

Published in final edited form as:

J VLSI Signal Process Syst Signal Image Video Technol. 2007 ; 49(1): 101–121.

Biosequence Similarity Search on the *Mercury* System

Praveen Krishnamurthy, Jeremy Buhler, Roger Chamberlain, Mark Franklin, Kwame Gyang, Arpith Jacob, and Joseph Lancaster

Department of Computer Science and Engineering Washington University in St. Louis

Abstract

Biosequence similarity search is an important application in modern molecular biology. Search algorithms aim to identify sets of sequences whose extensional similarity suggests a common evolutionary origin or function. The most widely used similarity search tool for biosequences is BLAST, a program designed to compare query sequences to a database. Here, we present the design of BLASTN, the version of BLAST that searches DNA sequences, on the *Mercury* system, an architecture that supports high-volume, high-throughput data movement off a data store and into reconfigurable hardware. An important component of application deployment on the *Mercury* system is the functional decomposition of the application onto both the reconfigurable hardware and the traditional processor. Both the *Mercury* BLASTN application design and its performance analysis are described.

1: Introduction

Computational search through large databases of DNA and protein sequence is a fundamental tool of modern molecular biology. Rapid advances in the speed and cost-effectiveness of DNA sequencing have led to an explosion in the rate at which new sequences, including entire mammalian genomes [35], are being generated. To understand the function and evolutionary history of an organism, biologists now seek to identify discrete biologically meaningful features in its genome sequence. A powerful approach to identify such features is *comparative annotation*, in which a *query sequence*, such as new genome, is compared to a large database of known biosequences. Database sequences exhibiting high similarity to the query, as measured by string edit distance [31], are hypothesized to derive from the same ancestral sequence as the query and in many cases to have the same biological function.

BLAST, the **B**asic **L**ocal **A**lignment **S**earch **T**ool [1], is the most widely used software for rapidly comparing a query sequence to a biosequence database. Although BLAST's algorithms are highly optimized for efficient similarity search, growth in the databases it uses is outpacing speed improvements in general-purpose computing hardware. For example, the National Center for Biological Information (NCBI) Genbank database grew exponentially between 1992 and 2003 with a doubling time of 12–16 months [24]. The problem is particularly acute for BLASTN, the BLAST variant used to compare DNA sequences, because each new genome sequenced from animals or higher plants produces between 10^8 and 10^{10} bytes of new DNA sequence.

One response to runaway growth in biosequence databases has been to distribute BLAST searches across multiple computers, each responsible for searching only part of a database. This approach requires both a substantial hardware investment and the ability to coordinate a

search across processors. An alternate approach that makes more parsimonious use of hardware is to build a specialized BLAST accelerator. By using an application-specific architecture and exploiting the high I/O bandwidth of modern storage systems, an accelerator can execute the BLAST algorithms much faster than a general-purpose CPU.

The *Mercury* system [9] is a prototype architecture that supports disk-based computation at very high data rates using reconfigurable hardware. Computing applications historically have been coded using the following paradigm: read input data into main memory with explicit I/O calls, compute on that data writing results back to main memory, and send the output from main memory with explicit I/O calls. In contrast, the Mercury system is built around the concept of continuous data flow. Data from disk(s) flow into the computational resource(s); one or more functions (often physically pipelined) are performed on the data; and the results flow to the intended destination. As the computational resources include reconfigurable hardware, application deployment requires hardware/software codesign. The Mercury system builds upon the work of Reidel [30] (active disks), Dally [11] (stream processors), and a host of work developed in the reconfigurable computing community.

This paper describes the re-engineering of the original BLASTN application for effective deployment on the Mercury system. We examine the existing application to explore its performance properties, propose a novel algorithmic optimization, prototype a number of critical components of the application, and evaluate the performance potential of the overall application running on the Mercury system.

2: System Architecture

The Mercury system (Figure 1) contains reconfigurable logic, associated with the disk controller, that provides computing capability in close proximity to the data flowing off the disk drive(s). Initial processing of the data occurs locally at the disk, prior to delivery to the processor. The reconfigurable logic is implemented via a Field-Programmable Gate Array (FPGA).

Application functionality is divided into two parts executing on the FPGA and the main processor, respectively. Application deployment therefore has the classic components of a hardware/software codesign problem, with the need to map application elements to multiple computational resources (i.e., FPGA and processor). A unique aspect of the Mercury system is that it was designed specifically to work well with high-volume data applications. The computational resource that is best suited to simpler, repetitive operations on a large data set is positioned closer to the data, while the resource best suited to more complex operations on smaller data volumes is (logically) farther away from the data.

The application set that is well matched to the Mercury system architecture is a pipeline that consumes a high data volume at its input, reduces that data volume to a smaller set, and performs higher-level processing on this smaller set. Our previous work has illustrated the use of the system for a number of text search applications [6,7,8,14,36,38]. BLASTN has properties that fit well with the Mercury system's capabilities.

While Figure 1 illustrates our vision of the system architecture, our prototyping work has so far been limited to a series of implementations that are progressively closer to, but do not yet exactly match, the architecture depicted in the figure. Our earliest prototypes used ATA drives [36,38] and were severely speed-limited by the disks. Our most recent prototypes are built using a set of 15,000 rpm Ultra320 SCSI drives organized in a RAID-0 configuration. On this configuration, we have demonstrated sustained read performance of over 800 MB/sec for continuous 500 GB reads. The prototype FPGA infrastructure is currently parallel to the disk

controller on the I/O bus, which limits throughput into the FPGA. We have, however, demonstrated sustained data throughput of over 700 MB/sec from the disk array into the FPGA [7,8].

In what follows, we refer to computations deployed in the FPGA as *firmware* and computations deployed on the processor as *software*. To facilitate the deployment of applications on the FPGA, we have developed a firmware socket interface that provides a consistent environment for the development of firmware application modules. Data from the disk array is delivered to the FPGA via the firmware socket, while outbound data from the reconfigurable logic is delivered into the main memory of the processor for access by software.

The current prototype system uses a Xilinx Virtex-II 6000 series FPGA, which provides 8,448 Configurable Logic Blocks (CLBs), 144 18 Kbit Block RAMs (BRAMs), 144 18×18 bit multipliers, and 1104 I/O pins. Each CLB is comprised of 4 slices and each slice provides two 4 input lookup tables (LUTs) for a total of 67,584 LUTs on chip.

3: Description of NCBI BLASTN

This section describes the open-source version of BLASTN distributed by the National Center for Biological Information (NCBI) and used by numerous biological research labs. As shown in Figure 2, BLASTN is functionally organized as a pipeline with three stages: word matching, ungapped extension, and gapped extension. The inputs to this pipeline are a query sequence and a database, each consisting of a string of DNA *bases*. A base is typically one of {A, C, G, T}, but other characters (a total of 15) are used to denote uncertainty about or special properties of certain bases. DNA sequences, including these special characters, can be represented using four bits per base; however, to minimize storage and I/O bandwidth, NCBI BLASTN stores its database using only two bits per base.

Each stage of BLASTN's pipeline implements progressively more sophisticated and more expensive computations to identify biologically meaningful similarities between query and database. In stage 1, BLASTN discovers *word matches* between query and database. A word match is a string of some fixed length w (hereafter called a " w -mer") that occurs in both query and database. Significantly similar sequences usually share a w -mer match for $w \approx 10$, though such matches also occur frequently by chance between unrelated sequences. Each word match is therefore filtered through stage 2, which tries to extend it into an *ungapped alignment* between query and database. An ungapped alignment may contain mismatched bases but consists primarily of matching base pairs. Ungapped alignments with too few matching base pairs are discarded, while the remainder are further filtered through stage 3, which extends them into *gapped alignments* that permit both mismatches and localized insertion or deletion of bases. In the final operation following the end of stage 3, gapped alignments with sufficiently many matching base pairs are reported to the user. The detailed algorithm for BLASTN stage 2 is described in [21], while that for stage 3 is a variant of the standard Smith-Waterman dynamic programming algorithm [31].

Although each stage of BLASTN is more compute-intensive than the last, each stage also discards a substantial fraction of its inputs. The volume of data that is processed at each stage therefore gradually decreases. Table 1 quantifies the data reduction at each stage of the pipeline¹. The match rate, p_i , represents the probability that an output from stage i is generated from an individual input to that stage. For stage 1, p_1 measures the number of matches per DNA base read from the database. Stages 1 and 2 are highly effective at reducing the data

¹Reduction measurements for NCBI BLASTN were taken in the same experiments used to generate the timings of Section 3.2.

volume to the next stage. Note that, as the query length increases, the rate at which matches are output from stage 1 into stage 2 also increases, raising the workload for stage 2.

In the performance predictions that follow, we will consider the throughput of individual stages of the pipeline as well as the throughput of the entire pipeline. To make throughputs comparable, they are normalized to be in units of input bases per second from the database. When executing on a single computational resource (i.e., software running on a single processor), the average compute time per input base can be expressed as $t_1 + p_1t_2 + p_1p_2t_3$, where t_i is the compute time for stage i for each input item (base, match, or alignment) to stage i . The normalized throughput is then $T_{put} = 1/(t_1 + p_1t_2 + p_1p_2t_3)$.

3.1: Details of BLASTN Stage 1

To facilitate later comparison with our firmware design, we now briefly describe the implementation of NCBI BLASTN's stage 1. This implementation uses a default word match length $w = 11$. Due to the speed advantages of comparing complete bytes at a time, discovery of 11-mer matches is implemented in two phases. BLASTN first checks two complete bytes of the database, containing 8 bases, against a lookup table constructed from the query. Only two-byte words occurring on full byte boundaries are checked. If the query contains the same 8-base word, BLASTN tries to extend this 8-base match to 11 bases by seeking additional matching residues on either side.

Two 11-mer matches that occur close to each other in both the query and database are likely to have arisen from the same underlying biological similarity. To avoid having later stages expend the effort to discover this similarity twice, NCBI BLASTN implements a *redundancy elimination* filter at the end of stage 2. The filter checks whether each new 11-mer match overlaps or is close to a previously observed match. If so, the new match is suppressed, since it would likely lead only to rediscovery of any feature found by the previous match.

3.2: Performance of NCBI BLASTN

To quantify the performance of NCBI BLASTN on a general-purpose CPU, we measured its execution time with default parameters on a 2.8 GHz Pentium 4 PC, with an L2 cache size of 512 KB and 1 GB of RAM, running Linux. We compared a database containing the mouse genome (1.16 Gbases after removing repetitive sequence) to queries of various lengths selected at random from the human genome. CPU time was measured separately for each of the three pipeline stages.

The length of a typical query sequence in BLASTN is application-dependent. For example, a short DNA sequence obtained in a single lab experiment may be only a few kilobases, while in genome-to-genome comparison, a query (one of the genomes) may be billions of bases long. A BLAST implementation should support the largest computationally feasible query length, both to accommodate long individual queries and to support the optimization of "query packing," in which multiple short queries are concatenated and processed in a single pass over the database. Conversely, queries longer than the maximum feasible length may be broken into pieces, each of which is processed in a separate pass.

In our experiments, we tested queries of 10 kbases, 25 kbases, 50 kbases, 100 kbases, and 1 Mbase, both to simulate different applications of BLASTN and to assess the impact of query length on the performance of our firmware implementation. One megabase is a reasonable upper bound on query size for NCBI BLASTN with standard parameters, since it generates 11-mer word matches by chance alone at a rate approaching one match for every base read from the database. Timings were averaged over at least 20 queries for each length, and each

query's running time was averaged over three identical runs of BLASTN. It should be noted that, given a query sequence of length n , BLASTN compares the database to both the sequence and its DNA reverse-complement, effectively doubling the query length. The performance numbers reported in this section and throughout the rest of the paper reflect such "double-stranded" queries.

Table 2 gives the distribution of times spent in each stage of NCBI BLASTN for various query sizes. Times are given with 95% confidence intervals. Time spent in stage 1 dominated that spent in later pipeline stages, while time spent in stage 3 was almost negligible. Although later stages are computationally more intensive, each stage is such an efficient filter that it discards most of its input, leaving later stages with comparatively little work.

From the measured running times of our experiments and the size of the mouse genome database, we computed the throughput (in Mbases from the database per second) achieved by NCBI BLASTN's pipeline for varying query sizes. The results are shown in the first row of Table 3. Throughput depends strongly on query length. To explain this observation, we used the predicted filtering efficiencies p_i for each pipeline stage and the distribution of running times by stage to estimate the average time spent to process each base in stage 1, each word match in stage 2, and each ungapped alignment in stage 3. These results are shown in the remaining rows of the table. While the overhead per input remains constant for stage 2 and actually decreases for stage 3, the cost per base in stage 1 grows linearly with query length. This cost growth derives from the linear increase in the expected number of matches per base that occur purely by chance, in the absence of any meaningful similarity.

The empirical performance of NCBI BLASTN's pipeline demonstrates that stage 1 is a performance bottleneck and therefore the first target for speedup in firmware.

4: Firmware Implementation of Stage 1

Our firmware implementation of stage 1 reflects the overall functionality of stage 1 in NCBI BLASTN but makes no attempt to implement this functionality using the same mechanisms. Our design decomposes stage 1 into 3 substages (Figure 3). The initial substage implements a prefilter using Bloom filters; the middle substage determines the query position of w -mers in the database that successfully pass through the Bloom filters (using hashing); and the final substage performs redundancy elimination.

4.1: Prefiltering using Bloom Filters

A Bloom filter [2] is a probabilistic algorithm to quickly test membership in a large set using multiple hash functions into a single array of bits. Bloom filters find many uses in networking and other applications [12]. Figure 4 illustrates a typical Bloom filter datapath. Programming the filter amounts to setting to '1' each of the bits of the memory locations obtained by the hash functions. Querying the Bloom filter yields a match when all the memory locations in the vector obtained from hashing the query contain '1'.

A Bloom filter yields no false negatives but does yield false positives at a rate f determined by the number of w -mers programmed into it and the length of its memory vector. The rate f can be modeled as $f = (1 - e^{-Nk/m})^k$, where N is the number of entries programmed into the filter (query size), m is the filter memory size in bits, and k is the number of hash functions. Figure 5 shows the false positive rate of a Bloom filter, as a function of memory size, for different query lengths. The number of hash functions, k , in Figure 5 is obtained as $k = \frac{m}{N} \ln(2)$. The false positive rate is obtained as $f = \frac{1}{2^k}$.

Bloom filters are more efficiently implemented in firmware than in software, as we can store the memory vectors on chip (using sets of block RAMs), calculate the hash functions in parallel, and look up the locations of the memory vector in parallel. However, as the number of ports on these block RAMs is finite, the hash functions are restricted to address only specific block RAMs.

4.2 Architecture of Bloom filters

In our implementation, each distinct memory vector made using a set of block RAMs is addressed using a unique hash function (see Figure 6, in which each ' m '-bit memory is constructed out of a set of block RAMs). The probability that a random bit in an ' m '-bit memory is set to 1 by a hash function is simply $\frac{1}{m}$. The probability that it is not set is therefore $1 - \frac{1}{m}$. Now since all N queries are programmed in each of the memories, the probability that a random bit is not set by any of the N queries is $(1 - \frac{1}{m})^N$. Hence, the probability that it is set is $(1 - (1 - \frac{1}{m})^N)$. A random query will trigger a false positive if it randomly collides with set bits in all the memories. In our design, the number of memories equals the number of hash functions k . Hence, the probability that a query triggers a match is obtained as $f = (1 - (1 - \frac{1}{m})^N)^k$, where m can also be interpreted as the address range of each hash function, N is the length of the query, and k is the number of hash functions used.

Because the Bloom filter implementation is the primary block RAM-intensive stage in our design, we dedicate 96 (about two-thirds) of the 144 block RAMs available on chip to it. The stage is designed to consume 16 bases every clock cycle and to operate at 133 MHz (yielding a potential throughput of 2 Gbases/s). To sustain this rate, we must process 16 w -mers every clock cycle and so require 16 identical copies of the Bloom filter.

We reduce the memory requirements for the Bloom filtering stage in half by using both of the ports of our dual-ported block RAMs. Furthermore, using the clock management circuitry on the FPGA, we double-clock the block RAMs and in effect make each block RAM quad-ported. This is illustrated in Figure 7. Four copies of the 4-way parallel Bloom filters are therefore sufficient to process all 16 w -mers (illustrated in Figure 8).

Reducing the number of Bloom filter copies needed helps to decrease the false positive rate of the queries, by dedicating more memory to each Bloom filter, or to process larger queries with the same false positive rate. These tradeoffs are illustrated in Figure 9, which shows the expected true match rate (solid line) and the overall match rate (including false positives) of stage 1a (dashed lines). This figure assumes that the input to stage 1a never stalls (i.e., 16 bases are available each clock cycle). The two double-clocking curves vary in how they utilize the additional effective memory. One doubles k , while the other doubles m . Doubling the clock rate of the block RAMs not only decreases the false positive rate for smaller query sizes but also helps us support larger queries than would otherwise be possible.

The maximum query size that can be supported on our prototype is partially determined by the rate of matches from stage 1a. Stage 1b (described below in Section 4.3) processes matches from this stage, and it is designed to support an input rate of approximately one match every clock cycle. Given an expected average input rate from the disk subsystem of 1.4 Gbases/s (700 MB/s \times 2 bases/byte) and a maximum ingest rate into stage 1a of 2 Gbases/s, 25 kbase double-stranded queries are reasonably supported.

As similarities can exist between the query and database sequences, there is a good chance that matches from stage 1a will be bursty. We maintain an on-chip queue of size 1000 w -mers to accommodate such bursts. In the event that this queue fills up, for example in the case where

long genes are conserved between sequences, we can either store the matches from stage 1a in off-chip DRAM or generate backpressure for the upstream stages. The performance implications of this backpressure are minimal, as assessed below.

4.3: Hash Lookup

The second substage of stage 1 uses a hash table to identify those w -mers from the database that actually occur in the query sequence. Each such w -mer must be mapped to its position or positions in the query. Note that, in contrast to NCBI BLASTN, we do not test only those w -mers falling on byte boundaries – every w -mer in the database is checked. The hash table is implemented in an external SRAM attached to the FPGA, since the latter's internal block RAMs are too limited in size to contain tables built from large query sequences.

The need to access a single external SRAM is a potential source of pipeline bottlenecks. Suppose that the SRAM can sustain only one read per clock cycle, which is a reasonable assumption at high FPGA clock speeds. The data reduction achieved by our Bloom filters is sufficient to ensure an average input rate to this substage of at most one w -mer per clock. However, if processing a w -mer were usually to entail multiple, serial accesses to the SRAM (e.g. to resolve hash collisions), we could not sustain even this modest data rate. Our design for this stage therefore seeks to dispose of the vast majority of w -mers with only a single SRAM lookup.

4.3.1: Near-Perfect Hashing—We have developed an FPGA-friendly approach to hash table design, called *near-perfect hashing*, that empirically approximates the properties of a perfect hash for our application. A perfect hash [32] for a set $k_1 \dots k_n$ of keys maps each key to a distinct table entry, with no collisions between keys in the set. Assuming we could find such a perfect hash function h of the w -mers in our query sequence, checking whether the query contains a given w -mer s would entail retrieving only a single entry $h(s)$ from SRAM; if the entry is empty or contains a key different from s , we know that s is not in the table.

Several methods exist [10,15,33] to generate perfect hash functions for applications (like ours) in which the table size is limited and so cannot be much larger than the number of keys. Unfortunately, to guarantee that a perfect hash exists for n keys, these approaches produce hash functions whose representation takes $\Theta(n \log n)$ bits. When n is on the order of tens of thousands, the number of on-chip bits required to compute hashes approaches a megabit, which would consume a significant fraction of our FPGA's capacity and would compete for space with the Bloom filters of stage 1a. We therefore sought a more area-efficient approach that nevertheless gives hashes with very few collisions.

Like many perfect hash generators, our approach uses the idea of *displacement hashing* [33]. To create a mapping from a set C of w -mers, each consisting of $2w$ bits, into a table of size 2^a for some fixed a , we construct a hash function h of the form

$$h(s) = A(s) \oplus \tau[B(s)],$$

where $A : \{0, 1\}^{2w} \rightarrow \{0, 1\}^a$ and $B : \{0, 1\}^{2w} \rightarrow \{0, 1\}^b$ are easily-computed functions of the key, \oplus is the bitwise XOR operation, and τ is a *displacement table* of 2^b small integers. The functions A and B are chosen so that the pairs $(A(s), B(s))$ are distinct for all keys $s \in C$. The displacement table can then be exploited to resolve collisions as follows: if w -mers s_1 and s_2 have $A(s_1) = A(s_2)$ but $B(s_1) \neq B(s_2)$, try to choose distinct values for $\tau[B(s_1)]$ and $\tau[B(s_2)]$ to ensure that $h(s_1) \neq h(s_2)$.

The hash functions A and B in our construction are chosen from the H_3 family described in [5,29], which is efficiently computable in hardware. An H_3 function g from p to q bits consists of q p -bit vectors $v_1 \dots v_p$. Let M be the $q \times p$ Boolean matrix whose rows are the v_i 's. Then for a p -bit input vector s , $g(s)$ is the vector-matrix product $s \cdot M$ over the mathematical field $(Z_2, \text{XOR}, \text{AND})$.

A novel aspect of our approach is that we choose M to be a random linear transformation of full rank; that is, the q rows of these matrices are linearly independent over Z_2^p . It can be shown that, if H_3 hash function M has rank q , then it maps exactly 2^{p-q} p -bit inputs to each of its 2^q possible outputs. Moreover, if we choose random A and B of full rank (a and b respectively), and $a + b \geq 2w$, then A and B have probability at least $1/4$ of mapping every $s \in C$ to a distinct pair $(A(s), B(s))$, as is required by our construction. Hence, a small number of trials is sufficient to find such A and B . Finally, we note that a random $p \times q$ Boolean matrix, $p \geq q$, with no zero rows has probability at least $1/4$ of being full rank, and we can test this for property efficiently using Gaussian elimination; hence, full rank H_3 functions are easy to generate. Proofs of the above properties may be found in [3].

Given suitable A and B , we are left with the problem of filling in the displacement table τ so as to resolve the collisions induced by the A hash alone. Unlike the abovementioned perfect hash generators, our τ has many fewer than 2^a entries, and these entries use fewer than a bits. For example, when $|C| \approx 50000$ (a 25 kbase query), we use a τ with 2048 entries of 8 bits each and so consume only 16 Kbits on chip. Such a τ is too small to guarantee existence of a perfect hash according to the analysis of, e.g., [33], so we settle for seeking values for the entries of τ that minimize the number collisions. Our scheme may not produce a perfect hash, whence comes the name “near-perfect.”

The following efficient greedy heuristic chooses values for the displacement table τ that empirically achieve few, and often no, collisions in our application. First, sort the 2^b values of B in decreasing order by the number of keys mapping to each. Let $\{\beta_1 \dots \beta_{2^b}\}$ be the resulting sequence of B -values. For each i from 1 to 2^b , greedily assign a value to $\tau[\beta_i]$ so as to minimize the number of collisions with all keys with B -values $\beta_1 \dots \beta_{i-1}$. Once τ has been filled in, we may heuristically improve it by trying to modify each entry $\tau[\beta_i]$ in turn to minimize collisions given the other values in τ . This cycle is repeated a fixed number of times, as computational costs permit.

To support 25 kbase queries with $w = 11$, we set $a = 18$, $b = 11$, and permit each entry of τ to be up to 8 bits. We therefore map w -mers to a table of 2^{18} entries in SRAM, using a modest 2^{14} bits of space on chip for τ . On a sample of thirty 25 kbase DNA query sequences drawn at random from the human genome, our heuristic almost always generated perfect hash functions, only occasionally producing one or two collisions. For comparison purposes, using the A hash alone with no displacement table produced an average of 7718 collisions over the same queries. On a 2 GHz AMD Opteron workstation, the time to generate h for these queries and construct an image of the resulting hash table was under 0.6 seconds. Additional empirical validation of near-perfect hashing may be found in [3]; theoretical performance guarantees for this strategy remain a topic for future work.

4.3.2: Hash Table Structure—The datapath of our hash table (stage 1b) is illustrated in Figure 10, and the control logic that manages this datapath is shown in Figure 11. Each 32-bit entry in our hash table stores a 16-bit query position and a *collision bit*, which is set if two or more distinct w -mers from the query map to that entry. Collisions occur if our hash function design does not yield a perfect hash. Unoccupied entries are marked by a special value in the query position field.

If a w -mer s maps to a table entry $h_1(s)$ with the collision bit unset, lookup proceeds as for a perfect hash, using only a single SRAM access. In case of a collision, the implementation performs a second check in a smaller *secondary* hash table (2^{16} entries in our implementation) similar to the primary table, which contains only those query w -mers that collide in the first table. The hash function used for the secondary table, h_2 , is chosen from the H_3 family, in similar fashion to A and B above. Because the primary hash function h_1 is near-perfect, the number of entries with collision bits set is small. Hence, accesses to the secondary table are rare, preserving our desired property of at most one SRAM lookup for most w -mers.

Query w -mers s_1 and s_2 that collide in both the primary and secondary tables are merged; that is, they are treated for hash lookup purposes as if they are two copies of the same w -mer. The duplicate facility described below ensures that probes of either s_1 or s_2 return all query positions at which either one occurs. Merging therefore creates false positive matches, which are removed by later pipeline stages, but does not create false negatives. Merging requires that two w -mers collide in both primary and secondary tables, which use independent hash functions, so it is extremely rare.

4.3.3: Resolving Duplicate Keys—When a w -mer occurs at two or more positions in the query, we wish to return not just one such position but all of them. We therefore extend each hash entry with a *duplicate flag* and a pointer into a separate *duplicate table*, also in SRAM. Any w -mer s with more than one occurrence has its duplicate flag set. One occurrence of s is stored in the table entry itself, while the remainder are stored in a contiguous range of 16-bit cells in the duplicate table. The first cell in the range, which is pointed to by the hash table entry, specifies the number of additional occurrences, while each additional cell gives the position of one occurrence.

We align the beginning of each range of cells in the duplicate table on a 32-bit boundary. Hence, a 14-bit pointer in the hash table entry is sufficient to support a table of 2^{15} 16-bit cells – enough for even a fairly degenerate 25 kbase query.

Accesses to the duplicate table compete with the hash table for SRAM bandwidth, so they increase the average number of probes per w -mer. However, a typical repeat-masked 25 kbase DNA sequence from the human genome has only a few hundred duplicate 11-mers, so the additional accesses to the table add only slightly to the overall load on SRAM.

There are two conditions under which a false positive result can come from stage 1b. The first, described earlier, is if there is a collision in both the primary and secondary hash tables. The second is when a false positive result from the Bloom filters collides in the primary hash table with a w -mer in the query. Our current implementation passes these false positives to later stages (where they are discovered and discarded). An alternative approach would be to perform an explicit test, comparing the w -mer in the database with the w -mer in the query. This would completely eliminate any false positives from stage 1. However, the additional load caused by false positives on later pipeline stages is small enough that it does not significantly impact performance, so removing them is not worth the additional design complexity it would entail.

4.4: Redundancy Filter

To avoid repeated generation of the same sequence alignment, NCBI BLASTN uses a redundancy filter to discard the w -mers which fall within the range that already has been inspected by extension in stage 2. This range can be defined as two w -mers which have an overlapping diagonal. Each w -mer is represented by an ordered pair (q_j, d_k) , where q_j and d_k are indices into the query and database, respectively. Define the *diagonal* of this w -mer to be $D = d_k - q_j$. To keep track of the redundant w -mers, a record for every diagonal is stored in a

data structure which contains the previous non-redundant database index. If d_k of the current w -mer is greater than the stored value, the w -mer is queued for further inspection. In NCBI BLASTN, the redundancy filter is updated after extension in stage 2. This method can eliminate more than strictly overlapping w -mers without decreasing sensitivity.

We use a slightly different heuristic which does not require feedback from stage 2 [28]. Instead of storing the highest endpoint of the database from the stage 2 extension, Mercury BLASTN implements the redundancy filter by finding all the truly redundant w -mers as well as the ones close to them. The end of the first non-overlapping w -mer out of stage 1b is stored for each diagonal and is updated every time a new non-overlapping w -mer is seen on that diagonal. The redundancy filter is also parameterizable to discard any w -mers on a diagonal that are within N bases of a previous diagonal, even if they do not overlap. Setting this parameter allows for a trade-off between filter strength and sensitivity.

Figure 12 illustrates an example query and database sequence plotted to form a matrix using a word length of 4, and a tail region of length $N = 2$. Non-redundant 4-mer matches are shown in bold, while redundant 4-mers are shown normal weight. The filter keeps a record of the first 4-mer in each diagonal and passes them to the later stages. 4-mers that are truly overlapping a previous 4-mer on that diagonal are discarded. This is shown on the figure as “True Overlap.” An example of a 4-mer that does not overlap the previous entry on that diagonal but is within the tail region (and therefore discarded) is shown in Figure 12 labeled “Tail Overlap.”

While the Mercury BLASTN redundancy filtering method is a less stringent filter than NCBI BLASTN’s heuristic, it does not require feedback from stage 2. We compared the two approaches empirically on comparisons of human and mouse genomic DNA and found that the simpler heuristic does not increase the number of matches emerging from stage 1 enough to create a performance bottleneck.

5: Performance Analysis

We assess the performance of our design in three phases. First, we assess the performance gain relative to software of stage 1 alone. Second, we assess the overall performance of a design that exploits the firmware implementation of stage 1 and continues to use software to implement stages 2 and 3. Finally, we discuss the benefit that can be gained from a (future) firmware implementation of stage 2 and provide performance targets for that design.

5.1: Word Matching (Stage 1) in Firmware

We have built several firmware prototypes of stage 1 that can consume 16 bases/clock. The latest runs at a clock rate of 133 MHz. The prototype is limited to a query size of 25 kbases. To process queries of size greater than 25 kbases, we pass the database through the firmware stage multiple times (say r), each pass consuming a fraction of the query ≤ 25 kbases. This results in an effective throughput of $\frac{1}{r}$ that of 25 kbase queries for larger queries. Note that this bound on the query length is due primarily to the number of available block RAMs on our current chip and is not a fundamental limit of the design itself.

To assist in the performance analysis task, we developed a detailed simulator that provides a cycle-accurate model of the stage 1 design. This simulation model is used to validate the analytical models used to generate Figure 9, and our assumption that we can process a match from stage 1a every clock cycle in stage 1b. We chose 3 of the design points from Figure 9 and executed simulations using a minimum of 30 different queries for each configuration. For each of these simulations we assumed that the input rate into stage 1 is 2.128 Gbases/s, which is its maximum ingest rate. Table 4 compares the results we obtain from these simulations to the

analytical predictions for a number of different parameters. We observe that the simulation performance predictions line up extremely well to the predictions using the analytic models.

While the mean time to process a match in stage 1b is slightly higher than the predicted value of 1, the input rate to this stage is such that the utilization remains < 100%. Also, the nature of genomic sequences is such that the assumption of uniform distribution of bases in sequences turns out to be pessimistic. The observed mean throughput of 2.128 Gbases/s (with extremely small standard deviations) builds further confidence in our analytic model results.

The raw throughput supported by our stage 1 implementation is about 2 Gbases/sec. Using the Mercury system infrastructure, which currently provides 700 MB/sec of input bandwidth, we can expect a data rate of 1.4 Gbases/sec into stage 1. Note that we use 4 bits per base, thereby eliminating potentially significant post-processing of masked sequences arising from NCBI BLASTN's use of 2 bits/base. Table 5 compares the performance of stage 1 in firmware to the software BLASTN.

As shown in Figure 8, a queue is present in the design to smooth the bursty nature of matches that are generated in stage 1a. The detailed simulation model described above was used to assess the usage of this queue, and also to assess the performance impact of bursty matches. Figures 13 through 15 plot the maximum queue length for 32 executions each of 3 system configurations. The 3 configurations shown are the same as those used for validation purposes above. Across the board, the maximum queue lengths are reasonable, with an overall maximum under 1600 matches. Note that should the queue fill, there is appropriate backpressure built into the design so that correctness is maintained. At issue is simply the performance impact that a full queue would have on overall throughput.

The bursty nature of the stage 1a match process has a negligible impact on the overall throughput performance. Observing the throughput predictions from simulation (see Table 4), the performance is quite stable (i.e., the standard deviations are quite low) and is not impacted significantly by the data variability.

5.2: Overall Performance of BLASTN on the Mercury System

We now consider *overall* pipeline performance. While the component pipeline stages above have been constructed individually and verified to work together, the complete system has yet to be integrated at full speed (we are currently undertaking this activity). As a result, the performance numbers that follow are model-based predictions.

When executing the application across multiple resources, the overall throughput is determined by the minimum throughput achieved on any one resource. Here, stage 1 executes in firmware, while stages 2 and 3 execute in software. The throughput is therefore

$$T_{put_{overall}} = \min \left(T_{put_1}, \frac{1}{p_1(t_2 + p_2 t_3)} \right),$$

where T_{put_1} , stage 1 throughput, is from Table 5; t_i , the time to perform stage i in software, is from Table 3; and p_i , the probability of an output from stage i , is from Table 1.

Table 6 compares the overall performance of Mercury BLASTN with that of NCBI BLASTN. Though we have shown significant speedup for stage 1 in firmware (refer to Table 5), the overall speedup is limited to a factor of 5 to 8. Overall performance is now limited by the

software-based stage 2. Hence, though we successfully deployed stage 1 in firmware with high throughput, the overall application still suffers from limitations imposed by the remaining pipeline stages.

If t_3 is the software compute time per match from stage 3 (from Table 3), the maximum pipeline throughput that can be sustained by stage 3 is $T_{put_3} = 1/p_1 p_2 t_3$ (as above, normalized to be in units of input bases per second from the database). For 1 Mbase query sizes, this rate is approximately 2 Gbases/sec, which matches the input rate supported by the firmware stage 1. Hence, stage 3 is unlikely to be a bottleneck to overall performance.

We next consider the overall performance impact of accelerating stage 2. This impact can be modeled as $T_{put_{overall}} = \min(T_{put_1}, T_{put_2}, T_{put_3})$, where T_{put_1} and T_{put_3} are as above and T_{put_2} is now $S_2 = p_1 t_2$. S_2 is a model input representing the speedup of a hypothetical firmware stage 2 implementation. This model determines the performance required of the stage 2 firmware in order to achieve a given overall pipeline throughput.

Figure 16 plots the throughput of the overall application, as a function of the stage 2 speedup S_2 , for various query sizes. By increasing the performance of the bottleneck stage 2, overall performance improves until the throughput reaches the limit imposed by stage 1, at which point it saturates.

Figure 17 plots the speedup, relative to a pure software implementation, of the entire application as a function of stage 2 speedup, again for various query sizes. If, as seems likely, we can achieve even modest speedup in a firmware stage 2, we predict that overall performance of Mercury vs. NCBI BLASTN will improve by two orders of magnitude.

6: Related Work

Biosequence similarity search is a fundamental task of modern biology. Several research groups have therefore implemented systems to accelerate similarity search, both in software alone and in hardware.

Several software tools exist that seek to accelerate BLASTN-like computations through algorithmic improvements. MegaBLAST [39] is used by NCBI as a faster alternative to BLASTN; it explicitly sacrifices substantial sensitivity relative to BLASTN in exchange for improved running time. The SSAHA [25] and BLAT [18] packages achieve higher throughput than BLASTN by requiring that the entire database be indexed offline before being used for searches. By eliminating the need to scan the database, these tools can achieve more than an order of magnitude speedup versus BLASTN; however, they must trade off between sensitivity and space for their indices and so in practice are less sensitive. In contrast, Mercury BLASTN aims for at least BLASTN-equivalent sensitivity.

Other software, such as DASH [19] and PatternHunter II [23], achieves both faster search and higher sensitivities compared to BLASTN using alternative forms of pattern matching and dynamic programming extension. DASH's reported speedup over BLASTN is less than 10-fold for queries of 1500 bases, and it is not clear how it performs at our query sizes, which are an order of magnitude larger. DASH's authors have also reported on a preliminary FPGA design for their algorithm [20]. PatternHunter II achieves a reported two-fold speedup relative to BLASTN, with substantially greater sensitivity, through judicious modification of its pattern-matching stage. We plan to implement similar improvements, based on our own studies of BLASTN-like pattern matching [4], in a future version of our system.

In hardware, numerous implementations of the Smith-Waterman dynamic programming algorithm have been reported in the literature, using both non-reconfigurable ASIC logic [13, 16] and reconfigurable logic [17,26,37]. These implementations focus on accelerating gapped alignment, which is heavily loaded in proteomic BLAST comparisons but takes only a small fraction of running time in genomic BLASTN computations. Our work instead focuses on accelerating the bottleneck stages of the BLASTN pipeline, which reduces the data sent to later stages to the point that Smith-Waterman acceleration is not necessary.

While one could in principle dispense with the pattern matching and ungapped extension stages of BLASTN given a sufficiently fast Smith-Waterman implementation, no such implementation is likely to be feasible with current hardware. The projected data rate of 1.4 Gbases/s shown for a 25 kbase query in Figure 16, if achieved by a Smith-Waterman implementation, would imply computation of around 10^{14} dynamic programming matrix cells per second. In contrast, existing FPGA implementations report rates of less than 10^{10} cells per second.

High-end commercial systems have been developed to accelerate or replace BLAST [27,34]. The Paracel GeneMatcher™ [27] relies on non-reconfigurable ASIC logic, which is inflexible in its application and cannot easily be updated to exploit technology improvements. In contrast, FPGA-based systems can be reprogrammed to tackle diverse applications and can be redeployed on newer, faster FPGAs with minimal additional design work. RDisk [22] is one such FPGA-based approach, which claims a 60 Mbases/sec throughput for stage 1 of BLAST using a single disk.

Two commercial products that do not rely on ASIC technology are BLASTMachine2™ from Paracel [27] and DeCypherBLAST™ from TimeLogic [34]. The highest-end 32-CPU Linux cluster BLASTMachine2™ performs BLASTN with a throughput of 2.93 Mbases/sec for a 2.8 Mbase query. Mercury BLASTN with only stage 1 implemented in firmware has a predicted throughput of 4.44 Mbases/sec for a 1 Mbase query. Hence, BLASTMachine2™ (with 32 nodes) has roughly twice the throughput of Mercury BLASTN (with 1 node).

The DeCypherBLAST™ solution uses an FPGA-based approach to improve the performance of BLASTN. This solution has throughput rate of 213 kbases/sec for a 16-Mbase query, which is comparable to that of Mercury BLASTN with only stage 1 in firmware, processing a query length of 1 Mbase.

7: Conclusions and Future Work

This paper presents the design of BLASTN, an important biosequence search application, for the Mercury system, an architecture that provides both FPGA and traditional processor computing resources and is optimized for disk-based, data-intensive applications. We constructed prototype application components for a firmware (FPGA-based) stage 1 of the BLASTN pipeline, including the addition of a Bloom filter-based prefilter, a firmware hash table, and a match redundancy eliminator.

We compared the performance of our firmware stage 1 implementation to that of NCBI BLASTN's software stage 1 implementation. We also estimated overall performance of Mercury BLASTN, both for the current version with only stage 1 in firmware and for a future version that will also deploy a firmware stage 2.

Because of the strong predicted impact of stage 2 speedups on overall application performance, we have proceeded with a firmware implementation of stage 2, which is reported in [21].

Currently, we are completing end-to-end deployment of BLASTN on the Mercury prototype. Each of the component pieces is now in place, functionally correct and running at speed on the reconfigurable hardware, and integration work is ongoing.

Acknowledgments

This work was supported by NSF Career grant DBI-0237902, NSF grants CNS-0313203, CCF-0427794 and CCR-0217334, and NIH/NGHRI grant 1 R42 HG003225-01.

References

1. Altschul SF, Madden TL, Schäffer AA, Zhang J, Zhang Z, Miller W, Lipman DJ. Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Research* 1997;25:3389–402. [PubMed: 9254694]
2. Bloom B. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM* May;1970 13(7):422–426.
3. Buhler, J. Technical report. Washington University; St. Louis: 2006. Mercury BLAST dictionaries: analysis and performance measurement.
4. Buhler J, Keich U, Sun Y. Designing seeds for similarity search in genomic DNA. *Journal of Computing and Systems Science* 2005;70:342–363.
5. Carter L, Wegman M. Universal classes of hashing functions. *Journal of Computer and System Sciences* 1979;18:143–154.
6. Chamberlain, R.; Cytron, R. Novel techniques for processing unstructured data sets. *Proc. of IEEE Aerospace Conf.*; March 2005;
7. Chamberlain, R.; Shands, B. Streaming data from disk store to application. *Proc. of 3rd Int'l Workshop on Storage Network Architecture and Parallel I/Os*; September 2005; p. 17-23.
8. Chamberlain, R.; Shands, B.; White, J. Achieving real data throughput for an FPGA co-processor on commodity server platforms. *Proc. of 1st Workshop on Building Block Engine Architectures for Computers and Networks*; October 2004;
9. Chamberlain, RD.; Cytron, RK.; Franklin, MA.; Indeck, RS. The *Mercury* system: Exploiting truly fast hardware for data search. *Proc. of Int'l Workshop on Storage Network Architecture and Parallel I/Os*; September 2003; p. 65-72.
10. Czech ZJ, Havas G, Majewski BS. Perfect hashing. *Theoretical Computer Science* 1997;182:1–143.
11. Dally, WJ., et al. Merrimac: Supercomputing with streams. *Proc. of Supercomputing Conf.*; November 2003;
12. Dharmapurikar S, Krishnamurthy P, Sproull T, Lockwood J. Deep packet inspection using parallel bloom filters. *IEEE Micro* 2004;24(1):52–61.
13. Singh, RK., et al. BioSCAN: a dynamically reconfigurable systolic array for biosequence analysis. *Proceedings of CERCS*; 1996.
14. Franklin, M.; Chamberlain, R.; Henrichs, M.; Shands, B.; White, J. An architecture for fast processing of large unstructured data sets. *Proc. of 22nd Int'l Conf. on Computer Design*; October 2004; p. 280-287.
15. Hagerup T, Miltersen PB, Pagh R. Deterministic dictionaries. *Journal of Algorithms* 2001;41:69–85.
16. Hirschberg, JD.; Hughley, R.; Karplus, K. Kestrel: a programmable array for sequence analysis. *Proceedings of IEEE International Conference on Application-Specific Systems, Architecture, and Processors*; 1996; p. 23-34.
17. Hoang, DT. Searching genetic databases on Splash 2. *IEEE Workshop on FPGAs for Custom Computing Machines*; 1993; p. 185-191.
18. Kent WJ. BLAT: the BLAST-like alignment tool. *Genome Research* 2002;12:656–64. [PubMed: 11932250]
19. Knowles, G.; Gardner-Stephen, P. DASH: localizing dynamic programming for order of magnitude faster, accurate sequence alignment. *Proceedings of the 3rd International IEEE Computer Society Computational Systems Bioinformatics Conference*; 2004; p. 732-35.

20. Knowles, G.; Gardner-Stephen, P. A new hardware architecture for genomic and proteomic sequence alignment. *Proc. of IEEE Computational Systems Bioinformatics Conf.*; 2004;
21. Lancaster, J.; Buhler, J.; Chamberlain, RD. Acceleration of ungapped extension in Mercury BLAST. *Proc. of 7th Workshop on Media and Streaming Processors*; November 2005;
22. Lavenier, D.; Guytant, S.; Derrien, S.; Rubin, S. A reconfigurable parallel disk system for filtering genomic banks. *ERSA'03, Engineering of Reconfigurable Systems and Algorithms*; 2003;
23. Li M, Ma B, Kisman D, Tromp J. Patternhunter II: highly sensitive and fast homology search. *Journal of Bioinformatics and Computational Biology* 2004;2:417–39.
24. National Center for Biological Information. Growth of GenBank. 2002.
<http://www.ncbi.nlm.nih.gov/Genbank/genbankstats.html>
25. Ning Z, Cox AJ, Mullikin JC. SSAHA: a fast search method for large DNA databases. *Genome Research* 2001;11:1725–9. [PubMed: 11591649]
26. Pappas, N. Master's thesis. Virginia Polytechnic Institute and State University; 2003. Searching biological sequence databases using distributed adaptive computing.
27. Paracel, Inc. <http://www.paracel.com>
28. Pevzner PA, Waterman MS. Multiple filtration and approximate pattern matching. *Algorithm-mica* 1995;13(12):135–154.
29. Ramakrishna MV, Fu E, Bahcekapili E. Efficient hardware hashing functions for high performance computers. *IEEE Transactions on Computers* 1997;46:1378–1381.
30. Reidel E, Faloutsos C, Gibson G, Nagle D. Active disks for large-scale data processing. *IEEE Computer June*;2001 34(6):68–74.
31. Smith TF, Waterman MS. Identification of common molecular subsequences. *Journal of Molecular Biology March*;1981 147(1):195–97. [PubMed: 7265238]
32. Sprugnoli R. Perfect hashing functions: a single probe retrieving method for static sets. *Communications of the ACM* 1977;20(11):841–850.
33. Tarjan RE, Yao ACC. Storing a sparse table. *Communications of the ACM* 1979;22(11):606–611.
34. TimeLogic Corporation. <http://www.timelogic.com>
35. Waterston RH, et al. Initial sequencing and comparative analysis of the mouse genome. *Nature* 2002;420:520–562. [PubMed: 12466850]
36. West, B.; Chamberlain, RD.; Indeck, RS.; Zhang, Q. An FPGA-based search engine for unstructured database. *Proc. of 2nd Workshop on Application Specific Processors*; December 2003; p. 25-32.
37. Yamaguchi, Y.; Maruyama, T.; Konagaya, A. High speed homology search with FPGAs. *Pacific Symposium on Biocomputing*; 2002; p. 271-282.
38. Zhang, Q.; Chamberlain, RD.; Indeck, RS.; West, B.; White, J. Massively parallel data mining using reconfigurable hardware: Approximate string matching. *Proc. of Workshop on Massively Parallel Processing*; April 2004;
39. Zhang Z, Schwartz S, Wagner L, Miller W. A greedy algorithm for aligning DNA sequences. *Journal of Computational Biology* 2000;7:203–14. [PubMed: 10890397]

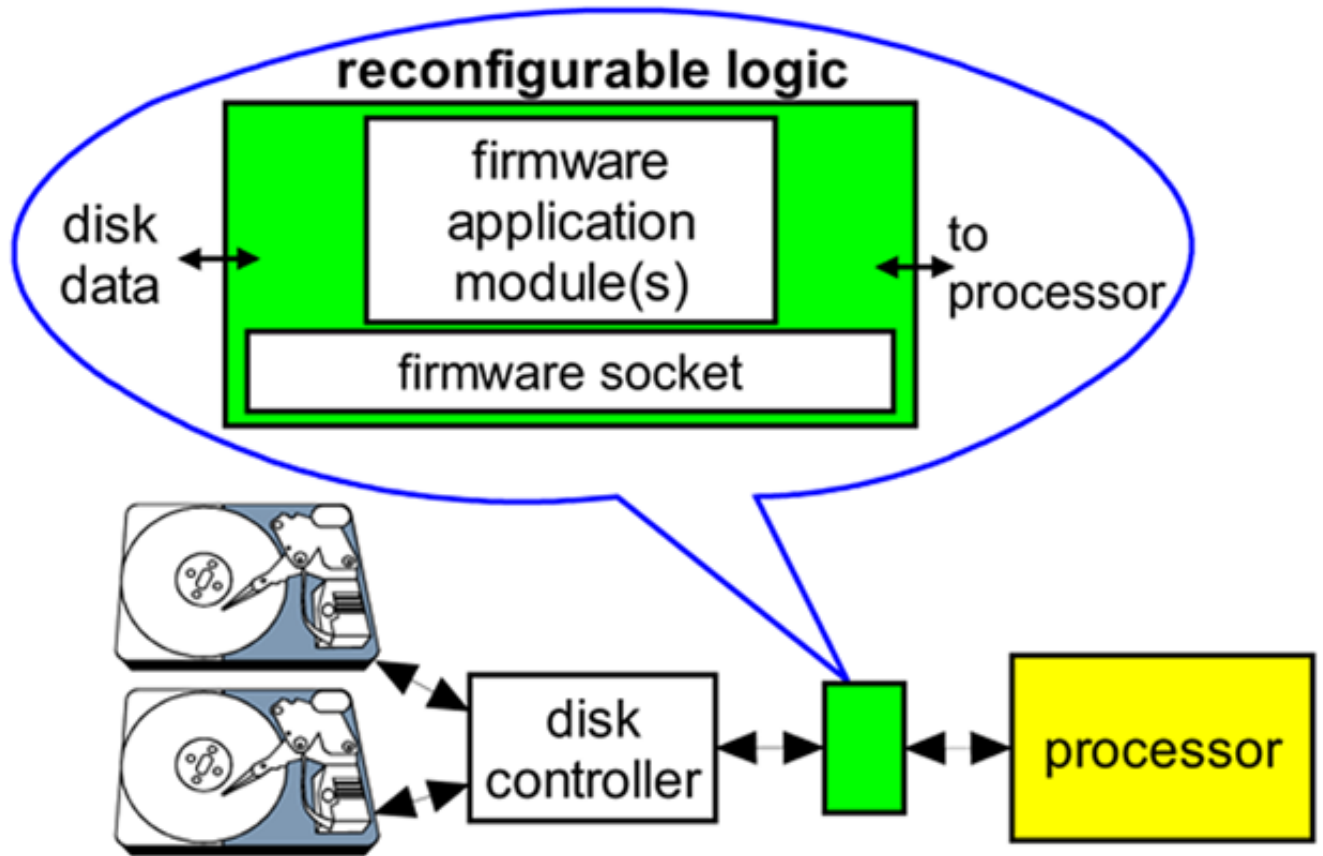


Figure 1.
Mercury system architecture




Figure 2.
Pipeline stages of NCBI BLAST algorithm




Figure 3.
Division of BLAST stage 1 (word matching) into 3 substages (1a: Bloom Filters, 1b: Hash Lookup, and 1c: Redundancy Eliminator)




Figure 4.
Typical Bloom filter functional diagram

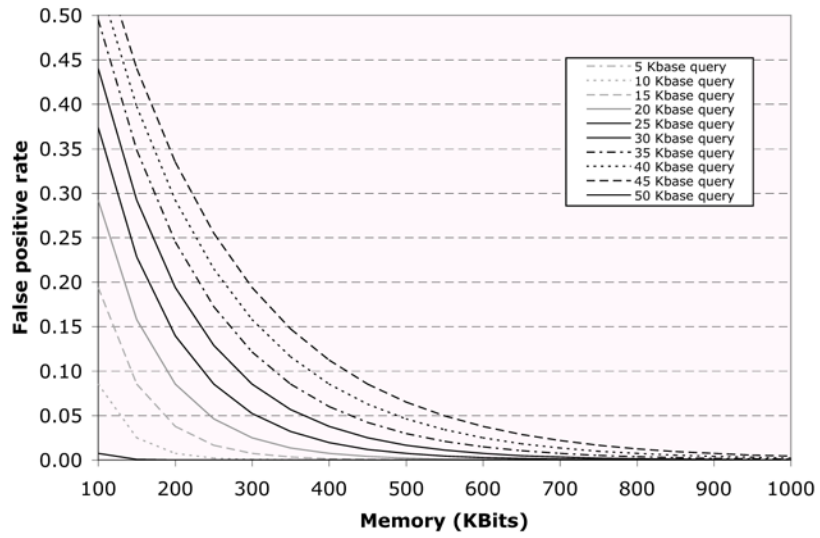


Figure 5.
Theoretical false positive rate of a Bloom filter vs. memory size for different query lengths




Figure 6.
Firmware Implementation of Bloom filter using block RAMs




Figure 7.
Four parallel Bloom filters

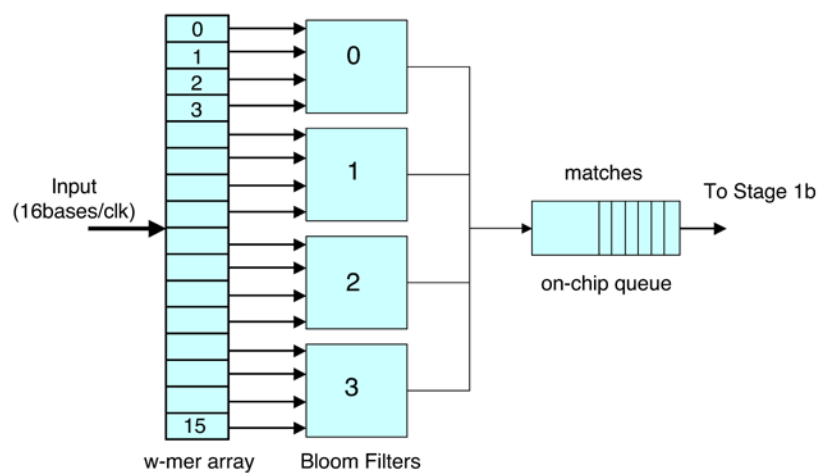


Figure 8.
Sixteen parallel Bloom filters




Figure 9.
Bloom filter output match rate vs. query size, NDC: no double clocking of block RAMs,
DC: double clocking of block RAMs

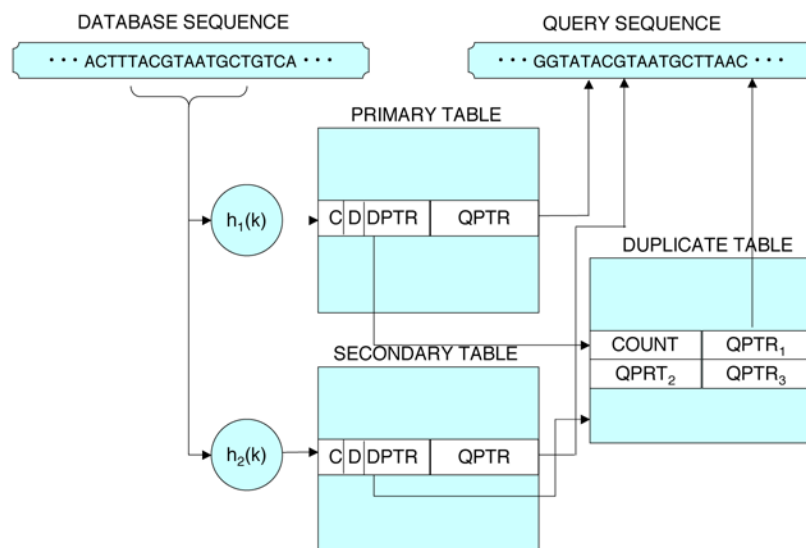


Figure 10.
Hash table (stage 1b) datapath

```

address =  $h_1(w\text{-mer})$ ; /* compute address for w-mer in primary table */
if (C == 1) { /* collision flag set */
    address =  $h_2(w\text{-mer})$ ; /* change to accessing secondary table */
}
if (D == 1) { /* duplicate flag set */
    address = DOFFSET + DPTR;
    return QPTRs; /* return all locations which match in query */
} else if (QPTR == NULL) { /* false positive identified */
    return NULL;
} else { /* unique match in query */
    return QPTR; /* return location in QPTR */
}

```

Figure 11.
Hash table control logic




Figure 12.
Redundancy filter example




Figure 13.
Maximum length of queue between stages 1a and 1b, query size = 12.5 kbases, $k=6$, $m=32Kb$




Figure 14.

Maximum length of queue between stages 1a and 1b, query size = 12.5 kbases, $k=6$, $m=64Kb$




Figure 15.

Maximum length of queue between stages 1a and 1b, query size = 20 kbases, $k=6$, $m=64\text{Kb}$




Figure 16.
Throughput of Mercury BLASTN with improved stage 2




Figure 17.
Speedup of Mercury BLASTN over NCBI BLASTN with improved stage 2

Table 1Match rates p across pipeline stages

Query Size (bases)	Stage 1 (p_1)	Stage 2 (p_2)	Stage 3 (p_3)
10 k	0.00858	0.0000550	0.320
25 k	0.0205	0.0000619	0.141
50 k	0.0411	0.0000189	0.194
100 k	0.0841	0.0000174	0.175
1 M	0.851	0.0000172	0.096

Table 2

Percentage of pipeline time spent in each stage of NCBI BLASTN

Query Size (bases)	Stage 1	Stage 2	Stage 3
10 k	86.53±1.33%	13.24±1.99%	0.23±0.02%
25 k	83.89±2.56%	15.88±4.40%	0.22±0.04%
50 k	82.63±2.94%	17.28±4.96%	0.09±0.01%
100 k	83.35±1.28%	16.58±2.17%	0.08±0.01%
1 M	85.39±3.34%	14.68±5.24%	0.03±0.01%

Table 3

Summary of performance results for software runs of NCBI BLASTN

Query Size	10 kbases	25 kbases	50 kbases	100 kbases	1 Mbase	Units
Throughput	67.0	29.2	14.9	8.76	0.648	Mbases/sec
Stage 1 (time per base, t_1)	0.0129	0.0287	0.0553	0.0951	1.32	μ sec/base
Stage 2 (time per match, t_2)	0.231	0.265	0.281	0.225	0.264	μ sec/match
Stage 3 (t_3)	71.3	60.4	81.8	58.9	34.4	μ sec/alignment

Table 4

Validation of analytic predictions using simulation

Configuration	Parameter	Analytic	Simulation
query=12.5kbases k=6, m=32Kbits	f	0.369	0.34
	f_r	0.81	0.82
	Process time stage 1b	≈ 1 clk	1.03
	Tput	2.128 Gbases/s	$2.128 \pm 2.9 \times 10^{-2}$ Gbases/s
query=12.5kbases k=6, m=64Kbits	f	0.016	0.014
	f_r	0.81	0.82
	Process time stage 1b	≈ 1 clk	1.07
	Tput	2.128 Gbases/s	$2.128 \pm 7.6 \times 10^{-6}$ Gbases/s
query=20.0kbases k=6, m=64Kbits	f	0.15	0.12
	f_r	0.69	0.71
	Process time stage 1b	≈ 1 clk	1.07
	Tput	2.128 Gbases/s	$2.128 \pm 7.0 \times 10^{-5}$ Gbases/s

 f : false positive rate from stage 1a f_r : %f removed by stage 1b

Table 5

Firmware vs. software stage 1 (throughput and speedup)

Query Size	10 kbases	25 kbases	50 kbases	100 kbases	1 Mbases	Units
NCBI BLASTN Stage 1 (T_{put_1})	77.4	34.8	18.1	10.5	0.771	Mbases/sec
Mercury BLASTN Stage 1 (T_{put_1})	1400	1400	700	350	35	Mbases/sec
Speedup (S_1)	18.1	40.2	38.7	33.4	45.4	

Table 6

Overall performance (throughput and speedup)

Query Size	10 kbases	25 kbases	50 kbases	100 kbases	1 Mbase	Units
NCBI BLASTN	67.0	29.2	14.9	8.76	0.657	Mbases/sec
Mercury BLASTN	497	181	86.0	52.6	4.44	Mbases/sec
Speedup	7.42	6.21	5.76	6.01	6.84	