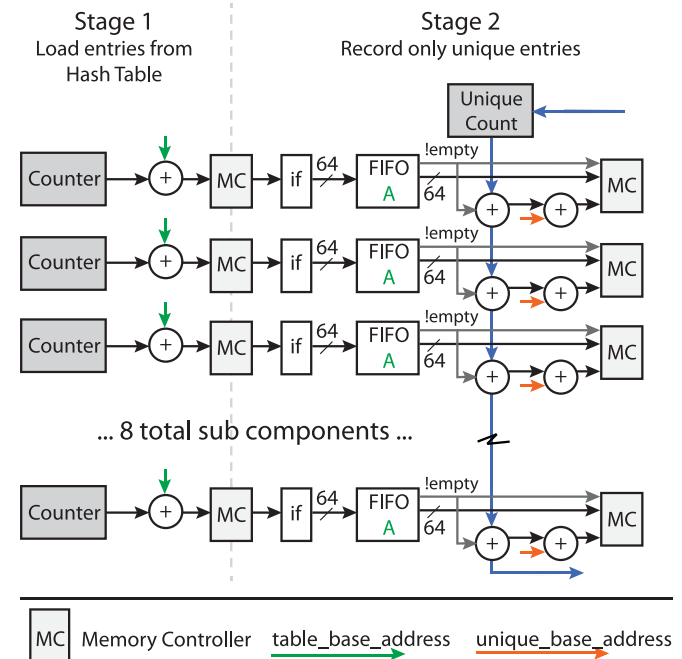


Fig. 8. Pipeline 1b runs through the hash table and collects all unique entries into a list of unique entries. Each of the eight subcomponents is responsible for loading 1/8th of the hashtable.

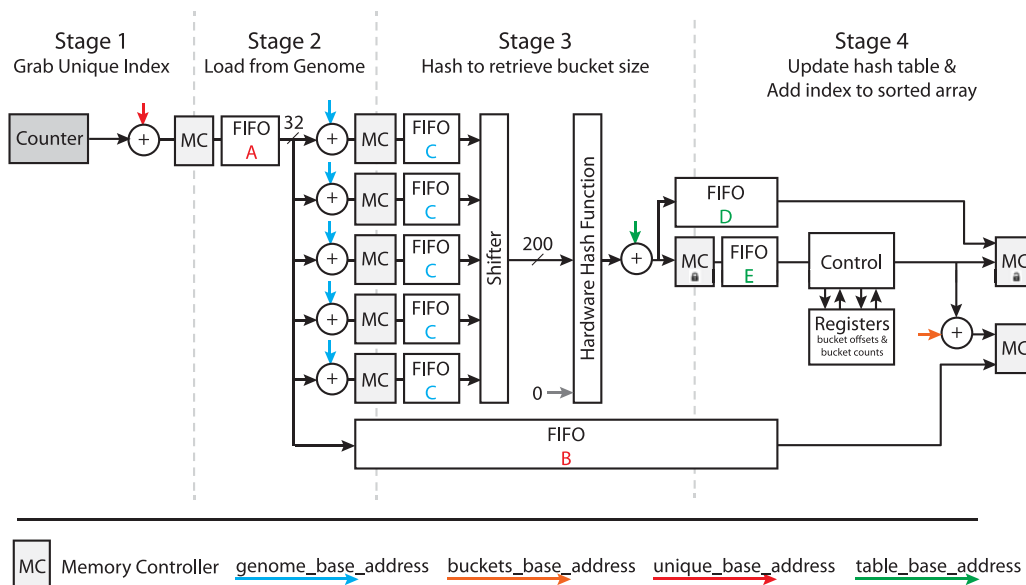


Fig. 9. Pipeline 2a/b: For the first pass, FIFO B and the MC port writing to the buckets array are disconnected. The first step is to load a unique reference genome index, which is used to load a 25 byte read from the genome in stage 2. In stage 3, the read is hashed to retrieve an index into the hash table. After the entry from the hash table is loaded, stage 4 does one of two things. If it is the 1st pass, it increments the count and saves the entry. If it is the 2nd pass, the control calculates the correct index into the buckets array to store the index. The hash table is updated with the location of the bucket in the buckets array and the number of indices for that bucket already stored.

keys that were just reseeded. If any of the keys collide when rehashed or if they fall into a taken spot in the hash table, the bucket is reseeded again. To reseed, the seed value to the hash function is simply incremented because a single bit change in the key or seed results in a waterfall change in the hash.

4.3.5 Pipeline 4 - Displace Singular Buckets

Once all large buckets containing two or more keys have been reseeded, the singular buckets containing just one key are placed into open spaces in the hash table. The open spaces in the hash table are those left empty after reassigning

all keys that fell into large buckets. The offset from their original location is recorded in the intermediate section of the table.

This process, shown in Fig. 11, must be done in order so that the keys receive small offset values. This is done by running through the hash table and bit array in order. The first entry in the hash table that has only a single key is matched with the first empty index in the bit array. The next single entry is matched with the next empty spot in the hash table, and so on, until the last single entry is matched with the last empty spot in the hash table. After this step is complete, the minimal perfect hash has been created. The only step left is to use the MPH to add the values to the hash table.

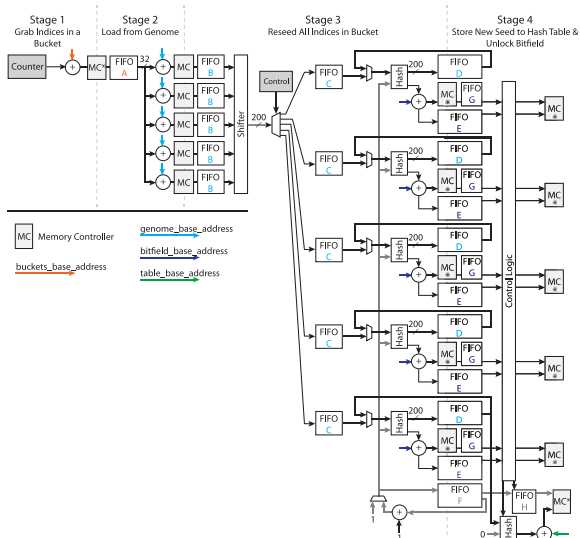


Fig. 10. Pipeline 3 runs multiple times, once for each size of bucket between 5 and 2. The bucket size controls the multiplexor that stores a read into one of the five FIFOs in bank C. Stage 4 performs a test and set operation on the bit array to ensure that each index is assigned a unique slot in the hash table.

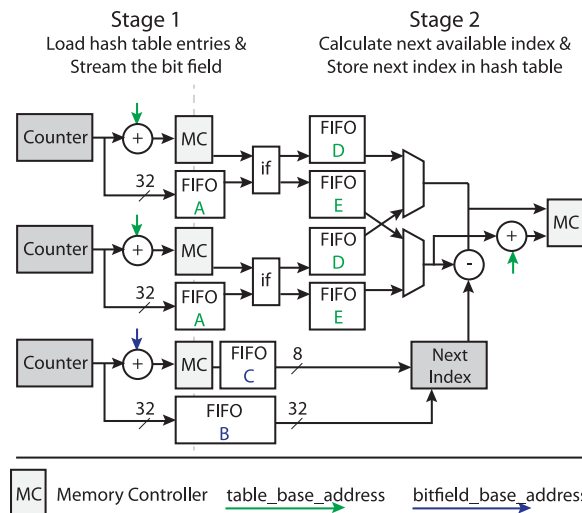


Fig. 11. Pipeline 4 reads multiple entries from the hash table at a time in stage 1, as not all entries are singular buckets. It also streams the bit array. An indexer calculates the next available index open in the hash table by looking for the next available bit in the bit array. The smallest index from the hash table in the bank of FIFOs D and E is then displaced and the intermediate value is written to memory.

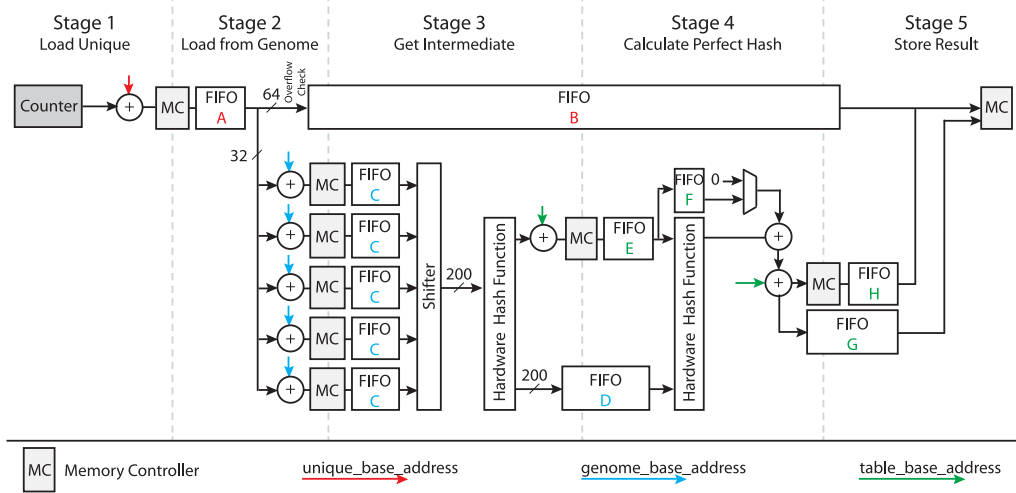


Fig. 12. Pipeline 5 streams in the entire set of keys (unique indices), loads the genome to retrieve the read, hashes to retrieve the intermediate value, rehashes, and finally stores the data from stage 2 into the table.

4.3.6 Pipeline 5 - Add Values to Table

Finally, the pipeline shown in Fig. 12 adds the unique indices and other data to the hash table. Unlike previous pipelines for minimal perfect hash table creation, this pipeline can be run on all four application engines. RAMPS partitions the list of unique keys in the following manner: AE0 adds keys 0, 4, 8, ..., AE1 adds keys 1, 5, 9, ..., and so on. The process of storing data in the table is similar to the process described in Fig. 1 and shares similarity with the hardware aligner.

4.4 Alignment via Hash Table Lookup

The alignment algorithm is simply to retrieve the correct alignment from the hash table. All possible exact match alignments have been preprocessed from the genome into the hash table. To retrieve an index of occurrence, the program simply attempts to lookup the read in the hash table, validating the index of occurrence is corrected by comparing the read to the genome at the retrieved index. Algorithm 3 formally describes this process. The five stages of the hardware pipeline for alignment are shown in Fig. 13. This pipeline is duplicated across all four application engines. The set of reads is split into four pieces, and each AE contains an identical pipeline to process its subset of the reads.

Fig. 13 shows the flow of data through the aligner's hardware pipeline. In stage 1 of the pipeline, a bank of four counters and the base address of the read data is used to calculate the address for a given read. These addresses are used to load short reads from four memory controller ports. After a memory latency of about 100 cycles, stage 2 hashes the short reads to obtain an index for retrieving a value from the intermediate table. In stage 3, the value from the intermediate table allows RAMPS to compute an index that is guaranteed to not collide with other entries in the table. The value is either a new seed value for the hash function, or an offset. The unique index is calculated and used to load the value from the hash table. In stage 4, the index in the genome loaded from the hash table is used to load the corresponding 100 base pairs from the reference genome. In stage 5, the reference genome is compared to the read. If they match, the index of occurrence in the genome and

number of occurrences are recorded in an output table in memory.

Algorithm 3. Pseudocode of Hash Table Lookup

```

1: procedure ALIGN(reads, genome, hashtable, results)
2:   for ( $i = 0; i < \text{length}(\text{reads}); i++$ ) do
3:      $r \leftarrow \text{reads}[i]$ ; ▷ Stage 1
4:      $h \leftarrow \text{hash}(r, \text{seed} = 0)$ ; ▷ Stage 2
5:      $\text{ivalue} \leftarrow \text{intermediateTable}[h]$ ;
6:     if (ivalue is an offset) then ▷ Stage 3
7:        $\text{index} \leftarrow \text{hash}(r, \text{seed} = 0) + \text{ivalue}$ ;
8:     else
9:        $\text{index} \leftarrow \text{hash}(r, \text{seed} = \text{ivalue})$ ;
10:    end if
11:     $\text{entry} \leftarrow \text{hashtable}[\text{index}]$ ;
12:     $\text{check} \leftarrow \text{genome}[\text{entry.index}]$ ; ▷ Stage 4
13:    if ( $r == \text{check}$ ) then ▷ Stage 5
14:       $\text{results}[i] \leftarrow \text{entry}$ ;
15:    else
16:       $\text{results}[i] \leftarrow \text{NULL}$ ;
17:    end if
18:  end for
19: end procedure

```

When the pipeline is full, each stage loads or stores data to its memory controller ports on every clock cycle. Much of the combinatorial logic, such as address calculations can occur at the same frequency. The exception is the hardware hash functions, each of which is a 34 stage pipeline. By allowing each portion of the pipeline to perform some part of the alignment on every clock cycle, we efficiently utilize resources and are able to achieve a theoretical throughput of 150 million reads per second per application engine (i.e., total of 600 million reads per second). In practical tests, the alignment speed was approximately 350 million reads per second.

In addition to the hardware aligner, which can only perform exact matches in its current implementation, the software aligner can be configured to split long reads into smaller parts, each part returning its own index of occurrence. If any indices match, a comparison is done between the entire read and the genome at the matching index

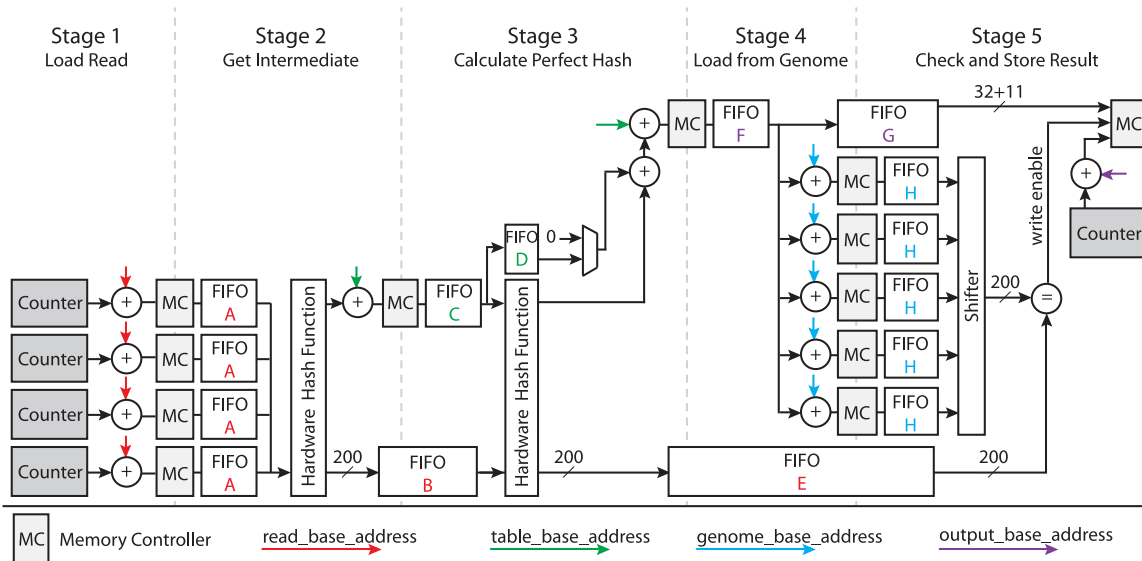


Fig. 13. RAMPS’s hardware hash table lookup pipeline is broken into five stages, shown here flowing left to right.

allowing for mismatches. This allows the ability of finding matching locations in the reference genome that could be possible locations of the read data. The software aligner takes the following arguments: read length (default = 25), key length (default = read length), and offset (default = key length).

The read length is the length of the read data in bytes. The aligner assumes that the read data is 8-byte aligned. For example, if using the default 25 byte read size, each read will occupy 32 bytes (the last 7 bytes are left empty) in order to align to an 8 byte boundary. The key length is the size of the key in bytes used for creating the hash table. The offset is used by the software aligner to perform more or less hashes in order to improve sensitivity. The RAMPS software aligner tool can run in one of three settings: exact matches, a medium sensitivity setting, and a higher sensitivity setting.

For example, using the medium sensitivity setting with the human genome (human g1k v37), the read length is the default of 25 bytes, while the hash table was created using 8 byte words from the reference genome (32 base pair). The technique used was to take three 32 bp chunks out of the 100 bp read, take the index that was in the majority out of the three chunks, and do a comparison between the genome at that index allowing for up to 5 mismatches. Though numbers will vary based upon the quality of read data used, this procedure was able to increase the number of reads aligned by about 20 percent up to 75 percent, whereas the number of reads containing exact alignments was 52 percent.

TABLE 1
Preprocessing Runtime Comparison of Popular Short Read Aligners

| Algorithm | Platform | Preprocessing | Memory (GB) |
|-------------|----------|---------------|-------------|
| Bowtie [13] | CPU | 4-5 hours | 16 |
| BWA [40] | CPU | 3 hours | 2.5 |
| RAMPS | CPU | 2-3 hours | 60 |
| RAMPS | Convey | 90 seconds | 60 |

5 RESULTS

5.1 Preprocessing Runtime

Most short read aligners perform preprocessing of the reference genome into a more suitable data structure that supports faster queries. Traditionally, it is faster to download the preprocessed data structure that someone else has created than to build your own. Using the index build times listed on various project websites and in papers, we compare the preprocessing time required before doing alignment to a certain reference genome.

Table 1 compares the preprocessing runtime and space requirement of RAMPS against the popular short read aligners Bowtie [13] and BWA [40]. RAMPS’s preprocessing time in hardware is orders of magnitudes faster than other approaches. While other hardware hash functions exist and hash tables with billions of entries have been created, RAMPS is comparatively much faster in hardware. RAMP’s preprocessing step requires 20 GB memory space for storing the resulting hash table, 20 GB for determining the unique keys, 10 GB for storing a list of indices, and 10 GB for storing a list of collisions. Table 2 compares the preprocessing runtime of RAMPS against other popular implementations of the minimal perfect hash algorithm.

It is important to note that the reported preprocessing runtime doesn’t include the time to copy data to co-processor memory and the time to reconfigure between creation personalities. Since this overhead is not dependent upon the

TABLE 2
Runtime Comparison of Popular Minimal Perfect Hash Algorithms

| Tool | Platform | Keys/second |
|--------------|----------------|-------------|
| CHD [33] | CPU | 770,000 |
| BPZ [33] | CPU | 910,000 |
| RAMPS | CPU | 260,000 |
| Botelho [32] | 14 CPU Cluster | 4,000,000 |
| RAMPS | Convey | 30,000,000 |

TABLE 3
Performance Comparison of Popular Short Read Aligners

| Tool | Platform | Speed (reads/s) | Memory (GB) |
|--------------|---------------|-----------------|-------------|
| MAQ [9] | CPU | 50 | 1.2 |
| SOAP | CPU | 70 | 14.7 |
| SOAP2 [11] | CPU | 2,000 | 5.4 |
| Bowtie [13] | CPU | 2,500 | 2.3 |
| BWA [40] | CPU | 10,000 | 3.5 |
| PerM [27] | Convey HC-1ex | 109,606 | 14.0 |
| SOAP3 [43] | GPU | 200,000 | 3.2 |
| BFAST [21] | CPU Cluster | 700,000 | 24.0 |
| BWA-Convey | Convey | 350,000 | - |
| BWT-GPU [19] | GPU | 400,000 | 10.0 |
| RAMPS | CPU | 800,000 | 23.3 |
| RAMPS | Convey HC-2 | 315,000,000 | 23.3 |

number of reads processed, this time gets lost asymptotically. Also, It should be noted that the preprocessing runtime is not an entirely fair comparison. In the sense that RAMPS performs the extra steps of removing duplicates from the genome, indirectly referencing the keys in the genome via index, and processed 2.8 billion keys. Both CHD and BPZ were being tested on a set of 20 million URLs.

5.2 Alignment Runtime

Our short read aligner's runtime is 895 milliseconds for processing 284,881,619 short read sequences from the NA06985 individual [41] and using human reference genome g1k v37 [31]. The timer starts before the reads begin streaming to the hardware pipeline, and ends after the last read's index and count are written to the result array. Table 3 compares the alignment runtime of RAMPS against other short read aligners using the alignment speeds from [42], [43], and other previous work.

RAMPS's runtime omits the one time costs associated with loading and preprocessing both the 22.5 GB hash table and the 780 MB reference genome. In addition, the stated runtime does not include the time it takes to load the reads from disk into memory. Convey's internal benchmarking puts the bandwidth from host memory to coprocessor memory at 2.4 GB/second. Since the aligner's pipeline can handle approximately 10 GB/second of read data, the host to coprocessor memory bandwidth is currently RAMPS's bottleneck. With an estimated world sequencing output of 15 quadrillion base pairs per year [4] (40 GB/s in compressed form), a single RAMPS machine would be able to process all the data in roughly seven days.

The work in [27] is the closest to our implementation. The authors developed a two-level hash-index based algorithm

on the Convey HC-1ex system. As shown in Table 3, RAMPS performs thousands times faster using 1 PE per FPGA than their proposed implementation, which uses 16 PEs per FPGA. While RAMPS utilizes multiple MCs per PE, the authors in [27] employ only one MC per PE where every PE aligns a different read which leads to higher probability of racing conditions and consequently a poor utilization memory bandwidth.

5.3 Hardware Usage

Another consideration was RAMPS's use of available hardware resources on the Xilinx Virtex 5 LX330 [44]. Each of the four FPGAs on the Convey system has 288 36Kbit of Block RAM, 207,360 Slice LUTs, and 207,360 Slice Registers. The data in Table 4 was collected from the map and place & route reports created after building the various bit files. It should be noted that a significant portion of the FPGA resources are used by Convey's optional memory components, such as the read order queue, crossbar switch, and write complete interface.

To estimate the amount of resources used by Convey's hardware interfaces to the memory controllers and application engine hub, a superfluous control personality, named Simpleton, was created. The Simpleton personality includes Conveys read order queue, crossbar switch, and write complete interfaces. It increments (reads and writes) a single address from memory on all the memory controller ports and nothing else. Table 4 shows the hardware usage of each pipeline. All designs utilize the crossbar switch and read order queue, but those pipelines that utilize the write complete interface are labeled. Additionally, the number of ports used for read operations and write operations are given; if one of the operations is tied to zero, some of the logic used for controlling read ordering is optimized out of the design. All memory controller write ports are utilized in the MPH creation pipelines 1-4 because these pipelines implement an erasing function in addition to their more complex pipelines.

RAMPS's alignment runtime, like many of the other short read aligners, is memory bound. The current design uses 12 out of the 16 available memory controller ports. Putting multiple pipelines on a single AE in order to use all available memory bandwidth would increase complexity and lead to only a marginal improvement of performance.

5.4 Alignment Percentage (Sensitivity)

RAMPS's hardware package is currently designed for a constant read length, though there are simple tweaks that could be used to support creating a hash table for read lengths less than 100 base pairs. RAMPS's software package includes

TABLE 4
Hardware Resource Usage Per Application Engine

| Personality | Block RAM | LUTs | Flip-Flops | MCs (Rd-Wr) | Write Complete |
|-----------------|-----------|------|------------|-------------|----------------|
| Simpleton | 20% | 32% | 39% | 16-16 | Yes |
| Aligner | 30% | 43% | 45% | 11-1 | No |
| MPH Create 1a/b | 32% | 68% | 72% | 12-16 | Yes |
| MPH Create 2a/b | 28% | 54% | 60% | 7-16 | Yes |
| MPH Create 3 | 44% | 81% | 85% | 8-16 | Yes |
| MPH Create 4 | 23% | 29% | 35% | 3-16 | No |
| MPH Create 5 | 28% | 43% | 45% | 8-1 | No |

TABLE 5
Alignment Comparison of Popular Short Read Aligners

| Tool | Platform | Max Alignment % |
|-------------|----------|-----------------|
| MAQ [9] | CPU | 93.2 |
| SOAP | CPU | 93.8 |
| SOAP2 [11] | CPU | 93.6 |
| Bowtie [13] | CPU | 91.7 |
| SOAP3 [43] | GPU | 96.8 |
| RAMPS | CPU | 75.4 |
| RAMPS | FPGA | 52.4 |

options to create hash tables based on any length of read that is a multiple of 4 (so as to align to the byte boundary). When performing alignments, long reads can be split into smaller parts, each part returning its own index of occurrence. If any indices match, a comparison is done between the entire read and the genome allowing for mismatches.

In Table 5, the read length is 25 bytes (100 base pairs). The hash table was composed of all 32 base pair words in the human genome, and alignment would divide each read into three parts: bytes 0-7, 8-15, 16-23. By reanalyzing the reads that didn't have exact matches, we can improve the alignment percentage beyond exact matches:

The hardware (FPGA) implementation performs only exact matches, while the software aligner allows up to five mismatches, though it is not guaranteed to find the best location or report all locations. By fundamentally changing the preprocessing time, we can encourage the use of reference genomes that more closely match the read data; e.g. by using a parent or relative's DNA. Combined with an increase in read quality, RAMPS can align more of the data.

6 CONCLUSION

RAMPS offers the bioinformatics community a new tool for fast exact match short read alignment with minimal preprocessing time. RAMPS is both the first known implementation of a generalized minimal perfect hash creation algorithm using FPGAs and an award winning, first place, exact match short read aligner [30].

Currently, the number of mismatches between the reference genome and read sequences occur from both imperfect read quality and genetic differences between individuals. The DNA of two individuals differs by roughly 0.1 percent, or about one base pair out of a thousand [30]. With minimal preprocessing, we fundamentally change the computational problem. Instead of using a generic human reference genome, people may now be able to use the DNA sequence of a blood family member as a reference in order to increase the percentage of exact matches. Higher quality sequencing machines already produce data sets that contain more than 60 percent exact matches. Thus, as quality improves and with the ability to use references that contain nearly identical DNA, exact match read aligners like RAMPS will be able to handle a larger share of the data.

6.1 Future Work

The alignment part of the problem disappears from the runtime when using approaches like RAMPS. Future work will be needed to pull reads from disk or the network at speeds

near 10 GB/s in order to keep such a hardware pipeline busy. While RAMPS offers the potential of increasing the speed of exact matches by orders of magnitude, work still needs to be done for inexact matching allowing mismatches and indels. One could also perform a comparative power study between the various aligners, though the results would likely be moot. FPGAs are known to be power efficient; their custom hardware means more energy is spent solving the problem and less energy is wasted. The RAMPS architecture spends such a small fraction of time performing alignments and runs at a clock frequency of 150 MHz, and would likely consume a fraction of the energy compared with other aligners.

Additionally, there are other application domains in which the RAMPS architecture could prove useful with minimal modification. In 2011, Google developed their own hash function, CityHash [45], to improve the speed of their hash table lookups at their data centers. A package similar to RAMPS, but allowing for variable length keys, would have been orders of magnitude faster. Minimal perfect hashing is an ideal answer for transforming URLs, which are normally static, into integers in a fixed range. Instead of using B-trees for databases numbering in the billions, minimal perfect hashing could provide solutions for new frontiers of database applications.

ACKNOWLEDGMENTS

This work has been partially supported by the National Science Foundation (NSF) under awards CNS-1116810 and CCF-1149539.

REFERENCES

- [1] (2012, Oct.). Illumina inc. sequencing portfolio [Online]. Available: <http://www.illumina.com/systems/sequencing.html>
- [2] A. Pollack. (2011). "DNA sequencing caught in deluge of data," *The New York Times* [Online]. Available: <http://www.nytimes.com/2011/12/01/business/dna-sequencing-caught-in-deluge-of-data.html>
- [3] (2012, Nov.). BGI. Beijing genomics institute [Online]. Available: <http://www.genomics.cn/en/index>
- [4] M. C. Schatz and B. Langmead, "The DNA data deluge," *IEEE Spectr.*, vol. 50, no. 7, pp. 28–33, Jul. 2013.
- [5] (2012, Nov.). A. Burke. DNA sequencing is now improving faster than Moore's law [Online]. Available: <http://www.forbes.com/sites/teconomy/2012/01/12/dna-sequencing-is-now-improving-faster-than-moores-law/>
- [6] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, 1965, pp. 114–117.
- [7] C. Nelson, K. Townsend, B. S. Rao, P. Jones, and J. Zambreno, "Shepard: A fast exact match short read aligner," in *Proc. Int. Conf. Formal Methods Models Codes.*, Jul. 2012, pp. 91–94.
- [8] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman, "Basic local alignment search tool," *J. Molecular Biol.*, vol. 215, no. 3, pp. 403–410, 1990.
- [9] H. Li, J. Ruan, and R. Durbin, "Mapping short DNA sequencing reads and calling variants using mapping quality scores," *Genome Res.*, vol. 18, no. 11, pp. 1851–1858, 2008.
- [10] T. F. Smith and M. S. Waterman, "Identification of common molecular subsequences," *J. Molecular Biol.*, vol. 147, no. 1, pp. 195–197, 1981.
- [11] R. Li, C. Yu, Y. Li, T.-W. Lam, S.-M. Yiu, K. Kristiansen, and J. Wang, "SOAP2: An improved ultrafast tool for short read alignment," *Bioinformatics*, vol. 25, no. 15, pp. 1966–1967, 2009.
- [12] H. Li and R. Durbin, "Fast and accurate short read alignment with burrows-wheeler transform," *Bioinformatics*, vol. 25, no. 14, pp. 1754–60, 2009.

- [13] B. Langmead, C. Trapnell, M. Pop, and S. L. Salzberg, "Ultrafast and memory-efficient alignment of short DNA sequences to the human genome," *Genome Biol.*, vol. 10, no. 3, pp. R25:1–R25:10. 2009.
- [14] P. Ferragina and G. Manzini, "Opportunistic data structures with applications," in *Proc. Annu. Symp. Found. Comput. Sci.*, 2000, pp. 390–398.
- [15] M. Burrows and D. J. Wheeler, "A block-sorting lossless data compression algorithm," Digital Equipment Corporation, Tech. Rep. 124, 1994.
- [16] A. Arbabi, M. Gholami, M. Varmazyar, and S. Daneshpajouh, "Fast CPU-based DNA exact sequence aligner," in *Proc. Int. Conf. Formal Methods Models Codes.*, 2012, pp. 95–98.
- [17] S. Aluru and N. Jammula, "A review of hardware acceleration for computational genomics," *IEEE Des. Test*, vol. 31, no. 1, pp. 19–30, 2014.
- [18] A. Aji, L. Zhang, and W. chun Feng, "GPU-RMAP: Accelerating short-read mapping on graphics processors," in *Proc. Int. Conf. Comput. Sci. Eng.*, 2010, pp. 168–175.
- [19] J. Torres, I. Espert, A. Dominguez, V. Garcia, I. Castello, J. Gimenez, and J. Blazquez, "Using GPUs for the exact alignment of short-read genetic sequences by means of the burrows-wheeler transform," *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, vol. 9, no. 4, pp. 1245–1256, Jul. 2012.
- [20] C.-M. Liu, T. K. F. Wong, E. Wu, R. Luo, S.-M. Yiu, Y. Li, B. Wang, C. Yu, X. Chu, K. Zhao, R. Li, and T. W. Lam, "SOAP3: Ultra-fast GPU-based parallel alignment tool for short reads." *Bioinformatics*, vol. 28, no. 6, pp. 878–879, 2012.
- [21] N. Homer, B. Merriman, and S. F. Nelson, "BFAST: An alignment tool for large scale genome resequencing," *PLoS ONE*, vol. 4, no. 11, Nov. 2009.
- [22] Convey Computer. (2011). Convey's new burrows-wheeler alignment delivers 15x increase in research efficiency [Online]. Available: <http://www.conveycomputer.com/files/3713/5095/9172/BWA.Final.NR.10.11.1-1.pdf>
- [23] Convey Computer. (2011). Convey computer smith-waterman personality [Online]. Available: <http://www.conveycomputer.com/files/1113/5085/5467/ConveySmithWaterman6202011.pdf>
- [24] H. Waidyasooriya and M. Hariyama, "Hardware-acceleration of short-read alignment based on the Burrows-Wheeler transform," *IEEE Trans. Parallel Distrib. Syst.*, Jun. 2015. Available: <http://www.computer.org/csdl/trans/td/preprint/07122348-abs.html>
- [25] H. M. Waidyasooriya, M. Hariyama, and M. Kameyama, "FPGA-accelerator for DNA sequence alignment based on an efficient data-dependent memory access scheme," in *Proc. Int. Symp. Highly-Efficient Accelerators Reconfigurable Technol.*, 2014, pp. 127–130.
- [26] E. Fernandez, J. Villarreal, S. Lonardi, and W. Najjar, "FHASt: FPGA-based acceleration of Bowtie in hardware," *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, vol. 12, no. 5, pp. 973–981, Sep. 2015.
- [27] G. Tan, C. Zhang, W. Tang, and N. Sun, "Accelerating irregular computation in massive short reads mapping on FPGA co-processor," *IEEE Trans. Parallel Distrib. Syst.*, Jun. 2015. Available: <http://www.computer.org/csdl/trans/td/preprint/07122363-abs.html>
- [28] J. Watson and F. Crick, "Reprint: Molecular structure of nucleic acids," *Ann. Internal Med.*, vol. 138, no. 7, pp. 581–582, 2003.
- [29] K. Mullis. (2012, Nov.). Polymerase chain reaction [Online]. Available: <http://www.karymullis.com/pcr.shtml>
- [30] S. A. Edwards, "MEMOCODE 2012 hardware/software codesign contest: DNA sequence aligner," in *IEEE/ACM Int. Conf. Formal Methods Models Codes.*, Mar 2012, pp. 85–90.
- [31] (2012, Mar.). 1000 genomes project: Human reference genome g1k v37 [Online]. Available: ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/reference/human%20g1k_v37.fasta.gz
- [32] F. C. Botelho and N. Ziviani, "Near-optimal space perfect hashing algorithms," Ph.D. dissertation, Federal Univ. Minas Gerais, Belo Horizonte, MG, Brazil, 2008.
- [33] D. Belazzougui, F. C. Botelho, and M. Dietzfelbinger, "Hash, displace, and compress," in *Proc. Eur. Symp. Algorithms*, 2009, pp. 682–693.
- [34] M. L. Fredman, J. Komlós, and E. Szemerédi, "Storing a sparse table with $O(1)$ worst case access time," in *Proc. 23rd Annu. Symp. Found. Comput. Sci.*, 1982, pp. 165–169.
- [35] K. Fredriksson and F. Nikitin, "Simple compression code supporting random access and fast string matching," in *Proc. 6th Int. Workshop Exp. Algorithms*, 2007, vol. 4525, pp. 203–216.
- [36] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Trans. Inf. Theory*, vol. 23, no. 3, pp. 337–343, May 1977.
- [37] B. Jenkins. (2012, Sep.). Spookyhash: A 128-bit noncryptographic hash [Online]. Available: <http://burtleburtle.net/bob/hash/spooky.html>
- [38] S. Hanov. (2011, Mar.). Throw away the keys: Easy, minimal perfect hashing [Online]. Available: <http://stevehanov.ca/blog/index.php?id=119>
- [39] T. Brewer, "Instruction set innovations for the Convey HC-1 computer," *IEEE Micro*, vol. 30, no. 2, pp. 70–79, Mar. 2010.
- [40] H. Li. (2012, Nov.). Manual reference pages - bwa [Online]. Available: <http://bio-bwa.sourceforge.net/bwa.shtml>
- [41] (2012, Nov.). Sequence read for individual NA06985 [Online]. Available: ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/data/NA06985/sequence_read/ERR050082.filt.fastq.gz
- [42] O. Knodel, T. Preusser, and R. Spallek, "Next-generation massively parallel short-read mapping on FPGAs," in *Proc. Int. Conf. Appl.-Specific Syst., Archit. Processors*, Sep. 2011, pp. 195–201.
- [43] T. Lam, T. Wong, Y. Li, P. U, and R. Li. SOAP3 alignment time [Online]. Available: <http://www.cs.hku.hk/2bwt-tools/soap3-dp/>
- [44] Xilinx. (2009, Feb.). Virtex 5 family overview [Online]. Available: http://www.xilinx.com/support/documentation/data_sheets/ds100.pdf
- [45] G. Pike and J. Alakuijala. (2011, Apr.). Introducing cityhash [Online]. Available: <http://google-opensource.blogspot.com/2011/04/introducing-cityhash.html>



Chad Nelson received the BS degree in computer engineering from the Iowa State University, in 2010, and the MS degree in computer engineering from the Iowa State University, in 2012. He is currently a Software Subject Matter Expert at the PreTalen. His research interests include CPU/GPU architecture, embedded systems, reconfigurable computing, and hardware/software co-design.



Kevin R. Townsend is currently working toward the PhD degree and joined Iowa State for the same, in 2011. He is the current expert on the Convey Computer (HC-2ex) the RCL group uses. He lead the team that won the 2014 Memocode competition, a one month application acceleration competition. He also played a critical role in winning the 2012 Memocode competition for the Iowa State University. He is currently researching methods to accelerate sparse matrix vector multiplication using FPGAs. After graduation, he will be joining Google. He is a student member of the IEEE.



Osama G. Attia received the BS degree in computer and systems engineering from the Mansoura University, Egypt, in 2010, and the MS degree in communication & information technology from the Nile University, Egypt, in 2012. He is currently working toward the the PhD degree in computer engineering and joined Iowa State University for the same in the Fall of 2012. His research interests include reconfigurable computing, hardware acceleration for graph processing algorithms, embedded systems, and wireless networks. He is a student member of the IEEE.



Phillip H. Jones received the BS degree in 1999 and the MS degree in electrical engineering from the University of Illinois, Urbana-Champaign, in 2002. He received the PhD degree in computer engineering from the Washington University, St. Louis, in 2008. He is currently an assistant professor in the Department of Electrical and Computer Engineering, Iowa State University, Ames, where he has been since 2008. His research interests are in adaptive computing systems, reconfigurable hardware, embedded systems,

and hardware architectures for application-specific acceleration. He is a member of the IEEE.



Joseph Zambreno received the BS degree summa cum laude in computer engineering, in 2001, and the MS degree in electrical and computer engineering, in 2002. Prior to joining the ISU, he was at the Northwestern University, Evanston, IL, where he graduated with the PhD degree in electrical and computer engineering in 2006. He has been with the Department of Electrical and Computer Engineering, Iowa State University since 2006, where he is currently an associate professor. His research interests

include computer architecture, compilers, embedded systems, reconfigurable computing, and hardware/software co-design, with a focus on run-time reconfigurable architectures and compiler techniques for software protection. He is a senior member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**