

Designing digital circuits for FPGAs using parallel genetic algorithms

Rizwan A. Ashraf¹, Francis Luna¹, Damian Dechev^{1,2}, Ronald F. DeMara¹

1: Department of Electrical Engineering and Computer Science

University of Central Florida

Orlando, FL, USA - 32816

2: Sandia National Laboratories

Scalable and Secure Systems Department

Livermore, CA

rizwan.ashraf@knights.ucf.edu, francis.luna@knights.ucf.edu, ddechev@sandia.gov, demara@ucf.edu

Keywords: FPGA, Synthesis, parallel genetic algorithms, evolvable hardware, evolutionary electronics

Abstract

Multicore processors are becoming common whereas current genetic algorithm-based implementation techniques for synthesizing Field Programmable Gate Array (FPGAs) circuits do not fully exploit this hardware trend. Genetic Algorithm (GA) based techniques are known to optimize multiple objectives, and automate the process of digital circuit design. In this paper, parallel GA algorithms are proposed for the synthesis of digital circuits for LUT-based FPGA architectures. Parallel modes of the GA such as Master-Slave and the Island model are compared to see which scheme results in better speedup and quicker convergence for effectively utilization of current multicore hardware. Speedup of about 5 over the sequential single-threaded implementation are achieved with both the schemes on six-core machine. Convergence is also found in fewer number of generations. The methods described here-in can be employed in evolvable hardware systems as well as FPGA CAD tools.

1. INTRODUCTION

The realm of developing electronic circuits via techniques based on evolutionary principles such as Genetic Algorithms is referred to as *Evolutionary Electronics* [22]. Electronic circuits such as amplifiers, analog and digital filters, digital circuits (e.g. combinatorial arithmetic circuits, parity circuits, sequential circuits, etc.) have been synthesized using such algorithms [1],[11],[22].¹ An evolutionary algorithm based design approach is an excellent tool for optimizing human-generated designs or synthesizing designs which meet multiple objectives of power constraints, size constraints in terms of number of gates required, and timing constraints in terms of circuit delay [6]. Essentially it automates the process of developing electronic circuits and is characterized by

its ability to search complex solution space. Further, electronic circuits for reconfigurable hardware such as Field Programmable Analog Arrays (FPAs) and Field Programmable Gate Arrays (FPGAs) can also be configured automatically [14], [19], [20]. This field, characterized by the use of reconfiguration techniques for hardware based on Genetic Algorithms, is known as *Evolvable Hardware* [21]. Adaptive systems can be realized using such techniques, which can adjust based on dynamics of the operating environment e.g. tolerate a failure by self-configuring a fault-tolerant design. The focus of this work is in the development of configurations for FPGAs.

Reconfigurable hardware such as FPGAs is an excellent platform for the application of the above mentioned techniques. A FPGA can be used to implement any given digital circuit and it can be reconfigured in runtime by using its dynamic reconfiguration feature [20], thus serving as an ideal platform for Evolvable Hardware systems. Its architecture is composed of reconfigurable logic and interconnect elements. Reconfigurable logic elements are based on SRAM-based Lookup Tables (LUTs). Multiple LUTs can be connected via programmable interconnects to implement a desired digital circuit. Adaptive systems based on this platform can utilize on-chip PowerPCs or use dedicated implementation in hardware to efficiently implement the genetic algorithm [5],[4]. In such systems, the performance of the genetic algorithm matters in terms of the time required to converge to a solution. The current trend in computing is pushing towards the use of multicore technology, which we intend to exploit in our implementation. Also with the introduction of FPGAs which tightly integrate on-chip multi-core processor with the reconfigurable logic [2] - there is a need to effectively utilize these parallel processing units for the above mentioned systems. Further, with the introduction of commodity multicore processors, the proposed technique can be effectively utilized for implementation in VLSI CAD tools [3].

Genetic algorithm based applications can be effectively parallelized to achieve significant speedups and to better utilize parallel processing units. Various attempts have been made to classify different models of parallel genetic algo-

¹Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

rithms [15],[4]. According to [15], genetic algorithms can be parallelized based on how fitness is evaluated, whether single or multiple populations are used, and whether GA operations like crossover and selection are performed locally in a sub-population or globally in the entire population. The authors also suggest that parallel genetic algorithms may perform better in terms of the solution found as compared to their sequential counterparts which may get trapped in the sub-optimal portion of the solution space. Thus, they have a better probability of getting out of this local optima as multiple sub-populations will tend to explore multiple portions of the search space. In this paper, the synchronous master-slave and island model of the parallel genetic algorithm as introduced in [15] are compared and contrasted for the proposed problem. We adopt these two models in our implementation.

The main contribution of this work is to design digital circuits for FPGA-based architectures using parallel genetic algorithms (GAs). The GA employed involves the use of a linear representation which can be readily employed for intrinsic evolution systems such as through direct manipulation of FPGA configuration bitstream as proposed in [16].

The paper is organized as follows. Section 2 describes the proposed methodology to parallelize the genetic algorithm and introduces how such techniques are used for the design of digital circuits for FPGAs. Section 3 presents the related works and highlights the unique contributions of this work. The sequential and parallel implementations are illustrated and contrasted in section 4. Experimental setup and the results are presented in section 5. The conclusions of this work are presented in section 6.

2. ALGORITHM

A genetic algorithm (GA) mimics evolutionary principles by maintaining multiple candidate solutions in the form of a *population*. Each candidate solution, referred to as an *individual*, describes a potential FPGA configuration and is initially generated purely randomly. The representation used will be described later. Each individual is ranked based on a *fitness* value, which is assigned according to the *fitness function*. The fitness function is used to quantify the correctness of a given candidate configuration. In this case it is described via the truth table of the desired digital function to be implemented on the FPGA. The genetic algorithm performs the operations of *crossover* and *mutation* on individuals according to user-specified probabilities, with the intent to increase the fitness of individuals. After application of these operators, the population for the next generation (iteration) is selected based on some selection scheme. The *selection* scheme is designed to guarantee the "survival of the fittest" i.e. the most fit individuals make it to the next generation. *Tournament-based* selection is chosen for this implementation, as described later. A constant-size population is maintained in this work (*finite-*

population GA). This process continues over multiple generations until an individual is found with the required *threshold fitness* or maximum number of generations t_{max} are achieved according to the described *exit criteria*. The sequential version of the algorithm as described above is summarized in Algorithm 1.

Algorithm 1 Genetic Algorithm

```

1:  $t := 0$ ;
2: Initialize Population  $P(0)$ ;
3: repeat
4:   Apply GA Operators (Mutation, Crossover)  $P'(t)$ ;
5:   Fitness Evaluation  $P'(t)$ ;
6:   Create new Population via Tournament-based Selection  $P(t+1)$ ;
7:    $t := t + 1$ ;
8: until Exit criteria (max Fitness achieved OR max Number of Generations  $t_{max}$  reached)

```

An individual, which models a FPGA configuration, is represented as shown in figure 1 and is adopted from [14]. The representation contains multiple Lookup Tables (LUTs) as shown. A fixed number of LUTs are selected for each design configuration depending on the complexity of the desired digital function. Each LUT has a function and four inputs associated with it. Each input of a LUT can be one of the inputs of the circuit or the outputs from one of the other preceding LUTs in the configuration e.g. LUT_α in the configuration $\in \{1, \dots, n, n+1, \dots, n+\alpha-1\}$ where n represents the inputs of the circuit whereas subsequent integers represents the LUTs in the configuration until LUT_α . This representation ensures feed-forward condition of the design is maintained, as only combinatorial circuits are evolved in this work. Each individual also encodes the outputs of the circuit. The outputs can be from any of the circuit inputs or from the outputs of the multiple LUTs in the FPGA configuration i.e. $output_\beta \in \{1, \dots, n, n+1, \dots, n+x\}$ where x is the maximum number of LUTs in the configuration.

The *fitness* of a candidate FPGA configuration which quantifies its correctness is calculated by exhaustively applying all the possible input combinations e.g. 8 test vectors are applied for a circuit with 3 inputs. For each input combination, the output(s) of the candidate individual is (are) evaluated and compared with the original required output(s) as described in the truth table of the digital function to be implemented. Bit-wise comparison is done in this case, incrementing the fitness value with each output line match. This is accumulated over all possible input combinations. As an example a circuit with n inputs and m outputs will have a maximum fitness of $2^n * m$ which indicates that every output line of the circuit matches the desired output for every possible input combination. Other fitness functions are also possible e.g. by taking the arithmetic difference between the desired and actual output(s) accumulated over all possible input combinations and which involve selective test vectors evaluation instead of exhaustive evaluation of all possible input combinations.




Figure 1. An individual: The representation of FPGA configuration in GA

The notion of population in genetic algorithms makes them an ideal candidate for parallelization. The classification of parallel versions of the GA into synchronous master-slave and island model is adopted in this work. The main classification is based on how the population is partitioned for various phases of the algorithm. For both classes, the fitness evaluation is done in parallel for a group of individuals. The GA operation of crossover is a binary operator thus it requires two individuals to be selected from the population at a user-specified *crossover-rate* as opposed to the mutation operator which requires only one individual. Similarly, the selection phase requires other individuals from the population e.g. tournament-based selection requires a pool of individuals to perform the competition based on fitness where most-fit individuals are selected for next generation. The master-slave employs a single population for its various operations as opposed to the island model which employs multiple independent sub-populations (demes). For example, the two individuals for the crossover operation can be randomly selected from the complete population in case of master-slave whereas in case of island model they are selected from within a sub-population (locally). Similarly, the pool of individuals (the size of which is user-specified) required for the tournament selection is formed from the complete population in case of master-slave whereas in case of the island model it is selected from within a sub-population. The selection pressure in case of island model is different as compared to the master-slave model as individuals in the master-slave model compete against the complete population. A single best individual is maintained for the whole population in case of master-slave model as opposed to the island model. The sub-populations involved in the island model each maintain their own best fit individuals and migration of best-fit individuals only take place if the best fit individual of a sub-population has a lower fitness value than any other sub-population's best fit individual. The frequency of these migrations may be as high as every generation. The salient features of the two schemes are highlighted in Table 1.

	Sync Master-Slave	Island Model
Number of Population(s)	Single	Multiple ρ
Number of Individuals in a Population	N	$\frac{N}{\rho}$
Crossover Operation	Global Population	Local Population
Selection Operation	Global Population	Local Population

Table 1. Highlights of the main differences among the Island and the Master-Slave parallel models with ρ parallel processors

3. RELATED WORK

Parallel Genetic Algorithms have been used for the design of digital [8] and analog circuits [1],[11]. GAs in general are known to solve multi-objective optimization problems and thus can be efficiently employed in the field of electronic circuit design [10] where often multiple constraints such as power, size, and timing have to be accomplished. It can also be used to automate the process of circuit synthesis which has useful applications in the field of evolvable hardware, where a design may need to be adapted at runtime to meet new requirements [7] or produce diverse designs which can be employed in a fault-tolerant system as proposed in [19]. In addition, GAs are also employed in VLSI layout tools for optimal placement and routing [13]. In this work, we focus on the design of digital circuits for the FPGA platform.

Design of analog circuits such as amplifiers and filters has been the focus of most works which employ parallel genetic algorithm [1],[11], as the fitness evaluation in this case is considered to be the most computationally expensive part. Synchronous master-slave implementation is employed in [12] and the workload of fitness evaluation is distributed off to slave nodes; fine-grain partitioning is done and a small number of configurations are evaluated by slave nodes at any given time, which potentially results in increased stall times. The experimentation for the above mentioned work is performed on a Beowulf cluster. Similarly, GA is parallelized in [1] to evolve simple VLSI circuits; coarse-grain parallelism is done using a shared memory programming model. Different implementations with a centralized single population and multiple distributed populations are done for the parallel GA in this case. The distributed scheme proposed in [1] involves infrequent communication among populations, whereas the island model proposed in our work involves communication at every generation and a possible migration is done if the sub-population has less fit individuals as compared to other competing sub-populations. Nearly linear speedups are reported for the different implementations in [1], though the results are limited and thus it is difficult to establish the generality of this work, e.g. speedups are only reported with 16 parallel computing units. Whereas, in our work, extensive experimentation with different number of threads is done, using both the synchronous master-slave and island model of the parallel GA.

In [8], digital circuits are designed using multi-expression programming. The representation employed in [8] is a variation of linear genetic programming, whereas our work employs a representation which is targeted for FPGA-based architectures. Asynchronous island model is employed in [8], where sub-populations are maintained on parallel machines which exchange individuals after a certain defined period (design time). The idea is to evolve multiple genetic programs in parallel on multiple processors. The computing nodes are connected in a ring topology and the message passing programming model is employed for communication between different nodes. Results show a considerable decrease in computational effort as compared to the non-parallel GA. In our work, as mentioned earlier, FPGA-based architectures are targeted, and comprehensive comparison is done between the synchronous master-slave and the island model. The island model employed in our work is different as compared to [8], as the best individual in our work may be communicated to sub-populations on every generation. Further, most of the implementations have been done on clusters of computers whilst using the MPI programming model [8],[12] and look to exploit the characteristics of a group of individuals by creating sub-populations which are evolved independently for many generations (this number is generally fixed) with little or no communication. Whereas, we have proposed the parallel models of the GA on a shared memory machine for targeting todays multicore processors.

4. IMPLEMENTATION

4.1. Sequential Implementation

The following sections present a bottom-up overview of the classes and functions that form the sequential program. These parts are reviewed in this section so as to see which portions could be made parallel and to see which parts might be problematic when they are made to work in parallel. We start from the smallest unit, the LUT object, and work our way up to the largest, a Generation. Finally, we review the main function to see how the algorithm works in this implementation.

4.1.1. LUT class

Each LUT object contains three vectors and its function type $\in \{NOT, AND, NAND, OR, NOR, XOR\}$. The vectors contain information about which LUT outputs are connected to the LUT inputs, the input binary values of these connections into the LUT, and the output value from the LUT.

4.1.2. CLB class

Each Configurable Logic Block (CLB) object contains multiple LUT objects. This class contains a vector of LUTs (i.e. $\{LUT_1, LUT_2, \dots, LUT_{CLB_max}\}$, where CLB_max is the number of LUT objects in a CLB). The size of this vector can be set by the user.

4.1.3. Individual class

Each individual (circuit) contains four vectors, a fitness value, and various functions. The first vector is of CLBs that the circuit uses. The other three are the inputs, outputs, and connections of the circuit. The main function from this class is the `CalculateFitness` function which goes through each LUT for each CLB and calculates each output value. It then compares these outputs with the expected value and assigns a fitness value to the individual by incrementing it's fitness, starting at zero, for every output that is correct.

4.1.4. Generation class

The generation object consists of a vector of Individuals, the generation number, and an index to keep track of the Individual having the maximum fitness. It also contains the functions for the genetic algorithm.

The `PerformCrossover` function, according to the crossover rate (probability of performing crossover), takes two individuals and randomly picks a crossover point, such that the boundaries of LUT objects are not violated. The LUT configurations before this point on Parent A and after this point on Parent B are copied to the offspring as shown in fig 2. If no crossover is to take place, the individual is just copied to the other individual.




Figure 2. The application of crossover operator between two individuals.

The `PerformMutation` function performs three types of mutation on one individual. The first type is functionality mutation where it takes the individual and for each selected LUT based on a user-defined mutation rate, randomly changes its function. The second type is interconnection mutation where it takes each selected LUT of the individual and randomly changes its input connections. Finally, it performs output line mutation where the output lines of the individual

are randomly assigned to either the inputs of the circuit or any of the LUTs in the configuration.

The Selection function randomly selects k (Tournament Size) individuals from both the parent population (current generation) and the offspring population (evolved individuals). From this selection, the best fit individual is picked to move on to the next generation. This is repeated until a new generation of the same size is formed.

The PerformElitism function copies the individual with maximum fitness in order to preserve it and maintain forward progress.

4.1.5. Main-The GA loop

The main function performs all the steps necessary for the implementation of the genetic algorithm. Firstly, individuals from the generation are chosen to have the genetic operators performed on them based on a user-defined probability. None of these operators directly modify the parent population; they simply modify a copy of the individual (offspring) which is then stored. For crossover, two random individuals are chosen from the parent population for crossover to perform on. After crossover and mutation take place, the fitness for the new offspring is calculated. This is repeated until all the individuals of the population have been evolved. Then selection is performed where individuals from the parent and offspring populations are chosen to form the new population. Then Elitism takes place. If there exists an individual with higher fitness than the elite individual, then that individual becomes the new elite and it replaces the old elite. If the highest fitness is less than the fitness of the elite individual, the elite individual is copied to the population, replacing a randomly chosen individual. The generation number is incremented and this procedure is then repeated until the desired fitness level is achieved. The program flow is outlined in fig 3.

4.2. Parallel Implementation

Understanding how the sequential program works allows us to parallelize the portions of it that would be most beneficial to the overall runtime. Two different models were used in the parallel versions of the program. The first, the master-slave model, has one master thread that calls the parallel functions. These functions operate on the entire population by partitioning the full population into sub-populations that each thread can operate on. The second, the island model, partitions the population in smaller sub-populations at the beginning of the GA loop. The functions used here are not parallel versions, they just operate on a smaller portion of the population. Each thread iterates through one instance of the GA loop independently of one another. These two models were used to determine if one achieved better speedup and performance than the other.




Figure 3. Program flow.




Figure 4. A generation with a population size of N

4.2.1. Master-Slave model

The portions that are made parallel in the master-slave model are within the GA loop, where the program spends most of its time running. More specifically the parts we parallelize are: the portion where genetic operators are performed, the Selection routine, and the portion that updates the maximum fitness. We did not want to change how the algorithm functioned and therefore only divide the population into smaller sub-populations that each thread will work on as

shown in fig. 4. These populations are divided evenly when genetic operators are performed, then combined once it is completed, divided once again when selection is performed, then combined, then divided again to search for the individual with highest fitness, updating the global value when a higher fit individual is found. This layout is shown in fig 5.




Figure 5. Master-Slave Population partitioning.

4.2.2. Island model

The GA loop is partitioned and made to run in parallel in the island model of the program. Each island (thread) calls the GA operators, selection, and determines the maximum fitness on their own sub-population. If any of the islands have a higher fit individual than the elite, that individual becomes the new elite and it gets distributed to the other islands at the end of a generation. This layout is shown in fig 6.




Figure 6. Island Population partitioning.

4.2.3. Library

A library that provided the following had to be chosen for the parallel versions of the program:

- concurrent vectors
- parallel loops
- locks

Since the sequential implementation used vectors throughout its implementation we needed to find a concurrent vector implementation that multiple threads can safely access at one time. We also needed an easy construct to parallelize the portions highlighted as well as an implementation of a *lock* which can be used to find the global max fitness. After some research we settled on Intel Threaded Building Blocks (TBB) [18] because it contained everything needed in an easy to use library. We have used constructs of `parallel_for`, `concurrent_vector`, and `mutex` for our parallel implementations.

The TBB `concurrent_vector` will replace any STL vector that will be accessed by more than one thread at a time. It guarantees that elements in the vector will never move until it is cleared which is needed in this implementation.

The TBB `spin_mutex` will be used to place a lock on the maximum calculated fitness and which individual it is. In order to reduce contention of the lock, it will only be grabbed when the new fitness of an individual is better than the current individual with maximum fitness. The `spin_mutex` version of a lock was chosen because the amount of time that the lock will be held is relatively short. Only two variables are required to be updated so there are very few operations that will be performed in the critical section. The overhead for other lock implementations will not be beneficial in our design. Also, this lock is not always acquired so contention will be low. It is only acquired when there is a higher max fitness than the one stored. After the lock is acquired, it checks this condition again to ensure no one has changed this value since before the lock was acquired and updates the values accordingly. This method of checking whether to grab the lock first, and then checking the conditions again, will greatly decrease the contention on that lock since it will only be acquired when needed as opposed to always grabbing the lock and then doing the comparison. This technique is mentioned in [9]. The following is an example from the code:

```

if (MaxCalculatedFitness < Individual(x)_Fitness)
{
    FitnessMutex.lock();
    if (MaxCalculatedFitness < Individual(x)_Fitness)
    {
        MaxCalculatedFitness = Individual(x)_Fitness;
        iMaxIndivi = x;
    }
    FitnessMutex.unlock();
}
  
```

5. EXPERIMENTS AND RESULTS

5.1. Experimental Setup

Experiments are done to evaluate the speedup and performance of the parallel implementations as compared to the sequential (single-thread) implementation as used in [17]. The genetic algorithm is supposed to start-off with completely random configurations as described earlier. Further, the GA operators are applied to configurations based on user-defined values. Similarly, the selection operation is performed on a pool (tournament size). We keep track of timestamps throughout the program in order to compute the speedups of the portions we parallelized and the entire program.

Experiments were conducted with population sizes of 120, 240, 480, and 960. The effect of population sizes on the speedups is to be observed. Each test was run 20 times for averaging. The objective of the experiments was to realize a 3-to-8 decoder configuration for a LUT-based FPGA architecture. Mutation and crossover rates were set to 0.007 and 0.60 for all the experiments and for both models. A fixed tournament size of 6 is used for all population size experiments. The experiments are run for a fixed number of generations (1000) to calculate the speedups. To measure GA performance, each test is run until a maximum fit individual is found and the number of generations it took to find this individual is recorded.

All tests were run on an Intel Xeon X5670 (6 cores) with 6GB RAM running Ubuntu 11.04 with Intel TBB version 3 update 5 [18]. The speedup and performance of the implementations are calculated with 2, 4, 6, 8, 12 and 16 threads. This effectively partitions the population size by the thread count into smaller subpopulations that each model uses.

5.2. Results

5.2.1. Speedup

Speedup over the sequential version of the algorithm was achieved, however the problem is observed to not be perfectly parallelizable. Figure 7 details the speedup achieved for the master-slave model with respect to varying thread count. The highest speedup achieved by the master-slave model is 5.01 with a population size of 960 running on 6 threads. The speedups for this model are fairly linear with the threshold being the physical number of cores the threads can run on, which in our setup was 6. After this limit is reached, the speedup is decreased as thread count increases. Also, larger problem sizes achieve higher speedup because the amount of computation in each thread outweighs the amount of communication in these cases. Figure 8 details the speedup achieved for the island model with respect to varying thread count. The highest speedup achieved by the island model is 5.04 with a population size of 960 running on 16 threads. Although the number of threads is larger than the number of available cores, a higher speedup is achieved than the master-

slave model. The speedup of the smaller populations begin to saturate while the speedup of larger populations continue to increase as thread count increases. This is unlike the master-slave model where speedups start to decrease after a certain threshold. Also, the speedups of all the population sizes are fairly close to one another with less variance between the population sizes than for the master-slave model.




Figure 7. Speedups achieved with sync master-slave model




Figure 8. Speedups achieved with Island model

5.2.2. GA Performance

Both models converged to find a perfect fit individual in all the runs within the maximum number of generations set as threshold for the experiments. The island model converged in fewer generations than the master-slave model for all population sizes as shown in figure 9.

6. CONCLUSION AND FUTURE WORK

Two models for parallelizing the genetic algorithm used for realizing configurations for LUT-based FPGA architectures were successfully realized. Results indicate speedups of approximately 5 are achieved for both the parallel modes of the genetic algorithm on a machine with six physical cores. In addition to achieving speedup over the sequential implementation, it was observed that using the island model for parallelizing this problem actually allowed the genetic algorithm to converge and find a maximum fit individual in fewer




Figure 9. Comparison of Convergence Properties of Master-Slave and Island Models

generations than the master-slave model. This results in the genetic algorithm running for less time.

Other areas of the algorithm could be explored in order to optimize them and make the parallel performance better. One such area is the fitness calculation which is inherently sequential, and could be improved by evaluating independent test vectors in parallel. Various parameter settings can also be tried to improve GA convergence properties.

REFERENCES

- [1] M. Davis, L. Liu, and J. Elias. VLSI circuit synthesis using a parallel genetic algorithm. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 104–109. IEEE, 1994.
- [2] K. DeHaven. Extensible Processing Platform Ideal Solution for a Wide Range of Embedded Systems. Technical report, Xilinx, 04 2010.
- [3] R. Drechsler. *Evolutionary algorithms for VLSI CAD*. Kluwer Academic Publishers, 1998.
- [4] S. E. Eklund. A massively parallel architecture for distributed genetic algorithms. *Parallel Comput.*, 30:647–676, May 2004.
- [5] P. Fernando, S. Katkoori, D. Keymeulen, R. Zebulum, and A. Stoica. Customizable fpga ip core implementation of a general-purpose genetic algorithm engine. *Evolutionary Computation, IEEE Transactions on*, 14(1):133 –149, feb. 2010.
- [6] F. Ferrandi, P. Lanzi, G. Palermo, C. Pilato, D. Sciuto, and A. Tumeo. An evolutionary approach to area-time optimization of fpga designs. In *Embedded Computer Systems: Architectures, Modeling and Simulation, 2007. IC-SAMOS 2007. International Conference on*, pages 145 –152, july 2007.
- [7] P. Haddow and G. Tufte. An evolvable hardware fpga for adaptive hardware. In *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, volume 1, pages 553 –560 vol.1, 2000.
- [8] F. Hadjam, C. Moraga, and M. Benmohamed. Cluster-based evolutionary design of digital circuits using all improved multi-expression programming. In *Proceedings of the 2007 GECCO conference companion on Genetic and evolutionary computation*, pages 2475–2482. ACM, 2007.
- [9] M. Heinrich and M. Chaudhuri. Ocean warning: avoid drowning. *SIGARCH Comput. Archit. News*, 31:30–32, June 2003.
- [10] T. Kalanova and J. Miller. Evolving more efficient digital circuits by allowing circuit layout evolution and multi-objective fitness. In *Evolvable Hardware, 1999. Proceedings of the First NASA/DoD Workshop on*, pages 54 –63, 1999.
- [11] J. Lohn, G. Haith, S. Colombano, and D. Stassinopoulos. Towards evolving electronic circuits for autonomous space applications. In *Aerospace Conference Proceedings, 2000 IEEE*, volume 5, pages 473–486. IEEE, 2000.
- [12] J. D. Lohn, S. P. Colombano, G. L. Haith, and D. Stassinopoulos. A Parallel Genetic Algorithm for Automated Electronic Circuit Design. Technical report, NASA, 2000.
- [13] P. Mazumder and E. M. Rudnick. *Genetic algorithms for VLSI design, layout & test automation*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1999.
- [14] J. F. Miller, D. Job, and V. K. Vassilev. Principles in the evolutionary design of digital circuits. *Genetic Programming and Evolvable Machines*, 1:7–35.
- [15] M. Nowostawski and R. Poli. Parallel genetic algorithm taxonomy. In L. C. Jain, editor, *KES*, pages 88–92. IEEE, 1999.
- [16] R. Oreifej, R. Al-Haddad, H. Tan, and R. DeMara. Layered approach to intrinsic evolvable hardware using direct bitstream manipulation of virtex ii pro devices. In *Field Programmable Logic and Applications, 2007. FPL 2007. International Conference on*, pages 299 –304, aug. 2007.
- [17] R. S. Oreifej, C. A. Sharma, and R. F. DeMara. Expediting ga-based evolution using group testing techniques for reconfigurable hardware. In *Reconfigurable Computing and FPGA's, 2006. ReConfig 2006. IEEE International Conference on*, pages 1 –8, sept. 2006.
- [18] J. Reinders. Intel threading building blocks, 2007.
- [19] A. Stoica, R. Zebulum, D. Keymeulen, R. Tawel, T. Daud, and A. Thakoor. Reconfigurable VLSI architectures for evolvable hardware: from experimental field programmable transistor arrays to evolution-oriented chips. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 9(1):227–232, 2001.
- [20] A. Upegui and E. Sanchez. Evolving hardware by dynamically reconfiguring xilinx fpgas. In *ICES'05*, pages 56–65, 2005.
- [21] X. Yao and T. Higuchi. Promises and challenges of evolvable hardware. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 29(1):87 –97, feb 1999.
- [22] R. S. Zebulum, M. A. C. Pacheco, and M. M. B. R. Vellasco. *Evolutionary Electronics: Automatic Design of Electronic Circuits and Systems by Genetic Algorithms*, volume 1. CRC Press, 2002.