

A RECONFIGURABLE ARCHITECTURE FOR THE PHYLOGENETIC LIKELIHOOD FUNCTION

*Nikolaos Alachiotis, Alexandros Stamatakis**

The Exelixis Lab
Dept. of Computer Science
Technische Universität München
email: {alachiot,stamatak}@in.tum.de

Euripides Sotiriades, Apostolos Dollas

Microprocessor and Hardware Lab
Dept. of Electronic and Computer Engineering
Technical University of Crete
email: {esot,dollas}@mhl.tuc.gr

ABSTRACT

As FPGA devices become larger, more coarse-grain modules coupled with large scale reconfigurable fabric become available, thus enabling new classes of applications to run efficiently, as compared to a general-purpose computer. This paper presents an architecture that benefits from the large number of DSP modules in Xilinx technology to implement massive floating point arithmetic. Our architecture computes the Phylogenetic Likelihood Function (PLF) which accounts for approximately 95% of total execution time in all state-of-the-art Maximum Likelihood (ML) based programs for reconstruction of evolutionary relationships. We validate and assess performance of our architecture against a highly optimized and parallelized software implementation of the PLF that is based on RAxML, which is considered to be one of the fastest and most accurate programs for phylogenetic inference. Both software and hardware implementations use double precision floating point arithmetic. The new architecture achieves speedups ranging from 1.6 up to 7.2 compared to a high-end 8-way dual-core general-purpose computer running the aforementioned highly optimized OpenMP-based multi-threaded version of the PLF.

1. INTRODUCTION

While the reconfigurable architecture community has been involved with bioinformatics, several interesting problems characterized by great computational demands came to light [1, 2]. New FPGA devices offer hardware resources that can be used to build powerful floating point arithmetic architectures. This feature now allows to deploy new generation FPGAs for novel classes of bioinformatics problems. One such challenging problem is the evaluation of the Phylogenetic Likelihood Function (PLF). Previous efforts to design reconfigurable architectures for this problem have been reported [3] but due to the wide number of methods available

and the biological data flood, the efficient computation of the PLF remains an open challenge.

The PLF is the most computationally intensive part of RAxML [4] and a plethora of other PLF-based codes such as GARLI, MrBayes, PAML, or PAUP. Those programs are widely used by biologists (over 20,000 citations according to Google Scholar). The input for PLF-based programs is a multiple sequence alignment, essentially an $n \times m$ data matrix, that contains n DNA sequences with a length of m nucleotide characters (m columns) each. More than 95% of overall execution time is spent to compute the PLF in the aforementioned programs. A phylogenetic tree is a binary tree structure that represents the evolutionary relationships among species. The tips (also called leaves or taxa) of the tree represent species alive today in contrast to internal (ancestral) nodes that represent species that have become extinct. Phylogenetic trees have many important applications in medical and biological research (see [5] for a summary). In this paper we present a new architecture which significantly extends an initial proof-of-concept design described in [5].

The initial architecture can only compute the PLF on fully balanced trees, but nonetheless yielded a significant performance boost. The more versatile architecture presented here, is capable to evaluate the PLF for any given tree topology at the same speed as the previous architecture, to exploit the intrinsic parallelism of the PLF in a more scalable manner, and to conduct so-called partial tree traversals which represent a fundamental mechanism current fast tree search algorithms. The proposed architecture yields exactly the same likelihood scores as the reference software. It has been fully post place and route simulated and executes several dozens of double-precision floating point operations during every clock cycle. Input/output issues have been taken into account to allow for adaptation to a modern platform. An accelerated hardware solution is desirable, since current biological analyses with RAxML on an IBM BlueGene/L supercomputer already require 2.25 million CPU hours and 89GB of main memory [6].

*Part of this work is funded under the auspices of the Emmy-Noether program by the German Science Foundation (DFG)

2. COMPUTING THE PHYLOGENETIC LIKELIHOOD FUNCTION

Felsenstein's pruning algorithm [7], described below, is the standard method to compute the PLF and hence the likelihood score for a given tree topology. The first step consists of determining a pair of child nodes i and j in the tree for which the likelihood vector at the common ancestor k ($1 \leq i, j, k \leq 2n - 2$) has not already been computed. The second step is to calculate the likelihood vector entries of the common ancestor (ancestral likelihood vector at k) and prune out the child nodes. These steps are executed recursively until the likelihood vector at the virtual root vr has been calculated and thus the pruning process has transformed the initial tree to only one node that is located at the virtual root. Phylogenetic trees under ML are unrooted for mathematical and computational reasons [7], but a virtual root vr can be placed into any branch of the tree to evaluate its likelihood score.

In order to compute the PLF on a tree one also needs the branch lengths and the remaining likelihood model parameters (see [5] for details). Given the branch lengths and model parameters, one initially needs to compute the entries for all ancestral likelihood vectors which are located at the inner nodes of the tree bottom up from the tips towards the virtual root. Every likelihood vector entry at position c ($c = 1 \dots m$) $\vec{L}(c)$ at the tips and at the inner nodes contains the four probabilities $P(A)$, $P(C)$, $P(G)$, $P(T)$ of observing a nucleotide A, C, G, or T at a specific column c of the input alignment. The probabilities at the tips (leaves) of the tree for which observed data (DNA sequences) is available are set to 1.0 for the observed nucleotide character at the respective position c , e.g., for the nucleotide A: $\vec{L}(c) = (1.0, 0.0, 0.0, 0.0)$.

Given, a parent node k and two child nodes i and j , their likelihood vectors $\vec{L}^{(i)}$ and $\vec{L}^{(j)}$, the respective branch lengths leading to the children b_i and b_j and the transition probability matrices $P(b_i)$, $P(b_j)$, the likelihood of observing an A at position c of the ancestral (parent) vector $\vec{L}_A^{(k)}(c)$ is computed as follows:

$$\vec{L}_A^{(k)}(c) = \left(\sum_{S=A}^T P_{AS}(b_i) \vec{L}_S^{(i)}(c) \right) \left(\sum_{S=A}^T P_{AS}(b_j) \vec{L}_S^{(j)}(c) \right) \quad (1)$$

This procedure is outlined in Figure 1. When the procedure reaches the root, the per-column likelihood $l(c)$ is computed as follows using the likelihood vector $\vec{L}^{(vr)}$ at the virtual root:

$$l(c) = \sum_{S=A}^T \pi_S L_S^{(vr)}(c) \quad (2)$$

The probabilities π_A through π_T are the prior probabilities (also called base frequencies) for observing A, C, G, or

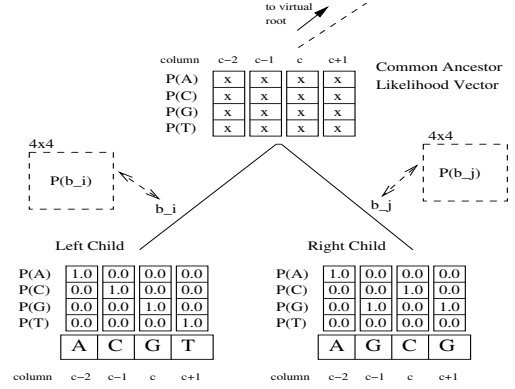


Fig. 1. Computation of ancestral likelihood vector entries.

T at vr and are typically drawn empirically from the input data. To compute the overall likelihood we then compute the product over all $l(c)$.

3. RELATED WORK

While there exist many tools and methods for phylogeny reconstruction, only few of them have been mapped to hardware. Bakos *et al* [8] map GRAPPA [9] to an FPGA which is based on gene order input data. Phylogenetic inference using gene order data is a discrete problem and does therefore not require floating point operations. Phylogenetic analyses that use gene order instead of DNA sequence input data are rarely used for real-world analyses at present. Also, a simple sequence-based method called UPGMA (Unweighted Pair Group Method with Arithmetic Mean) has been mapped to hardware [10]. UPGMA is one of the most simple tree reconstruction methods and is currently not used for real-world phylogenetic analyses. Mak and Lam [3] report the mapping of a reduced floating point precision PLF implementation to FPGAs that is based on the simple Jukes-Cantor (JC69 [11]) model of nucleotide substitution. The most commonly used and most complex model is the GTR (General Time Reversible) model of nucleotide substitution which we implement in our architecture.

4. ARCHITECTURE

We propose a master-worker architecture that consists of two main units: the Target Pair Unit (TPU , master) and the Computational Basic Core (CBC , worker). The TPU (master) performs the first of the previously described algorithmic steps while the CBC (worker) unit performs the second. The TPU executes the tree traversal steps of the pruning algorithm on the contents of a local memory which is used to hold information about the nodes of the tree, i.e., the tree topology. The TPU tracks down which tips, inner nodes or combination thereof should be combined for computing the




Fig. 2. Architecture overview.




Fig. 3. Architecture of the Target Pair Unit (TPU).

entries of an ancestral likelihood vector. The information required by the worker to locate the tip sequences or the likelihood vectors in the external or internal memories in order to start calculating the ancestral likelihood vector is provided via shared registers. Both the *TPU* as well as the *CBC* have access to these registers. Once the *TPU* has written the required addresses and selection bits, it sends a start signal to the *CBC* which indicates that there are valid data available in the registers. This means that the *CBC* can start its calculation process. At the same time the *TPU* switches to stand-by mode. The *TPU* waits for the *CBC* to calculate the ancestral likelihood vector and write back the address and selection bits of the memory position (to the shared registers) where the newly computed ancestral likelihood vector has been stored. The *CBC* then announces that the calculation process has been completed by sending a respective signal to the *TPU*. The *TPU* will then update its local memory accordingly using the information from the shared registers and will determine the next pair of tips and/or nodes for which an ancestral vector needs to be computed.

We denote this design as master-worker architecture because only the *TPU* can initiate computations on the *CBC*. The general architectural scheme of the proposed design is illustrated in Figure 2.

4.1. The Target Pair Unit (TPU)

As already mentioned the *TPU* executes Felsenstein’s pruning algorithm on the contents of its local memory. The top-level architecture of this unit is outlined in Figure 3. The unit consists of three structural components, a *Control Unit* implemented as FSM (Finite State Machine) and a local memory. The *Examine Line* and *Control Unit* components jointly implement the pruning algorithm that determines the pairs of tips and/or nodes to be combined.




Fig. 4. Layout of a tree topology memory line.

The *Store Tree* unit is used to initialize the contents of the local memory. It reads in the input tree to be evaluated that is encoded as an integer array. The integer array contains the node depths (distances from *vr*) in depth-first order from the virtual root. For example, the tree illustrated in Figure 5 (top right) is given by the integer array: 1 (depth of A), 2 (D), 3 (C), 3 (B). The *Create New Line* unit is used to retrieve the information that has been written to the shared registers by the *CBC* which describes the ancestral vector that has just been calculated. The fields in the memory lines of the local memory which holds the tree structure are outlined in Figure 4. The *valid* field indicates whether the memory line contains useful information for the remainder of the pruning process. The *depth* field contains the depth of the tip or node that the line describes. The remaining fields hold the necessary information for addressing the nucleotide sequences or the ancestral likelihood vectors. The address field is an index to the first nucleotide of a DNA sequence (tip case) or the first likelihood vector entry at an inner (ancestral) node. Both nucleotide sequences and ancestral likelihood vectors are stored contiguously in the external and the internal memory respectively. The address field contains a memory address. The *ext_int_sel* (external/internal selection) field specifies whether this address refers to internal (value set to 1) or external memory (value set to 0). Since internal memory is organized into two parts, the *lft_right_sel* (left/right selection) field is used to select between these two parts (0-left, 1-right). Two snapshots of the local memory that illustrate how the *TPU* works are depicted in Figure 5. For the four-taxon tree shown in Figure 5 (top right), the master unit initializes the memory as shown by the table in the top left corner. Every tip of the tree has been labeled with a capital letter. There is also one number for every node (tips, internal nodes) that indicates the distance to the virtual root. The component *Examine Line* reads the memory from the top and stores the contents of fields: address, *ext_int_sel*, and *lft_right_sel* in the shared registers. Then, the *Control Unit* triggers the *CBC* to initiate the calculation. When the *CBC* has calculated the respective likelihood vectors it updates the shared registers with the address and selection bits of the ancestral vector which has been labeled by E in the second tree of Figure 5 (lower part). Then, the *Create New Line* unit updates the memory and the contents of the memory are used in the subsequent step of the pruning algorithm. The memory update sequence is thus equivalent to the pruning steps in the tree.

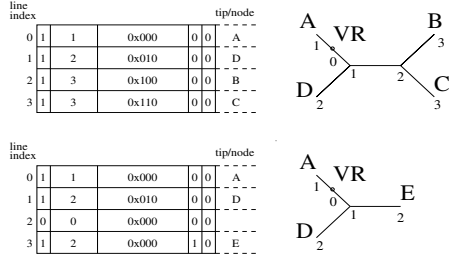


Fig. 5. Two consecutive snapshots of the tree topology memory and the respective trees.

4.2. The Computational Basic Core (CBC)

Initially, the *CBC* idles while the *TPU* determines the nodes to be combined. When the *CBC* is triggered by the *TPU* to start the calculations, it fetches the data at the position indicated by the addresses and selection bits in the shared registers and starts the calculation of the likelihood vectors. An overview of the *CBC* architecture is provided in Figure 6. It consists of the Control FSM, the *Basic Cell Array*, the *Fetch Units* (FIFOs), internal memories and the *Likelihood Score Unit*. The *Fetch Units* have been specifically designed to hide the latency of linear external memory accesses to tip sequences (evidently the latency can not be hidden for the first accesses to the tip address) from the *Basic Cells*. Since the likelihood vectors are long, i.e., typically $m \geq 1,000$, the latency for the access to the first datum of an array is negligible while we can achieve infinitely large burst ability.

The *Basic Cell Array* consists of ten *Basic Cells* which perform the additions and multiplications given in Equation 1. All *Basic Cells* work in parallel on a different column of the sequence alignment. According to the node pair provided by the *TPU* the appropriate data are pre-fetched and stored in the *Fetch Units* or accessed directly in internal memories. The resulting ancestral likelihood vector is then written to the internal memories. Once the likelihood vector of the virtual root has been computed, the *Basic Cells* are used again to calculate the per-column likelihood scores (see Equation 2) and the product over the per-column scores $l(c)$ is then computed by the *Likelihood Score Unit*.

Each *Basic Cell* is arranged as a tree of double precision floating point adders and multipliers, which operate in parallel as shown in Figure 7. The *Basic Cells* are fully pipelined with a total pipeline depth of 58 cycles. Each *Basic Cell* evaluates one likelihood value per cycle yielding one likelihood vector entry every 4 cycles. Because the operations required for the calculation of the likelihood vector at the virtual root are slightly different, the *Basic Cell* also contains a vector of pipeline registers and a 4 to 1 multiplexer in order to perform the appropriate operations when the respective *mode* signal is set.

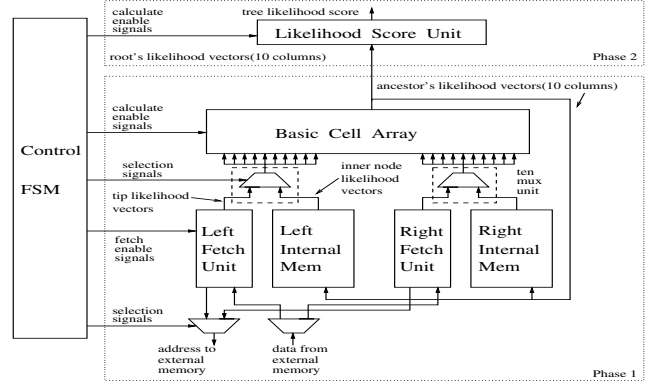


Fig. 6. Overview of the Computational Basic Core (CBC).

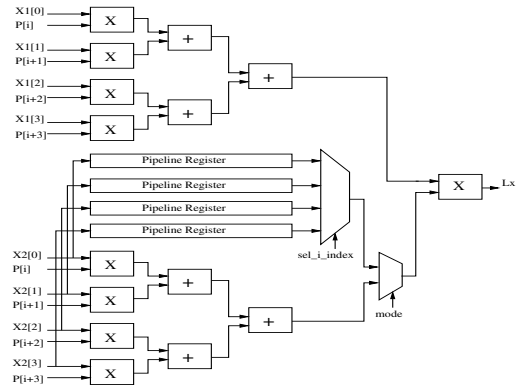


Fig. 7. Basic Cell Architecture.




Fig. 8. FPGA vs Sun x4600 multi-threaded execution times.

5. SYSTEM EVALUATION & PERFORMANCE

Extensive post place and route simulations were conducted to verify the functionality of the proposed architecture. The input data sets contained trees with 4, 8, 16, 32, 64, 128, 256, and 512 taxa (tips) and a length of 1,000 nucleotides each. In order to conduct a fair performance comparison between the hardware and software implementations, we designed

# taxa	1 Thr.	2 Thr.	4 Thr.	8 Thr.	16 Thr.
4	7.04	3.92	2.40	2.16	3.77
8	5.76	3.32	2.14	2.18	3.08
16	5.22	3.04	1.95	1.62	2.84
32	5.19	2.90	1.86	1.66	2.92
64	5.49	3.00	1.86	1.64	2.89
128	5.50	3.20	2.18	1.73	2.27
256	5.44	3.20	2.33	1.94	2.34
512	5.41	3.18	2.28	2.01	2.96

Table 1. Hardware versus multi-core software speedups

and optimized (based on 8 years of programming experience with phylogenetic inference software) a light-weight software version to compute the PLF that omits the overhead of the standard RAxML open-source distribution. We compiled this light-weight version (available at <http://www.kramer.in.tum.de/exelixis/software>) using the Intel icc compiler (v 10.1, optimization option -O3) that generates faster code than gcc for RAxML and also parallelized the PLF with OpenMP. We executed the program with 1, 2, 4, 8, and 16 threads on a high-end SUN x4600 system with 64 GB of main memory and 8 dual core AMD Opteron processors running at 2.6 GHz. In order to obtain accurate software timing results, we measured the time required by a loop that executes 20,000 full tree traversals. The proposed architecture was mapped to a Xilinx V5 SX240T FPGA. The design with 10 parallel basic cells uses 87% of the slice LUTs, 94% of the BRAMs, and 93% of the DSP48Es on this device.

The clock speed that was measured for the design amounts to 101 MHz (static timing report of the Xilinx Tools, ADVANCED 1.53 speed file). Figure 8 provides the projected FPGA execution times and actual software execution times on the Sun x4600 for 1, 2, 8, and 16 threads as a plot over input tree size for the respective data sets. As can be observed in the plot, the execution of the software implementation with 16-threads is slower than with 8-threads because of a lack of scalability caused by an unfavorable communication to computation ratio. For all datasets, FPGA performance is better than software performance on the multi-core system. While the execution times are within the millisecond range, in real application scenarios, search algorithms will invoke likelihood computations millions of times to conduct ML estimates of model parameters and to search for the best ML tree topology which is an NP-complete optimization problem.

6. CONCLUSION & FUTURE WORK

A new architecture for the PLF was presented. The speedups of 1.6–7.04 are modest, though within the typical order of magnitude for double precision floating point kernel implementations on FPGAs. Moreover, performance is compared in a fair way to a highly optimized and OpenMP-parallelized

code on a high-end multi-core machine. We will work towards redesigning the critical path in order to increase the clock speed, and plan for mapping the design to actual hardware and integrate it with software for phylogenetic inference. We will further extend the current hardware design by modules to calculate the likelihood score of very large trees (with respect to the number of taxa n), which require a scaling mechanism to avoid numerical underflow, as well as by modules for supporting the branch length optimization process which is conducted via an iterative Newton-Raphson procedure.

7. REFERENCES

- [1] K. Puttegowda, W. Worek, N. Pappas, A. Dandapani, P. Athanas, and A. Dickerman, "A run-time reconfigurable system for gene-sequence searching," in *Proc. 16th Int. Conf. on VLSI Design*, 2003, pp. 561–566.
- [2] F. Xia, Y. Dou, and J. Xu, "FPGA-Based Accelerators for BLAST Families with Multi-Seeds Detection and Parallel Extension," in *Proc. of 2nd Int. Conf. on Bioinformatics and Biomedical Engineering*, 2008, pp. 58–62.
- [3] T. Mak and K. Lam, "Embedded computation of maximum-likelihood phylogeny inference using platform FPGA," in *Proc. IEEE Computational Systems Bioinformatics Conference*, 2004, pp. 512–514.
- [4] A. Stamatakis, "RAxML-VI-HPC: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models," *Bioinformatics*, vol. 22, no. 21, pp. 2688–2690, 2006.
- [5] N. Alachiotis, E. Sotiriades, A. Dollas, and A. Stamatakis, "Exploring FPGAs for Accelerating the Phylogenetic Likelihood Function," in *Proc. of HICOMB2009*, Rome, Italy, May 2009.
- [6] M. Ott, J. Zola, S. Aluru, A.D. Johnson, D. Janies, and A. Stamatakis, "Large-scale phylogenetic analysis on current HPC architectures," in *Scientific Programming*, 2008, pp. 255–270.
- [7] J. Felsenstein, "Evolutionary trees from DNA sequences: a maximum likelihood approach," *J. Mol. Evol.*, vol. 17, pp. 368–376, 1981.
- [8] J. Bakos, "FPGA Acceleration of Gene Rearrangement Analysis," in *Proc. 15th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, 2007, pp. 85–94.
- [9] B. Moret, J. Tang, L. Wang, and T. Warnow, "Steps toward accurate reconstructions of phylogenies from gene-order data," *J. Comp. Syst. Sci.*, vol. 65, no. 3, pp. 508–525, 2002.
- [10] J. Davis, S. Akella, and P. Waddell, "Accelerating phylogenetics computing on the desktop: experiments with executing UPGMA in programmable logic," in *Proc. 26th Annual Int. Conf. of the IEEE Engineering in Medicine and Biology Society*, vol. 2, 2004.
- [11] T. Jukes and C. Cantor, *Evolution of protein molecules*. Academic Press, New York, 1969, ch. III, pp. 21–132.