

# Scalability Studies of the BLASTn Scan and Ungapped Extension Functions \*

Siddhartha Datta      Ron Sass  
Electrical & Computer Engineering Department  
University of North Carolina at Charlotte  
9201 University Blvd., Charlotte, NC 28223  
{skdatta, rsass}@uncc.edu

## Abstract

*BLASTn is a ubiquitous tool used for large scale DNA analysis. Detailed profiling tests reveal that the most computationally intensive sections of the BLASTn algorithm are the scan and ungapped extension stages consuming approximately 94% of its total execution time. Exploiting the inherent parallel nature of an FPGA is the fundamental advantage of porting BLASTn onto it. The aim of this paper is two-fold. First, we build upon our prior work with BLAST to develop a hardware/software co-design that is composed of multiple cores that can be scaled in two dimensions. The ungapped extension and a second dimension are new in this work. Second, we use this non-trivial example to explore spatially scalable designs. To provide the ability to move the design to a future generation chip, a mathematical model of performance that incorporates all of the system design parameters and the user's preference (high throughput vs low latency) is developed. We demonstrate here that the model correctly predicts the optimal ratio between the two dimensions on a Xilinx Virtex-4 and measures four to five times faster performance figures as compared to a state of the art general purpose processor.*

## 1 Introduction

For the last two decades, applications have benefited from frequency scaling. Every new generation of ICs offer both more transistors and faster clock rates. Increasing the clock rate allows every sequential application to execute faster on newer devices. Even though it is expected that the number of transistors will continue to increase at a dramatic rate, the system clock frequency in newer devices will grow much slower than in the past. Rather than relying on gains in frequency, future applications will need to be

designed such that they are spatially scalable to utilize the additional transistors that each new technology generation provides. This is especially true of FPGA-based applications. While many designs consist of simple cores that can be easily replicated, this does not always result in a spatially scalable design whose performance grows linearly.

BLASTn is an excellent example of this phenomenon. Running on a single core, this critical bioinformatics application is compute-bound. However, the application quickly becomes I/O-bound by increasing the number of computational cores. This is because each instance needs a large look-up table that is stored in main memory and all the cores are using the same I/O subsystem to read databases. Thus, multiple cores are contending for space in the cache, a portion of the fixed bandwidth to memory, and bandwidth to the I/O subsystem. An FPGA core that accelerates BLASTn cannot be arbitrarily replicated because the off-chip bandwidth is fixed and, ultimately, will limit performance.

This paper describes a parametrized implementation of NCBI BLASTn for Platform FPGAs. Our goal is two-fold: we extend our prior BLASTn core and then use the new core to explore a spatially scalable design. Specifically, the BLAST implementation described here adds another dimension to our prior work[5] that reduces latency (as well as improves throughput). This new core also introduces ungapped extensions to the basic hardware core — a function that was previously implemented in software. Using this new core, a spatially scalable design has been implemented. Based on the number of BLAST cores, their organization, and the off-chip memory bandwidth, this paper describes a mathematical formulation of the design's performance. This allowed us to develop a novel algorithm that — given the characteristics of an FPGA device, the board's memory bandwidth, and user preferences — calculates the optimal parameters for the design. This approach highlights a new level of scalability and portability: as larger chips become available and I/O technology evolves, one simply enters the characteristics of a new system and the algorithm determines the design's generic parameters.

\*This project was supported in part by the National Science Foundation under NSF Grants CNS 06-52468 (EHS) and CNS 04-10790 (CRI). The opinions expressed are those of the authors and not necessarily those of the Foundation.

## 2 Background and Related Work

The BLAST application was originally developed at Washington University, but in 1990 [2, 1], development of the public domain code shifted to the National Center for Biotechnology Information (NCBI). The application, which is actually a collection of programs, is used to compare an unknown genomic sequence, called a *query* against an existing *subject* genomic database to identify high similarity regions between them. The BLAST software has been developed in five flavors based on the nature of the subject database and query. These are: BLASTn, BLASTp, BLASTx, TBLASTn, TBLASTx. BLASTn is the program used to make local searches between a nucleotide based query and subject database. Nucleotides comprise of four different letters (*monomers*): A (*Adenine*), C (*Cytosine*), G (*Guanine*) and T (*Thymine*), represented using 2 bits.

The algorithm is divided into four major steps. The first step involves creation of the query lookup table and overflow table. The query is run through once forming all eight lettered words (*w-mers*). The offsets of these words are stored in the query lookup table. Each word is converted into a 16 bit hexadecimal number which is the location in the query lookup table where its offset is stored. If a certain word is repeated in the query, its offsets are stored in an overflow table. The offset value of that word in the lookup table is updated to a pointer to the overflow table. Since BLASTn is a similarity search rather than exact matching, the maximum number of repetitions considered per word is set to eight. The size of the lookup table and the overflow table is  $[2^{2(\text{WordLength})} \times 32]$  bits.

The second step or the scan stage identifies exact matching words in the database and the query sequence. These exact matches are called hits and their score value is eight. The stride length is set to four letters or eight bits. The scan stage is measured to consume  $\approx 78\%$  of the total execution time. The offsets of the matched words in the subject database and query are sent to the ungapped extension stage. where each hit is extended to the left and right by the stride length number of letters. Their scores are incremented by one for every exact match found. All hits that score greater than a threshold score set by the user are sent to the gapped extension stage. Depending upon the user's selection, gapped extension is done using either the Needleman-Wunsch algorithm[13] or the Smith-Waterman algorithm[15]. The ungapped extension is measured to consume another 12% to 16% of the total execution time. The complexity of the algorithm is:

$$O(N_{Seq} \times (N_H + N_G)) \quad (1)$$

where  $N_{Seq} \gg N_H \gg N_G$

where  $N_{Seq}$  is the number of sequences in the database,  $N_H$  is the number of hits found during the scan stage and  $N_G$  is

the number of hits sent for gapped extension.

Biological sequence searching using FPGAs and systolic arrays [9, 8, 4, 6] has been an active area of research for over a decade. A lot of research has been done in implementing the Smith-Waterman [15, 22] and Needleman-Wunsch [13] algorithm over hardware platforms to enhance the performance of search algorithms. The bottleneck for BLASTn resides in the first stage where the hits are located. This issue has been addressed by Mercury BLAST [11, 3]. More work has also been done in implementing the entire application on the FPGA like Tree-BLAST [7] and the TUC implementation [17, 16]. Another interesting design was proposed by [21] to port the entire application on chip. RCBLAST [12] proposed an index based searching scheme for BLASTn. Commercial companies, like Time Logic [20] and SGI/Mitronics [19] sell the Decypher machine and RASC Appliance respectively, market high speedup numbers over present multi core processors, but provide little detail about their designs. Recently, the IBM Cell broadband Engine[18] has also been used to speedup BLASTn.

All the previous implementations have shown significant speedups compared to software; however, most do so by compromising fidelity. Mercury BLASTn [3] claims 98% to 99% accuracy in results. Tree-BLAST [7] reports extra alignments (although reporting the same *E-values*) and [10] considers two repetitions of any word in the query, instead of eight. While one could argue that BLAST is a heuristic and 100% compatibility with NCBI BLAST is unnecessary, it is difficult to convince biologists of the same. The typical user would have no idea whether the differences are statistically significant. Thus, in addition to "better performance", complete consistency with NCBI BLAST is important.

## 3 Design

The goal of this design is two-fold. First, we augment our previous high-throughput design [5] to include a latency reducing single query option. We also add ungapped extension to the hardware core. Second, we build a bandwidth-aware scalable design that allows us to mathematically relate the bandwidth requirements to the size of various generics. Hardware implementations are based on the current 2.2.20 version of NCBI BLAST software.

### 3.1 Modified Query Lookup and Overflow Table

The inclusion of ungapped extension to the scan hardware core required modifications to be made to the query lookup and overflow tables. Four letters, each to the right, left and their count in the query are grouped into 32 bit words and written to the offset's next element in the query



**Table 1. Time taken to execute the Scan and Ungapped Extension functions for sequences of env\_nt.nt (3GB) Database**

Database Sequence Length	Hardware Core on FPGA	AMD Opteron (2GHz) One Core	Number of hits
638 Bytes	112 $\mu$ s	121 $\mu$ s	41
6378 Bytes	573 $\mu$ s	750 $\mu$ s	399
9138 Bytes	613 $\mu$ s	1639 $\mu$ s	743
13359 Bytes	583 $\mu$ s	1149 $\mu$ s	590
14480 Bytes	631 $\mu$ s	1987 $\mu$ s	779

and wait for an acknowledge, measured to be seven clock cycles.  $T_{FLT}$  is the time taken to read query offset and Left/Right Data from the query lookup table.  $T_{FOT}$  is the time taken to read eight overflow query offsets and their corresponding Left/Right Data.  $T_{UHW}$  is the time taken to fetch the query offsets from the DDR2 and compute the ungapped extension score.

$$\max(T_{UHW}) = (2T_I) + T_{FLT} + T_{FOT} = 48$$

$$\therefore \max(T_{HW}) = (N_{Seq}((S_{Len}/4) + 48))\text{clock cycles} \quad (3)$$

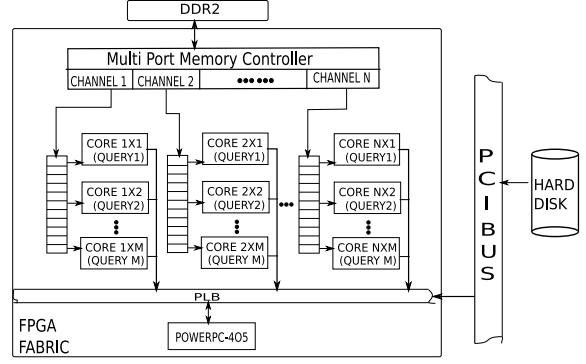
Since the ungapped extension is run during the scan in hardware, the complexity of the algorithm is reduced to  $O(N_{Seq} \times (N_G))$ . The performance of the software execution of NCBI `blast_nascan` and `nt_word_finder` functions run on a 2GHz AMD Opteron (single core) with 128KB L1 cache, 1MB L2 cache and 8GB DDR2 are as follows:  $T_{SSW}$  is the time taken to run scan function in software.  $T_L$  is the time taken to update the `FOR` loop and incrementing the subject offset counter.  $T_S$  is the time taken to shift every byte of the subject database to the index register which is compared to the query lookup table for locating hits. Let  $T_{HC}$  be the time taken to make a hit check and update the offset data structures in case of a hit. These instructions are measured to take four clock cycles. All cases assume 100% L1 cache hits and the entire database sequence and query is pre-buffered into the L2 cache.

$$T_{SSW} = (S_{Len}(T_L + T_S + T_{HC})) = 4S_{Len} \quad (4)$$

$T_U$  is the time taken to calculate ungapped extension scores.  $T_{UL}$  is the time to make loop comparisons over the number of hits found.  $T_{SA}$  and  $T_{QA}$  is the time taken to access four letters to the left and right of the hit in the subject database and query respectively. Each of these instructions take one clock cycle to execute.  $T_C$  is the time taken to make eight comparisons and update the score.  $NH$  is the total number of hits found in the scan function.

$$T_U = (NH(T_{UL} + T_{SA} + T_{QA} + T_C)) = 13NH \quad (5)$$

$$\therefore \min(T_{SW}) = N_{Seq}(13NH + 4S_{Len})\text{clock cycles} \quad (6)$$



**Figure 3. Design of the System**

Since, the AMD Opteron is clocked at 2GHz and the FPGA core at 100MHz, from equations 3 and 6:

$$\therefore \min(\text{Speedup}) \propto (2.6NH + 0.8S_{Len}) \quad (7)$$

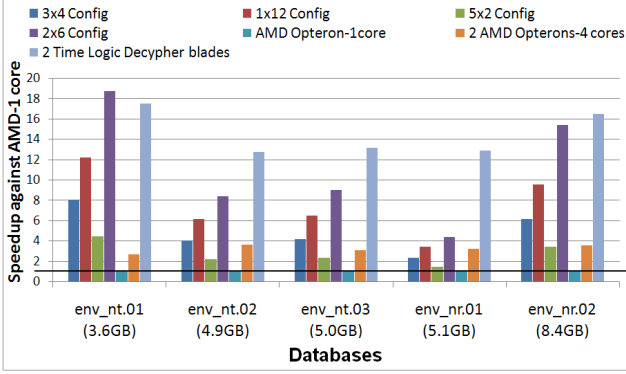
According to 7, the implementation of the scan and ungapped extension of the FPGA is theoretically about 1.2 – 1.5 times faster and the rate increases with the increase in number of hits. Table 1 lists the experiments' results which validates this point.

### 3.3 Scalable Design Model

In [5], various queries' hit index tables were loaded into multiple cores connected across a single direct channel of the Multi Port Memory Controller. This resulted in incomplete utilization of all its ports, but high throughput due to hits located across several queries in parallel. ( $\text{latency of system} \propto (1/\text{throughput of system})$ ). Sequences of the database streamed over multiple channels from the DDR2 in parallel results in low latency but the number of cores loaded with different queries' Hit Index Tables is reduced, reducing the throughput. For a XC4VFX60 FPGA, the channel's bandwidth ( $ch\_bw$ ) is measured to be 1.6GB/s when data is read continuously over one channel (continuous mode). However, when reads occur simultaneously over multiple channels, the burst mode bandwidth is measured to be 0.98GB/s. With two ( $n$ ) Hit Index Tables in the core, data is consumed at a rate of 0.4 GB/s. The mathematical model for designing an optimal ( $N \times M$ ) grid of cores with  $N$  channels for streaming in the sequences of the database over  $M$  queries for any platform FPGA is as follows:

$$N = \begin{cases} 1 & \text{-- max throughput} \\ num\_ch & \text{-- min latency} \\ \left\lfloor \frac{ch\_bw}{0.2n} \right\rfloor & \text{-- optimal performance} \end{cases} \quad (8)$$

where  $num\_ch$  is the total number of channels available. Each Hit Index Table utilizes four 18Kb BRAM blocks.



**Figure 4. Speedup Numbers against one core of the AMD Optron**

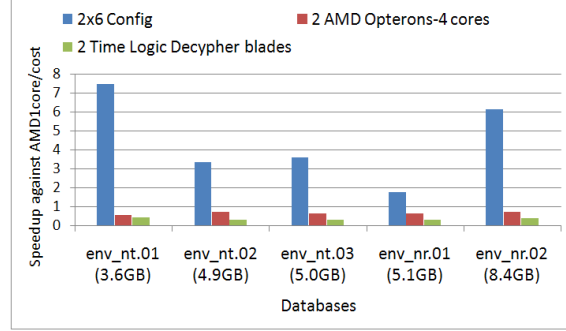
Each lookup queue and output FIFO utilizes one 18Kb BRAM block each. The subject FIFO utilizes two 18Kb BRAM blocks. For adding every query, four lookup queues, two output FIFOs and  $n$  number of Hit Index Tables are added to the system. For every channel, one subject FIFO is added to the system. The DDR2 wrapper consumes twenty-one 18Kb BRAM blocks and eight 18Kb BRAM blocks for every channel added to the system. A 32KB and 64KB cache for the PPC consumes sixteen and thirty-two 18Kb BRAM blocks respectively. The trimode macc FIFO and pci bridge consume two 18Kb BRAM blocks each. Let  $B$  be the total number of 18Kb BRAM blocks available on the chip.

$$M = \begin{cases} \left\lfloor \frac{B - (10N + 41)}{(4n + 6)N} \right\rfloor - 32\text{KB Cache} \\ \left\lfloor \frac{B - (10N + 57)}{(4n + 6)N} \right\rfloor - 64\text{KB Cache} \end{cases} \quad (9)$$

Hence, using equations 8 and 9, depending upon any system's specifications, its optimal  $N \times M$  grid size can be determined. Figure 3 demonstrates the layout of the system.

## 4 Experimental Setup and Evaluation


To evaluate the effectiveness of the proposed technique, a set of experiments was designed to test the absolute rate of computation and scalability of the FPGA accelerated implementation. A Xilinx ml410 evaluation board embedded with the XC4VFX60 FPGA is used for the tests. Two Powerpc-405 processors are fabricated on the FPGA. Linux 2.6.26 is run on one of the processors at 300MHz. The BLASTn executables, formatted databases and formatted queries are stored in a WD1600AAJS 160GB hard disk. The interface to the hard disk is through PCI and the PCI-to-PLB bridge (Figure 3). Four configurations of cores ( $1 \times 12$ ,



**Figure 5. Performance with respect to cost of equipment**

$2 \times 6$ ,  $3 \times 4$  and  $5 \times 2$ ) on the Virtex-4 were successfully synthesized and tested with the BLASTn algorithm running on Linux. All the four configurations' synthesis reports show 99% on-chip memory and 98% slice utilization. 12 queries of various sizes, ranging from 0.77KBases to 136KBases, were searched over the `env.nt` and `env.nr` databases, downloaded from the NCBI website. For performance comparison purposes, these tests were run on a server with 2 dual core AMD Optrons (2GHz) with 128KB L1 cache, 1MB L2 cache, 8GB DDR2, 1TB HDD (SATA interface) and two blades of TimeLogic's Decypher machine. The results of the tests are shown in Figure 4. The bar chart indicates speedups measured with different configurations on the FPGA, the AMD Optron server and the Decypher machine over a single core's performance of the AMD Optron server. The best performance on the FPGA was achieved for the  $2 \times 6$  configuration of cores, validating our mathematical model (equations 8, 9) described in section 3. The Decypher machine though is measured to be about twice as fast as the  $2 \times 6$  configuration of cores on the FPGA. The cost of the Decypher machine is quoted at \$42000, the AMD Optron server is quoted at \$3000 and the Xilinx ml410 board is quoted at \$2500. The performance/\$ results shown in Figure 5 plot the ratio of speedups attained against a single core of the AMD Optron for the same tests described in Figure 4. The FPGA's best implementation priced 16 times lower than the cost of the Decypher machine proves to be an inexpensive yet one of the fastest existing solutions to the BLASTn problem having complete consistency with the NCBI results.

Figure 6 shows the measured relations between spatial computing and their execution time taken. The scan and un-gapped extension functions are measured to utilize on an average 3% and the communication overhead takes 1% of the total execution time when run on the FPGA, as opposed to an average 87% and 5% respectively on the AMD Optron server. This indicates the compute-bound problem in software is converted to be I/O bound by increasing the number



**Figure 6. Spatial Computing vs Time**

of cores and introducing parallelism in the hardware implementation. The majority of the total execution time of the algorithm is taken to read data from the hard drive to the RAM. Hence, with an increase in the number of cores, more physical area on the chip is utilized but it provides lower execution latencies.

## 5 Conclusion

This paper describes a BLASTn accelerator core that performs the scan and ungapped extension functions. The core, used in a spatially scalable design, uses a mathematical model of performance and a novel algorithm to determine the optimal parameters for a given FPGA system. A variety of configurations were tested and the model correctly predicted their performance of the configuration. The resulting FPGA design has been shown to be very competitive; it outperforms a typical AMD Opteron compute server in terms of throughput and performance/\$.

The design described here runs on single FPGA node. However, this node is just one of 64 nodes in a high-performance computing cluster[14]. The layout of the single node system was designed so that the replicated cores could, in the future, be extended to span all of the FPGAs in this cluster. This, combined with parallel disk access and the high speed network, will change the bandwidth numbers and introduce new variables into the mathematical formulation. The next step will be to extend the spatial scalability study to include this multi-FPGA platform.

## References

[1] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman. Basic local alignment search tool. *J Mol Biol*, 215:403–410, October 1990.

[2] S. F. Altschul, T. L. Madden, A. A. Schäffer, J. Zhang, Z. Zhang, W. Miller, and D. J. Lipman. Gapped blast and psi-blast: a new generation of protein database search programs. *Nucleic Acids Res*, 25(17):3389–3402, September 1997.

[3] J. D. Buhler, J. M. Lancaster, A. C. Jacob, and R. D. Chamberlain. Mercury blastn: Faster dna sequence comparison

using a streaming hardware architecture. In *Reconfigurable Systems Summer Institute*, July 2007.

[4] R. H. Computer and R. Hughey. Massively parallel biosequence analysis. Technical report, 1993.

[5] S. Datta, P. Beeraka, and R. Sass. RC-BLASTn: Implementation and Evaluation of the BLASTn Scan Function. In *FCCM '09*, April 2009.

[6] D. H. Department, D. T. Hoang, and D. P. Lopresti. Fpga implementation of systolic sequence alignment. In *International Workshop on Field Programmable Logic and Applications*, 1992.

[7] M. C. Herbordt, J. Model, B. Sukhwani, Y. Gu, and T. VanCourt. Single pass streaming blast on fpgas. *Parallel Computing*, 33(10-11):741–756, 2007.

[8] D. Hoang. Searching genetic databases on splash 2. In *FPGAs for Custom Computing Machines, 1993. Proceedings. IEEE Workshop*, pages 185–191, Apr 1993.

[9] R. Hughey. Parallel sequence comparison and alignment. *Application-Specific Systems, Architectures and Processors, IEEE International Conference on*, 0:137, 1995.

[10] S. Kasap, K. Benkrid, and Y. Liu. Design and implementation of an fpga-based core for gapped blast sequence alignment for the two-hit method. volume 16, pages 443–452, 2008.

[11] P. Krishnamurthy, J. Buhler, R. Chamberlain, M. Franklin, K. Gyang, A. Jacob, and J. Lancaster. Biosequence similarity search on the mercury system. *J. VLSI Signal Process. Syst.*, 49(1):101–121, 2007.

[12] K. Muriki, K. D. Underwood, and R. Sass. Rc-blast: Towards a portable, cost-effective open source hardware implementation. In *IPDPS '05- Workshop 7*, page 196.2, 2005.

[13] S. B. Needleman and C. D. Wunsch. A general method applicable to the search for similarities in the amino-acid sequence of two proteins. *Journal of Molecular Biology*, 48:443–453, 1972.

[14] R. Sass, W. V. Kritikos, A. G. Schmidt, S. Beeravolu, and P. Beeraka. Reconfigurable computing cluster (rcc) project: Investigating the feasibility of fpga-based petascale computing. In *FCCM*, pages 127–140, 2007.

[15] T. Smith and M. S. Waterman. Identification of common molecular subsequences. *Journal of Molecular Biology*, 147:195–197, 1981.

[16] E. Sotiriades and A. Dollas. Design space exploration for the blast algorithm implementation. In *FCCM '07*, pages 323–326, 2007.

[17] E. Sotiriades and A. Dollas. A general reconfigurable architecture for the blast algorithm. *J. VLSI Signal Process. Syst.*, 48(3):189–208, 2007.

[18] A. Wirawan, K. C. Keong, and B. Schmidt. Parallel dna sequence alignment on the cell broadband engine. Lecture Notes in Computer Science, pages 1249–1256, 2007.

[19] www.sgi.com/pdfs/4010.pdf. Sgi products.

[20] www.timelogic.com. Time logic biocomputing solutions.

[21] F. Xia, Y. Dou, and J. Xu. Families of fpga-based accelerators for blast algorithm with multi-seeds detection and parallel extension. volume 13 of *Communications in Computer and Information Science*, pages 43–57. Springer, 2008.

[22] C. W. Yu, K. H. Kwong, K. H. Lee, and P. H. W. Leong. A smith-waterman systolic cell. In *FPL 2003*, pages 375–384. Springer, 2003.