

Fault Tolerance of the Input/Output Ports in Massively Defective Multicore Processor Chips

Piotr Zajac^{1,2}, Jacques Henri Collet¹, Jean Arlat¹, and Yves Crouzet¹

¹ LAAS-CNRS, Université de Toulouse,
7, avenue du Colonel Roche, 31077 Toulouse Cedex 04, France

² Department of Microelectronics and Computer Science,
Technical University of Lodz, Al. Politechniki 11, 93-590 Lodz, Poland
{pzajac, jacques.collet, jean.arlat, yves.crouzet}@laas.fr

Abstract

This paper addresses the fault tolerance issues concerning the input-output ports (IOPs) of future multicore chips built up using massively defective nanotechnologies. Recall that the IOPs are critical to system dependability as they constitute bottlenecks for all communications between the chip and external resources. Various levels of modular redundancy in the IOPs are being considered for which we calculate the probability to maintain correct operation. We also calculate the cost attached to the proposed protective designs for the IOP, in terms of circuitry overhead.

Keywords: Ultra-large-scale integration, nanotechnology, multicore processor architecture, fault tolerance.

1 Introduction

As witnessed by recent announcements, new design paradigms are being considered by silicon technology manufacturers as alternatives to the device downsizing and increase frequency races that is anticipated to soon lead to a dead end for what concern power dissipation. Indeed, manufacturers are increasingly considering large-scale multiprocessor chips (LSMC) [also often referred to as chip multiprocessor (CMP)] and even grid architectures for future processor chips [1], in place of SMP-derived architectures. Similarly, the availability of a large number of basic processing cores, say tens or several hundreds is also assumed for the future chips implemented with nanoscale technologies emerging from molecular electronics. Reliance on redundant constructs is being considered to overcome the non-deterministic and failure-prone behaviors of the underlying devices [2]. The reason for this evolution is simply that replication of one or several basic blocks is the simplest manner (and perhaps the only one) to control the physical complexity of large systems.

Anticipating with the need to cope with massively defective nanotechnologies, in a previous work, we have studied the impact on yield and resilience that can be achieved by

exploiting the multiplicity and connectivity offered by the nodes within regular grid architectures to support on-line reconfiguration strategies [3, 4].

In practice, several issues have still to be solved for such architectures to be readily operational. Indeed, the beneficial evolution in the physical layer generates number of problems in the other layers. These concerns relate system programmability (for instance parallelization and dynamic allocation), execution control and communications. Recent efforts have proposed software-based solutions that modify the application software to circumvent faulty cores in a LSMC with no hardware cost and performance overhead for the fault-free cores [5]. In this paper we concentrate on communication issues. The related problems concern two levels of the LSMC architectures:

- intra-chip: communication between the cores within the chip;
- inter-chip: communication between the chip and its environment via dedicated input-output ports (IOPs).

For what concerns *intra-chip communication*, it is known for decades that the increase of the number of cores in a multiprocessor system (typically more than 32 cores) cannot be envisaged as part of SMP architectures, because the shared memory bus becomes inexorably a communication bottleneck, which slows down the transactions with the memory (e.g., see a discussion of this problem in [6]). Of course, it is always possible to consider architectures with several buses or based on non-blocking crossbars, but this only provides a very partial answer to such communication problems. In such a context, one better understands the motivation for providing core arrays, and for instance the 80-core grid by Intel [1]), even if regular arrays have also specific issues, as for instance the rapid increase of the latency access to remote data within the grid. In the sequel, we concentrate on grid architectures to support the implementation of general-purpose processing chip (GPP).

For what concerns communications via IOPs, the key problem is that increasing the number of cores inevitably increases the need and intensity of the communication with the environment which may result in a communication bottleneck if only one single IOP is considered irrespective of the number of cores. Accordingly, the obvious approach is to consider multi-port chip architectures. Another important issue is the protection of these IOPs when considering chips made up with massively defective nanotechnologies. The investigation of mechanisms for tolerating the impact of faulty cores (e.g., see [7-9]) is out of the scope of this paper.

The rest of the paper is organized into 5 sections. Section 2 briefly describes the kind of target architecture considered in the paper. In Section 3, we analyze the tolerance of faults in the IOPs. We consider the application of R -modular redundancy in the design of the IOP and the use of external tests to select r fault-free modules out of R in each IOP. Section 4, addresses the tolerance of faults in the vicinity of each IOP. Finally, in Section 5, we estimate the attached cost in terms of circuitry overhead to protect the IOPs and their vicinity. In particular, this section depicts the scalability issues in defective technologies. Finally, some concluding remarks are drawn in Section 6.

2 The Target On-chip Grid Architectures

The specificity of the considered on-chip grid architectures is that downsizing enables increasing the core number, but scaled down technologies are becoming increasingly faulty or vulnerable to radiation effects [10-12]. Moreover, non deterministic and unreliable behaviors are anticipated for both C-MOS and non C-MOS nanoscale technologies device [13]. Thus, there is a significant probability for the nodes in the grid and for the IOPs to be faulty.

However, tolerating faults in the IOPs or tolerating faults affecting the cores or interconnects within the grid is not at all the same problem. As we have previously proposed in [3, 4], mitigating the impact of a fraction of defective nodes in a grid can be achieved by stopping all communications with the defective elements (that are identified by mutual tests), and in discovering the suitable communication routes in the grid which circumvent the defective core nodes. This does not apply for the IOPs: they cannot be by-passed as they implement the necessary communication between the chip and the environment.

Increasing the number of IOPs would not be much useful: on one hand, it would reduce the risk of saturating the traffic on each port, but, on the other hand, there would be an increased risk for having faulty IOPs as we augment their number to balance the traffic! Indeed, none IOP must fail. For example, in a 4-IOP chip, if one assumes the probability for an IOP to be fault-free to be 0.8, then the probability that all IOPs be non-defective could be estimated as $(0.8)^4 \approx 0.41$. Consequently, there would be a

significant risk that the majority of chips would not operate correctly if no protection mechanisms were implemented.

Consequently, the IOPs and their vicinity (e.g., the first of level of connected core nodes) are definitely the most critical zones of a LSMC using a massively defective technology.

For sake of clarity, Figure 1 depicts the type of grid architecture that we are typically considering in this work. It features a 2D array with 7x9 nodes and 4 IOPs (labeled N, E, S, W) located in the middle of each edge. Each node is made up with a processing core and a router, respectively represented by a square and a circle. The solid lines between two routers denote the interconnects. This grid is massively defective, as it contains 14 faulty cores out of 59. The IOP are also possibly faulty.




Figure 1: Example of 4-IOP grid including 14 defective cores

Note that in this example, the IOPs are located on the edges of the chip. However, flip chip bonding has been today generalized and devices are assembled in a “face down” orientation to a package carrier (substrate) by using a flip chip bonder. It results that the IOPs can be placed anywhere in the grid.

3 Fault Tolerance in the IOPs


Several complementary strategies are possible for the direct protection of the IOPs:

- Fault avoidance: For instance, one may think of using hardened technology at the component level. However, such an approach is complex when it is necessary to mix several technologies on the same chip.
- Fault tolerance: Applying R -modular redundancy and possibly majority voting at the level of each IOP.

In the case of fault tolerance, it is important to distinguish the types of faults considered: manufacturing defects or operational faults (most likely, transient disturbances).

Concerning manufacturing defects, an obvious solution consists in testing the modules with an external controller, and in selecting one of them through multiplexing and demultiplexing circuitry (MD).

Figure 2 shows a possible R -redundant IOP (denoted RIOP) featuring 3 modules M_i ($i = 1, 2, 3$) — thus, $R=3$, 4 links (labeled n, e, s, w) for connection to the 2D-grid and a link (labeled X_{tern}) connecting to the environment.



Note: b_1 and b_2 are two digital inputs to select one of the 3 M_i modules, based on the results of external tests. Bold lines show the links that are activated when module M_3 is selected.

Figure 2: RIOP design with $R=3$ redundant I/O modules

At start up, a chip could be simply validated as operation if all RIOPs are operational, considering that a RIOP is operational if at least one of its modules is deemed as fault-free following the external verification/diagnosis test.

For what concerns transient faults in operation, different techniques should be applied. Replicated executions at runtime are mandatory, possibly with level 2 redundancy for dual check and level 3 for error masking.

Accordingly, we rather suggest validating the grid if at least 3 out of R modules are fault-free at start-up in *each* RIOP. Thus, the probability to validate a chip according to this criterion is simply:

$$P_{W, IOP} = \left[\sum_{i=3}^R \binom{R}{i} (1 - p_{f,M})^i p_{f,M}^{R-i} \right]^{N_{IOP}} \quad (1)$$

where N_{IOP} is the number of RIOPs and $p_{f,M}$ the permanent-fault probability of each M_i module.

A key point of this approach is that the fault probability of the MD circuitry (see Figure 2) should be much smaller (and even *negligible*) compared to that of each M_i module which can comprise several million transistors.

Various sets of probabilities $P_{W, IOP}$ are plotted in Figure 3 for a 4-port chip (i.e., $N_{IOP} = 4$). Note that each curve is indexed with a label r/R where: r (the required number of

fault-free modules (M_i) in each RIOP — see Figure 2) is always 3, and R denotes the redundancy level of the RIOP.

For instance, point A in Figure 3 (coordinates X_A and Y_A) shows that for: i) a redundancy $R = 6$ and ii) a fault probability for each of the M_i modules $p_{f,M} = X_A = 0.2$, the probability $P_{W, IOP}$ that at least $r = 3$ of the $R = 6$ M_i modules in each RIOP are fault-free is approximately $Y_A = 0.93$.




Figure 3: Validation probability of a 4-port chip, for $r = 3$ and $R = 5, 6, 7$, as a function of module fault probability $p_{f,M}$

Note that it is easy from Figure 3 to deduce the operation probability for any other number of RIOPs N_{IOP} . Because of the structure of Eq. 1, simply shift any point of ordinate Y in Figure 3 to the new ordinate $Y^{N_{IOP}/4}$. For instance, in a chip with 8 RIOPs, the validation probability for $p_{f,M} = 0.2$ and $R=6$ simply results in shifting point A from ordinate 0.93 to the new ordinate $(0.93)^2 = 0.865$.

4 Fault Tolerance Around each RIOP

To be consistent with the notation used in the previous section, in the sequel we keep using the term RIOP. However, the discussion and results in this section apply also, should one consider non redundant IOPs.

As a first approximation, it can be assumed the communication bandwidth to be proportional to the number of fault-free nodes directly adjacent to a RIOP. Thus, we consider increasing the RIOP connectivity n_C to protect the bandwidth against the failure of the nodes in the vicinity of the RIOP.

Examples of the local modification of the grid topology around each RIOP aimed at tolerating faulty adjacent nodes are shown in Figure 4. Note that the connectivity of the nodes adjacent to the RIOP is not changed to keep the design simple.




Figure 4: Examples of local changes of the grid topology around a RIOP

Ultimately, it is necessary to sort and validate the chips versus the number of good adjacent nodes around each IOP. The simple expression giving the probability $P_L(k, n_c, p_{f,N})$ that at least k adjacent neighbors (out of n_c) are NOT defective around each IOP reads:

$$P_L(k, n_c, p_f) = \left[\sum_{i=k}^{n_c} \binom{n_c}{i} (1 - p_{f,N})^i p_{f,N}^{n_c-i} \right]^{N_{IO}} \quad (2)$$

where $p_{f,N}$ is the node fault probability.

Various sets of probabilities $P_L(k, n_c, p_{f,N})$ are plotted in Figure 5 for a 4-port architecture (i.e., $N_{IOP} = 4$). Again, each curve is identified by a label k/n_c , where: k denotes the minimal number of fault-free nodes adjacent to the IOP and n_c is the IOP connectivity.




Figure 5: Probability for a 4-port chip that each port is linked to a minimal number of fault free nodes, versus the node fault probability $p_{f,N}$.

We analyze the cases $k/n_c = 3/\{4,6,8\}$ and $k/n_c = 4/\{6,8\}$ in accordance to the topologies depicted in Figure 4. For instance, point C in Figure 5 means: When the IOP connectivity is $n_c = 8$, the probability is about 0.95 to have at least $k=3$ fault-free nodes adjacent to each IOP (out of 8) when the node fault probability $p_{f,N}$ is 0.25.

5 Protection cost

The two sets of figures may be combined to estimate the overhead (in terms of silicon area) necessary to protect the IOPs and their environment. For instance, points A (in Figure 3) and B (in Figure 5) jointly mean: When the fault probabilities of the IOPs and nodes are $p_{f,M} = p_{f,N} = 0.2$ and when the grid has 4 ports, each one connected to 6 nodes, the probability that 3 out of 6 I/O modules work in each RIOP is $P_{W,IOP} = 0.93$ (Figure 3) and the probability that at least 3 out of 6 direct adjacent neighbors of each RIOP work is approximately $P_L(k, n_c, p_{f,N}) = 0.93$ (Figure 5). Thus, the probability to produce a chip fulfilling these constraints is approximately $P_{W,IOP} \times P_L(k, n_c, p_{f,N}) \approx 0.86$.

The price to pay for protecting the IOPs, i.e., the fraction of additional silicon used to implement the RIOPs, is approximately:

$$Q = \frac{(R-1) N_{IO} \cdot A_{IO}}{N \cdot A} \quad (3)$$

where N is the number of core nodes, A_{IO} denotes the size of a I/O module (Mi) and A the size of a node. In the sequel, we will assume that fault probabilities for a Mi module and a node are proportional to their respective size.

Considering expression (3), we have calculated the chip area overhead as a function of the relation A_{IO}/A . As an example, we have analyzed a chip with the following parameters: $N = 300$ nodes, $N_{IOP} = 4$ RIOPs, each having $n_c = 8$ adjacent nodes (see Figure 4-c). In addition, we have considered the following validation criteria (VC):

- VC1) To protect communication bandwidth of each RIOP, at least 3 neighboring nodes (out of 8) must be fault-free.
- VC2) The validation yield (combining the probability values obtained from Figure 3 and Figure 5) must be such that: $P_{W,IOP} \times P_L(k, n_c, p_{f,N}) \geq 80\%$.

The values of the IOP redundancy R that matches the VCs and the resulting chip area overhead are depicted in Figure 6, for two values of $p_{f,N}$, namely 0.2 and 0.3. We have chosen these values because we previously showed that reconfiguration techniques are able to maintain sufficient communication in the array for this range of values [3, 4]. The vertical axis displays the chip area overhead and the horizontal axis represents the ratio A_{IO}/A .

For both curves, labels R identify the minimal redundancy level to be implemented to protect the ports in accordance to the set of VCs stated above.




Figure 6: Chip area overhead vs. the ratio of IOP and node sizes as a function of node fault probability ($p_{f,N} = 0.2; 0.3$)

Concerning VC1, Figure 5 shows that, when $p_{f,N} = 0.2$, the probability to have 3 fault free nodes out of 8 (curve labeled 3/8) is very close to 1. Thus, the condition $P_{W,IOP} \times P_L(k, n_c, p_{f,N}) \geq 80\%$ (VC2) reduces to $P_{W,IOP} \geq 0.8$. The minimal redundancy of the IOPs is thus determined from Figure 3. For instance, when $p_{f,M} = 0.3$, it can be seen that the redundancy R necessary to fulfill criteria VC2 is $R \geq 7$. This is reflected in Figure 6 by considering the curves for the X axis value $p_{f,M}/p_{f,N} = 1.5$, when $p_{f,N} = 0.2$.

Let us now consider point B in Figure 6. The meaning is that if the fault probability of a node $p_{f,N} = 0.2$ and if a RIOP module is 1.7 times larger than a node, then 13.5% of chip area is devoted to protect the ports. In general, we can see that the size of the RIOP module should be preferably smaller than a node to keep the area overhead below 7%. If the RIOP module is larger than a node, the overhead rapidly becomes prohibitive; it reaches 20% for $A_{IOP}/A=2$. Surprisingly, the differences between the two curves $p_{f,N}=0.2$ and $p_{f,N}=0.3$ are quite small.

6 Conclusion

In this paper, we have studied how protecting the dedicated input-output ports (IOPs) in multiport grid chips, as well the attached cost in terms of circuitry overhead.

Protecting the IOPs is a critical issue in chips developed using massively defective nanotechnologies, because these are necessary paths for all communications with the environment, whereas the protection of the processing or communications inside the grid is relatively easy, simply by moving around the defective zones. In other words, there is a natural redundancy in the grid which is missing in the IOPs and which must be implemented.

The results depicted in Figure 3 and Figure 5 reveal the price to pay, in terms of redundancy and IOP connectivity,

to maintain the operation probability of all ports above a user-defined threshold. We have also shown that the IOP module size has a crucial impact on the chip area overhead necessary to protect the ports and should be as small as possible in comparison with the node size (see Figure 6).

Acknowledgement

This work was supported in part by the Network of Excellence *ReSIST: Resilience for Survivability in IST* [www.resist-noe.org] of the European Union FP6 IST Program (contract: 026764).

References

- [1] S. Vangal *et al.*, "An 80-Tile 1.28 TFLOPS Network-on-Chip in 65nm CMOS," in *Proc. IEEE International Solid-State Circuits Conference (ISSCC-2007)*, San Francisco, CA, USA, 2007, pp. 98-99 & 589, (IEEE CS Press).
- [2] W. Robinett, G. S. Snider, P. J. Kuekes and R. S. Williams, "Computing with a Trillion of Crummy Components," *Com. of the ACM*, vol. 50, no. 9, pp. 35-39, September 2007.
- [3] P. Zajac, J. H. Collet, J. Arlat and Y. Crouzet, "Resilience through Self-Configuration in Future Massively Defective Nanochips," in *Supplemental Volume of the 37th Annual IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN-2007) - Workshop on Dependable and Secure Nanocomputing*, Edinburgh, Scotland, UK, 2007, pp. 266-271.
- [4] P. Zajac and J. H. Collet, "Production Yield and Self-Configuration in the Future Massively Defective Nanochips," in *Proc. 22nd IEEE Int. Symp. on Defect and Fault-Tolerance in VLSI Systems (DFT'07)*, Roma, Italy, 2007, pp. 197-205, (IEEE CS Press).
- [5] A. Meixner and D. J. Sorin, "Detouring: Translating Software to Circumvent Hard Faults in Simple Cores," in *Proc. 38th Annual IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN-2008)*, Anchorage, AK, USA, 2008 (To appear in June), (IEEE CS Press).
- [6] W. Hlayhel, D. Litaize, L. Fesquet and J. H. Collet, "Optical Bus versus Electronic Bus for Address-Transactions in Future SMP Architectures," in *Proc. 7th IEEE Int. Conf. on Parallel Architectures and Compilation Techniques (PACT'98)*, Paris, France, 1998, pp. 22-29, (IEEE CS Press).
- [7] M. A. Breuer, S. K. Gupta and T. M. Mak, "Defect and Error Tolerance in the Presence of Massive Numbers of Defects," *IEEE Design & Test of Computers*, vol. 21, no. 3, pp. 216- 227, May-June 2004.

- [8] S. Murali, T. Theocharides, N. Vijaykrishnan, M. J. Irwin, L. Benini and G. D. Micheli, "Analysis of Error Recovery Schemes for Networks on Chips," *IEEE Design & Test of Computers*, vol. 22, no. 5, pp. 434- 442, Sept.-Oct. 2005.
- [9] J. C. Smolens, B. T. Gold, B. Falsafi and J. C. Hoe, "Reunion: Complexity-Effective Multicore Redundancy," in *Proc. 39th Annual IEEE/ACM Int. Symposium on Microarchitecture*, Orlando, FL, USA, 2006, pp. 223-234, (IEEE CS Press).
- [10] Y. Zorian, "Nanoscale Design & Test Challenges," *IEEE Computer*, vol. 38, no. 2, pp. 36-39, February 2005.
- [11] L. Anghel, R. Leveugle and P. Vanhauwaert, "Evaluation of SET and SEU Effects at Multiple Abstraction Levels," in *Proc. 11th IEEE Int. On-Line Testing Symposium (IOLTS-2005)*, Saint Raphaël, France, 2005, pp. 309-312, (IEEE CS Press).
- [12] M. Zhang, T. M. Mak, J. Tschanz, K. S. Kim, N. Seifert and D. Lu, "Design for Resilience to Soft Errors and Variations," in *Proc. 13th IEEE Int. On-Line Testing Symposium (IOLTS-2007)*, Heraklion, Crete, Greece, 2007, pp. 23-28, (IEEE CS Press).
- [13] L. N. Chakrapani, P. Korkmaz, B. E. S. Akgul and K. V. Palem, "Probabilistic System-on-a-Chip Architectures," *ACM Transactions on Design Automation of Electronic Systems*, vol. 12, no. 3, pp.1-28, August 2007.