The Design of an SRAM-Based Field-Programmable Gate Array, Part II: Circuit Design and Layout

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. In Part I, we described the high-level architectural design of an SRAM-programmable FPGA. This paper, Part II, will address the circuit design issues through to the physical layout. We consider area-speed trade-offs in the design of the logic block circuits and in the connections between the logic and the routing structure.**

All commercial FPGA designs are done using full-custom hand layout to obtain absolute minimum die sizes. This is both labor and time-intensive. We propose a design style with a *mini-tile* that contains a portion of all the components in the logic tile, resulting in less full-custom effort. The mini-tile is replicated in a 4×4 array to create a macro tile. The mini-tile is optimized for layout density and speed, and is customized in the array by adding appropriate vias. This technique also permits easily changing the hard-wired connections in the logic block architecture, and the segmentation length distribution in the routing architecture.

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

I. Introduction

THE design and implementation of Field-Programmable Gate Array (FPGA) technology is rarely described in the literature because much of the information is proprietary. In Part I [1], we describe some of the prior work and show the high-level architectural decisions used to select the logic block and routing architecture for the design of a high-performance FPGA. A symmetric array of hard-wired 4-input logic blocks with a segmented routing architecture was selected. Figure 1 illustrates such an array.

This paper reports on the circuit design and layout considerations that were made during the implementation of our prototype chip. In Section II we discuss the various circuit design issues that must be considered. A novel layout style for FPGAs is presented in Section III. The area, speed, and performance of the chip are given in Section IV. Finally, Section V provides some conclusions.

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

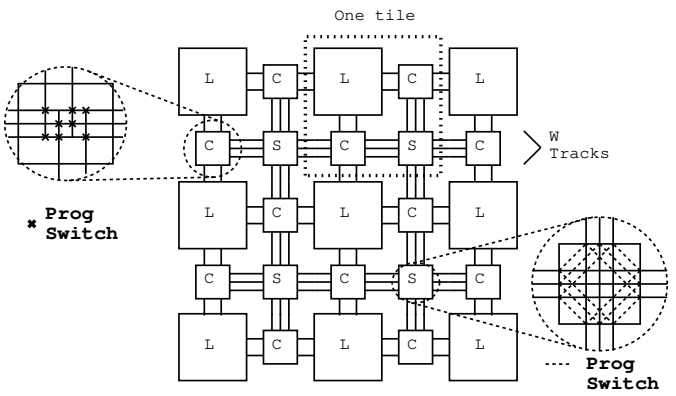

Fig. 1. Architectural Definitions

II. Circuit Design

In this section we discuss the circuit-level design of the FPGA architecture described in Part I [1]. The technology used in this design is a $1.2\mu m$ two-level metal n-well CMOS process.

Figure 2 illustrates the signal path starting at the input of an L3-4.2 lookup table, which was selected in Section II of Part I [1]. Following the lookup table is an output stage that contains a multiplexer to select the latched or unlatched version of the output before the signal is driven through a connection block (C block) onto the routing tracks. The signal then goes through one switch block (S block), or a number of S blocks, before it enters another C block where it will reach the input of the next logic block.

In the following sections we discuss issues that arose in the design of this path. Simulation with HSPICE [2] was used to evaluate different circuit options.

A. Multiplexers versus Switches

The C block is used to connect signals in the routing tracks to the input pins of the logic blocks, and can be implemented either using multiplexers or point-to-point switches. Figure 3 illustrates the use of a multiplexer and Figure 4 shows a connection block using one switch for each track-to-pin connection. The boxes labelled SRAM are the programming bits. In UTFPGA1 [3] we used multiplexers for the inputs and switches for the outputs. The reasoning was that since only one wire in the channel can be connected to an input pin, there is no need for the flexibility provided by switches. This reduces the number of bits

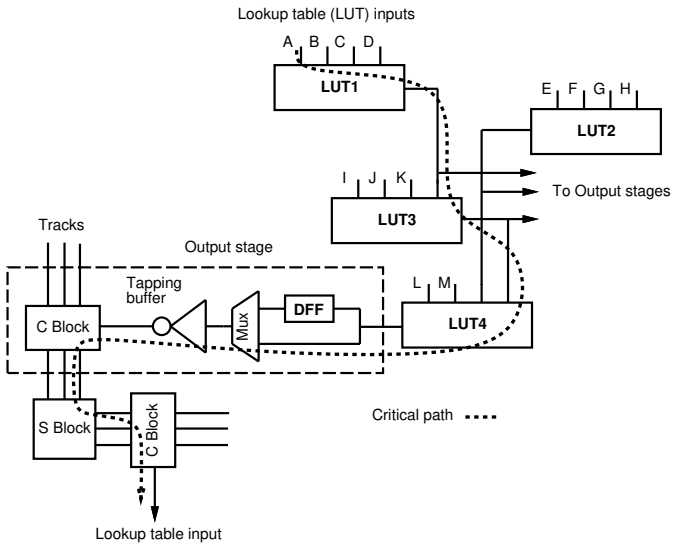

Fig. 2. Architectural Definitions

needed to select the input connection and the total number of transistors needed. For the output, having a switch from a pin to each track provides fanout capability.

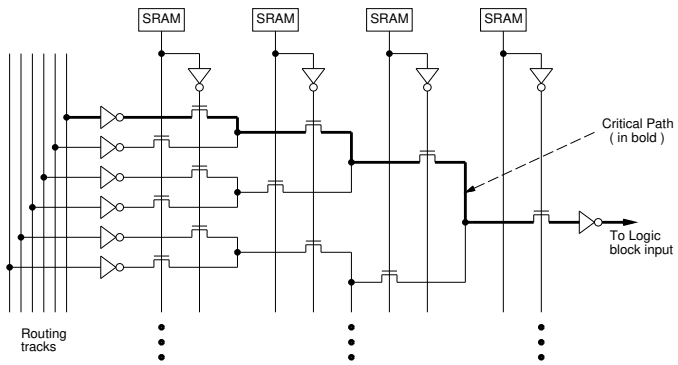

Fig. 3. Connection block input using a multiplexer.

Table I compares the two possibilities on the basis of transistor count. The buffer inverters from the routing tracks are not counted. In LEGO, every C block is connected to two logic blocks with four inputs on each side. For the multiplexer technique, eight multiplexers with 30 transistors each are required to connect to the 16 tracks. This assumes that n-channel switches are used instead of full complementary switches, and that every track can be connected to every logic block pin. Four programming bits are needed to program each multiplexer as well as 32 inverters (4 inverters/mux $\times$ 8 muxes) to generate the complementary inputs. For the pass transistor technique, eight sets of 16 switches are required to connect to the 16 tracks. Each switch will require one programming bit. For either technique, each programming bit will require five transistors. The circuit for the programming cells is discussed in Section III-A. Unit-sized transistors are used in the programming cells and the transistors that would be used in the switch or multiplexers are about three unit transistors in area. The result is that the pass-transistor technique is

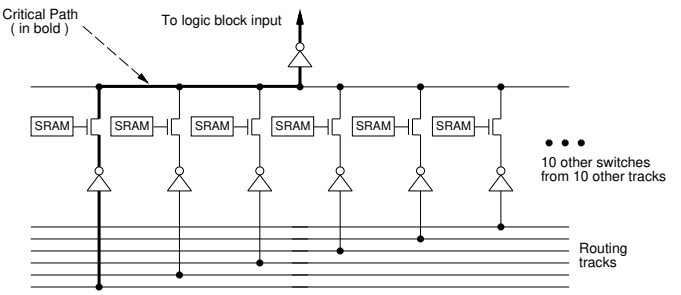

Fig. 4. Connection block input using pass transistors.

about 1024 unit transistors in area, compared to 944 for the multiplexer.

Although using a multiplexer results in less area, it has a significant performance impact because the signal must go through a number of series transistors. In LEGO, we chose to use switches at both the inputs and the outputs of the logic block. The price is that a significant fraction of the area is devoted to the C blocks.

B. Isolation

Another important issue is the capacitance at the C block inputs. With switches connected as shown in Figure 5(a), the capacitance seen from the routing track depends on whether the switch is on or off. If the switch is off, the contribution to the capacitance is due to the switch transistor itself. If the switch is on, then any of the circuitry on the other side of switch will contribute to the load seen by the driving signal. This means that the loading will be a function of the number of C-block switches that are turned on. This is related to the problem noted in the design of the Triptych [4] logic block, and their solution was to use gate logic instead of switch logic. The problem in LEGO can be reduced by adding a buffer between the track and each switch input, which makes the load independent of how many switches are turned on. However, this still has one load per switch input. Further isolation can be provided by using the configuration shown in Figure 5(b). In LEGO, each connection block connects the tracks to two sets of logic block inputs, so an additional optimization can be done as shown in Figure 6, which reduces the number of buffers required by one half, also reducing the capacitance on the track.

C. Lookup Table Design

For the LUT design, we have shown in Section II of Part I [1], that hard-wired connections can significantly improve performance. At the implementation level, the actual placement of the hard-wired connections is also important.

Figure 7 shows a part of the multiplexer tree in the LUT. It shows that the load of input S is only one transistor and one inverter, while the load of input P is eight transistors and an inverter. Input S is the least loaded input and also the input that determines the final output of the multiplexer. Therefore, when connecting the output of a LUT to the input of the next LUT, it is best to use the inputs that are closest to the output. In this way, the delay through

TABLE I
AREA COMPARISON FOR THE TWO POSSIBLE C BLOCK TECHNIQUES.

Method	(Programming Bits) × (Transistors/Bit)	Inverters × 2	(Connection Transistors) × Scale	Total Transistors
Pass transistor	(8 × 16) × 5	0	(8 × 16) × 3	1024
Multiplexer	(8 × 4) × 5	32 × 2	(8 × 30) × 3	944

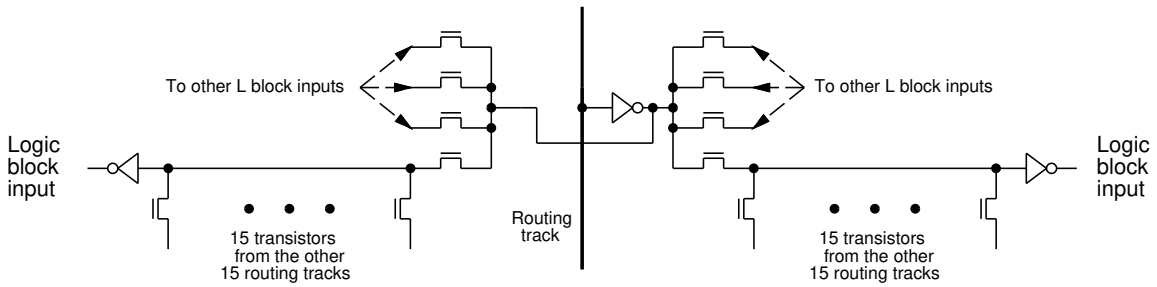

Fig. 6. Track isolation using shared buffer.

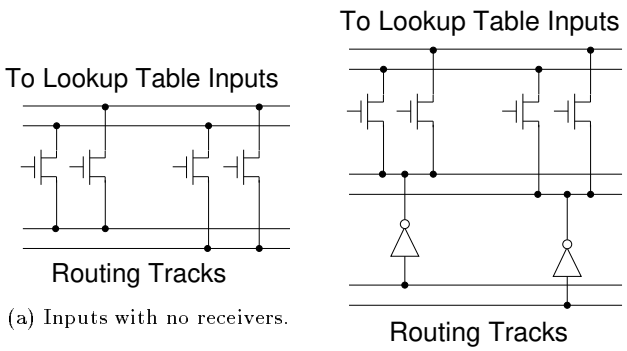

(a) Inputs with no receivers.

(b) Inputs with isolation.

Fig. 5. C-block input isolation options.

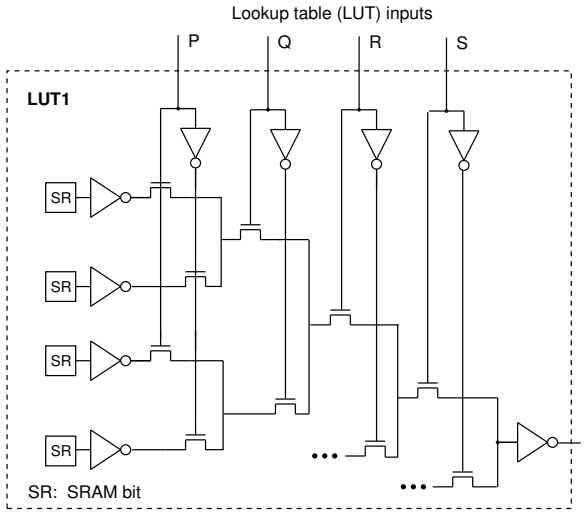

Fig. 7. Top quarter of LUT multiplexer tree.

the LUTs can be partially overlapped. Figure 2 shows the hard-wired connections used in LEGO. Note that the output of LUT2 is connected to input R of LUT4 (using the input labels of Figure 7), while the output of LUT3 is connected to input S. This is because the output of LUT3 will arrive the latest, so it should be connected to the fastest input, which is input S.

D. Full Complementary Versus n-channel Only Switches

Throughout the array there are switches in the LUT, the C block, and the S block, so the design of the switches will have a significant impact on the area and speed of the array. The most obvious area gain can be made by using a single n-channel pass transistor instead of a full transmission gate. In this section, we summarize the important considerations when deciding between the two alternatives.

By using a pass transistor, the area benefit is not just because there are half the transistors, but also because of a savings in well areas, which can be significant depending on whether the switches are scattered enough that individual wells must be provided for each transmission gate. As well, internal nodes of n-channel pass transistor circuits may not require contacts, saving area and reducing junction capacitance. There is also a savings in routing for the complementary clock or providing the local clock inverter if that technique is used.

For signal integrity, the pass transistor causes a V_t voltage drop when passing highs and has a slower rise time towards the end of a low to high transition. These effects can be partially compensated by lowering the switching threshold of the succeeding gates. Unfortunately, the reduced high also means that the pullup transistors in succeeding gates will not be fully turned off, resulting in static power dissipation. The problems become even greater as supply voltages move towards $3V$ because the worst case V_t voltage drop could still be around $1.2V$ to $1.3V$ ($V_t = 0.7V$ nominal at room temperature) significantly degrading the noise margin. This means that pass transistors are not viable unless the gate voltage of the pass transistor is bootstrapped

to above the supply.

As for speed, consider the construction of a transmission gate by the addition of a p-channel transistor that is 2.5 times the size of the n-channel transistor. The resistance changes by 0.5 but the junction capacitance increases by 3.5 or more, depending on the use of contacts. This means that the RC delay is at least $0.5 \times 3.5 = 1.75$ times larger so that the n-channel pass transistor would normally be faster, although careful simulations would always be required to make sure that this is true.

In LEGO, which was designed for $5V$ operation, we chose to use pass transistors because we could get a better propagation delay and reduced area. To compensate for the V_t drop we used simulations to adjust the switching point of the gates following the switches and to make sure that the propagation delay for rising and falling signals was balanced.

E. Power-up and Programming Protection

A difficult problem occurs between the time after chip power-up but before the programming has been downloaded. At this point the configuration bits have not been programmed and it is possible that there are conflicts between drivers on the routing tracks. This problem has been considered previously. In UTFPGA1 [3] the output drivers are disabled using a tristate buffer as shown in Figure 8(a). The buffers will be disabled during power up and programming. This design, however, impacts the speed because the disabling transistors of the tristate buffer will be in the critical path. Figure 8(b) shows the method used in LEGO. The output routing switch is controlled not only by a configuration bit, but also by another signal, shown as **Progb**, which is a global signal activated during power-up that disables all connections to the track, preventing the hazard.

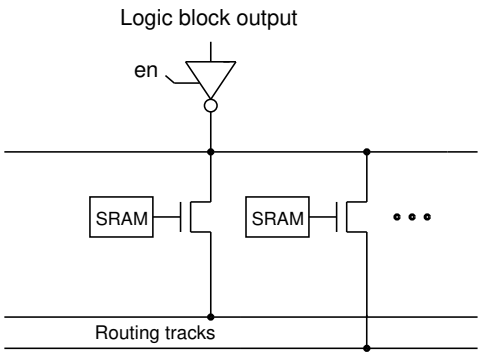

(a) Using a tristate buffer.

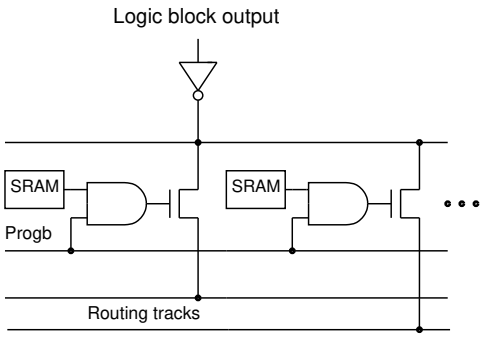

(b) Using a buffer and the **progb** signal.

Fig. 8. Hazard prevention alternatives.

F. Transistor Sizing and Signal Transition Balancing

Transistor sizing is a major concern for achieving high performance. Small transistors limit the current while large transistors have a higher capacitive load. Therefore, some optimization is required. In Triptych [4], the method of Logical Effort [5] was used to assist in transistor sizing and buffer insertion for speed. We chose to use actual simulation because we felt it was easier to balance the rise and fall times. We began by simulating the components of the critical path individually but later tuned the design by simulating the entire critical path at once.

Our initial assumption was that the n-channel pass transistors used for the S block switches and C block outputs had to be at least as large as the transistors of the drivers so as to reduce the effects of the switch sizes on the critical path delay. This allowed us to choose an initial driver size. We then tuned the design by simulating with different switch sizes. An example result is shown in Figure 9, which shows the effect of varying the logic block output driver size on the total delay of the critical path. It can be seen that the rise and fall times are equal at about $37mm$ for the size of the n-channel driver pulldown. In LEGO we actually used $34mm$ because in our initial simulations we used nine tracks in the channels, which was later changed to 16 tracks, and thus a higher load.

It should also be noted that the tuning was done by starting at the end of the critical path. This gives a better idea of the loads to be driven as sizes are fixed.

Figure 10 shows the critical path through the logic block section, and Figure 11 shows the critical path in routing using the tuned transistor sizes. From simulation, the total delay is about $4.1ns$ along the critical path broken into about $2.7ns$ for the logic block and $1.4ns$ for the routing.

III. Physical Design and Layout

Besides the concerns for optimizing the speed and area of the circuits, we also considered some of the issues related to the overall physical design and layout of the chip that would be used to implement the FPGA circuit. In this section we will describe the programming architecture of LEGO and some of the layout issues.

A. Programming Architecture

Three techniques for programming the configuration bits were identified. The simplest is to connect all of the bits into a long shift register. This requires only one pin to input the programming data. Although this is the slowest method, the speed of programming was not of concern in this project. Of greater concern is that each bit must contain a master-slave section because it behaves as a shift register, and hence each bit is very large. An alternative, which requires smaller bit cells, is to organize the program-

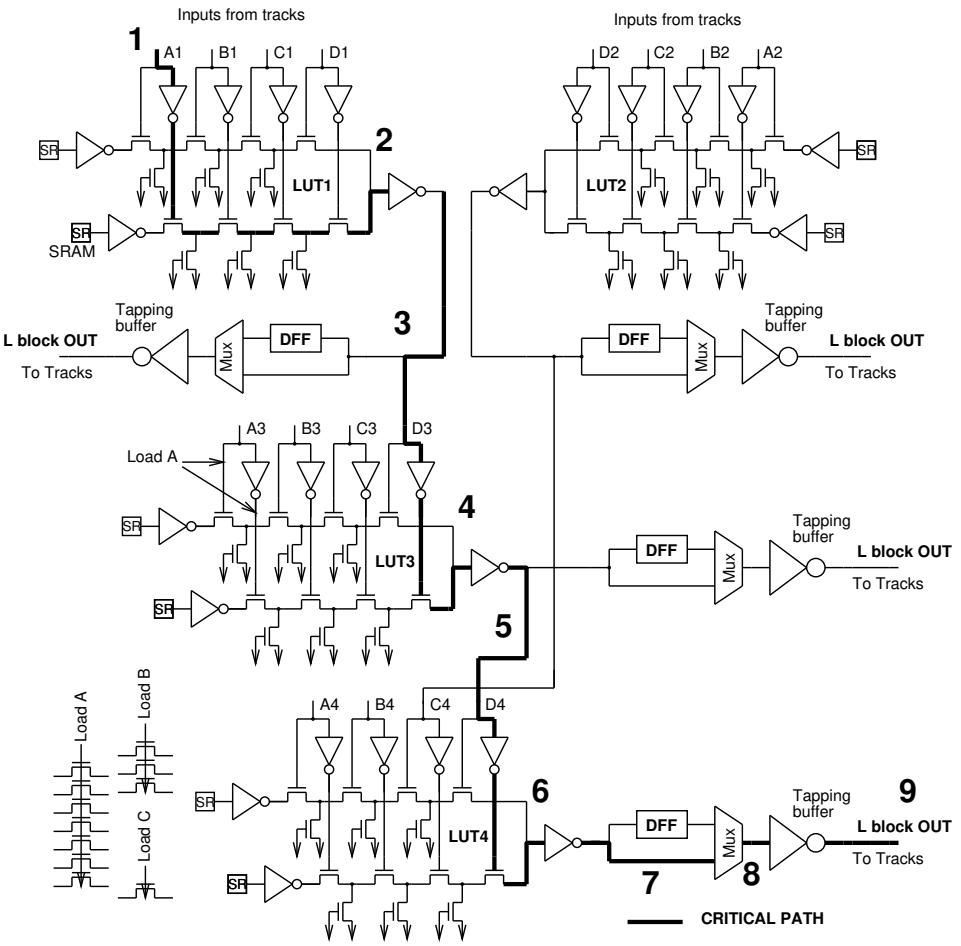

Fig. 10. Details of critical path through the logic block.

ming bits into rows and columns as in a regular memory. This will require significantly more pins for the parallel data as well as for selecting the word to be programmed.

A compromise of the serial and parallel techniques was used in LEGO and is shown in Figure 12. This serial-parallel technique uses a master-slave shift register to serially load a row of bits. The vertical shift register is initialized with the top bit on and the remainder off. This bit is shifted after each row is programmed to select the next row. After a row of bits has been loaded in the top shift register, the whole row of data is copied into the selected row in one clock pulse. The static RAMs that store the data bits do not have to be master-slave because the data is not shifted beyond this point. This halves the size of a programming bit as compared to the scheme that uses a single shift chain. This technique still leaves the option of programming any row, as in the fully parallel method, but it requires fewer programming pins.

The regular structure of the serial-parallel architecture also fits well with the layout methodology used for LEGO, which is described next.

B. The Mini-tile Layout Style

We wanted to find a better layout style than just trying to attack the problem purely as a full-custom design. When doing VLSI layout, it is desirable to find some regularity in the topology to make the task easier. In an FPGA, the first level of regularity is the arraying of the tiles as shown in Figure 1, where a tile includes an L block, an S block, and two C blocks. This partitioning was used for UTFPGA1 and in LEGO. However, even the full custom layout of a tile is a significant task, and it is desirable to find even more regularity. In UTFPGA1 [3] we attempted to use the same C block layout for both the vertical and horizontal channels but this resulted in a very poor tile floorplan because it was difficult to pack the pieces together. For LEGO we wanted a solution that had a finer granularity.

Instead, we divided the tile into 16 *mini-tiles*. Each mini-tile contains one sixteenth of the L block, S block, and C blocks. The tile[1] is implemented by creating a 4×4 array of mini-tiles, and adding vias at the appropriate locations to customize the local function of each mini-tile. The flip-flops, muxes to select the latched or combinational outputs, and output tapping buffers are at the edge of the array, since there are only four of each. These are shown as part of the output stage in Figure 2.

The motivation for this topology comes from the topology used for PLAs [6], where the goal is to make the overall

[1] We will use *tile* to denote the complete tile, and *mini-tile* to denote the smaller building block piece.

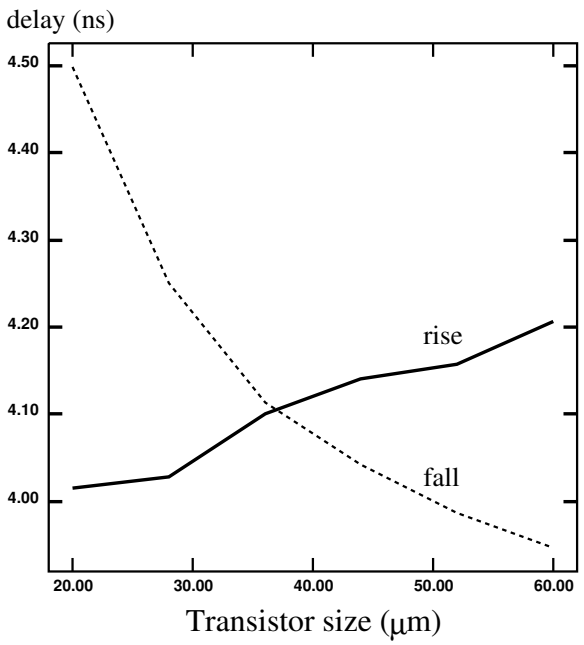

Fig. 9. An example of varying switch sizes for transistor sizing.

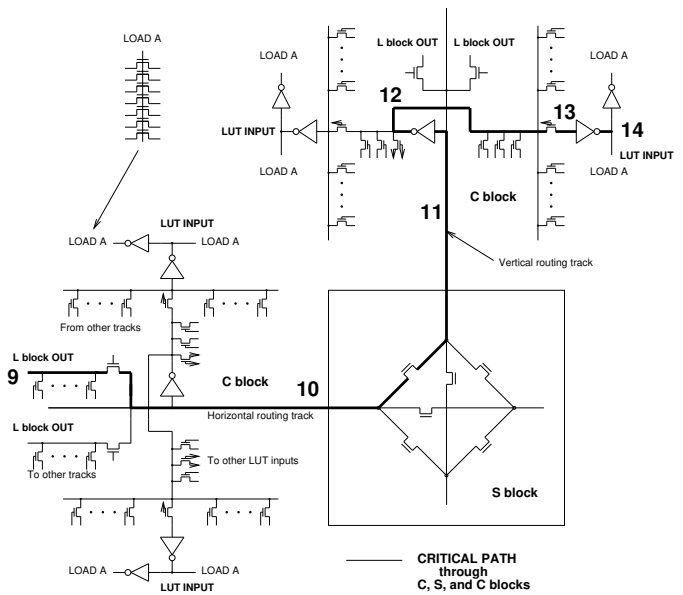

Fig. 11. Details of critical path through the routing network.

structure regular and have irregularity restricted to the final programming. Figure 13 shows the details of the mini-tile.

The major problem with dividing the functionality of the tile across 16 mini-tiles is the routing needed to interconnect the mini-tiles. It is difficult to organize the routes correctly, and it is likely that the extra wires will cost more in area than a non-mini-tile approach. The area issue is addressed further in Section IV. Figure 13 shows all the necessary wiring. Examples of via placement are shown by the Xs. The horizontal and vertical tracks of Figure 1 are shown in heavier lines intersecting at the switch block (S block).

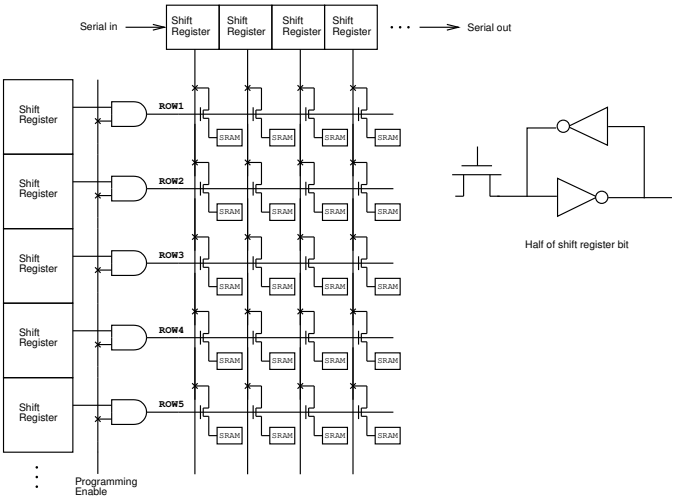

Fig. 12. Serial-parallel programming of the configuration bits.

The choice of how many tracks to use depends on the actual size of the overall array of tiles in a chip and on the number of mini-tiles. For larger FPGAs more tracks would be needed, and in the case of LEGO, the number will be a multiple of four because of the mini-tile array size.

Each mini-tile contains one quarter of a mux-tree, which is really a 4-to-1 mux that has the SRAM bits at its inputs and a separate 2-to-1 mux. The three mux interconnect tracks going through the quarter mux-tree are used to connect between the mini-tiles to form the full 4-input lookup tables. These tracks are segmented in the mini-tile to allow sharing of the tracks and can be connected with a piece of metal as needed during final configuration. One of the 2-to-1 muxes is not required when hooking up the full lookup table.

In addition, the outputs of the LUTs are routed on the *LUT connection* tracks to the edge of the tile for connection to the D flip-flops and to the driver to the tracks. The output of the driver is routed back into the mini-tiles along a *Lblk to C blk inputs* track for connection to the C block for output. The remaining vertical wires are to distribute signals to the mini-tiles along the vertical dimension of the array.

Driver *DRI2* is the inverter shown at the inputs to the LUTs in Figure 7. Driver *DRI1* drives the input from the C block to the LUT inputs.

The drivers marked *DRI3* and *DRI4* are the track isolation drivers as shown in Figure 6. The transistors at the outputs of these drivers are the switches described in Section II-A. In this partitioning, the C block includes connections to the south- and east-neighbour tiles so some of the wires in each mini-tile are used to route signals to the LUT inputs in those neighbours.

Using the mini-tile methodology we sacrifice the potential for global layout optimization across the whole tile in exchange for a simpler full-custom layout task. However, we could afford to spend significant effort on the layout of the mini-tile. The floorplan of the mini-tile is shown in Figure 14. The vertical dimension is limited in pitch by

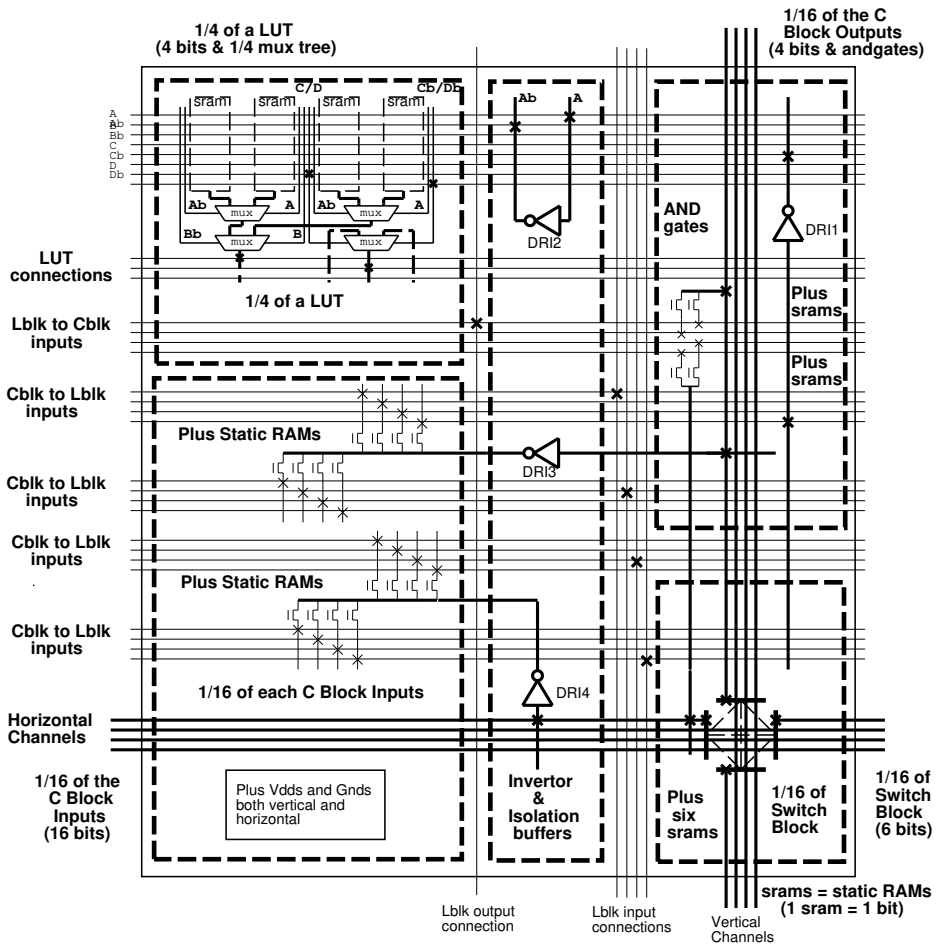

Fig. 13. Details of the mini-tile (1/16 of a logic tile).

the number of horizontal metal 2 tracks. In the horizontal dimension, there are still some visible internal vertical routing channels that do not go over any active area. If we had a third metal routing layer so that the vertical routing could go over active area, then the horizontal dimension could be reduced by about 25%. With this area removed, the layout efficiency would be quite good as most of the metal routing would be on top of active area, which cannot be compacted much. Since we did not attempt a full custom layout of the tile we do not have a quantitative measure of our layout efficiency. However, we feel that by using this methodology there was not a great sacrifice in overall density, and shortening the time to do the layout was more important.

An added benefit that has not been explored at this point is that with this ability to configure the tile after the hard layout work is done, it is possible to try other hard-wired interconnections for the LUTs and other channel segmentation distributions. This may be of benefit if our work in non-homogeneous logic blocks suggests that a mix of tiles with different types of hard-wired connections is beneficial.

C. Other Layout Issues

Our experience with UTFPGA1 showed that doing a full-custom layout of a complete tile is extremely difficult, and distributing power was one of the major problems. With the mini-tile methodology, it is easy to have a rectangular floorplan for the tile, and this makes it easier to distribute power throughout the array.

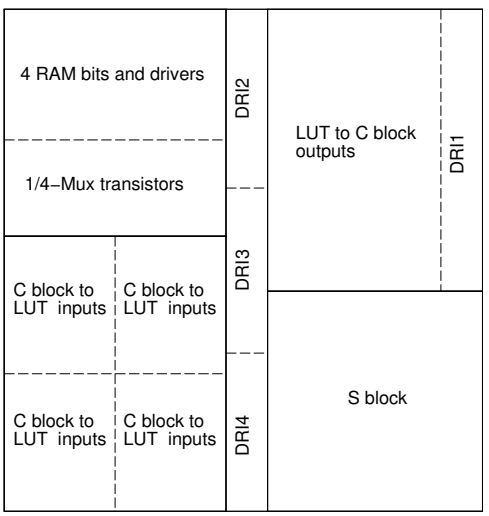

Fig. 14. Floorplan of mini-tile.

As described in Section III of Part I [1], LEGO uses a

segmented routing scheme. The segments repeat at intervals of one, two, and three tiles. To achieve this with a single tile, and without requiring customization of the interconnect between the tiles, the scheme shown in Figure 15 was used. This scheme will work as long as the number of tracks for a repeat interval is the same as the repeat interval. In LEGO, there are four tracks of length two and three tracks of length three.

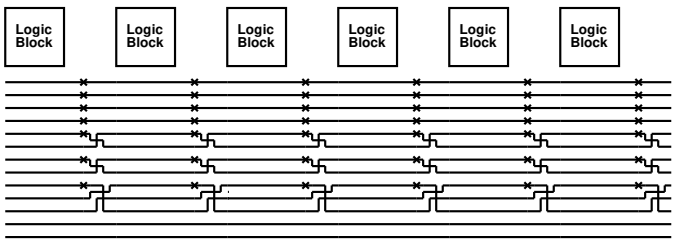

Fig. 15. Segmentation scheme for use in a tile layout style.

IV. Results and Performance

The goal of this project was to design an FPGA with high performance without greatly sacrificing area. In this section we analyze how well this goal was met. We first provide area measurements and speed measurements from simulation. Subsequently, measured results from a fabricated test chip are provided.

A. Area and Speed

The size of each tile, which contains four 4-input lookup tables, four D flip-flops, and connections for 16 tracks, is $1240\mu m$ by $1184\mu m$. By way of comparison, the Xilinx 3000 tile in the same feature size technology ($1.2\mu m$ CMOS), which contains the equivalent of two four-input lookup tables, two D flip-flops and connections for eight tracks (about half the functionality of the LEGO tile) is roughly $500\mu m$ by $400\mu m$. This is about a factor of four smaller, which comes from the fact that there is twice as much logic and routing in the LEGO tile, the mini-tile strategy, and design decisions that sacrificed area for speed. In particular, the large size drivers and the switches (rather than a multiplexer) used in the C block take up significant area. Each LEGO tile requires 492 programming bits.

From simulation, we measure a delay of $4.1ns$ for the critical path shown in Figure 10 and Figure 11. The measured delay for LEGO is simulated time for nominal conditions, not worst-case.

An attempt to compare LEGO with commercial FPGAs is shown in Table II. It is only provided as a point of reference since doing more complicated circuits across a number of architectures would be difficult. The circuit used is shown in Figure 16. The figure and some of the numbers in Table II are taken from some performance benchmarks published by Crosspoint Solutions [7]. The numbers in the *published* column are the reported numbers assuming that the output goes to a pad driving a $50pf$ load. The *adjusted* numbers subtract $25ns$ as an attempt to take out the effect of the output pad[2]. This is a gross approximation as we do not have the numbers for the delays due to only the logic. However, it can be seen that LEGO compares well even though the others are implemented in higher performance technologies. Much of this is because of the larger functionality of the LEGO logic block. The LEGO number is from an HSPICE simulation under nominal conditions. The other numbers are for worst case process conditions.

TABLE II
Performance comparisons for the ATN benchmark.

FPGA	Published (ns)	Adjusted (ns)
Crosspoint CP20420-1	37.4	12.4
Crosspoint CP20420-0	54.2	29.2
Actel A1280-1	70.0	45.0
Xilinx XC4005-5	41.4	16.4
LEGO	n/a	11.5

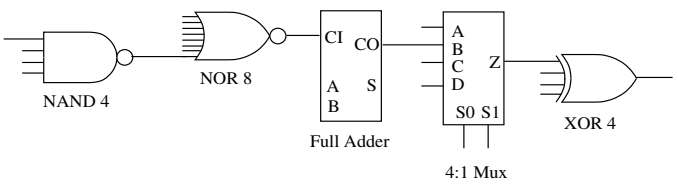

Fig. 16. The ATN benchmark circuit.

B. Fabrication Results

A chip comprising two LEGO tiles was fabricated and the photomicrograph is shown in Figure 17. The extra wires that exit the column and rows were either attached to pads or "wrapped" around the tiles to connect to the other side of the channel. This chip has 50 pads, including 12 power and ground, 23 signal pins into the FPGA core, eight programming pins, five control pins, and two test pins.

The testing of the chip showed that the logic blocks were not fully functional, and the cause has not been determined, though we suspect that having only nominal parameters to simulate the pass transistor logic may have lead to some threshold problems. However, we were able to make reliable connections using the programmable routing.

A smaller test chip was also fabricated. It contained only four of the mini-tiles, which was enough to build one of the four-input lookup tables. From this chip, we were able to make delay measurements that included one part of the logic block. Based on these measurements, we observed a pad to pad delay through one S block switch transistor of $19ns$, where the output was unloaded. The pad to pad delay through part of the critical path shown in Figure 10 and Figure 11 was measured. The measured path exits

[2]The Crosspoint benchmarks erroneously used data for a Xilinx 4005-7 part, which does not exist. The Xilinx data has been multiplied by 5/7 to reflect the speed of a 4005-5.

Fig. 17. Photomicrograph of the LEGO test chip.

the logic block through the tapping buffer at node 3 of Figure 11, rather than node 6. The measured delay was $26ns$, giving a delay of about $7ns$ when the pad delay is eliminated. The simulated delay for this path is about $3ns$. Speed measurements were done using a tester and have an error of $\pm 1ns$. We attribute most of the difference to the fact that the simulations are based on nominal parameters.

V. Conclusions

We have described the low-level circuit design and layout methodology used in the design of a high-performance FPGA. The important issues in circuit design that affect the critical path delay have been shown. We have focused on achieving a high-speed design while keeping in mind the area tradeoffs. The most critical issues are the design of the switches and minimizing the capacitance of the routing network. As well as the circuit issues, we have also presented a layout methodology for FPGAs that reduces the amount of custom layout required without sacrificing excessive density.

Our results have shown that the LEGO design compares favorably with existing commercial FPGAs. The most important contribution of our work is that we have gained a greater understanding of the real design issues in an FPGA circuit.

As technology continues to scale, supply voltages will decrease, and the number of metal layers will increase. The electrical properties of the new technologies will require that new tradeoffs be considered in the circuits and layouts of FPGAs. These properties will also impact the architectures of both the logic blocks and the routing. It is clear that both the architectural and circuit-level aspects of FPGA design must continue to evolve.

VI. Acknowledgments

This research was supported by a grant from the MICRONET Network of Excellence. The authors would like to thank Jaro Pristupa for his support of our CAD and VLSI labs, Paul Chow, (no relation) for help with the testing, and Ken Martin for his insights into circuit design using pass transistors. The Canadian Microelectronic Corporation provided fabrication access using the facilities of Nortel. Their help is always appreciated.

We also acknowledge the valuable feedback provided by the referees.

References

[1] 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 I: Architecture," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. XX, no. YY, pp. xxx–yyy, Month 1999.

[2] Meta-Software, Inc., 1300 White Oaks Road, Campbell, CA 95008, *HSPICE User's Manual*.

[3] 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.

[4] 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.

[5] Ivan E. Sutherland and Robert F. Sproull, "Logical effort: Designing for speed on the back of an envelope," in *Proceedings of the 1991 University of California at Santa Cruz Conference on Advanced Research in VLSI*, Mar. 1991, pp. 1–16.

[6] Carver Mead and Lynn Conway, *Introduction to VLSI Systems*, Addison-Wesley, 1980.

[7] "Crosspoint solutions, performance benchmarks, rev 1.01," Aug. 1992.

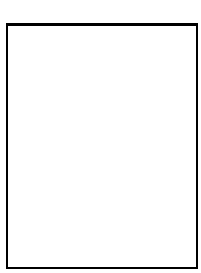

Paul Chow (S'79–M'83) received the B.A.Sc. degree with honours in Engineering Science, and the M.A.Sc. and Ph.D. degrees in Electrical Engineering from the University of Toronto, Toronto, Ont., in 1977, 1979 and 1984, respectively.

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.

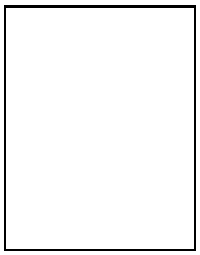

Soon Ong Seo received the B.A.Sc. degree with honours in Electrical Engineering from the University of Ottawa, Ottawa, Ont. in 1985, and the M.A.Sc. degree from the University of Toronto, Toronto, Ont., in 1994.

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.

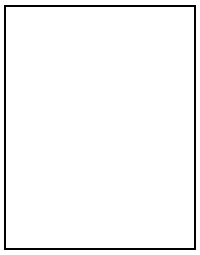

Jonathan Rose is an Associate Professor of Electrical and Computer Engineering at the University of Toronto, and an NSERC University Research Fellow.

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.

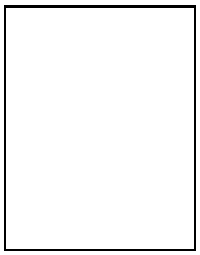

Kevin Chung received the B.A.Sc., M.A.Sc., and Ph.D. degree at the Department of Electrical and Computer Engineering of the University of Toronto in 1986, 1988 and 1994 respectively.

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.

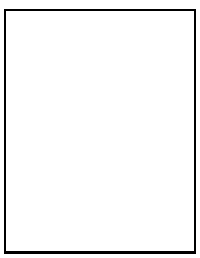

Gerard Páez-Monzón (BEE,Villanova University, PA-USA '79; DEA in Computer Systems, Université Pierre et Marie Curie, Paris-France '83; Ph.D. in Computer Science, UPMC, Paris-France '86). He is a Professor with the Engineering School of the Universidad de Los Andes-Venezuela. He was a visiting scholar 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.

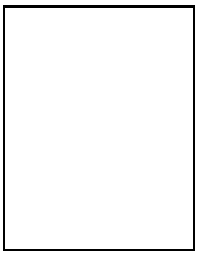

Immanuel Rahardja obtained his Electrical Engineering Bachelor degree from the University of Toronto in 1992. He worked at Xilinx from 1992 to 1996 on various projects exploring and evaluating different FPGA architectures. He currently works at a startup company, Aristo Technology, developing the graphical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