

Classical Versus Evolved Fault Tolerance: Comparing Metrics and Performance

Pauline C Haddow, Morten Hartmann and Asbjørn Djupdal

Complex, Reconfigurable, Adaptive and Bio-inspired Hardware (CRAB)Lab
Dept. of Computer and Information Science
Norwegian University of Science and Technology
(pauline,mortehar,djupdal)@idi.ntnu.no

Abstract

Triple-modular redundancy has been a successful fault tolerant approach, improving reliability in digital systems for over 40 years. The expansion of space exploration during this period has increased the need for reliable systems in dynamic environments. Even with the invent of reconfigurable technology, correction of failures is difficult as the access window to reconfigure devices is getting smaller and smaller as space exploration moves further and further afield. Evolvable Hardware, applying artificial evolution to the design of digital and analogue hardware, is an emerging field that shows much promise within the area of fault tolerant design in dynamic environments. However, applying traditional metrics to these non traditional techniques is far from straight forward. This paper considers reliability metrics for traditional and evolved fault tolerant circuits and experiments are conducted on simple multiplier circuits.

1 Introduction

The need for fault tolerance is an important issue of modern electronic design. High density chips increase the possibility of failing compo-

nents and the complexity of design increases the probability of human errors. Space exploration in unknown and dynamic environments places a greater pressure on fault tolerance in terms of the faults themselves and the cost of repair. The need for fault tolerant designs is stated amongst the long term grand challenges of the International Technology Roadmap for Semiconductors, ITRS (2005) [3].

The introduction of reconfigurable technology has, in some ways, reduced the need for fault tolerance as the reconfigurable nature of these devices enables runtime correction, assuming a fault detection mechanism is available. However, when one considers space applications, runtime correction is not so easy. Reconfiguration requires uploading the new configuration when access windows permit access to the device. For small satellites, a 10 minute access window may be available 2/3 times a day with transfer rates of perhaps only 10 to 19kbits/sec [7]. With, for example, Xilinx's virtex II's 10Mbit configuration stream, even with the possibility of partial reconfiguration and techniques such as data compression and splitting and fusing of configuration data, fault tolerance still needs to be addressed to reduce the need to reconfigure.

The field of evolvable hardware, EHW, where

artificial evolution is applied to the design of electronic circuits, is a promising field in the area of fault tolerant design for dynamic environments. However, how good are the results being produced? One major problem is that this field is still in its infancy, creating relatively simple designs and evaluating these designs using internal metrics e.g. number of generations, which have no meaning at all in terms of traditional designs.

In this paper we address the fault tolerant metric reliability, using two variations of this metric. The first R_{trad} , as used in traditional design evaluation, is a measure of the probability that a system will not fail under specified conditions. This reliability measure is, of course, a measure of the probability that a circuit is 100% functional and says nothing about how dis-functional the circuit is when it is not 100% functional. The second reliability measure applied R_{ehw} , most often applied as the fitness measure in design by evolution — see section 2, represents how correct the solution is, on average, over a spectrum of faults. All experiments are conducted on simulations of a simple electronic circuit — a 2 bit multiplier, and the traditional designs evaluated are based on a non-redundant multiplier and a multiplier design with triple modular redundancy (TMR) [8].

The paper provides an introduction to Evolvable hardware and, in particular, evolving fault tolerant designs in section 2. Section 3 provides information relating to the simulator designed for the evolution of fault tolerant designs and also applied to fault testing of the traditional circuits. The experiments conducted, their results and analysis of these results are given in section 4. Finally section 5 draws some conclusions from this work.

2 Evolvable Hardware

The principles of artificial evolution are based on Darwin's theory of evolution by natural selection. His theory was first adapted to the field of artificial evolution in 1975 [2]. The main features of natural evolution consisting of reproduction by cloning, mutation and crossover through the exchange of and alteration of genetic material are included in today's evolutionary algorithms. However, an artificial selection mechanism was introduced, which unlike that in nature, steers artificial evolution towards a given goal i.e. a solution to a given problem.

The application of evolutionary techniques to hardware design is termed evolvable hardware (EHW) [9, 6], the main goal being to replace traditional design methods with evolutionary techniques for given hardware applications. These applications are either not achievable using traditional methods or would benefit from an evolutionary approach. A number of algorithms have been developed for evolutionary design and many have been applied to evolvable hardware. To highlight some of the main concepts behind the concept of an evolutionary algorithm, a genetic algorithm is described in section 2.1. It should be noted, however, that the evolutionary algorithm applied in this work is termed Cartesian genetic programming [5].

2.1 Genetic Algorithms

The evolutionary process is a dynamic process where at a given point in time we have what is termed in biology a generation: a population of individuals for the given species. Biologically, a generation change is not an event but a constantly unfolding process. However, in artificial evolution, we introduce new generations through the application of genetic operators to selected individuals within the current population to form a new population of individuals i.e.

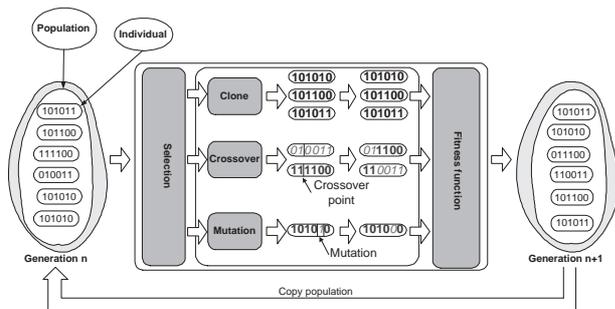


Figure 1: A Genetic Algorithm

a new generation. Unlike biological evolution, in artificial evolution we have a concrete goal, that is to create a correct solution i.e. a functionally correct circuit which solves a given problem. The goal is specified in terms of the fitness measure. This measure may define the functionality of the circuit i.e. how correct a given circuit solution is. However, it may also take into account other features such as area usage and power consumption and weight these different factors. It is the purpose of the selection mechanism to steer evolution towards the goal described by the fitness measure.

Figure 1 illustrates the process of a genetic algorithm. The set of individuals — circuit solutions, make up the population in the current generation. Individuals are selected and genetic operators such cloning (copying), crossover (swapping information between individuals) or mutation (inverting a bit) may be applied before fitness is calculated and the individuals are placed in the new population. When the new population is full then the next generation can begin. The process continues until an individual in the population achieves 100% fitness or the simulator is stopped after a certain number of generations.

2.2 Evolving Fault-tolerant Designs

In the process of evolving a design, the mutation operator may be said to act as noise in the evolution process thus, in some cases, giving rise to a certain amount of implicit fault tolerance in the evolved design. As such, even with 0% faults applied, a certain amount of fault tolerance may be present in the evolved design. In this work, evolved designs are both implicitly and explicitly designed for fault tolerance.

Explicit fault tolerance is achieved by testing individuals during evolution with various fault scenarios and evaluating how well the individual functions under these scenarios. These results, R_{ehw} , are used to provide a fitness measure for the individual. The goal is that this fitness measure combined with the selection method applied are suitably defined so as to drive evolution towards finding good fault tolerant solutions.

R_{ehw} provides evolution with information, not only on the correct solutions but on how good the non correct solutions are. As such, the R_{ehw} measure provides more fine grained information to the evolutionary process about the quality of different solutions than that which would be achievable with R_{trad} .

3 Simulator

3.1 Experimental Setup

All experiments were conducted on simulations of 2 bit multiplier designs where faults were applied as worst case scenarios — inversion of the output, and applied on a fault rate per gate basis. Each point in the graphs represents the reliability measure under investigation applied to a circuit which has been tested with 1000 fault scenarios at the given fault rate. Fault rates investigated ranged from 0.0 to 0.2.

For the evolved circuits, a population of 20 randomly generated individuals was given to the

evolution process. The selection method applied was tournament selection with elitism [4], crossover was applied at a rate of 0.2 and mutation at a rate of 0.05. In these experiments, mutation was applied at the gate level. A mutated gate either has one of its input connections rewired or its type changed. The evolution of a single solution was allowed to continue for a maximum of 100,000 generations and the maximum size of an individual (the genome) was 80 gates.

The goal of evolving a fault tolerant circuit solution can be interpreted as the evolution of a circuit that successfully produces the correct mapping between input and output vectors despite faults. The target behavior of the multiplier was specified by a truth table. Each individual of the evolution process was evaluated by applying the complete set of input vectors and calculating the hamming distance between the output of the circuit to the target truth-table. The fitness was then the average of this result for the 10 fault scenarios. The fitness result was normalised between 0 and 1, 0 being no correct outputs and 1 being all correct outputs. The evolutionary process was repeated 20 times to achieve 20 individuals. For a given fault rate, the results plotted represent the best reliability result (either in terms of R_{trad} or R_{ehw}) of the 20 evolved circuits, tested under the 1000 fault scenarios.

One important consideration when applying faults as faults per gate is that the bigger the circuit i.e. more gates present, the larger the challenge one faces with respect to reliability. As such, the goal was to apply a traditional designed multiplier using a minimum number of gates to the experiments. A 9 gate multiplier was thus implemented and applied as the traditional design in the simulator.

To incorporate traditional fault tolerance, TMR was applied to the circuit. It was chosen to apply TMR to the multiplier and not to

the voter. This choice was based on the fact that applying TMR to the voter would drastically increase the number of gates in the circuit but there would still be the problem of three outputs instead of one.

3.2 Evolved circuit Representation

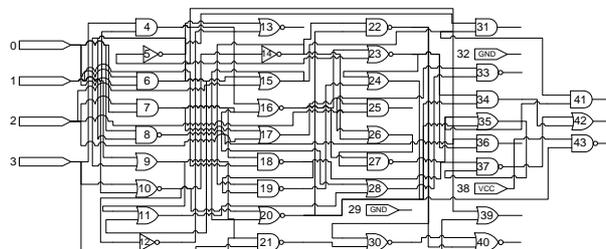


Figure 2: Example Circuit

what	label	type	input A	input B
input	0			
input	1			
input	2			
input	3			
gate	4	AND	3	1
gate	5	NOT	4	
gate	6	AND	0	0
gate	7	AND	2	0
.				
.				
gate	38	VCC		
gate	39	OR	4	9
gate	40	NOR	30	22
gate, output	41	AND	15	27
gate, output	42	OR	37	34
gate, output	43	NAND	38	24

Figure 3: Symbolic Netlist of the Example Circuit

Both evolved and traditional circuits were expressed in the simulator using a symbolic netlist representation. Figure 3 provides a symbolic description of the circuit illustrated in figure 2. Each row represents a gate, with its type and from where to connect its two inputs. Connections can only be made to lower numbered gates (or the circuit inputs). This simple scheme assures a pure feed-forward network and thus combinatorial circuits. The last gates in the netlist are connected to the circuit output.

4 Experiments and Results

The goal of these experiments was to understand and compare the effect of using the two metrics : reliability traditional R_{trad} and reliability evolvable R_{ehw} , with respect to digital circuits — in particular a 2 bit multiplier.

4.1 Reliability of Evolved Circuits

In earlier work on evolving reliable circuits [1], R_{ehw} was the reliability metric applied. Faults were applied both at the inputs and output to a gate providing the possibility that a fault would not necessarily result in a gate failure. Fault scenarios applied during evolution were based on the fault rate per gate under investigation. The best individuals from 20 evolutionary runs were then each tested 1000 times using the same fault rate as applied during evolution. R_{ehw} was calculated for each of the 20 results and the average as well as the best and worst results were plotted.

These experiments are repeated, herein, but with a worst case fault scenario — inverting the output to cause a gate failure. The results are displayed in figure 4 as R_{ehw} . It should be noted that in this case R_{ehw} represents the best result of the 20 individuals. The reason for choosing the best result was to illustrate what evolution is capable of rather than what it achieves on average. As shown, evolution achieves circuits with over 70% reliability even with 20% faults in a worst-case fault scenario.

These figures may, of course, be said to be misleading, in terms of a more traditional view of reliability. As such, it was important to evaluate these circuits in terms of R_{trad} . As such, R_{trad} was calculated, similar to R_{ehw} , as the best result for the 20 individuals under the 1000 fault scenarios — see R_{trad} in figure 4. As can be seen from the results, at a fault rate of 0.04 and greater, evolution does not achieve any circuits that are 100% functional.

The break at 0.04 would indicate that at this point the task of achieving 100% correctness becomes too hard for evolution. However, as indicated by R_{ehw} , evolution manages to retain a reasonable sub-optimal solution, degrading gracefully from around 90% R_{ehw} at 0.04 to around 75% at 0.2. It should be noted that these results where the fault rate used during evolution matches the fault rate applied in the 1000 fault scenarios is termed “increasing fault rate” herein.

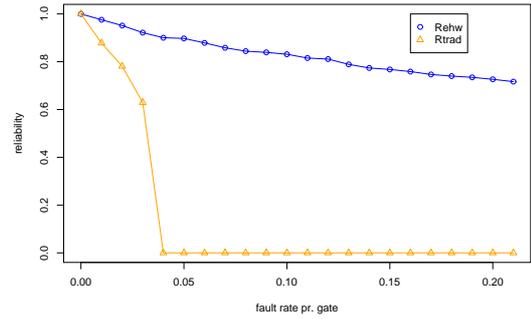


Figure 4: Circuits Evolved with Increasing Fault Rates

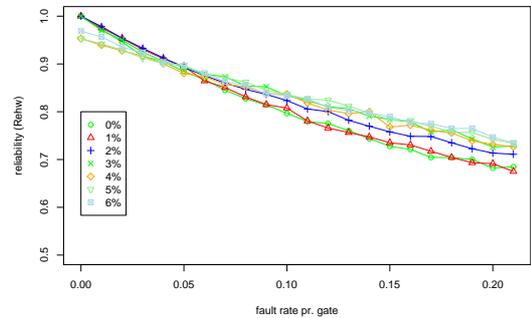


Figure 5: R_{ehw} : Circuits Evolved with Fixed Fault Rates (0 to 0.06)

Looking again at R_{trad} in figure 4 made us wonder if we could exploit the better results

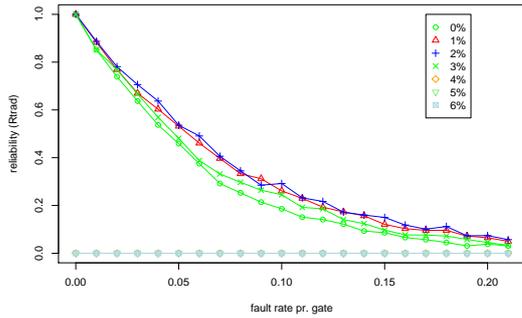


Figure 6: Rtrad: Circuits Evolved with Fixed Fault Rates (0 to 0.06)

achieved at fault rates from 0 to 0.03. How reliable were these circuits at higher fault rates than those fault rates that the circuits were exposed to under evolution? To investigate either side of the breakpoint it was chosen to investigate fixed fault rates of 0 to 0.06 i.e. 7 sets of experiments. That is, the evolved circuits from each set of experiments were tested with 1000 fault scenarios for each of the fault rates from 0.0 to 0.2. Both R_{ehw} and R_{trad} were measured, as illustrated in figures 5 and 6 respectively. For the case of R_{ehw} , as shown, circuits evolved with lower fault rates perform well below fault rates of 0.04. On the other hand circuits evolved with higher fault rates perform well above fault rates of 0.04. It should be noted that the scale on the y-axis has been increased, compared to other graphs so that the results are more visible. There is a substantial improvement in R_{trad} when fixed fault rates of 0.0 to 0.04 are applied. One may assume that fault rates of 0.05 and 0.06 provide too difficult a task for evolution and, even at a fault rate of 0.00, the best evolved circuits for 0.05 and 0.06 are not able to produce any 100% functional circuits from the 1000 fault scenarios. The best results from fixed fault rates for both R_{ehw} and R_{trad} are displayed in figure 7.

Comparing the fixed fault rate solution (“Best

of 0-6%”) with that of increasing fault rates (“Evolved with given fault rate”) in figure 8 one can see that there is little difference in R_{ehw} . However, R_{trad} is substantially improved — see figure 9. Even at a fault rate of 0.2 some 100% functional solutions may be found.

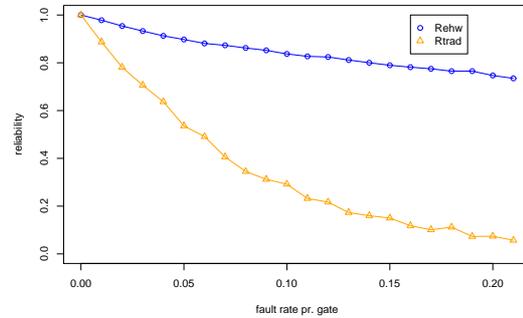


Figure 7: Circuits Evolved with Fixed Fault Rates (0 to 0.06)

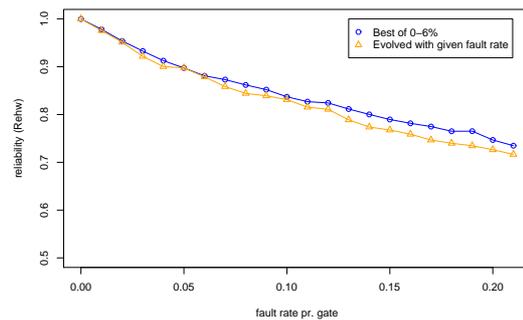


Figure 8: Rrehw: Evolved Fixed vs Evolved Increasing

4.2 Reliability of Traditional Circuits

To calculate R_{ehw} for the traditional circuit, 1000 fault scenarios were applied and their average calculated for each of the fault rates from 0 to 0.2. Figure 10 illustrates the results for both

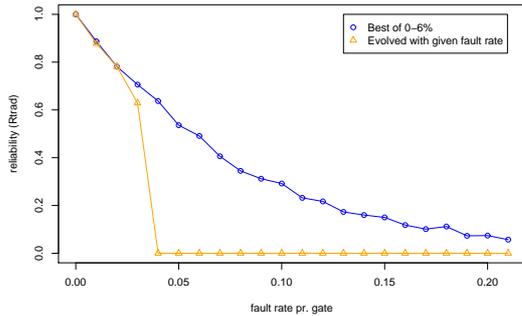


Figure 9: Rtrad: Evolved Fixed vs Evolved Increasing

the traditional and traditional with TMR version of the circuit. Unfortunately, but not unexpected, the weighting of the 16 vulnerable gates in the voter compared to the small 9 gate multiplier modules is a big disadvantage to TMR as a methodology. Also the sheer size of the circuit makes the average number of occurring faults relatively large due to faults being generated based on gate reliability. The traditional circuits without redundancy are, therefore, more reliable than those with TMR.

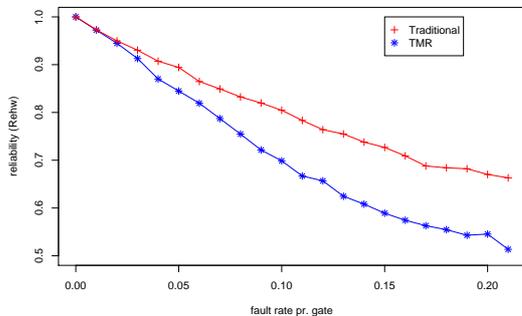


Figure 10: Rehw: Traditional Circuits

Figure 11 illustrates R_{trad} for both the traditional and TMR version of the tradition circuit. Statistical estimates which illustrate the

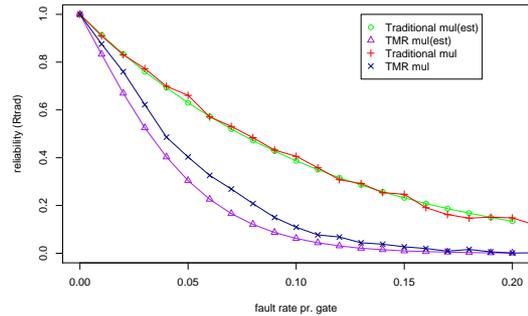


Figure 11: Rtrad: Traditional Circuits

validity of the results received are also plotted. It should be noted that the TMR multiplier results deviate from the statistical estimates. The reason for this is that the statistical estimates do not take into account a number of features inherent in TMR. One such feature is the fact that two faults may occur on a TMR triple gate without resulting in gate failure. As such, the measured reliability will be higher than the estimated results. It may, also, be noticed that the traditional circuit with TMR performs, as in the R_{ehw} case, poorer than the non redundant version.

Similar to the evolved circuits, it may be seen that R_{ehw} results give a more positive impression of reliability than the R_{trad} results.

4.3 Traditional vs Evolved

Rehw for the traditional circuits is somewhat poorer than the evolved results — see figure 12, indicating that the evolved solutions perhaps exhibit more graceful degradation. On the other hand, one can see from figure 13 that the evolved circuits are less reliable, with respect to R_{trad} , than the traditional solutions without TMR. As stated, in the evaluation phase the evolving circuits are not being optimised against a specific fault scenario. Instead a sub-optimal solution

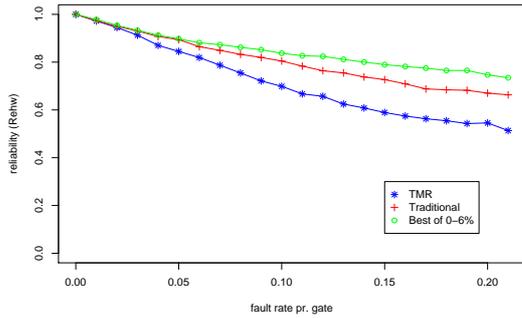


Figure 12: Rehw: Traditional vs Evolved

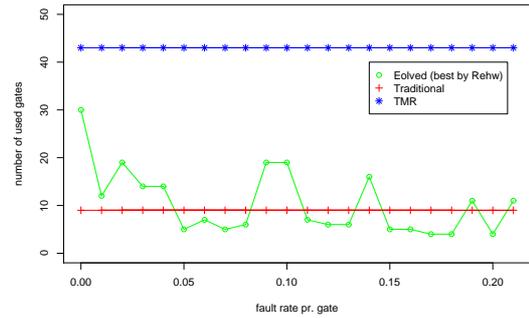


Figure 14: Rehw: Number of Gates versus Fault Rate (per gate)

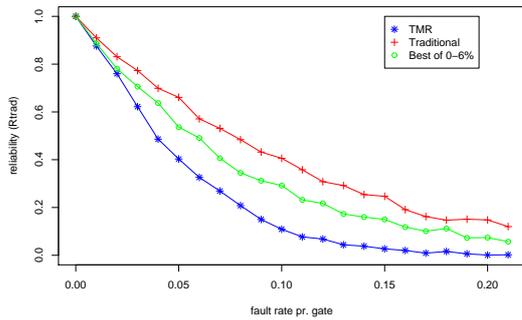


Figure 13: Rtrad: Traditional vs Evolved

is being sought with respect to individual fault scenarios which gives rise to a more optimal solution for the set of fault scenarios. Further, the evolved results achieve a better R_{trad} than the traditional circuits with TMR.

4.4 Gate Usage

Figure 14 illustrates gate usage for traditional, traditional with TMR and evolved with fixed fault rate circuits. Each point in the evolved with fixed fault rate experiments represents the best result, with respect to R_{ehw} , of the 7 result sets for each fault rate. Although the results are somewhat unstable in nature, there is a tendency for evolution to reduce the size of the circuits to

moremorte

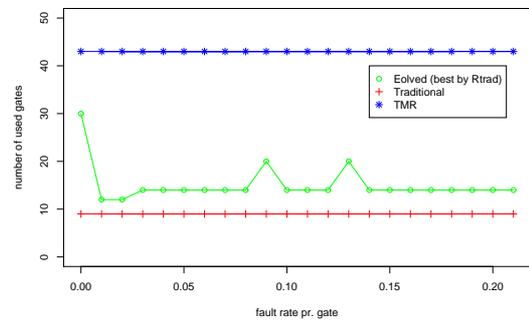


Figure 15: Rtrad: Number of Gates versus Fault Rate (per gate)

achieve better R_{ehw} values. Looking at the gate count for R_{trad} —see figure 15, more stability can be seen in the gate count for the evolved circuits. This is due to the fact that, for most of the points in the graph, it is the same circuit of the 20 evolved circuits which achieves the best R_{trad} .

5 Conclusion

This paper has addressed the issue of applying reliability metrics to both evolved and tradi-

tional circuit solutions and compared the performance of circuits with regards to those metrics in terms of a worst case fault scenario applied on a per gate basis.

When considering the results it is worth noting that the fault scenarios used herein are hard to solve. Redundance techniques, in general, can cope with single faults. The traditional circuit with TMR, for instance, would handle any single fault not occurring in the voter element. The faults herein are, however, applied as the probability of a gate failing on a per gate basis. All circuits are tested 1000 times in arbitrary fault scenarios. Several of these scenarios will present the circuit with a large number of failing gates, often including the output gates. Several of the generated scenarios may, thus, be fault scenarios that are impossible to solve, regardless of the circuit architecture whether evolved or designed using traditional techniques.

Evolved circuits, herein, are evolved with the goal of improving R_{ehw} in the presence of faults. On the other hand, R_{trad} is a design metric tuned to traditional design methods. As a consequence it is not so surprising that evolved circuits are more reliable in terms of R_{ehw} and traditional circuits are more reliable in terms of R_{trad} .

The circuits used in the experiments conducted may be said to be small toy problems. For larger problems observed results may differ. This is especially true in the case of TMR, as the technique is not suited for problems of the current size. Further investigations would, therefore benefit from scaling up the application. Scaling up solutions is, however, currently one of the greater challenges faced by the evolvable hardware community as a whole.

The field of evolvable hardware is still a field, far from maturity, and applying evolvable techniques is not a straight forward science. As shown, applying fixed fault rates during evolution led to greatly improved reliability, espe-

cially in terms of R_{trad} . One might expect that one could use R_{trad} to drive evolution instead of R_{ehw} and thus improve R_{trad} for evolved designs. However, R_{trad} does not provide sufficiently fine grained circuit feedback so as to separate potential solutions and, therefore, limits evolution's search for a solution to a rougher search.

Another topic that needs to be addressed in the future is the suitability of R_{ehw} and R_{trad} for different applications. Whilst R_{trad} has obvious advantages as a metric, R_{ehw} is also beneficial in certain cases such as in circuit evolution. Other cases may include those where partially correct outputs would still be of greater value than a completely incorrect output. This could, for instance, be correctness for certain input vectors in a multiplier or reduced resolution or precision in general. Also, as technology scales, production defects and lifetime faults are likely to become an increasing problem. As such, one can expect R_{trad} to fall as it is going to be harder and harder to tolerate faults 100%. It may, therefore, become important to not only investigate R_{trad} but also R_{ehw} in traditional designs to study circuits degradation in the presence of faults.

Another interesting fact observed was that, in certain cases, evolution was seen to have high tolerance to faults with up to 3 times the number of gates of the traditional non-TMR solution. In these cases a question arises. What type of redundancy is evolution using so as not to experience the negative effect of a large number of gates experienced by the TMR solution? Although this is not a general result, these cases are interesting in their own right and worth further investigation to find out whether we can improve traditional redundancy techniques based on architectural observations from evolved circuits.

References

- [1] M. Hartmann and P. C. Haddow. Evolution of fault-tolerant and noise-robust digital designs. *IEE Proc. -Comput. Digit. Tech.*, 151(4):287–294, 2004.
- [2] J. Holland. *Adaption in Natural and Artificial Systems*. The University of Michigan Press, 1975.
- [3] International Roadmap Committee. Executive summary. In *International Technology Roadmap for Semiconductors*. 2005. <http://public.itrs.net>.
- [4] J. Koza. *Genetic Programming*. The MIT Press, 1993.
- [5] J. F. Miller, D. Job, and V. K. Vassilev. Principles in the evolutionary design of digital circuits – part i. *Journal of Genetic Programming and Evolvable Machines*, 1(1):8–35, 2000.
- [6] M. Sipper, E. Sanchez, D. Mange, M. Tomassini, A. Prez-Uribe, and A. Stauffer. A phylogenetic, ontogenetic, and epigenetic view of bio-inspired hardware systems. *IEEE Transactions on Evolutionary Computation*, 1(1):83–97, April 1997.
- [7] T. Vladmirova, X. Wu, K. Sidibeh, D. Bernhart, and A.-H. Jallad. Enabling technologies for distributed picosatellite missions in leo. In *Proc. First NASA/ESA Conference on Adaptive Hardware and Systems (AHS'06)*, pages 330–337. IEEE, 2006.
- [8] J. von Neumann. Probabalistic logics and synthesis of reliable organisms from unreliable components. *Automata Studies of; Annals of Mathematical Studies*, (34):43–98, 1956.
- [9] X. Yao and T. Higuchi. Promises and challenges of evolvable hardware. In I. T. H. et al, editor, *Evolvable Systems: From Biology to Hardware. First Int. Conf., ICES 96*, volume 1259 of *Lecture Notes in Computer Science*. Springer, 1996.