The Design of an SRAM-Based Field-Programmable Gate Array, Part I: Architecture

Paul Chow, Soon Ong Seo, Jonathan Rose, Kevin Chung, Gerard Páez-Monzón and Immanuel Rahardja

Abstract—**Field-Programmable Gate Arrays (FPGAs) are now widely used for the implementation of digital systems and many commercial architectures are available. Although the literature and data books contain detailed descriptions of these architectures, there is very little information on how the high-level architecture was chosen and no information on the circuit-level or physical design of the devices. This paper, Part I, describes the high-level architectural design of an SRAM-programmable FPGA. Part II will address the circuit design issues through to the physical layout.**

The logic block and routing architecture of the FPGA was determined through experimentation with benchmark circuits and custom-built CAD tools. The resulting logic block is an asymmetric tree of four-input lookup tables that are hard-wired together and a segmented routing architecture with a carefully chosen segment length distribution.

Keywords— **field-programmable gate arrays, FPGA, FPGA architecture, SRAM programmable**

I. Introduction

FIELD -Programmable Gate Array (FPGA) technology permits the design of many different complex digital circuits using a single off-the-shelf device [1]. The time-to-market pressures and low financial risk has made FPGAs and Complex Programmable Logic Devices (CPLDs) an increasingly popular vehicle for prototyping, and in many cases, actual production. There are many different architectures now available from over fifteen silicon vendors. Although the databooks and related publications usually describe the architecture in detail, there is little information on *how* the architecture was chosen, and no information on the circuit-level or physical layout level of the design. The contribution of this paper is to describe how a suitable architecture was chosen for an SRAM-programmable FPGA. A subsequent paper [2] examines the circuit and layout issues encountered when implementing that FPGA. Our primary goal was to produce a high-speed architecture and circuit with a reasonable logic density.

FPGAs were first introduced in 1986 by Xilinx using a memory-based programming technology [3]. Since then there have been many new commercial architectures [4], [5], [6], [7], [8] and several new programming technologies including two types of anti-fuse [9], [10], [11] and floating gate transistors, which are UV and electrically erasable [12], [13], [14].

A few non-commercial FPGA architectures have been reported for which the design details are more readily available. The Triptych FPGA [15], [16] matches the physical structure of the routing architecture to the fanin/fanout nature of the structure of digital logic by using short connections to the nearest neighbours. Segmented routing channels are used between the columns to provide for nets with fanout greater than one. No discussion is given about how the segmentation length distribution was selected. This routing architecture does not allow the arbitrary point-to-point routing available in general FPGA structures. The logic block implements logical functions using a multiplexer-based three-input lookup table followed by a master-slave D-latch, and can also be used for routing. Initial results show potential implementation efficiencies in terms of area using this structure. The Montage FPGA [17], [16] is a version of the Triptych architecture modified to support asynchronous circuits and interfacing separately-clocked synchronous circuits. This is achieved by the addition of an arbiter unit and a clocking scheme that allows two possible clocks or makes the latches transparent.

Earlier work at the University of Toronto resulted in the implementation of an architecture (UTFPGA1) using three cascaded four-input logic blocks and segmented routing [18]. UTFPGA1 used information from previous architectural studies but there was very little transistor-level optimization (for speed), and little time was spent on layout optimization. This was a first attempt that provided some insight into the problems faced in the design and layout of an FPGA. An earlier version of the work presented here appeared in [19].

Modern trends in computer architecture have shown the importance of considering both the compiler technology and the hardware technology at the same time, when designing for high performance [20]. This is also true in the design of an FPGA, where it is important that the CAD tools collaborate with the architecture of the FPGA. At the University of Toronto, we have developed a number of custom CAD tools that have been used to perform *experimental* architectural studies. In this paper, we will describe two important architectural innovations that lead to higher-speed FPGAs, and influenced the architecture of our second-generation FPGA. In Part II [2], we will

Paul Chow and Jonathan Rose are with the Department of Electrical and Computer Engineering at the University of Toronto. Soon Seo is now with ATI Technologies. Kevin Chung is now with Xilinx, Inc. Gerard Páez-Monzón is with CEMISID-Universidad de Los Andes-Venezuela and Immanuel Rahardja is now with Aristo Technology Inc.

Paul Chow can be reached at pc@eecg.toronto.edu

also introduce a novel layout style for FPGAs that builds the basic tile using a set of identical mini-tiles and customization by the addition of vias. This allows us to still achieve reasonable density while significantly reducing the time spent doing custom layout. These features have all been implemented in our test chip, called LEGO (**L**ogic that's **E**rasable and **G**reatly **O**ptimized).

In Section II we show how various hard-wired interconnect topologies between lookup tables can form larger logic blocks and influence delay, and how to select an appropriate topology. In Section III, the routing architecture is derived and we show that longer dedicated wire lengths in the routing architecture are beneficial. Section IV provides a summary of our results and some conclusions.

II. Logic Block Architecture

The speed and density gap between FPGAs and mask-programmable gate arrays (MPGAs) is mostly due to the routing structures used to connect logic components in each technology. In MPGAs, the logic elements are connected with mask-programmed metal wires. In FPGAs, logic block pins are connected using field-programmable switches. Regardless of the type of programmable switch used in the FPGA (e.g. static RAM-controlled pass transistors [4] or anti-fuses [21], or floating-gate based switching [14]), the capacitance, resistance, and size of programmable connections makes them much slower and larger than a simple metal wire.

One way to improve the speed and density of FPGAs is to replace some of the programmable connections between basic logic blocks with hard-wired connections, which are simple metal wires [22]. By using hard-wired links to construct more coarse-grained logic blocks (called Hard-wired Logic Blocks, or HLBs) from several basic blocks, the delay and size of circuits can be reduced. For example, Figure 1(a) shows a network of basic blocks with five programmable connections in the routing along the critical path (which we define as the combinational path with largest number of logic levels) from block 1 to block 5, and nine programmable connections in total. Suppose that three of the basic blocks are hard-wired together to create a hard-wired logic block with a fast three-block path as shown in Figure 1(b). The resulting circuit in Figure 1(c) has only two slow programmable links instead of five along the critical path. This is a sizable reduction in routing delay. Also, the total number of programmable connections has been reduced from nine to four and this may lead to a reduction in routing area.

A. The Hard-Wired Logic Block Design Space

Given this general notion of hard-wired logic blocks, it is clear that there are many possible ways to implement it. Assuming that there is only one type of basic block that is connected in a given HLB, the choices lie in the topology of the interconnection between the blocks and the size of the basic block itself.

Each choice will provide different delay and area for a given circuit. In general, the more hard-wired links in a hard-wired logic block, the faster will be the circuits implemented using that logic block. However, a greater number of hard-wired connections provide less connection flexibility and may lead to lower density because basic blocks are wasted. In this section we give a brief description of the experiments performed to determine a good choice for the HLB topology. We omit the work done to determine the choice of basic block, which the interested reader can find in [23], along with a more detailed description of these experiments.

The basic block we choose to work with is the 4-input lookup table (hereafter referred to as a 4-LUT) as previous studies have indicated that this is a reasonable choice for both speed and density [24], [25], [26], [27], [28] and as such has been adopted by several commercial vendors [4], [5], [6].

We restricted the choice of topology to be a tree, as circuits often are tree-like in nature. Figure 2 illustrates several of the topologies that were investigated. The naming convention of these structures is as follows: it begins with the letter "L", then the height of the HLB (in basic blocks), then a dash ("–") followed by a listing of sizes of the subtrees from a pre-order traversal of the canonical HLB tree. Each subtree size is separated by a dot (".") with the restriction that leaf inputs and single-LUT subtrees are not listed.

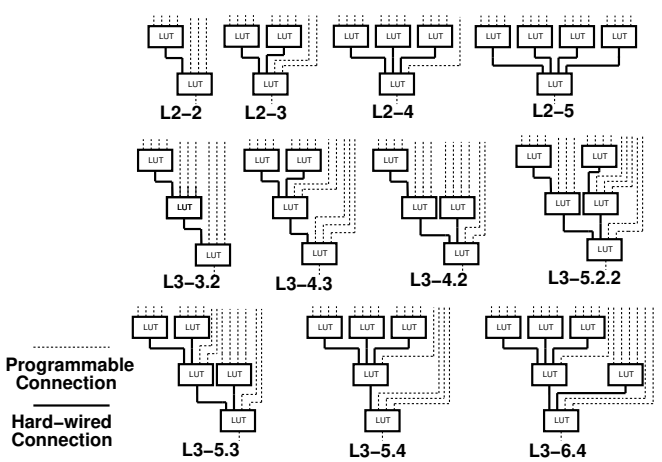

Fig. 2. Example of some of the tree topologies investigated.

In addition to the HLB topologies illustrated in Figure 2, all possible tree topologies with nine or fewer 4-LUTs and three or fewer LUT levels were explored. Initial experiments indicated that although HLBs with more than nine 4-LUTs provide greater speedups they require too much area to be considered practical.

An important architectural assumption is that each HLB has a buffer on the output of each LUT basic block that is accessible to the programmable routing. This direct access has two important advantages: delay is reduced because an output can be accessed without propagating it through downstream logic blocks, and density is increased because unrelated pieces of logic can be packed together in a multi-output HLB.

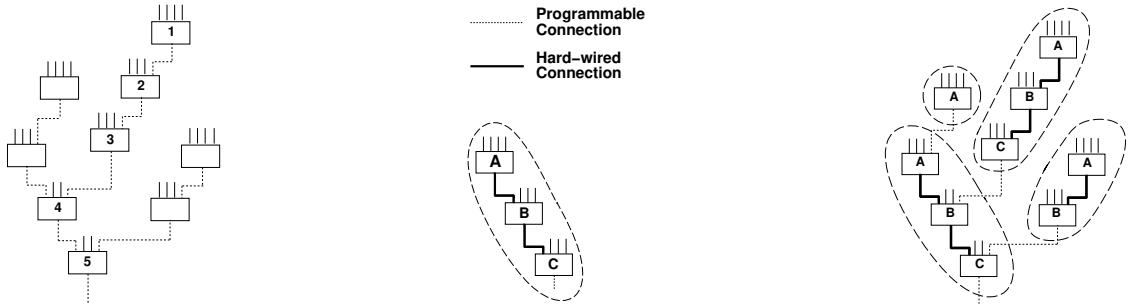

(a) Network of Basic Blocks (b) An HLB (c) Faster HLB Network

Fig. 1. Hard-wired connections and logic blocks.

The goal of the experiments described below is to select from among the many possible hard-wired 4-LUT architectures one that produces high system speeds, at a reasonable cost in area.

B. Experimental Method

To evaluate the HLB topologies, a set of 15 benchmark combinational circuits from the MCNC suite [29] were each "implemented" as a set of FPGAs with each HLB topology. The area and delay of each implemented circuit is then calculated using area and delay models, and the results are averaged over all circuits for each HLB topology. By "implementation" we mean that each circuit is processed through a suite of CAD tools that take it from the logic level (equations) through to the physical placement upon the hypothetical FPGA and the global routing on that FPGA. Note that a custom logic synthesis tool, called TEMPT, was developed to allow synthesis into the hard-wired structures [30].

The focus of this paper is not the details of these experiments but the results. As such the description of the area and delay models and the CAD tools needed to provide the implementation are omitted but can be found in [23]. The only detail necessary for the discussion below is the fact that the CAD implementation stream has two modes: one in which the resulting FPGA circuit is optimized for speed (sometimes at great cost in area), and the other in which the circuit is optimized for area at the cost of some speed. Our method of selecting an appropriate logic block for LEGO is to inspect the results of both experiments using both optimizations and to deduce a reasonable compromise.

C. Experimental Results

Table I provides the summarized results of the implementation of the benchmark circuits for those HLB topologies that appear the most promising. Table I gives the highest-speed (lowest critical path delay) topologies for each *size* of HLB, where size is measured by the number of 4-LUTs in the HLB, when the circuits are optimized for both speed and area.

The first column of Table I lists the number of 4-LUTs in the HLB. If there are two rows for a given number of 4-LUTs, the second entry is an alternative HLB that shows a slower speed but at a lower area cost. The second column gives the label of the HLB topology, as described above. The third and fifth columns give the normalized speedup of the HLB with respect to L1 (a 4-LUT FPGA with no hard-wired links), for the speed and area-optimized implementations. The speed is calculated as the geometric average of the inverse of the critical path delay for each circuit. It is then normalized with respect to L1. The fourth and sixth columns of Table I give the normalized area with respect to L1, for both the speed and area cases. The un-normalized absolute values for speed and area are given in the L1 row, assuming a $1.2\mu m$ CMOS process. Absolute areas come from actual layouts of cells used in an area model. Absolute times come from a delay model, which incorporates delay data derived from the simulation of the cell layouts.

The speed-optimized results show that the hard-wired links can increase system speed from 14% up to 50% compared to a flat non-hard-wired 4-LUT when the circuits are speed optimized. However, the area cost of this speed is an increase of 19% to 70% in circuit size. Some of this increase is caused by wasteful logic synthesis tools that sacrifice area for speed, but it also comes from the fundamental speed-area trade-off with hard-wired links as discussed above.

The area-optimized results show that several topologies are actually *more* area-efficient than the flat 4-input lookup table. All topologies provide a gain in speed, although less than that of the speed optimized case.

Our criteria in choosing a block is to select one that gives a good combination of speed and good density. From Table I, we reject those blocks that have too high an area penalty: those in the last four rows. Of the remaining blocks, L3-5.2, L2-5, L3-4.2, and L2-4 provide the best speed up. However, the L3-4.2 block proved superior to the others in the area-optimized case, and so it was selected. Of note is that this block provided the best logic density overall, while also providing significant speed up. The L3-4.2 HLB is illustrated in Figure 2.

It is tempting to use programmable multiplexers in place of the hard-wired connections to allow either the direct link or to access an input. We do not consider their use in this study, for two reasons: First, we are interested in finding the fastest topologies, and the multiplexers do not improve the speed; Secondly, each extra input to the logic block is

TABLE I
SPEED AND AREA OF HLBS-BASED FPGAS, SPEED-OPTIMIZED CIRCUITS

Number of LUTs in HLB	HLB Topology	Speed-Optimized Circuits		Area-Optimized Circuits	
		Speed Factor	Area Factor	Speed Factor	Area Factor
1	L1	1 (35 MHz)	1 (15.3 × 106 μm^2)	1 (22 MHz)	1 (6.7 × 106 μm^2)
2	L2-2	1.14	1.19	1.17	0.93
3	L2-3	1.26	1.21	1.17	0.96
4	L2-4	1.34	1.32	1.09	1.00
4	L3-4.2	1.28	1.31	1.27	0.93
5	L2-5	1.42	1.41	1.03	1.10
5	L3-5.2	1.37	1.37	1.24	1.04
6	L3-6.5	1.43	1.65	1.04	1.23
7	L3-7.3	1.43	1.66	1.19	1.28
8	L3-8.3.2	1.47	1.61	1.25	1.32
9	L3-9.4.2	1.50	1.70	1.24	1.42

expensive in terms of area, and part of the purpose of the hard-wired connection is to save that area.

III. ROUTING ARCHITECTURE

The FPGA routing architecture is the most important determining factor of system speed and logic density because the programmable switches, which are pass transistors driven by SRAM cells, have significant resistance and capacitance and require large area. The routing architecture is the manner in which wire segments and switches are organized.

In the original Xilinx 2000 and 3000 architectures [3], [31] a very simple architecture was employed: most wire segments spanned only the length (or width) of one logic block, and had switches at each end. One way to improve the speed of connections that travel long distances is to provide segments that span multiple logic blocks without being switched, as illustrated in Figure 3. In this way, a segment with appropriate length for a connection can be selected, which results in less resistance along the path. The general notion of segmented routing was first introduced in [9].

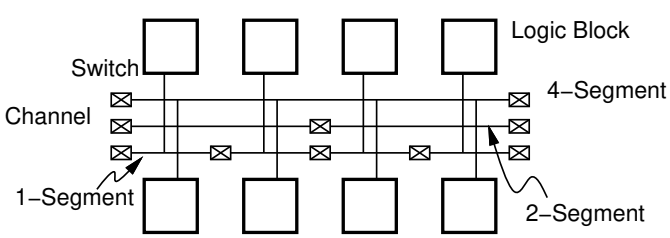

Fig. 3. Segmented Routing Architecture

A key architectural question is the determination of the number of segments of each length, called the *segmentation length distribution*. The greater number of longer segments, the greater the likelihood that long and otherwise time-consuming connections will be routed with fewer series switches, resulting in a faster circuit. An excessive number of longer segments, however, will mean that some of the long segments will have to be used for short connections, resulting in the waste of part of the segment. This waste leads to a decrease in logic density. Thus it is important to find a segmentation distribution that improves speed, but does not waste too much area.

There is a second important architectural feature of segmented routing architectures besides the distribution: a segment is called internally *populated* if it is possible to make connections from the middle of a segment to logic blocks or other routing segments. The advantage of unpopulated segments is that it has less parasitic switch capacitance connected to the segment, which makes it faster. The disadvantage is that the reduction in routing flexibility (without population there cannot be internal fanout) may result in the need for more tracks and so a loss of logic density.

In this section we briefly summarize a study on segmented routing architecture distribution and population that was used as the basis for decisions on the routing architecture for LEGO. The basic approach was experimental, similar to the one described in the previous section: several benchmark circuits are "implemented" as FPGAs, each with a different segmentation distribution. These experiments required several CAD tools, including a detailed router specifically designed for FPGAs [32]. The number of tracks needed to complete the route using segments of only length one is called the minimum. When the channel includes segments of other lengths, the number of tracks required above the minimum is measured.

A. Definitions and Experimental Method

Figure 4 illustrates the broad architectural features of the LEGO FPGA. The logic block (L) has pins that are connected to routing channels in a structure called a *connection block* (C). The *switch block* (S) provides connectivity among its attached channels. A channel consists of C and S blocks. Each channel contains W routing tracks.

We assume that each track consists of only one type of segment length. We will consider only segments of length one, two, and three, which will be referred to as 1-segments, 2-segments and 3-segments.

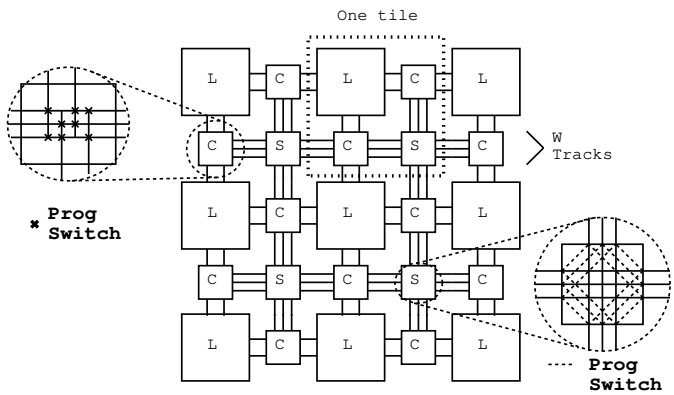

Fig. 4. Architectural Definitions

Let D_1 be the fraction of the W tracks that contain 1-segments. Similarly D_2 and D_3 are the fractions for 2- and 3-segments so that $D_1 + D_2 + D_3 = 1$.

Concerning the issue of the internal population of the segments, we divide it into two parts: whether or not the connection blocks internal to the segment (i.e., those not at its ends) should be populated, and whether or not the internal switch blocks should be populated. The following notation indicates the level of population: CB_u indicates that all unpopulatable connection blocks are unpopulated, while CB_p means that they are populated. Similarly, for switch block population we use SB_u and SB_p. Thus there are four combinations of population. Notice that for 2-segments, only the middle switch block can be depopulated, and for 3-segments, there are two switch blocks and one connection block that can be de-populated.

B. Experimental Results and Decisions

The following experiments were performed: for all possible values of D_1, D_2 and D_3, using all four population combinations, the number of excess tracks needed over the minimum (for $D_1 = 1$) were measured and averaged for five benchmark circuits. We addressed the questions of population and distribution using these experiments. Note that the logic block used in these experiments was a single 4-input LUT, with a D flip-flop. While this was not the block chosen for LEGO, we believe the general results discussed below will hold for most logic blocks.

B.1 Population of Segments

The first question we address is that of population of the S and C blocks. Figure 5 is a plot, for $D_2 = 0$, of the average number of excess tracks versus D1 for the five circuits. The four curves are the different combinations of population for the switch and connection block. Notice that as D_1 decreases from 1, D_3 increases from 0 ($D_3 = 1 - D_1$). Since this provides a greater number of longer segments with lower flexibility, we expect the number of tracks needed to successfully route will increase, and this is borne out in the figure.

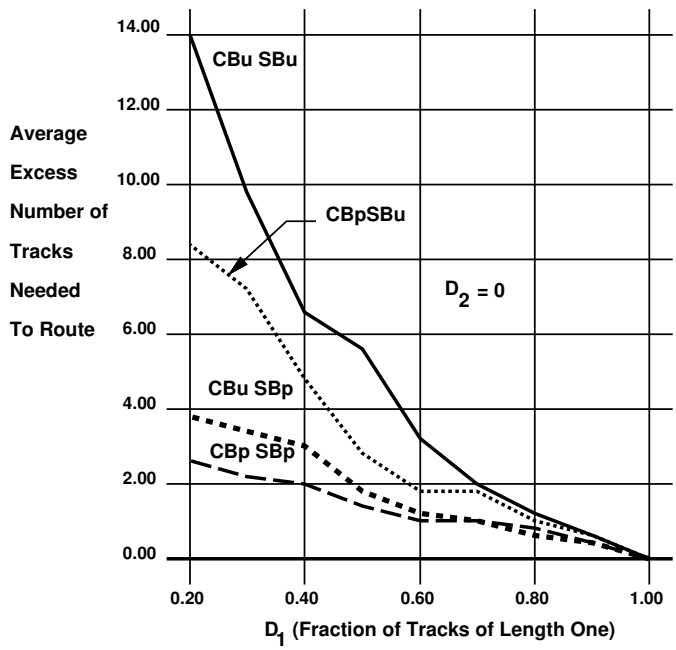

Fig. 5. Population Comparison

The data show that both cases in which the switch block is un-populated (the upper two curves) result in significantly more excess tracks. For this reason we decided to populate the switch block in LEGO. On the other hand, the bottom two curves (with the switch block populated) illustrate that only a minor increase in tracks is experienced when the connection block is de-populated. Since there is an advantage in speed for de-population (less capacitance on the track) we decided to de-populate the connection block in LEGO.

B.2 Segmentation Distribution

Table II gives the average number of excess tracks for many possible values of D_1, D_2 and D_3. Note that since $D_1 + D_2 + D_3 = 1$ there are only two independent variables, which we chose as D_2 and D_3 in the table. The rows vary D_2 and the columns vary D_3. This table is for architectures with both the C blocks un-populated and the S blocks populated, as discussed above.

Table II illustrates that even with 80% length three segments, that only about four extra tracks on average are required. (The minimum number of absolute tracks required for the five circuits ranged from 9 to 15.) We decided that we were willing to tolerate only about one extra track per channel above the minimum, so that the area cost of the higher speed tracks would be small. Thus architectures with D2 in the range of 0.0 to 0.6 and D3 in the range of 0.1 to 0.4 (those shown in boxes in Table II) would reflect such a choice. The faster architectures are those with higher values of D_3.

In Part II [2] we give the reasons that lead to the choice of 16 tracks per channel in LEGO. For the segmentation

TABLE II

Average excess tracks for all distributions, population CB_uSB_p. Absolute Average Number of Tracks = 12, Standard Deviation = 2.2.

D_2 \ D_3	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
0.0	0.0	0.4	0.6	1.0	1.2	1.8	3.0	3.4	3.8
0.1	0.2	0.6	1.0	1.2	1.6	2.2	2.8	3.4	
0.2	0.6	0.8	1.0	1.0	1.8	3.0	3.2		
0.3	0.6	0.6	1.2	1.8	2.2	2.4			
0.4	0.6	1.0	1.0	2.0	2.6				
0.5	0.8	1.2	1.4	2.0					
0.6	1.2	1.6	2.0						
0.7	1.4	1.8							
0.8	1.8								

distribution, there is an additional constraint that arises from the one-tile style of layout that was used: there must be an even number of 2-segment tracks, and the number of 3-segment tracks must be a multiple of three. Taking this constraint and the above distribution into account, the following segmentation distribution was used in LEGO: nine of the 16 tracks are 1-segments, four tracks are 2-segments, and three are 3-segments. This corresponds to a segmentation distribution in which $D_1 = 0.56$, $D_2 = 0.25$ and $D_3 = 0.19$, which fits within the architectural range indicated above.

IV. Conclusions

We have described the high-level architectural decisions used to select the logic block and routing architecture for the design of a high-performance FPGA. We developed the architecture using an experimental process in which custom-built CAD tools are used to implement benchmark circuits on candidate architectures. This method has the danger that the tools may have unfair algorithmic biases towards certain architectures, resulting in "incorrect" architectural decisions. We have made an honest effort, however, to use the highest quality algorithms possible. This method is the only practical way of dealing with the enormous complexity of FPGA architecture development because theoretical analysis can never account for all of the conflicting constraints that must be considered in a real design. In addition, this kind of approach produces an architecture that is tuned to the capabilities of the tools, which is key to the success of an FPGA.

The architecture will be a symmetric array of the L3-4.2 HLB logic blocks, with a segmented routing architecture employing the distribution and population described in Section III. Figure 4 illustrates such an array. An FPGA with these characteristics was designed and fabricated. The circuit and layout issues are described in Part II [2].

As technology continues to scale, the area of the logic will decrease but the relative effect of the delays due to wiring will increase. This will create further demands on the routing architecture, especially as the amount of logic available on a single FPGA will also increase, resulting in larger systems being built. Architecturally, there will continue to be a need for innovation in new logic and routing structures, and any new studies must always match the capabilities of the design and implementation tools to the architecture.

It is particularly important that a good set of benchmark circuits be available, especially to the academic community. With the trend towards larger systems-on-a-chip, studies like ours can only have merit when they use circuits that reflect this trend. We hope that the community will work towards improving the publicly available benchmark suites such as the one provided by MCNC [29]. It is essential that the industrial community participate in this effort as they possess such circuits and stand to gain the most from the research that uses them.

In the future, we will explore other aspects in architecture and design, including the incorporation of large blocks of memory, more complex routing structures for high-speed interconnect, and the global distribution of routing tracks.

V. Acknowledgments

This research was supported by a grant from the MICRONET Network of Excellence. We also acknowledge the valuable feedback provided by the referees.

References

[1] S. Brown, R. Francis, J. Rose, and Z. Vranesic, *Field-Programmable Gate Arrays*, Kluwer Academic Publishers, May 1992.

[2] Paul Chow, Soon Ong Seo, Jonathan Rose, Kevin Chung, Gerard Páez-Monzón, and Immanuel Rahardja, "The design of an SRAM-based field-programmable gate array, Part II: Circuit Design and Layout," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. XX, no. YY, pp. xxx–yyy, Month 1999.

[3] W. Carter, K. Duong, R. H. Freeman, H. Hsieh, J. Y. Ja, J. E. Mahoney, L. T. Ngo, and S. L. Sze, "A user programmable reconfigurable gate array," in *Proceedings of the 1986 Custom Integrated Circuits Conference (CICC-86)*, May 1986, pp. 233–235.

[4] H. Hsieh, W. Carter, J. Y. Ja, E. Cheung, S. Schreifels, C. Erickson, P. Freidin, and L. Tinkey, "Third-generation architecture

boosts speed and density of field-programmable gate arrays," in *Proceedings of the 1990 Custom Integrated Circuits Conference (CICC-90)*, 1990, pp. 31.2.1–31.2.7.

[5] Barry K. Britton, Dwight D. Hill, William Oswald Nam-Sung Woo, and Satwant Singh, "Optimized reconfigurable cell array architecture for high-performance field-programmable gate arrays," in *Proceedings of the 1993 Custom Integrated Circuits Conference*, 1993, pp. 7.2.1– 7.2.5.

[6] Richard Cliff et al., "A dual granularity and globally interconnected architecture for a programmable logic device," in *Proceedings of the 1993 Custom Integrated Circuits Conference*, 1993, pp. 7.3.1– 7.3.5.

[7] D.E. Smith, "Intel's FLEXlogic FPGA architecture," in *Compcon Spring '93*, 1993, pp. 378–384.

[8] D. Tavana, W. Yee, S. Young, and B. Fawcett, "Logic block and routing considerations for a new SRAM-based FPGA architecture," in *Proceedings of the 1995 Custom Integrated Circuits Conference*, 1995, pp. 511–514.

[9] Abbas El Gamal, Jonathan Greene, Justin Reyneri, Eric Rogoyski, Khaled A. El-ayat, and Amr Mohsen, "An architecture for electrically configurable gate arrays," *IEEE Journal of Solid-State Circuits*, vol. 24, no. 2, pp. 394–398, Apr. 1989.

[10] J. Birkner et al., "A very-high-speed field-programmable gate array using metal-to-metal antifuse programmable elements," *Microelectronics Journal*, vol. 23, pp. 561–568, 1992.

[11] David Marple and Larry Cooke, "An MPGA compatible FPGA architecture," in *Proceedings of the 1992 Custom Integrated Circuits Conference*. IEEE, 1992, pp. 4.2.1–4.2.4.

[12] S.C. Wong, H.C. So, J.H. Ou, and J. Costello, "A 5000-gate CMOS EPLD with multiple logic and interconnect arrays," in *Proceedings of the 1989 Custom Integrated Circuits Conference*, 1989, pp. 5.8.1 – 5.8.4.

[13] A. Gupta, V. Aggarwal, R. Patel, P. Chalasani, D. Chu, P. Seeni, P. Liu, J. Wu, and G. Kaat, "A user configurable gate array using CMOS-EPROM technology," in *Proceedings of the 1990 Custom Integrated Circuits Conference (CICC-90)*, 1990, pp. 31.7.1–31.7.4.

[14] R. Patel et al., "A 90.7MHz - 2.5 million transistors CMOS PLD with JTAG boundary scan and in-system programmability," in *Proceedings of the 1995 Custom Integrated Circuits Conference*, 1995, pp. 507–510.

[15] Carl Ebeling, Gaetano Borriello, Scott A. Hauck, David Song, and Elizabeth A. Walkup, "TRIPTYCH: a new FPGA architecture," in *FPGAs*, Will Moore and Wayne Luk, Eds., chapter 3.1, pp. 75–90. Abingdon EE&CS Books, 15 Harcourt Way, Abingdon OX14 1NV, England, 1991, Presented at the Oxford 1991 International Workshop on Field Programmable Logic and Applications.

[16] G. Borriello, C. Ebeling, S. A. Hauck, and S. Burns, "The triptych FPGA architecture," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 3, no. 4, pp. 491–500, Dec. 1995.

[17] Scott Hauck, Gaetano Borriello, Steven Burns, and Carl Ebeling, "MONTAGE: An FPGA for synchronous and asynchronous circuits," in *Proceedings of the Second International Workshop on Field-Programmable Logic and Applications*, Vienna, Austria, Sept. 1992.

[18] Paul Chow, Soon Ong Seo, Dennis Au, Terrence Choy, Bahram Fallah, David Lewis, Cherry Li, and Jonathan Rose, "A 1.2μm CMOS FPGA using cascaded logic blocks and segmented routing," in *FPGAs*, Will Moore and Wayne Luk, Eds., chapter 3.2, pp. 91–102. Abingdon EE&CS Books, 15 Harcourt Way, Abingdon OX14 1NV, England, 1991, Presented at the Oxford 1991 International Workshop on Field Programmable Logic and Applications.

[19] Paul Chow, Soon Ong Seo, Kevin Chung, Gerard Paez, and Jonathan Rose, "A high-speed FPGA using programmable mini-tiles," in *Symposium on Integrated Systems, previously the Conference on Advanced Research in VLSI*, Mar. 1993, pp. 103–122.

[20] John L. Hennessy and David A. Patterson, *Computer Architecture: A Quantitative Approach, Second Edition*, Morgan Kaufmann Publishers, Inc., 1995.

[21] M. Ahrens, A. El Gamal, D. Galbraith, J. Greene, S. Kaptanoglu, et al., "An FPGA family optimized for high densities and reduced routing delay," in *Proceedings of the 1990 Custom Integrated Circuits Conference (CICC-90)*, 1990, pp. 31.5.1 – 31.5.4.

[22] Kevin Chung, Satwant Singh, Jonathan Rose, and Paul Chow, "Using hierarchical logic blocks to improve the speed of field-programmable gate arrays," in *FPGAs*, Will Moore and Wayne Luk, Eds., chapter 3.3, pp. 103–113. Abingdon EE&CS Books, 15 Harcourt Way, Abingdon OX14 1NV, England, 1991, Presented at the Oxford 1991 International Workshop on Field Programmable Logic and Applications.

[23] K. Chung, *Architecture and Synthesis of Field-Programmable Gate Arrays with Hard-wired Connections*, Ph.D. thesis, Department of Electrical and Computer Engineering, University of Toronto, 1994.

[24] Jonathan Rose, Robert J. Francis, Paul Chow, and David Lewis, "The effect of logic block complexity on area of programmable gate arrays," in *Custom Integrated Circuits Conference*. IEEE, May 1989, pp. 5.3.1–5.3.5.

[25] Jonathan Rose, Robert J. Francis, David Lewis, and Paul Chow, "Architecture of field-programmable gate arrays: The effect of logic block functionality on area efficiency," *IEEE Journal of Solid-State Circuits*, vol. 25, no. 5, pp. 1217–1225, Oct. 1990.

[26] Satwant Singh, "The effect of logic block architecture on the speed of field-programmable gate arrays," M.A.Sc. thesis, University of Toronto, Department of Electrical Engineering, Toronto, Ontario, M5S 3G4, 1991.

[27] Satwant Singh, Jonathan Rose, Paul Chow, and David Lewis, "The effect of logic block architecture on FPGA performance," *IEEE Journal of Solid-State Circuits*, vol. 27, no. 3, pp. 281–287, Mar. 1992.

[28] Jack L. Kouloheris and Abbas El Gamal, "PLA-based FPGA area versus cell granularity," in *Proceedings of the 1992 Custom Integrated Circuits Conference*. IEEE, 1992, pp. 4.3.1–4.3.4.

[29] Microelectronics Centre of North Carolina (MCNC), *Logic Synthesis and Optimization Benchmarks User Guide, Version 3.0*, Jan. 1991.

[30] K. Chung and J. Rose, "TEMPT: Technology mapping for the exploration of FPGA architectures with hard-wired connections," in *Proceedings of the 29th Design Automation Conference (DAC-29)*, 1992, pp. 361–367.

[31] Hung-Cheng Hsieh, Khue Duong, Jason Y. Ja, Roy Kanazawa, Luan T. Ngo, Liane G. Tinkey, William S. Carter, and Ross H. Freeman, "A second generation user-programmable gate array," in *Custom Integrated Circuits Conference*. IEEE, 1987, pp. 515–521.

[32] S. Brown, J.S. Rose, and Z. Vranesic, "A detailed router for field programmable gate arrays," *IEEE Transactions on Computer-Aided Design of Circuits and Systems*, vol. 11, no. 5, pp. 620–628, May 1992.

Paul Chow (S'79–M'83) received the B.A.Sc. degree with honours in Engineering Science, and the M.A.Sc. and Ph.D. degrees in Electrical Engineering from the University of Toronto, Toronto, Ont., in 1977, 1979 and 1984, respectively.

In 1984 he joined the Computer Systems Laboratory at Stanford University, Stanford, CA, as a Research Associate, where he was a major contributor to the MIPS-X project. In January 1988, he joined the Department of Electrical and Computer Engineering at the University of Toronto where he is now an Associate Professor. He spent the 1995-1996 year at ATI Technologies Inc., in Thornhill, Ont., doing IC Design and working on new CAD flows. His research interests include high performance computer architectures, architectures for programmable digital signal processors, VLSI systems design, and field-programmable gate array architectures and applications. More details about his research can be found at www.eecg.toronto.edu/~pc.

He is currently the Chair of the Technical Advisory Committee for the Canadian Microelectronics Corporation.

Soon Ong Seo received the B.A.Sc. degree with honours in Electrical Engineering from the University of Ottawa, Ottawa, Ont. in 1985, and the M.A.Sc. degree from the University of Toronto, Toronto, Ont., in 1994.

In 1992 he joined ATI Technologies Inc., Thornhill, Ont., as a Design Engineer, where he is involved in the design and support of many graphics chips. His interests include high performance architecture and design, graphics, and field-programmable gate array design and architecture.

Jonathan Rose is an Associate Professor of Electrical and Computer Engineering at the University of Toronto, and an NSERC University Research Fellow.

He received the Ph.D. degree in Electrical Engineering in 1986 from the University of Toronto. From 1986 to 1989, he was a Research Associate in the Computer Systems Laboratory at Stanford University. In 1989, he joined the faculty of the University of Toronto. He spent the 1995-1996 year as a Senior Research Scientist at Xilinx, Inc., in San Jose, CA, working on a next-generation FPGA architecture.

He is the co-founder of the ACM FPGA Symposium, and remains part of that Symposium on its steering and program committees. He has worked for Bell-Northern Research and a number of FPGA companies on a consulting basis. A paper co-authored with Steve Brown won a distinguished paper award at the 1990 ICCAD Conference. His research covers all aspects of FPGAs including architecture, CAD, Field-Programmable Systems, and graphics and vision applications of rapid prototyping systems.

Kevin Chung received the B.A.Sc., M.A.Sc., and Ph.D. degree at the Department of Electrical and Computer Engineering of the University of Toronto in 1986, 1988 and 1994 respectively.

He is currently working as a Senior Software Engineer at Xilinx, Inc. at the Xilinx Toronto Development Centre in Toronto and has been with Xilinx since 1994. Previously, he was a Software Engineer at Data I/O Corporation from 1993 to 1994.

His main interests are in Field-Programmable Gate Array (FPGA) architectures and CAD algorithms for FPGAs.

Gerard Páez-Monzón (BEE,Villanova University, PA-USA '79; DEA in Computer Systems, Université Pierre et Marie Curie, Paris-France '83; Ph.D. in Computer Science, UPMC, Paris-France '86). He is a Professor with the Engineering School of the Universidad de Los Andes-Venezuela. He was a visiting scholar at the University of Toronto, Canada '91-92 where he worked with the FPGA group. Páez-Monzón is a principal member of the Research Center CEMISID of the Universidad de Los Andes in the areas of Microprocessor Architecture, Microprogramming, Floating-Point Units Design, FPGA, and VLSI Design. He is the author of a Computer Architecture book and a number of Computer Science research papers. He is currently on leave at National Semiconductor Corporation/Cyrix-West as a Microprocessor Architect working with the Architecture Group in x86 architecture designs in Santa Clara, CA. USA.

Immanuel Rahardja obtained his Electrical Engineering Bachelor degree from the University of Toronto in 1992. He worked at Xilinx from 1992 to 1996 on various projects exploring and evaluating different FPGA architectures. He currently works at a startup company, Aristo Technology, developing the graphical user interface software for a set of block-level system IC planning and implementation design tools.