Algorithms for Efficient Runtime Fault Recovery on Diverse FPGA Architectures

John Lach
UCLA EE Department
jlach@icsl.ucla.edu

William H. Mangione-Smith
UCLA EE Department
billms@ee.ucla.edu

Miodrag Potkonjak
UCLA CS Department
miodrag@cs.ucla.edu

Abstract

The inherent redundancy and in-the-field reconfiguration capabilities of field programmable gate arrays (FPGAs) provide alternatives to integrated circuit redundancy-based fault recovery techniques. An algorithm for efficient runtime recovery from permanent logic faults in the Xilinx 4000 architecture has been expanded to include interconnect fault recovery and has been applied to a diverse set of FPGA architectures. The post-fault-detection system downtime is minimized, and the end user need not have access to computer-aided design (CAD) tools, making the algorithm completely transparent to system users. Although some architectural features allow for a more efficient implementation, high levels of fault recovery with low timing and resource overhead can be achieved on these diverse architectures.

1. Introduction

Taking advantage of the flexible and inherently redundant nature of FPGAs, a low overhead fault recovery algorithm has been developed capable of recovering from faults at runtime with minimal system downtime and no end user CAD tool requirements. Assuming detection[1], localization, and diagnosis[2] of a fault, a configuration of the design can be loaded that does not utilize the faulty resource(s). The alternate configurations are previously generated by the CAD tools at design-time and are available in memory. The proper configuration is then activated based on the location of the faults. The previously prepared configuration need only be applied to the device, thereby not requiring substantial system downtime or the end user to have access to CAD tools. Therefore, the algorithm, and even the very existence of an FPGA in the system, remains transparent to the user. Previous algorithms achieved such runtime fault recovery exclusively for logic faults [7], but the algorithm has been expanded to include interconnect faults and has been applied to a variety of FPGA architectures.

1.1. General algorithm

We propose partitioning the physical design into a set of tiles. Each tile is composed of a set of physical resources (i.e. logic blocks and interconnect), an interface specification which denotes the connectivity to neighboring tiles, and a netlist. Logic and local interconnect reliability is achieved by providing multiple configurations of each tile, each of which does not use certain resources within the tile. Furthermore, by using immutable tile interfaces, the effects of swapping a tile configuration do not propagate to other tiles, thereby making each tile independent and reducing the storage overhead[3]. Any paths that cross tile boundaries can be made reliable by reserving other inter-tile interconnect to be used as spares.

---

[1] Many FPGA testing techniques are currently available for both logic and interconnect [1-6].

[2] Any fault that occurs can be considered permanent (the resource will no longer be used). Therefore, diagnosis is not entirely necessary, but it may be helpful in identifying non-permanent faults that can be corrected.

[3] Tile independence requires the system to generate and store individual tile information and its corresponding instances. However, no inter-tile information need be generated or stored, thus reducing CAD tool effort and storage requirements.

This approach has three main benefits compared to redundancy-based fault recovery techniques: very low overhead, the option for runtime management, and flexibility. The overhead required to implement this fine-grained approach, which can be measured in both physical resources on the FPGA (logic blocks, I/O blocks, and interconnect) and circuit performance, is extremely low compared to redundancy. This is due primarily to the inherent redundancy in FPGAs, as opposed to the introduction of redundant elements into fixed designs that is required for reliability and yield enhancement [8-10]. Runtime management can be a very valuable feature of a system, particularly for mission-critical applications. This fault recovery approach handles runtime problems on-line, minimizing the amount of system downtime and eliminating the need for outside intervention. The flexibility that this approach provides allows for application-specific solutions. The degree of fault recovery can be adjusted to meet timing and resource constraints or estimated logic block and interconnect reliability.

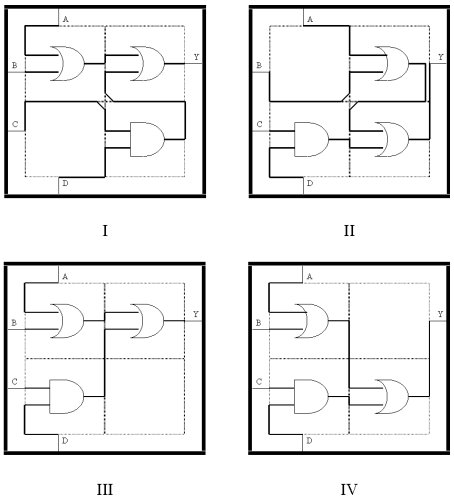

Figure 1. Logic fault recovery

1.2. Example

Consider the Boolean function $Y=(A \wedge B) \wedge (C \vee D)$, which might be implemented in a tile containing four logic blocks as shown in the Figure 1. This partitioning contains one spare logic block, which is available if a fault should be detected in one of the occupied logic blocks. Upon detecting such a fault, an alternate configuration of the tile is activated which does not rely on the faulty logic block. Each implementation is interchangeable with the original, as the interface between the tile and the surrounding areas of the design is fixed and the individual configurations implement the same function.

Recovering from local interconnect faults can be handled in much the same manner. Interconnect can be set aside as unused in each tile configuration[4], and when a fault disables an interconnect line an instance not utilizing that line can be activated.

Recovering from faults to global and overlapped segmented interconnect requires a different approach, as much of that interconnect crosses tile boundaries, thus eliminating the independent nature of each tile. Ignoring tile boundaries, interconnect can be set aside that acts as backup for used global and overlapped segmented interconnect. However, the backup interconnect must be able to fulfill the connections of the failed interconnect without requiring an alteration to the affected tiles. Figure 2 shows how a global line (dotted) can act as a backup for several segmented interconnect lines (solid) in the Xilinx XC4000 family. Upon a segmented line sustaining a fault, the backup global line can be activated. This requires only the programming of the proper connections and does not affect the logic in any way.

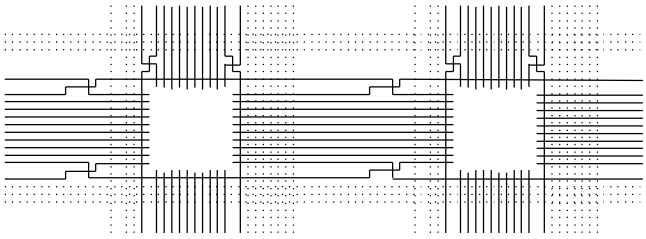

Figure 2. Inter-tile interconnect fault recovery

2. FPGA architectures

The implementation of the algorithm varies depending on the target FPGA architecture. The following sections describe the implementation on three architectures:

[4] Throughout the configuration generation for logic reliability, most local interconnect is unused in at least one instance, making it unnecessary to generate many additional instances for local interconnect reliability.

Sanders' context switching reconfigurable computing (CSRC) technology [11], Xilinx's XC4000 family [12], and Altera's Flex 10k series [13].

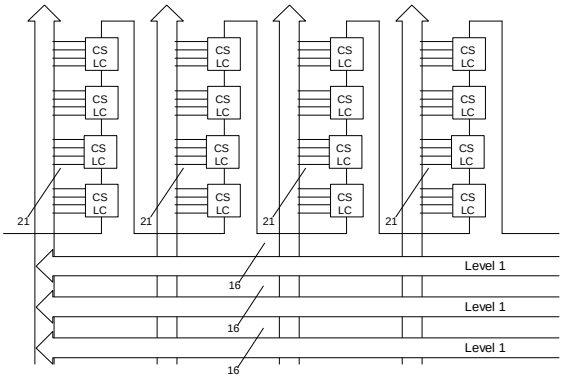

Figure 3. CSLA and Level 1 routing

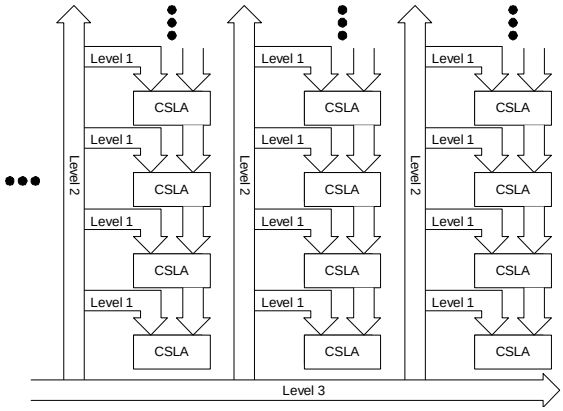

Figure 4. Data pipes and Level 3 routing

2.1. Sanders CSRC

The CSRC device is based on new architectural techniques that exploit dynamic reconfiguration via context switching. Each context is identical and is composed of logic and interconnect arranged in a hierarchical structure. At the lowest level of logic is the context switching logic cell (CSLC). Sixteen CSLCs comprise one context switching logic array (CSLA), as shown in Figure 3. Each logic cell has a carry-in and carry-out. The sixteen CSLCs are sectioned into four groups that are connected and driven by Level 1 routing. The next level of logic and interconnect is the 16-bit data pipe which is composed of CSLAs and connected by Level 2 routing as shown in Figure 4. Each context is composed of a set of data pipes that are connected by Level 3 routing, completing the highest level of logic and routing.

Developing a fault recovery approach for the CSRC device can be done in two ways: 1) looking at each context as an independent entity, and 2) allowing functionality belonging to one context to be transferred to another in the face of a fault. Analysis reveals that the latter is less desirable. Moving a section of a pipeline, for example, to another context would require switching to the other context in mid-stream before switching back again to complete the context's original function. The CSRC technology is capable of performing such a task, but the performance cost would be great. For example, leaving an entire context free to serve as a spare to be programmed and activated in the case of a fault to a portion of an active context is undesirable. In general, such an approach would require a tremendous amount of overhead (X spare contexts for every Y active contexts[5]) and only be able to recover from X faults throughout the lifetime of the system. Fault recovery in FPGAs can be implemented efficiently because of the inherent redundancy (just as contexts are redundant), but the larger the redundant block, the higher the resource overhead and the smaller number of tolerable faults. Allocating or reserving spare contexts in the CSRC device would be analogous to simply adding redundant FPGAs in a non-CS system.

Looking at each context as an independent entity, a much finer grained approach, helps to alleviate these problems. For example, looking at smaller blocks of redundancy (CSLAs or even CSLCs) creates the opportunity for lower resource overhead and a smaller performance impact. CAD tools rarely map logic to maximum density, leaving many CSLAs and CSLCs unused. Using these as redundant blocks eliminates the effective resource overhead. Additional unused resources can, and often should, be added to raise the number of tolerable faults, thus creating resource overhead. Both the naturally unused and additional redundant

[5] Sanders' CSRC device has four contexts. With three active contexts and one spare, resource overhead would be 25% while being capable of tolerating only one fault.

blocks must also be distributed for easier fault-recovery and a smaller performance impact, but the overhead is still significantly reduced from a larger redundant block approach.

This approach also raises the number of recoverable faults and creates a more efficient and realistic fault model. Very rarely would a fault occur that destroys a large portion of the chip. If such a situation arose, it would be unlikely that enough of the chip would remain functional, thus rendering any on-chip fault recovery algorithm useless. Most faults that would occur are single faults that affect a small segment of memory (e.g. LUT), logic (e.g. multiplexor or flip-flop), or interconnect at any level. Such a fault model dictates the use of smaller redundant blocks, as one faulty wire or LUT should not render an entire context faulty.

Breaking the CSRC device contexts into independent fault recovery blocks (tiling) can be done in a number of ways. Selecting the most efficient may be application dependent. The pipeline nature of the interconnect makes it difficult to have a single CSLC be redundant for the others in its array (see Figure 3). If a logic cell in the middle of a 4-cell pipe portion should fail, it may not be possible for the CAD tool to route the proper signals to reach the redundant cell. Therefore, it becomes necessary to look at a slightly larger block, i.e. the 4-cell pipe portion. Each portion has the same connectivity, making it possible for one pipe portion to be redundant for any other in the array. Thus, each array has one 4-cell pipe portion that is redundant for the other three, and three other configurations are generated that would be instantiated upon a failure to one of the active pipe portions.

Almost all types of failures can be tolerated at this level. Faults to the cells, the cell pins, and the wires leading from the Level 1 routing to the cell pins can all be tolerated, as each can be attributed to a specific pipe portion. Simply switching to an array configuration not utilizing the pipe portion containing the faulty hardware recovers from the fault. Figure 5 shows how a CSLA may be reconfigured if a fault were to occur in the routing from Level 1 to the logic cell pin wires in pipe portion three (or anywhere within the third pipe portion).

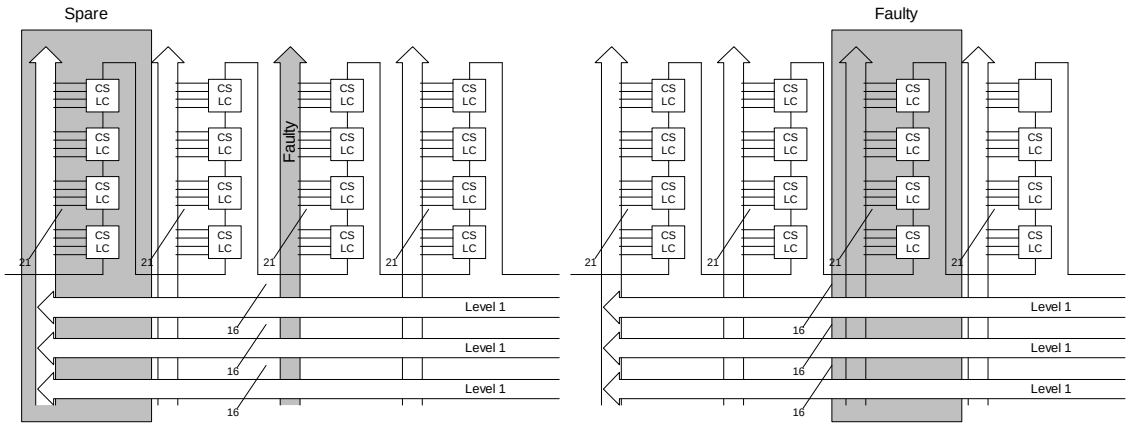

Figure 5. Original and recovered configuration after an internal CSLA fault

One problem with this tiling approach concerns inter-pipe interconnect. This interconnect includes the Level 1 routing and the carry lines, as such lines cannot be attributed to a single pipe portion. The connections among Level 1 routing are plentiful and flexible enough that the CAD tools can find acceptable routing if a Level 1 line should fail, but configurations must be ready at runtime that do not necessitate in the field CAD use. Therefore, configurations must be generated at design-time that do not make use of each Level 1 line in at least one configuration, just as each pipe-portion is unused in at least one configuration.

Recovering from faults to the carry line is more difficult, as such a fault potentially renders two pipe portions inoperative. The backend CAD tool deals with the carry lines when there is a break in the pipe (i.e. a middle pipe portion is faulty and, therefore, unused) during its design-time configuration generation. However, if the carry-out of one pipe portion and the carry-in of another are faulty, an inter-pipe portion problem arises. The CAD tool may be able to

implement the array's functionality without using the carry line (i.e. each carry line unused in at least one configuration). If not, two pipe portions may have to be set aside for the array, or such a fault may be considered intolerable within the array[6].

This problem can be resolved through a second tiling approach that involves looking at entire CSLAs as the redundant block size. This can be achieved on the CSRC device by using the same steps as the previous approach but applying them to larger areas: 4-cell pipe portions become CSLAs, Level 1 routing becomes Level 2, etc. Configurations can be generated that leave one CSLA to be redundant for the active arrays in its 16-bit data pipe. The problems that existed for the first approach exist again at this next level, and they can be dealt with in the same manner. Level 2 interconnect can be reserved as unused in various contexts, just as Level 1 lines were, and the carry lines in and out of the arrays pose the same problem as before and must be dealt with similarly. But, this approach solves the carry line problem at the logic cell level. If a fault occurs in a logic cell carry line, the entire array can be disabled as faulty. This approach also is more efficient if there are highly correlated faults which may render entire arrays faulty. The approach can also be taken up to the next level with pipes of CSLAs being the redundant block size with the Level 3 routing and the carry lines in and out of the pipe posing the same problems. But again, redundancy at this next level solves the carry line problem from the previous level and recovers from correlated faults that may disable entire pipes.

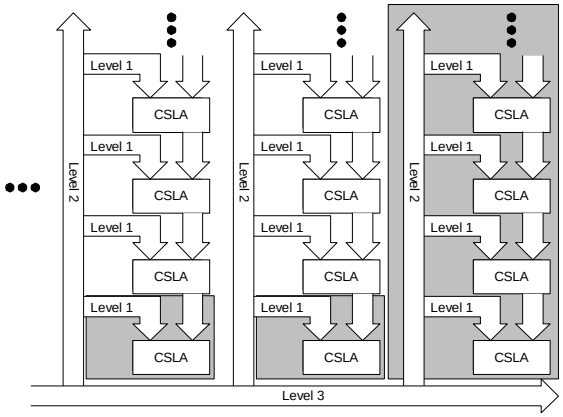

Figure 6. Hierarchical redundancy (shaded areas are spares)

Choosing the proper level of redundancy may depend on many factors, including the application, the desired level of fault recovery (the number and types of faults to be tolerated), the amount of spare resources, and the acceptable amount of overhead (timing and area). Often times, the best technique will be a fractal-like approach combining all of the above levels. Within each logic array, there can be a redundant 4-cell pipe portion. Within each data pipe, there can be a redundant CSLA, and within each context, there can be a redundant data pipe (see Figure 6). The CSRC device is inherently self-similar and hierarchical in design and is therefore well suited for such an approach. This emphasizes the flexible nature of the general approach. That is, different levels of fault recovery and various acceptable amounts of overhead can be achieved for specific applications within the same general algorithm. Therefore, a wider array of fault models can be accommodated. Isolated single faults can be handled within individual arrays, as there is no need to render large sections of the chip useless for a small fault. Conversely, the higher levels can more efficiently tolerate a large number of correlated faults that may render entire arrays or even entire data pipes inoperable.

2.2. Xilinx XC4000

Implementing the same algorithm on the Xilinx XC4000 family requires many alterations, but the general algorithm remains unchanged. The main architectural features that require the alterations are the non-fractal and non-hierarchical nature of the family and the wide use of

[6] Considering such a fault intolerable under the given approach is a reasonable concession. The four carry lines occupy a relatively small area and, therefore, are less susceptible to faults than the other interconnect or logic which occupy a much larger area of the die.

Figure 7. PREP 5 before and after tiling with one tile congfiguration identified

segmented, overlapping interconnect. The former requires that the tiling algorithm be altered, and the latter requires the use of the inter-tile interconnect approach described in Section 1.2.

The logic in the 4000 family is not broken up into cells, arrays, and pipes as the CSRC device. Instead, there is simply a general array of configurable logic blocks (CLBs) which, along with the local interconnect, can be tiled into smaller groups of CLBs. Figure 7 shows an example of a small design (from the PREP benchmark set) implemented on the 4000 architecture. The first shows the original layout, and the second shows the design after tiling and with one configuration for one tile identified with two spare CLBs.

The placement and shape of the tiles are determined by the following three key factors listed in decreasing order of importance: amount of interconnect across the tile interface, tile logic density, and tile size. Tiling lines are drawn across areas with little inter-tile interconnect to ease the interface locking process and minimize the performance degradation. The logic density of each tile must allow some unused logic for redundancy and should be flexible and malleable to enable various configuration possibilities. Tile size is also a factor, as large tiles may incur large overhead and low fault recovery levels as described in Section 2.1. If the first tiling attempt does not meet the user area or fault recovery specifications, the algorithm must repeat and find a different partition. Once tile lines are drawn, the independent tile implementation becomes quite similar to that for the CSRC device. Instances of each tile are generated at design-time that leave a portion of the tile (CLBs and interconnect) unused. When a fault occurs, a tile instance can be activated that does not utilize the faulty resource.

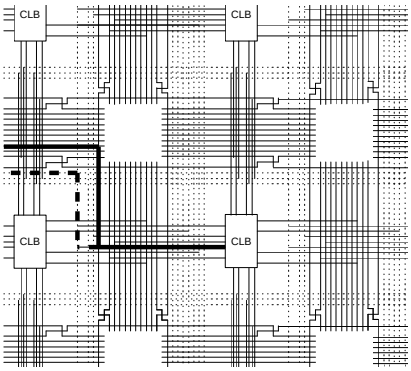

Figure 8. Inter-tile interconnect fault recovery on the Xilinx XC4000

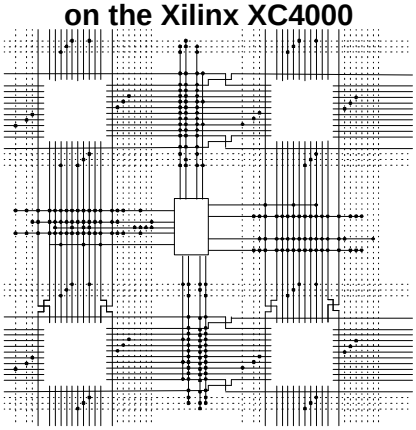

Figure 9. Sample of interconnect connections available on XC4000 family

The second change for implementation on the 4000 family involves the inter-tile interconnect. The 4000 architecture contains many lines that cross tile boundaries, and many are segmented and overlapped. Figure 8 shows an example of how the use of the inter-tile interconnect algorithm from Section 1.2 could be implemented on the 4000 architecture. The solid lines represent the segmented interconnect, and the dotted are the global lines. The solid bold line shows the signal before a fault is detected on one of the segmented lines, and the dotted line reflect a possible recovery from that fault using the global line

The solution works theoretically on this architecture, but the current implementations would not allow the algorithm to work in practice. Although lines could be made available to backup inter-tile lines, the connections do not currently exist for such an implementation, preventing the signals from making the proper routing alterations (see Figure 9). The general architecture supports the approach, but not all connections were made possible in the implementation. There would not be a significant cost increase for adding possible connections, making the eventual implementation of the algorithm on the architecture a possibility.

2.3. Altera Flex 10k

The Altera Flex 10k family is more similar to the CSRC architecture than the Xilinx XC4000 family. The hierarchical structure returns, and the interconnect is more contained. Therefore, the implementation on this architecture is similar to that on the CSRC device.

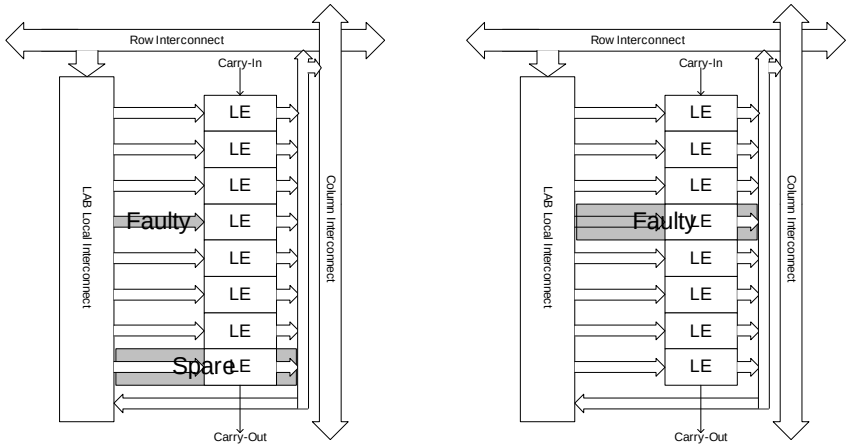

Figure 10. Original and recovered configuration after an internal LAB fault

The unit analogous to the CSLC is the logic element (LE), eight of which comprise the logic for a logic array block (LAB), which is analogous to the CSLA. The local interconnect within a LAB is also quite contained, allowing for a redundancy similar to that used in the CSLA, as it is structured like the Level 1 interconnect. One LE can be redundant for the others within the group of eight, and local interconnect can be set aside in each instance of the LAB generated at design-time by the CAD tools. Figure 10 shows the internal structure of a LAB and how the block can recover from an LE (or associated interconnect) fault.

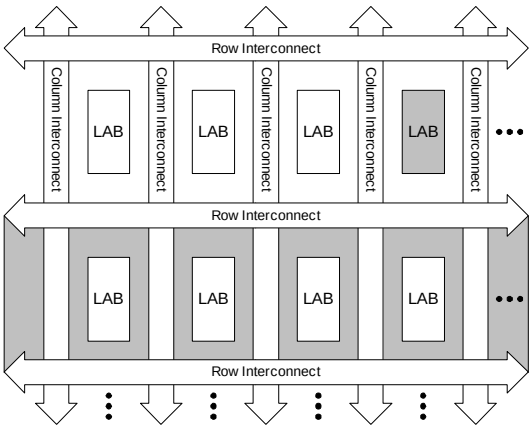

Figure 11. Hierarchical redundancy (shaded areas are spares)

The hierarchical structure that makes for an efficient implementation on the Flex 10k family as was done for the CSRC device also creates the same problems that existed for the CSRC device. LEs have carry lines that cannot easily be made fault-tolerant within a single LAB. Fortunately, Flex 10k similarities to the CSRC device allow for similar solutions to be available. Therefore, the next level of redundant block size must be inspected, beginning the fractal-type approach for this architecture.

Groups of LABs form logic arrays, similar to the pipe level of the CSRC device, and the row or column interconnect is analogous to the Level 2 routing. One LAB can be redundant for the others within its pipe, and configurations can be generated for each LAB that may become faulty. Carry lines may flow in and out of pipes,

potentially requiring that entire logic arrays also have a spare on the device, in the same way that pipes could have a spare on the CSRC device.

Using these levels of redundant blocks in combination on the Flex 10k architecture (Figure 11 reveals the hierarchy) can yield the same benefits as those described for the CSRC architecture, including reduced area and timing overhead and increased levels of fault recovery.

3. Experimental results

The key factor to consider when evaluating this approach is the additional reliability provided the system by its implementation. However, overhead in terms of area (physical resources) and timing must be considered to reveal its efficiency compared to traditional redundancy-based approaches to fault recovery. If a higher reliability is necessary, area and timing overhead may suffer. Conversely, if area and/or timing is a strict system constraint, reliability must be sacrificed. The approach presented in the paper is flexible with respect to this reliability/overhead tradeoff.

For the Xilinx XC4000 family (the architecture least friendly to the implementation of the technique), the proposed approach and optimization algorithms were applied to nine MCNC designs, with various logic and interconnect densities. Table 1 shows the timing and area metrics for the designs before and after the application of the fault recovery approach. This overhead is flexible, and the values in Table 1 only reveal one possible implementation instance. This instance was used for the reliability calculations displayed in Tables 2 and 3.

Table 1. Timing and area overhead

Design	Slowest - Fastest / Fastest	Final – Original / Original
9sym	0.21	.072
c499	0.25	.026
c880	0.17	.049
duke2	0.42	.077
rd84	0.45	.041
planet1	0.39	.056
styr	0.28	.039
s9234	0.41	.063
sand	0.26	.102

Table 2 shows reliability improvements for the MCNC benchmarks under the uniform random fault model. The first column indicates the assumed probability (p) that a physical element (logic or interconnect) is fault free. The next two columns show the probability that the original and fault recoverable designs of a particular benchmark are functioning properly. Table 3 shows the reliability for the same set of designs (original and tiled) with four different cluster variability factors, μ, assuming a 90% fault free probability for physical resources.

Table 2. Reliability of the original and tiled designs against resource reliability

p	.900		.950		.990		.999		.9999	
Design	Orig.	Tiled	Orig.	Tiled	Orig.	Tiled	Orig.	Tiled	Orig.	Tiled
9sym	0.78	10.99	7.54	36.53	47.5	91.8	91.7	99.5	96.8	100
c499	0.01	1.06	2.08	24.38	32.2	83.3	85.2	98.3	96.6	100
c880	0.00	0.39	1.17	20.61	28.7	85.2	84.8	98.1	97.1	100
duke2	0.01	0.46	1.82	20.74	29.6	84.7	85.2	98.4	96.5	100
rd84	4.68	25.09	18.01	48.56	58.9	91.3	92.6	99.2	97.8	100
planet1	0.01	3.51	7.48	21.19	28.8	89.6	85.7	98.0	96.1	100
styr	0.01	1.89	1.63	20.41	32.3	87.5	88.1	97.9	97.5	100
s9234	0.00	0.00	0.01	2.06	9.8	75.8	75.6	98.3	96.2	100
sand	0.02	0.99	1.19	1.63	31.4	85.2	83.5	98.6	96.9	100

Table 3. Reliability of original and tiled designs using Stapper's correlated failure model [14] with resource reliability of 90% and a variable μ

	1		5		20	
	Orig.	Tiled	Orig.	Tiled	Orig.	Tiled
9sym	40.69	46.80	18.53	26.07	5.70	19.31
c499	37.64	44.33	12.94	23.53	1.79	9.43
c880	36.34	42.71	11.90	21.97	1.35	7.28
duke2	37.44	44.14	12.61	16.84	1.65	9.30
rd84	43.16	49.86	23.99	36.66	11.70	34.32
planet1	37.51	44.14	12.68	16.84	1.68	9.30
styr	38.29	44.27	14.04	25.74	2.36	10.01
s9234	34.52	41.86	8.52	18.46	0.41	3.46
sand	37.96	44.20	13.46	20.87	2.05	9.62

The same evaluations can be performed on the Sanders CSRC and Altera Flex 10k architectures. The analysis in Sections 2.1 and 2.3 reveals that the hierarchical nature and minimal segmented overlapped interconnect of the CSRC architecture and the Altera's Flex 10k family create even greater opportunities for efficient implementation. Therefore, the above results reveal a worst case analysis in terms of applying the approach to the diverse architectures examined here.

4. Conclusions

Efficiently implementing a runtime fault recovery algorithm on a variety of architectures shows the flexible nature of the algorithm and reveals architectural features that enhance or hinder the approach. Runtime implementation of the algorithm on each architecture is straightforward and minimizes system downtime, and the approach minimizes the amount of memory and CAD tool effort required. Experimental results reveal that the area and time overhead remain low and that the fault recoverability is high even for the Xilinx architecture, the most difficult on which to implement the approach.

5. References

[1] A. L. Burress, P. K. Lala, "On-line Testable Logic Design for FPGA Implementation", *International Test Conference*, 471-8, 1997.

[2] W. Feng, W. K. Huang, F. Lombardi, "Structural Testing of Programmable Interconnects", *Journal of Microelectronic Systems Integration*, vol. 5, no. 3, 129-44, Sept. 1997.

[3] C. Metra *et al.*, "Novel Technique for Testing FPGAs", *Design, Automation and Test in Europe*, 89-94, 1998.

[4] M. Nicolaidis, "On-Line Testing for VLSI: State of the Art and Trends", *Integration, The VLSI Journal*, vol. 26, no. 1-2, 197-209, Dec. 1998.

[5] M. Renovell *et al.*, "RAM-Based FPGAs: A Test Approach for the Configurable Logic", *Design, Automation and Test in Europe*, 82-8, 1998.

[6] M. Renovell *et al.*, "Testing the Interconnect of RAM-Based FPGAs", *IEEE Design & Test of Computers*, vol. 15, no. 1, 45-50, Jan.-March 1998.

[7] J. Lach, W.H. Mangione-Smith, M. Potkonjak, "Low Overhead Fault-Tolerant FPGA Systems", *IEEE Transactions on VLSI Systems*, vol. 6, no. 2, 212-21, 1998.

[8] G. A. Allan, A. J. Walton, "Automated Redundant Via Placement for Increased Yield and Reliability", *Proceedings of the SPIE - The International Society for Optical Engineering*, vol. 3216 (Microelectronic Manufacturing Yield, Reliability, and Failure Analysis III), 114-25, 1997.

[9] F. Distante, M. G. Sami, R. Stefanelli, "Harvesting Through Array Partitioning: A Solution to Achieve Defect Tolerance", *IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, 261-9, 1997.

[10] I. Koren, Z. Koren, "Defect Tolerance in VLSI Circuits: Techniques and Yield Analysis", *Proceedings of the IEEE*, vol. 86, no. 9, 1819-38, Sept. 1998.

[11] S.M. Scalera, J.R. Vázquez, "The Design and Implementation of a Context Switching FPGA", *6th IEEE Symposium on FPGA-Based Custom Computing Machines*, 1998.

[12] Xilinx, The Programmable Logic Data Book, San Jose, CA, 1996.

[13] Altera, *Data Book*, San Jose, CA, 1996.

[14] C. H. Stapper, "A New Statistical Approach for Fault-Tolerant VLSI Systems", *The Twenty Second International Symposium on Fault-Tolerant Computing*, 356-65, 1992.