08371 Abstracts Collection
Fault-Tolerant Distributed Algorithms on VLSI Chips
— Dagstuhl Seminar —

Bernadette Charron-Bost[1], Shlomi Dolev[2], Jo Ebergen[3] and Ulrich Schmid[4]

[1] Ecole Polytechnique - Palaiseau, FR
[2] Ben Gurion University, IL
dolev@CS.bgu.ac.il
[3] Sun Microsystems - Menlo Park, US
jo.ebergen@sun.com
[4] TU Wien, AT
s@ecs.tuwien.ac.at

Abstract. From September the 7th, 2008 to September the 10th, 2008 the Dagstuhl Seminar 08371 "Fault-Tolerant Distributed Algorithms on VLSI Chips " was held in Schloss Dagstuhl – Leibniz Center for Informatics. The seminar was devoted to exploring whether the wealth of existing fault-tolerant distributed algorithms research can be utilized for meeting the challenges of future-generation VLSI chips. During the seminar, several participants from both the VLSI and distributed algorithms' discipline, presented their current research, and ongoing work and possibilities for collaboration were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Fault-tolerant distributed algorithms, fault tolerance, VLSI systems-on-chip, synchronous vs. asynchronous circuits, digital logic, specifications.

08371 Executive Summary – Fault-Tolerant Distributed Algorithms on VLSI Chips

Bernadette Charron-Bost, Shlomi Dolev, Jo Ebergen and Ulrich Schmid

The Dagstuhl seminar 08371 on Fault-Tolerant Distributed Algorithms on VLSI Chips was devoted to exploring whether the wealth of existing fault-tolerant distributed algorithms research can be utilized for meeting the challenges of future-generation VLSI chips. Participants from both the distributed fault-tolerant algorithms community, interested in this emerging application domain, and from

the VLSI systems-on-chip and digital design community, interested in well-founded system-level approaches to fault-tolerance, surveyed the current state-of-the-art and tried to identify possibilities to work together. The seminar clearly achieved its purpose: It became apparent that most existing research in Distributed Algorithms is too heavy-weight for being immediately applied in the "core" VLSI design context, where power, area etc. are scarce resources. At the same time, however, it was recognized that emerging trends like large multicore chips and increasingly critical applications create new and promising application domains for fault-tolerant distributed algorithms. We are convinced that the very fruitful cross-community interactions that took place during the Dagstuhl seminar will contribute to new research activities in those areas.

Keywords: Fault-tolerant distributed algorithms, fault tolerance, VLSI systems-on-chip, synchronous vs. asynchronous circuits, digital logic, specifications

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1927

Distributed Computing: Time and Faults

Bernadette Charron-Bost (Ecole Polytechnique - Palaiseau, FR)

In this talk, I discuss the main features of fault-tolerant distributed computing, and present classical models of distributed systems in the context of failures.

Keywords: Time, clock, causality, failure models, synchrony degree, consensus and agreement problems

Self-Stabilization Copes with Soft Errors

Shlomi Dolev (Ben Gurion University, IL)

Self-stabilization is a property of a system that ensures automatic recovery in the sense that the system starts operating as it should from arbitrary state. It turned out that to achieve self-stabilization one needs to ensure that the Hardware is stabilizing otherwise the stabilizing algorithm/software/protocol will not be executed.

The talk starts describing ways to analyze the influence of soft-errors and describe ways to design Hamming preserving circuits; then a description of ways to verify and enforce stabilization of a processor are described; at last stabilization preserving compiler is described.

Keywords: Self-stabilization, soft-errors, hardware, compiler, error-correcting

See also: Shlomi Dolev, Yinnon A. Haviv: Stabilization Enabling Technology.

See also: SSS 2006: 1-15

See also: Shlomi Dolev, Yinnon A. Haviv: Self-Stabilizing Microprocessor: Analyzing and Overcoming Soft Errors. IEEE Trans. Computers 55(4): 385-399 (2006)

See also: Shlomi Dolev, Yinnon A. Haviv, Mooly Sagiv: Self-stabilization Preserving Compiler. Self-Stabilizing Systems 2005: 81-95

Introduction to VLSI

Jo Ebergen (Sun Microsystems - Menlo Park, US)

This set of slides presents a brief tutorial on VLSI. Topics include CMOS basics, wire scaling and its implications, flops and latches, timing, soft errors, and variations in parameters.

Synchrony and Asynchrony in VLSI

Jo Ebergen (Sun Microsystems - Menlo Park, US)

This set of slides presents a brief introduction to synchrony and asynchrony in VLSI.

The slides cover the following topics: synchronous versus asynchronous circuits – as in clocked versus clockless circuits –, clock distribution, synchronizers and arbiters, reliability of synchronizers, and examples of clockless circuits.

Probability Distribution of Some Time Characteristics of Fault-Tolerant Systems

Sergey Frenkel (Russian Academy of Sciences - Moscow, RU)

This talk discusses a problem of estimation of time, which can be spent to discover a fault effect at output of a system represented as a combination of interacted finite state machines (FSM) under random inputs when this effect occurred in its inner states. This topic deals with both "latency" of the fault detection and so-called "Self-healing". Self-healing enables the system to continue functioning correctly on the event of the failure to determine the errors and to recover from them.

One the one hand, the shorter the latency the better the system can recover, that means the reliability increases. On the other hand, the increase in dependability from self-healing is proportional to the probability that a self-healing will prevent the failure Therefore, both these time-distributed phenomena are very considerable aspects of high-reliable systems design, and computation of probability distribution function of time to self-healing needs for fair prediction of fault-tolerant systems reliability.

In order to be computed, the computation probability distribution function of fault detection latency and probability of self-healing ability needs in some joint consideration of the fault and fault-free system models at a given specific level of the system modeling, at FSM level, in particular, that is at rather early design stages.

The fault detection latency is modeled as a time to absorption for a product of fault-free and faulty (i.e. corrupted by a fault) systems of interacting Markov chains. A notion of a "product" of two models (of fault-free and faulty, correspondingly) is convenient in many cases for such modeling.

The self-healing is analyzed in terms of coupling of Markov chains describing faulty and fault-free FSMs under independent random input binary signals.

This work is an attempt to consider possible ways of such modeling, using a generalized view on some probabilistic approaches to analysis of self-healing paradigm-based fault-tolerant computer systems. A model of the FSMs product behavior under random input vectors as a Markov chain is considered. The state space of the Markov chain model can be constructed as a direct product of fault-free/faulty states, as well as Kronecker product of the states. An approach to FSM decomposition into a network of smaller component FSMs for latency reduction is also outlined, as well as a possibility to use the Markov chain model for its fault latency analysis.

These models can be a base of a tool of fault-tolerant systems design, dealing with such fault tolerance aspects as fault detection latency (for permanent faults) and self-healing in presence of some transient faults (Soft Upset Errors in particular).

Keywords: Fault detection latency, self-healing, FSM

See also: I. Levin, A. Matrosova, S. Ostanin: Survivable Self-checking Sequential Circuits, Proceedings of the DFT 01, p. 395, 2001.

See also: S. Frenkel, A. Pechinkin, V. Chaplygin, I. Levin: A mathematical Tool for Support of Fault-Tolerant Embedded Systems Design, ERCIM/DECOS Dependable Smart Systems: Research, Industrial Applications, Standardization, Certification and Education. Workshop on “Dependable Embedded Systems”, Lübeck, Germany, 2007.

Probability Distribution of Some Time Characteristics of Fault-Tolerant Systems

Sergey Frenkel (Russian Academy of Sciences - Moscow, RU)

Fault detection latency and self-healing phenomena are very considerable aspects of high-reliable systems design. The computation of both probability distribution function of fault detection latency and probability of self-healing ability needs in some joint consideration of the fault and fault-free system models at a given specific level of the system modeling. Finite state machine (FSM) is very

popular model of a computer system behavior at rather high levels of the system design. A notion of a "product" of the models of fault-free and faulty FSMs can be convenient for testability and fault-tolerant features modeling. However, currently this model was applied the only to a single FSM.

In this talk some approaches to analysis of fault detection latency for FSM decomposed into a network of smaller component FSMs, are outlined. A possibility to use the Markov chain model for its fault detection latency and self-healing analysis is shown.

These models can be a base of a tool of fault-tolerant systems design, namely, in order to compute the probabilistic distribution functions of both fault detection latency (for permanent faults) and time to self-healing in presence of some transient faults (Soft Upset Errors in particular).

Keywords: Fault detection latency, self-healing, finite state machine, Markov Chains

Implications of VLSI Fault Models and Distributed Systems Failure Models – A hardware designer's view

Gottfried Fuchs (TU Wien, AT)

The fault and failure models as well as their semantics within the VLSI and the distributed systems/algorithms community are quite different. Pointing out the mismatch of those fault respectively failure models is the main part of this work. The impact of the implemented failure model in terms of hardware effort and system complexity will be shown on different VLSI implementations of distributed algorithms.

However, still, there are a lot of open questions left mostly related to the coverage analysis of hardware implemented fault-tolerant algorithms.

Keywords: VLSI, fault model, distributed system, failure model

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1924

Search for a formal VLSI model

Matthias Függer (TU Wien, AT)

Formally stating properties of VLSI circuits and proving their correctness is of major interest when designing highly dependable systems. In this talk I give an overview on a system model, that very well suits VLSI and at the same time allows to derive (correctness, worst case performance, impossibility) properties by analogous means as in distributed systems proofs.

Keywords: System model, VLSI, formalization

Arbitration-Free Synchronization

Leslie Lamport (Microsoft Corp. - Mountain View, US)

Implementing traditional forms of multiprocess synchronization requires a hardware arbiter. Here, we consider what kind of synchronization is achievable without arbitration. Several kinds of simple arbiter-free registers are defined and shown to have equal power, and the class of synchronization problems solvable with such registers is characterized. More powerful forms of arbiter-free communication primitives are described. However, the problem of characterizing the most general form of arbiter-free synchronization remains unsolved.

Keywords: Arbiter, marked graphs, multiprocess synchronization

Full Paper:
http://research.microsoft.com/users/lamport/pubs/pubs.html#arbiter-free

Fault Tolerance in Reconfigurable Fabrics

Rajit Manohar (Cornell University, US)

Reconfigurable fabrics by their very nature seem to be well-suited for handling faults. I will provide an overview of reconfigurable fabrics, and describe some practical challenges in making them robust to faults. I will also describe a number of techniques that have been used to increase the fault-tolerance of reconfigurable systems and, in particular, some of the issues that arise in the context of asynchronous reconfigurable fabrics.

Keywords: Reconfigurable fabrics, FPGAs, fault-tolerance, asynchronous circuits

Robustness and soft-error tolerance of asynchronous VLSI

Alain Martin (CalTech - Pasadena, US)

I will briefly describe a special approach to asynchronous logic called "quasi-delay-insensitive" or QDI, which among all varieties of asynchronous logics, has the weakest delay assumptions: QDI circuits are "almost entirely" insensitive to delay variations. I will argue that QDI circuits are by themselves very robust to all parameter variations (so-called PVT variations) and also to all effects of soft errors that affect timing.

I will then briefly describe how QDI circuits can be made tolerant to soft errors that flip a bit (single event upsets). I will mention associated issues with arbiters and memories. If time permits I will describe the three levels of distributed programming models that make it possible to compile a distributed algorithm into a QDI circuit.

Keywords: Message-passing, handshake protocols, VLSI, asynchronous, QDI, arbiter, soft-error tolerance

Full Paper:
http://www.async.caltech.edu/Pubs/PDF/async_softerror.pdf

See also: http://www.async.caltech.edu/Pubs/PDF/2006_ieee_soc.pdf

Fault-Tolerance in Biological Systems

Chris Myers (Univ. of Utah - Salt Lake City, US)

This talk will briefly describe some examples of how biological systems deal with noise and uncertainty while remaining fault-tolerant. The major example discussed will be a synthetic genetic Muller C-element, and a potential application for this genetic circuit.

Keywords: Genetic circuits, stochastic behavior, synthetic biology, fault-tolerance

Error and Fault Tolerance in VLSI

Lirida Naviner (ENST - Paris, FR)

This talk presents an overview on digital VLSI design aspects related to reliability.

It begins with some basic concepts/definitions currently used by hardware community and the impact of new (nano) technologies on fault/tolerance of circuits. Then are presented the challenges for hardware designers to guarantee reliable processing on unreliable devices.

These are defect and fault tolerance techniques but also methods and tools to evaluate reliability.

The emphasis is given onto the solutions of architectural level, which allow a relative independence with respect to the hardware technology and seem suitable to complement fault tolerant mechanisms in higher and lower levels.

Keywords: Reliability, Nanoelectronics, Fault tolerance, Defect tolerance, Reliability analysis

Methods and Metrics for Reliability Assessment

Lirida Naviner (ENST - Paris, FR)

This paper deals with digital VLSI design aspects related to reliability. The focus is on the problem of reliability evaluation in combinational logic circuits.We present some methods for this evaluation that can be easily integrated in a traditional design flow. Also we describe suitable metrics for performance estimation of concurrent error detection schemes.

Keywords: Reliability, fault tolerance, combinational logic

Joint work of: Alves de Barros-Naviner, Lirida; Naviner, Jean-François; Teixeira Franco, Denis; Correia de Vasconcelos, Mai

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1925

Power Aware Soft Error Robust Low Power Fault Tolerance Techniques

Dhiraj Pradhan (University of Bristol, GB)

This talk will present some power aware novel memory and processor designs that achieves soft error protection and fault-tolerance.

Multiple Event Upsets Aware FPGAs Using Protected Schemes

Dhiraj Pradhan (University of Bristol, GB)

Multiple upsets would be available in SRAM-based FPGAs which utilizes SRAM in different parts to implement circuit configuration and to implement circuit data. Moreover, configuration bits of SRAM-based FPGAs are more sensible to upsets compared to circuit data due to significant number of SRAM bits. In this paper, a new protected Configurable Logic Block (CLB) and FPGA architecture are proposed which utilize multiple error correction (DEC) and multiple error detection. This is achieved by the incorporation of recently proposed coding technique Matrix codes [1] inside the FPGA. The power and area analysis of the proposed techniques show that these methods are more efficient than the traditional schemes such as duplication with comparison and TMR circuit design in the FPGAs.

Keywords: FPGA, SEUs, ECC, Reliability, MTTF

Joint work of: Argyrides, Costas; Pradhan, Dhiraj K.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1926

Distributed Algorithms and VLSI - An Appetizer

Ulrich Schmid (TU Wien, AT)

In order to demonstrate that (1) distributed algorithms research can indeed be successfully applied in the VLSI context, and (2) that new research challenges for distributed algorithms emanate from this application domain, we present an overview of our DARTS fault-tolerant distributed clock generation approach.

Starting out from a simple Byzantine fault-tolerant distributed clock synchronization algorithm, we first describe the modifications required for a direct VLSI implementation using asynchronous digital logic. The need for adaptions originates from two peculiarities of hardware implementations, namely, (i) fine-grained parallel asynchronous computations, which undermines the concept of atomic steps common to all distributed computing models, and (ii) very limited resources, which makes even apparently simple operations prohibitively costly. We proceed with sketching our correctness proof, which is based on a low-level modeling framework based on continuous time that integrates the state- and event-based view of binary signals. We conclude with a glimpse on our fully functional FPGA and ASIC prototype implementation.

Keywords: Fault-tolerant distributed algorithms, VLSI, clock generation, SoCs

Detecting Faults and Hardware Bugs Using Dynamic Verification

Daniel J. Sorin (Duke University - Durham, US)

We are using dynamic verification to detect hardware errors due to transient and permanent faults and hardware design bugs. The ultimate goal is to build a fully self-checking system-on-chip.

Keywords: Fault-tolerance, architecture, dynamic verification

Error Containment in the Presence of Metastability

Andreas Steininger (TU Wien, AT)

Error containment is an important concept in fault tolerant system design, and techniques like voting are applied to mask erroneous outputs, thus preventing their propagation. In this presentation we will use the example of DARTS, a fault-tolerant distributed clock generation scheme in hardware, to demonstrate that metastability is a substantial threat to error containment. We will illustrate how metastability can originate and propagate such that a single fault may upset the system. The main conclusion is that modeling efforts on all design levels are definitely required in order to mitigate and quantify the deteriorating effect of metastability on system dependability.

Keywords: Metastability, fault tolerance, clock generation

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1923

From Manual Proofs to Model Checking

Helmut Veith (TU Darmstadt, DE)

Concurrent systems are notoriously error-prone and hard to analyze, both for human engineers and verification tools.

In my talk, I will discuss the verification methodologies at hand, and show how abstraction makes model checking applicable to infinite state systems and parameterized systems.

Keywords: Model checking, verification, formal specification

Fault Tolerant Pulse Synchronization

Jennifer L. Welch (Texas A&M University, US)

I will provide some background on pulse synchronization and then describe some preliminary ideas for adding Byzantine fault-tolerance to one previously known algorithm for the problem.

Keywords: Fault tolerance, Byzantine faults, pulse synchronization

Design of Asynchronous Arbiters using Petri Nets

Alex Yakovlev (University of Newcastle, GB)

In this talk I am planning to discuss how a wide range of arbitration circuits can be designed, systematically, with the use of Petri nets. Petri nets provide a convenient stepping stone between the conceptual understanding of the arbitration algorithm and its hardware implementation.

Keywords: Arbiters in asynchronous circuits