

# Building Reliable Systems-on-Chip in Nanoscale Technologies

Andreas Steininger(1), Horst Zimmermann(2), Axel Jantsch(3), Michael Hofbauer (2) , Ulrich Schmid (1), Kurt Schweiger (2), Varadan Savulimedu Veeravalli (1)

(1) Vienna University of Technology, Institute of Computer Engineering

(2) Vienna University of Technology, Institute of Electrodynamics, Microwave and Circuit Engineering

(3) Vienna University of Technology, Institute of Computer Technology

**Abstract:** Modern application-specific integrated circuits (ASICs) contain complete systems on a single die, composed of many processing elements that communicate over a dedicated router-based on-chip network. As systems-on-chip comprise billions of transistors with feature sizes in the range of 10nm, reliable operation cannot be established without carefully engineered support at all levels, from technology to the circuit- and the system-layer. This article surveys contributions of research groups at TU Vienna to this field. At lower levels of abstraction, they range from the generation of fault models for simulation that closely match reality and are at the same time efficient to use, to circuit-level radiation-tolerance techniques. At the level of on-chip networks, novel fault-tolerant routing algorithms are being developed together with architectural techniques to isolate faulty parts while keeping the healthy parts connected and active.

The article will briefly portray the associated research activities and summarize their most relevant results.

**Keywords:** system-on-chip, network-on-chip, fault tolerance, radiation tolerance, fault model, fault-tolerant routing

## Der Entwurf zuverlässiger Systems-on-Chip in Nanotechnologien

**Kurzfassung:** Moderne integrierte Bausteine (ASICs) beinhalten auf einem einzigen Chip ganze Systeme, bestehend aus einer Vielzahl an Funktionsblöcken, die über eigene Router-basierte „on-chip“-Netzwerke kommunizieren. Der zuverlässige Betrieb eines solchen Milliarden an Transistoren mit Feature-Size in Bereich von 10nm umfassenden Systems kann nur durch sorgfältig ausgelegte Maßnahmen auf allen Ebenen, von der Technologie über das Schaltungsdesign bis hin zur Systemebene, gewährleistet werden. Der vorliegende Artikel gibt einen Überblick über die diesbezüglichen Beiträge der Forschergruppen an der TU Wien. Auf den unteren Abstraktionsebenen reichen diese von der Erstellung möglichst wirklichkeitsgetreuer Fehlermodellen für die Simulation, die dennoch handhabbar bleiben, bis hin zu schaltungstechnischen Maßnahmen zur Erhöhung der Strahlungsfestigkeit. Auf der Ebene der on-chip-Netzwerke werden neuartige fehlertolerante Routing-Algorithmen in Kombination mit Architekturmaßnahmen entwickelt, die fehlerhafte Bereiche isolieren während funktionierende Teile verbunden und aktiv bleiben.

Der Artikel umreißt die entsprechenden Forschungsaktivitäten und skizziert ihre wichtigsten Ergebnisse.

**Schlagwörter:** System-on-Chip, Network-on-Chip, Fehlertoleranz, Strahlungsfestigkeit, Fehlermodell, fehlertolerantes Routing

## 1. Introduction

The advances in ASIC fabrication technology nowadays allow to accommodate a complete system on a single die, comprising billions of transistors, with feature sizes in the range of 10nm. While, without doubt, this progress fuels the unprecedented technological advances that modern computer systems deliver to our society, it can only be sustained if we manage to keep removing the roadblocks along the way. One of these is decreasing reliability. Even though the fault rate per transistor could be kept relatively constant over the years in spite of the dramatically reduced feature sizes, supply voltages and critical charges, the sheer amount of transistors per die (or system) ultimately makes its failure more likely [1]. So fault tolerance comes into the focus. Unfortunately, however, traditional fault tolerance techniques developed and refined for application domains like aerospace over decades cannot usually be applied directly, as they are prohibitively expensive or become ineffective with newer chip technologies. Moreover, the architectural structure of the systems has changed as well: Traditional architectures were comprised of locally well separated “discrete” function blocks (separate chips) that interacted via a shared bus. Modern systems on chip (SoC) integrate all required functional units on a single die, communicating over a mesh-based network that employs routers at the cross-points. These networks on chip (NoCs) scale much better than a bus but add routing issues to the picture.

Consequently, there is a need for continued research on fault tolerance, and TU Vienna is taking an active role here. This article surveys some of the related activities of selected working groups. More specifically, Section 2 will survey a project dedicated to elaborating an accurate analog simulation model for radiation-induced faults in digital circuits, which constitutes an important tool for the development and assessment of novel fault tolerance techniques. Related approaches will be addressed in Section 3, where we present the architecture and some design highlights of an ASIC that is specifically tailored to exploring radiation fault propagation and masking effects. Section 4 is dedicated to the system-level communication aspect and will review some approaches for protecting packet payload and control information in NoCs. Finally, Section 5 will conclude the paper.

## 2. Fault Modeling

When designing a fault tolerance strategy, it is obviously crucial to have a clear understanding of the faults that need to be tolerated. Otherwise, the resulting system will either be overdesigned and consequently too expensive, too slow, too power-hungry, etc., or it will ultimately not exhibit the desired reliability. Although this seems quite evident, there is still a substantial lack of good fault models that are realistic and easy to use. The very simple and popular single bit-flip model, e.g., fails to consider multi-bit upsets caused by a single ionizing particle, which increasingly plague nanometer VLSI technology. Physical-level simulation tools like TCAD, on the other hand, allow the consideration of every detail of a specific particle hit, but due to excessive computational demands, their use is limited to a scope of a few transistors only; in addition, the choice of suitable simulation parameters remains an open problem.

With this motivation, we have started a joint project with the vision of developing high-level fault models that remain computationally feasible for a larger design scope, but still faithfully reflect the physical behavior of digital circuits under fault conditions. A cornerstone of this research project

FATAL<sup>1</sup> are fault models for radiation-induced faults caused by ionizing particles, which primarily originate in high-energy cosmic rays hitting Nitrogen or Oxygen atoms in the atmosphere.

An ionizing particle hitting the junction of a reverse-biased transistor creates free charge along its track [2]. This charge causes an undesired current in the affected transistor, which, in turn, causes a transient voltage pulse, a so called single-event transient (SET) in the circuit. This SET finally has the potential to flip a bit in a storage cell (single-event upset, SEU), thus causing erroneous behavior of the circuit.

However, there are many masking effects involved that potentially prevent the generation of an SET. Their effectiveness depends on the amount of charge actually created (in turn dependent on particle energy, angle, hit location,...), the current paths taken by the charge carriers, the conversion from these currents to voltages (in turn dependent on effective impedances and capacitances according to circuit technology, state, etc.) and finally the timing of the particle impact relative to the circuit operation. At the same time, SETs can multiply, e.g., at forks in their propagation path. Therefore it is mandatory to really understand all involved masking and propagation effects [12] and have a useful model at hand [11].

Our aim is to develop an accurate analog-level electrical simulation model (SPICE) for SET generation in modern VLSI technologies, since at this level simulation of larger circuit parts is doable. In fact, analog-level simulation has established itself as the method of choice for modelling SET effects in circuits. The tricky part here, however, is the actual representation of the electrical effects generated by a radiation-induced fault. According to the state of the art, SET generation is modeled by a current source that creates a double-exponential current pulse. However, as we found several indications that this model may not reflect all relevant physical effects, especially for newer circuit technologies, our aim was to improve it. To this end we resorted to physical-level simulation (TCAD), which allowed us a detailed view of the actual charge flows within the silicon. The major problem here was to find the right technology parameters for the TCAD model, as some of these (like doping profiles) are considered confidential by the library vendors and hence simply not available. To be able to calibrate our TCAD model and to verify it against reality, we thus had to resort to physical experiments. Figure 1 summarizes the overall structure of the FATAL project and its various models.

---

<sup>1</sup> The project „Fault-Tolerant Asynchronous Logic (FATAL)” was a joint FWF basic research project (P21694) of TU Vienna’s Institute of Computer Engineering and the Institute of Electrodynamics, Microwave and Circuit Engineering. It started in October 2009 and ended in March 2014; its SET-related part is now continued in the FWF basic research project EASET (P26435).

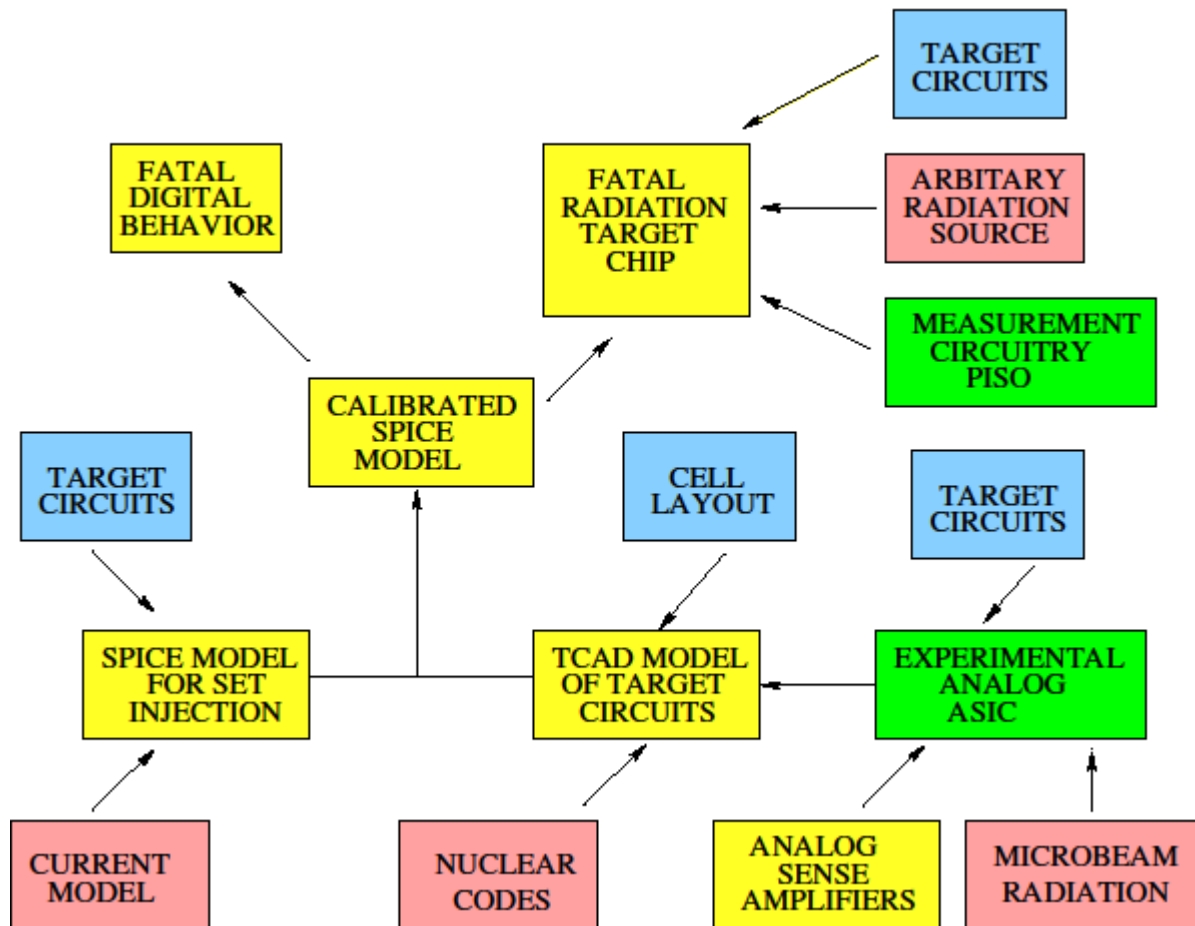


Figure 1: Overall structure of the FATAL project and its various models.

For our experiments, we built a special target ASIC that comprises fundamental digital logic structures like an inverter chain or a Muller pipeline as radiation targets, enhanced with analog amplifiers that facilitate the accurate measurement of SET pulse shapes. Thanks to their low intrusiveness and high bandwidth, the on-chip amplifiers allowed us to convey the analog SET signals generated at relevant nodes of the digital target circuits to a high frequency oscilloscope with very low distortion (see Figure 2). To selectively trigger SETs, we exposed this ASIC to well-controlled radiation in microbeam facilities<sup>2</sup>: They allow to target, with high resolution, a specific location of the ASIC (namely, a specific transistor) from a specific direction, with particles of a specific type (gold particles, alpha particles, ...), and with an energy chosen such that they create a noticeable effect. In these radiation experiments, we could, for the first time ever, directly observe the shapes of SETs generated in real VLSI circuits [5], [6].

<sup>2</sup> We performed radiation experiments at GSI in Darmstadt, PTB in Braunschweig and SNAKE in Munich.

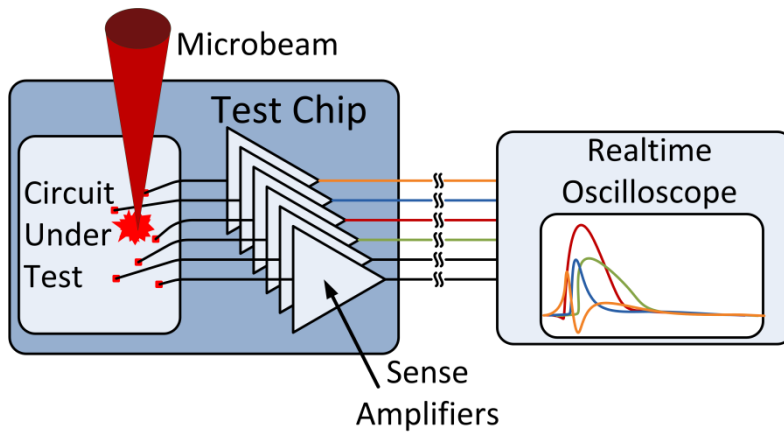


Figure 2: Simplified illustration of the measurement setup including the target ASIC with the analog amplifiers.

Figure 3 shows a glimpse of the wealth of results we gained so far [5], namely, the dependence of the SET pulse shape on the power supply voltage in digital inverters after a particle hit in the sensitive part of the inverter, in the case where a '0' is applied at the input. In this case, the nMOS transistor is in the off-state and therefore sensitive to particle hits. It becomes apparent that the fall-time of the SET observed at the output is almost independent of the supply voltage of the inverter, whereas the rise-time increases strongly for decreasing supply voltages. By using a fully digital measurement approach, only the total pulse widths could have been measured. Furthermore, small SETs may not be visible in the digital world, but are of course captured by our analog measurement approach. This additional information is not only mandatory for calibrating our TCAD model, but also helps to get a better understanding of the underlying effects.

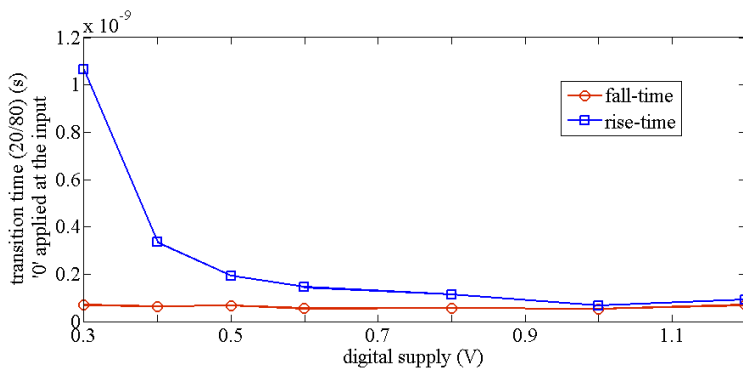


Figure 3: Sample result from the SET measurements in the FATAL project [5].

Based on these and other results we are now in the position to develop a new SPICE modelling approach for SETs in the scope of our follow-up FWF basic research project EASET<sup>3</sup>.

### 3. Fault-Tolerant Circuit Design

A complementary line of our research is dedicated to the design of a purely digital ASIC for long-term radiation experiments, which are dedicated to gather meaningful statistics on the SET generation and

<sup>3</sup> The project „Accelerator-based Experimental Analysis and Simulation Modeling of Single-Event-Transients in VLSI Circuits (EASET)“ is a joint FWF basic research project (P25218) of TU Vienna’s Institute of Computer Engineering and the Institute of Electrodynamics, Microwave and Circuit Engineering. It started in April 2014.

propagation process in various target circuits. Unlike the mixed-signal ASICs described in Section 2, which can only be used in microbeam experiments, the circuit described below can be employed in any radiation environment, including realistic ones like high-atmosphere flights. Besides gathering real statistical data, we will use this circuit also for validating digital radiation fault models: Such models allow fault rate and -effect predictions for much larger circuits than the ones that are amenable to SPICE simulations.

Designing such a circuit leads to challenging problems. First, observing fault effects on many target structures (each in turn comprising multiple circuit nodes) for a significant amount of time causes the need for “streaming” the observations from the target ASIC to a host computer more or less in real time. Here, a manageable data rate can only be obtained by preprocessing the recordings on the target ASIC and transferring condensed information (like upset counts) only. Figure 4 gives an overview of our proposed ASIC architecture. One can see the target circuits in the center, surrounded by the preprocessing infrastructure (mainly SET counters), and finally parallel/serial converters to transfer the counter values to the host over a serial (low footprint) interface.

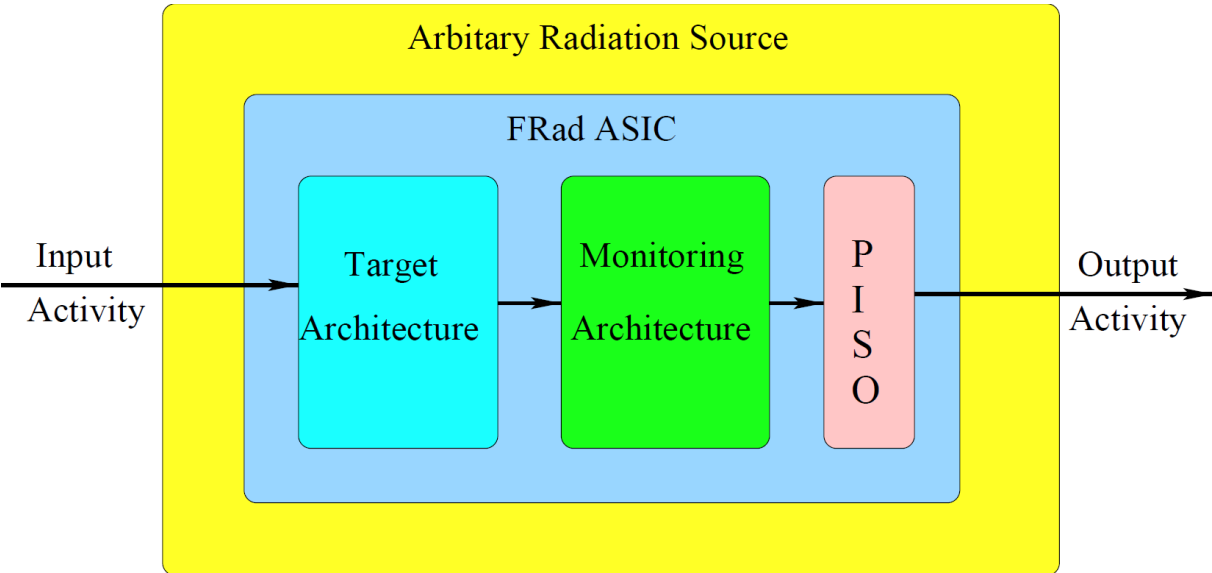


Figure 4: Basic structure of the fully digital SET monitoring ASIC

The on-chip infrastructure for preprocessing and communication, however, will, for uncontrolled radiation sources, necessarily be exposed to the same radiation as the target structures, hence must be radiation tolerant. When designing this infrastructure, we soon realized that the available radiation hardening techniques (like transistor sizing) are not useful for our purpose [9], as they increase the area of the infrastructure to the point where the actual target structures’ area becomes negligible. We would hence experience far more (undesired) radiation particle hits in the infrastructure than in the targets. Our final solution is therefore primarily based on replicated (asynchronous) difference counters tracking the number of SETs in the target circuits, implemented as Muller pipelines or LFSRs (see Figure 5). Instead of hardening those, they are replicated in a very efficient manner that also allows us to regard those as additional, secondary target circuits [3] [4]. So if one Muller pipeline is hit, we can precisely detect this fact and still come up with a correct SET count for the primary target.

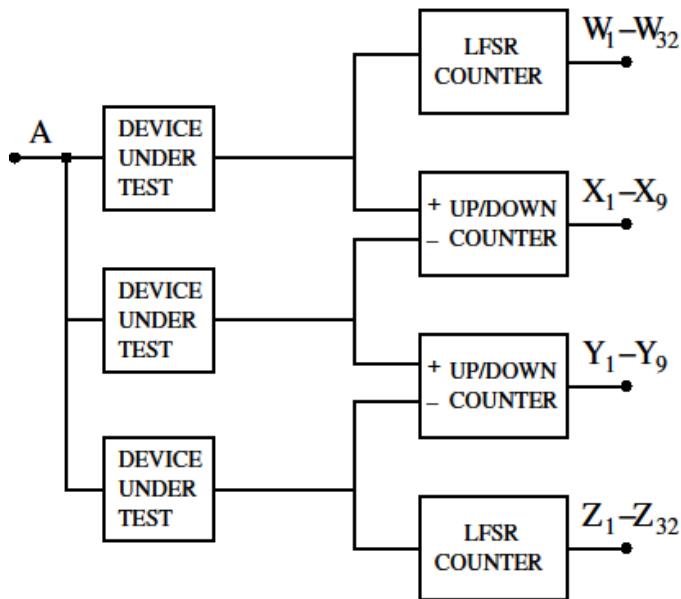


Figure 5: Redundant implementation of the counters.

In addition to this architectural concept, we also implemented and evaluated (via SPICE-level simulation involving our SET generation models) several methods for protecting combinational [9] and sequential [8] gates against radiation, methods for studying fault propagation along selected paths [11], as well as methods for measuring the length on an SET with a purely digital measurement architecture [7].

Our next steps will be to layout and fabricate the ASIC, and conduct the envisioned experiments<sup>4</sup>. Considering the unique infrastructure provided by this circuit, we expect to reveal interesting insights in fault generation rates as well as masking, multiplication and propagation effects that have not been accessible to a systematic experimental evaluation so far.

#### 4. NoC routing & fault isolation

Since 2000 Networks-on-Chip (NoC) have emerged as the interconnect structure of choice for homogeneous and heterogeneous SoCs with many cores, custom hardware processing units and subsystems. Although a large variety of NoC designs have been proposed, their salient features include packet based communication with a network of routers that forward packets in a pipelined fashion. A great variety notwithstanding, the predominant NoC architecture is based on a 2-dimensional mesh network with dimension order routing and wormhole switching. The deterministic *dimension order routing* technique, sometimes also called XY routing, first routes a packet along the X-dimension. When the packet has arrived at the target column of the mesh, i.e. when the difference in the X-dimension to the destination's X-coordinate is 0, the packet is routed along the Y-dimension until the destination node absorbs it. *Wormhole switching* breaks the packet into a sequence of flow control units, also known as flits. The flits move through the network like a worm, with typically only 2-8 flits buffered in each router. When the header flit is blocked due to traffic congestion, the trailing flits are stopped as well in their respective router buffers.

With ever decreasing feature size the challenges to produce 100% fault free chips with billions of transistors and kilometers of wires steadily increase due to process variability, aging effects and the

<sup>4</sup> We are grateful for the funding from TU Vienna that we received for this purpose in the framework of the innovative project „Robust Nanoscale Logic Devices“.

growing importance of radiation-induced transient errors [13] [14] [15]. After memory structures which are routinely protected by error codes, NoCs are typically the largest subsystem on-chip and are essential for the system's smooth operation. Hence, they require protection as well. Consequently, researchers have been quick to develop techniques to protect NoCs against transient and permanent faults at the link level, the network level, and the end-to-end level. Hundreds of proposals have been put forward, explored and evaluated [16]. At the link level, from one router to its neighbor, traditional error detecting and correcting codes are typically employed with emphasis on low cost hardware implementation of the encoder and decoder structures. The most sensitive aspect is the added delay and two cycles additional delay at each link traversal due to encoding and decoding amounts to a penalty of ten cycles across a network with five hops, which is usually unacceptable. An alternative is end-to-end error protection which requires encoding and decoding only once at the source node and at the destination node, respectively. While this works well for the payload, it does not suffice for the control information of the packet, which has to be parsed by every router on the path to allow for correct delivery. Different strategies are therefore used for protecting the control information and the payload of packets. The payload can be safely encoded at the source node, blindly forwarded by all routers on the path, and decoded at the destination node. The control information can be protected by redundant wires and circuitry in the links and the routers to avoid encoding and all unnecessary delay. This strategy is sensible because the control information constitutes only a small fraction of the data information, typically 10-15% for small packets and less than 1% for large packets.

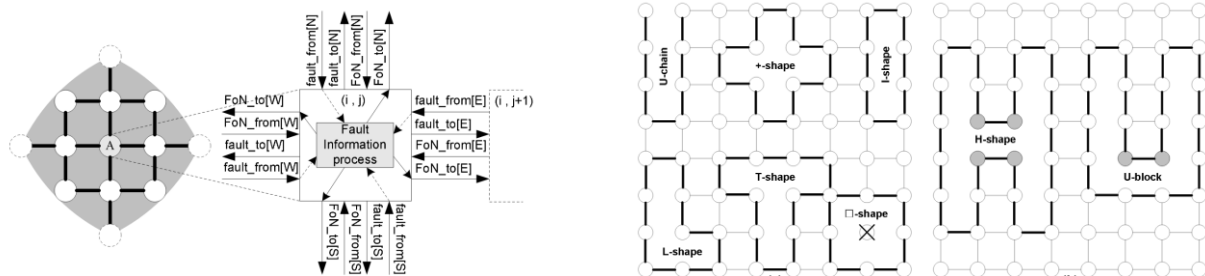


Figure 1. Fault on Neighbor technique: Exchange of fault information between routers (left) and the fault shapes that can be tolerated (right) [18].

Error codes protect well against transient faults but are not very effective for permanent faults because when they accumulate over time they tend to overwhelm the limited error correction capabilities. When error codes detect a fault and trigger retransmission, this is very costly and cannot be tolerated for permanent faults. Fault tolerant routing algorithms have been developed to avoid damaged links and routers [17]. We developed Fault on Neighbor [18] which is a particularly low cost fault tolerant routing algorithm utilizing only local information. After a faulty link or router is identified, this information is propagated to its neighbors, that route around faults by fumbling along the edge of the faulty region like a blind person who can only use the local information available to her hands. By equipping the algorithm with a simple memory and backtracking even complicated concave shapes can be handled (Figure 1 right). Because of the simplicity of the algorithm the area and energy overhead is less than 20% compared to a non-fault tolerant router, but it is limited because only local information is used. Alternative routing algorithms, that acquire and use global information, can find the new optimal route through a network after a fault is identified. To this end we developed a Q-Learning based algorithm that propagates fault information slowly through the whole network such that, after a learning period, all routing nodes have global information about the faults in the network and thus find globally optimal routes through the faulty network [19] [20]. Figure 2 compares the learning behavior of a flat and a hierarchical Q-Learning algorithm and shows



that in a 64 node network the learning period in a 2 level hierarchical learning algorithm takes a few hundred cycles to the new optimal routes after a faulty router has been detected.

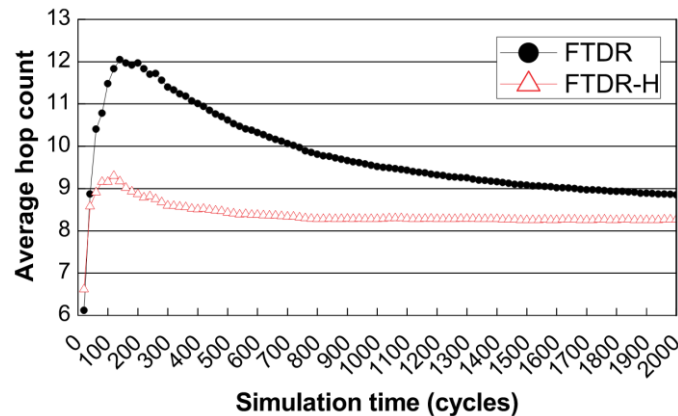


Figure 2. Learning behavior of a Q-Learning based routing algorithm [20].

As witnessed by a wealth of published work on the subject [16], many other techniques have been developed to address specific failures at the different levels of the communication stack. In particular, realistic fault models and complete solutions for fault tolerant NoCs are subjects of intense attention of the research community.

## 5. Summary & Conclusion

The examples of research activities presented in this paper reveal that several working groups at TU Vienna complement each other in devising innovative solutions for overcoming the roadblocks expected to be seen with tomorrow's nanoscale technologies. The topics that are covered, partly in close cooperation in joint projects, span from fault modeling over radiation-tolerant circuit design to reliable NoC design.

With respect to modeling radiation-induced transient faults, the current problem is that device-level physical models are very accurate but often lack the ability to model real circuits, partly due to excessive demands in computational power, and partly due to confidentiality issues with relevant technology parameters. On the other hand, analog-level circuit simulation is very precise in handling circuits of reasonable size (even technology models are available at that level), but usually lack efficient and realistic SET generation models. Our research therefore targets such models, primarily by means of combining simulation and experimental analysis.

With respect to radiation-tolerant circuit design, we can offer innovative approaches covering architectural (replication-based) concepts as well as circuit level measures. These approaches are currently leveraged in an ASIC design, which will serve as a target for further radiation experiments.

Concerning communication within large digital circuits, we contributed approaches for efficiently protecting payload as well as control information in NoC-based communication infrastructures.

In summary, these aspects cover a substantial portion of the problem space "reliable circuit design". Nevertheless, much more research is needed on both the above topics and additional ones, like fault-tolerant clocking of SoCs, robust timing domain interfacing, self-repair, etc., which are at the heart of our future activities.

## References

- [1] A. Dixit and A. Wood  
The impact of new technology on soft error rates  
*Proc. IEEE International Reliability Physics Symposium (IRPS'11)*, pages 5B.4.1 –5B.4.7, April 2011.
- [2] R. Baumann  
Radiation-induced soft errors in advanced semiconductor technologies  
*IEEE Transactions on Device and Materials Reliability*, 5(3):305–316, Sept. 2005.
- [3] V. S. Veeravalli, T. Polzer, A. Steininger, and U. Schmid  
Architecture and design analysis of a digital single event transient/upset measurement chip.  
*Proceedings 15th Euromicro Symposium on Digital System Design: Architectures, Methods and Tools (DSD'12)*, pages 8–17, sep 2012.
- [4] V. S. Veeravalli, T. Polzer, U. Schmid, A. Steininger, M. Hofbauer, K. Schweiger, H. Dietrich, K. Schneider-Hornstein, H. Zimmermann, K. Voss, B. Merk, M. Hajek  
An infrastructure for accurate characterization of single-event transients in digital circuits  
*Microprocessors and Microsystems*, 37 (2013), pp 772 - 791.
- [5] M. Hofbauer, K. Schweiger, H. Zimmermann, U. Giesen, F. Langner, U. Schmid, A. Steininger  
Supply Voltage Dependent On-chip Single-Event Transient Pulse Shape Measurements in 90 nm Bulk CMOS Under Alpha Irradiation  
*IEEE Transactions on Nuclear Science*, vol 60 (2013), pp 2640 - 2646.
- [6] M. Hofbauer, K. Schweiger, H. Dietrich, H. Zimmermann, K.O. Voss, B. Merk, U. Schmid, A. Steininger  
Pulse Shape Measurements by On-chip Sense Amplifiers of Single Event Transients Propagating through a 90 nm Bulk CMOS Inverter Chain  
*IEEE Transactions on Nuclear Science*, vol 59 (2012), pp 2778 - 2784.
- [7] V. S. Veeravalli, A. Steininger, U. Schmid  
Measuring SET Pulsewidths in Logic Gates using Digital Infrastructure  
*15th International Symposium & Exhibit on Quality Electronic Design (2014)*  
ISBN: 978-1-4799-3946-6; 7 S.
- [8] V. S. Veeravalli, A. Steininger  
Efficient Radiation-Hardening of a Muller C-Element  
*2012 Single Event Effects Symposium (SEE 2012)*
- [9] V. S. Veeravalli, A. Steininger  
Radiation-Tolerant Combinational Gates - An Implementation Based Comparison  
*15th IEEE International Conference on Design and Diagnostics of Electronic Circuits and Systems (DDECS 2012)*, pp. 115 - 120.
- [10] V. S. Veeravalli, A. Steininger  
Performance of Radiation Hardening Techniques under Voltage and Temperature Variations  
*2013 IEEE Aerospace Conference*, 6p
- [11] R. Najvirt, V. S. Veeravalli, A. Steininger  
Particle Strikes in C-Gates: Relevance of SET Shapes  
*2nd Workshop on Manufacturable and Dependable Multicore Architectures at Nanoscale*, (2013), 4p
- [12] A. Steininger, V. S. Veeravalli, D. Alexandrescu, E. Costenaro, L. Anghel  
Exploring the State Dependent SET Sensitivity of Asynchronous Logic - The Muller-Pipeline Example  
*32nd IEEE International Conference on Computer Design (ICCD)*, (2014), ISBN: 978-1-4799-6492-5; Paper-Nr. 69, 7p
- [13] Borkar, S., "Designing reliable systems from unreliable components: the challenges of transistor variability and degradation", *Micro*, IEEE, vol. 25, pp. 10-16, no. 6, November-December 2005.
- [14] Kuhn, K.J., Giles, M.D., Becher, D., Kolar, P., Kornfeld, A., Kotlyar, R., Ma, S.T., Maheshwari, A., and Mudanai, S., *Electron Devices*, IEEE Transactions on, "Process Technology Variation", vol. 58, no. 8, pp. 2197-2208, 10.1109/TED.2011.2121913, 0018-9383, Aug 2011.
- [15] John Keane and Chris Kim, "An Odometer for CPUs", *IEEE Spectrum*, vol. 48, pp. 26-31, no. 5, May 2011.
- [16] Martin Radetzki, Chaochao Feng, Xueqian Zhao, and Axel Jantsch, "Methods for Fault Tolerance in Networks-on-Chip", *ACM Computing Surveys*, vol. 46, no. 1, 0360-0300, pp. 8:1--8:38, 8, 38, 10.1145/2522968.2522976, ACM, New York, NY, USA, July 2013.
- [17] *Routing Algorithms in Networks-on-Chip*, edited by Maurizio Palesi and Masoud Daneshtalab, Springer, 2014.
- [18] Chaochao Feng, Zhonghai Lu, Axel Jantsch, Jinwen Li, and Minxuan Zhang, "FoN: Fault-on-Neighbor aware Routing Algorithm for Networks-on-Chip", *International SOC Conference*, Las Vegas, Nevada, September 2010.
- [19] Chaochao Feng, Zhonghai Lu, Axel Jantsch, Jinwen Li, and Minxuan Zhang, "A Reconfigurable Fault-tolerant Deflection Routing Algorithm Based on Reinforcement Learning for Networks-on-Chip", *Proceedings of the International Workshop on Network on Chip Architectures (NoCArc)*, Atlanta, Gorgia, November 2010.
- [20] Chaochao Feng, Zhonghai Lu, Axel Jantsch, Minxuan Zhang, and Zuo Cheng Xing, "Addressing Transient and Permanent Faults in NoC With Efficient Fault-Tolerant Deflection Router", *IEEE Transactions on Very Large Scale Integration Systems (TVLSI)*, vol. 21, no. 6, pp. 1053-1066, , June 2013.