


IAF0530 (MSc)
IAF9530 (PhD)

Dependability and fault tolerance

Gert Jervan
Department of Computer Systems
Tallinn University of Technology (TTÜ)




General Information

- Contents:
Dependability and fault tolerance
www.pld.ttu.ee/IAF0530
- Lecturer & Examiner:
Gert Jervan
ICT-527 620 2261
gert.jervan@ttu.ee
www.pld.ttu.ee/~gerje

© Gert Jervan

2




Gert Jervan

- MSc from TTÜ in 1998
 - Exchange student at TIMA Labs (Grenoble, France), Fraunhofer Institute (Dresden, Germany), Linköping University (Sweden)
- PhD from Linköping University (Sweden) in 2005
- Senior research fellow at TTÜ since 2005, professor since 2012
- Vice-Dean for Research at the Faculty of IT (2012), Dean (2013)
- Published more than 80 papers at international conferences and journals
- Organized many international conferences and coordinated several research projects, incl. 7-year project CEBE (Centre for Integrated Electronic Systems and Biomedical Engineering)

© Gert Jervan

3




Course Plan

- 16 occasions, á 1,5 hours
Mondays 10:00-11:30
- Longer sessions during the second half of the semester (will be announced separately)
- 7-10 Lectures with discussion sessions. No meeting on March 27, April 24 (Tentatively). Always check the course homepage!!
- Individual project work
- Oral exam (discussion)

© Gert Jervan

4




Individual work

- Reading
- Writing
- Presenting
- The course requires weekly reading and participation in discussions
 - All missing assignments have to be compensated during the exam

© Gert Jervan

5




Reading

- **Various papers (on the course homepage)**
www.pld.ttu.ee/IAF0530
- Textbooks
- Incident/accident reports
- Web pages

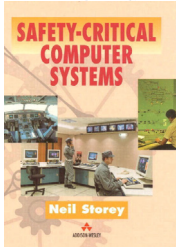
© Gert Jervan

6




Textbooks

- Safety-Critical Computer Systems
 - Neil Storey
 - Addison Wesley, 1996.
 - An introductory text which provides overview of safety related aspects and methods in computer systems development.
 - Available in the TTÜ library



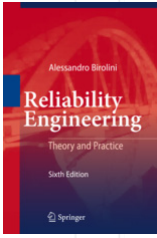
© Gert Jervan

7




Textbooks

- Reliability Engineering: Theory and Practice.
 - Alessandro Birolini
 - Springer
 - 2014 (7th ed.) 2010 (6th ed.), 2007 (5th ed.)
 - This book shows how to build in, evaluate, and demonstrate reliability & availability of components, equipment, systems. It presents the state-of-the-art of reliability engineering, both in theory and practice
 - TTÜ library has several copies of the latest edition.



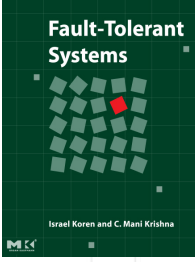
© Gert Jervan

8



Textbooks


- Fault-Tolerant Systems
 - Israel Koren and C. Mani Krishna
 - Morgan-Kaufman Publishers, 2007



This book covers comprehensively the design of fault-tolerant hardware and software, use of fault-tolerance techniques to improve manufacturing yields and design and analysis of networks. Additionally it includes material on methods to protect against threats to encryption subsystems used for security purposes.


© Gert Jervan

9



Textbooks


- Fault-Tolerant Design
 - Elena Dubrova
 - Springer, 2013



- This textbook serves as an introduction to fault-tolerance, intended for upper-division undergraduate students, graduate-level students and practicing engineers in need of an overview of the field. Readers will develop skills in modeling and evaluating fault-tolerant architectures in terms of reliability, availability and safety. They will gain a thorough understanding of fault tolerant computers, including both the theory of how to design and evaluate them and the practical knowledge of achieving fault-tolerance in electronic, communication and software systems. Coverage includes fault-tolerance techniques through hardware, software, information and time redundancy. The content is designed to be highly accessible, including numerous examples and exercises.

© Gert Jervan

10




Case Studies

- The exact format will be announced during the second lecture (and it depends of the number of students we will have)
- Topic categories:
 - Accident analysis
 - System safety analysis
 - Literature survey
 - Something else (implementation, tool study, etc.)
 - Requires prior ack.

Literature and sample (!) topics on the webpage

© Gert Jervan

11

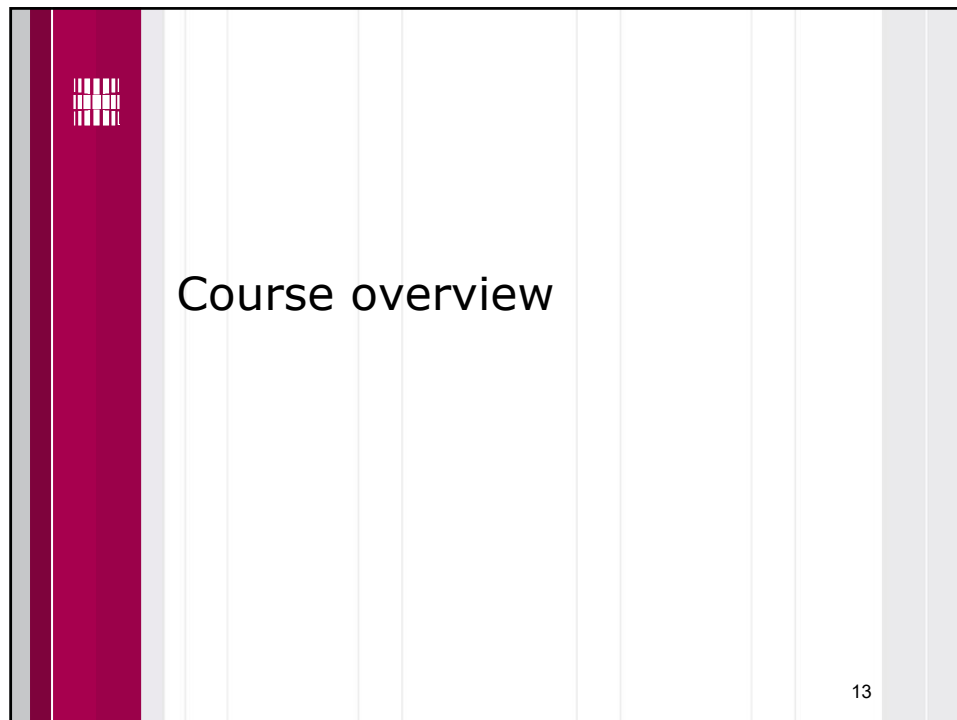


Case Studies

- Some examples (from 2016):
 - Estimating availability of the KSI service.
 - Dependability and Fault Tolerance of PaaS
 - Real-time Transport Protocol security considerations in Source-Specific Multicast topology
 - Fault tolerance on Cryptography
 - Automatic train protection systems
 - Software Fault Injection Methods
 - Safety and reliability of autonomous vehicle technologies
 - Evolution of Fault Tolerance in PostgreSQL
 - Self-checking network-on-chip layout design
 - Verified compilation
 - Fault tolerance in wireless systems
 - Critical Information Infrastructure vulnerability analysis methods

© Gert Jervan

12


A presentation slide with a white background and a dark blue vertical bar on the left. The bar contains a small white grid icon. The title "Course Overview" is centered in a large, dark blue font. Below the title is a bulleted list. The number "14" is in the bottom right corner.

© Gert Jervan

Course Overview

- Reliability: increasing concern
 - Historical
 - High reliability in computers was needed in critical applications: space missions, telephone switching, process control, medical applications etc.
 - Contemporary
 - Extraordinary dependence on computers: on-line banking, commerce, cars, planes, communications etc. Emergence of internet-of-things.
 - Hardware is increasingly more fault-prone (complexity, technology, environment)
 - Software is increasingly more complex
 - Things simply do not work without special reliability measures

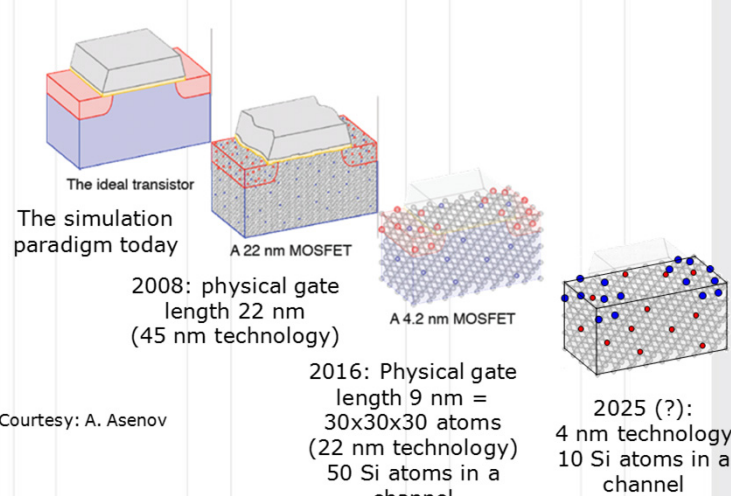




Moore's Law

- Growth rate
 2x transistor & clock speed every 2 years over 50 years
 10x every 6-7 years
This won't last for long...
- Dramatically more complex algorithms previously not feasible
 - Dramatically more realistic video games and graphics animation (e.g. Playstation 4, Xbox 360 Kinect, Nintendo Wii)
 - 1 Mb/s DSL to 10 Mb/s Cable to 2.4 Gb/s Fiber to Homes.
 - 2G to 3G to 4G to 5G wireless communications
 - MPEG-1 to MPEG-2 to MPEG-4 to H.264 video compression
 - 480 x 270 (0.13 million pixels) NTSC to 1920x1080 (2 megapixels) HDTV resolution to 4K UHD 3840 x 2160 (8.3 megapixels)

© Gert Jervan



Scaling

The ideal transistor

The simulation paradigm today

A 22 nm MOSFET

2008: physical gate length 22 nm (45 nm technology)

A 4.2 nm MOSFET

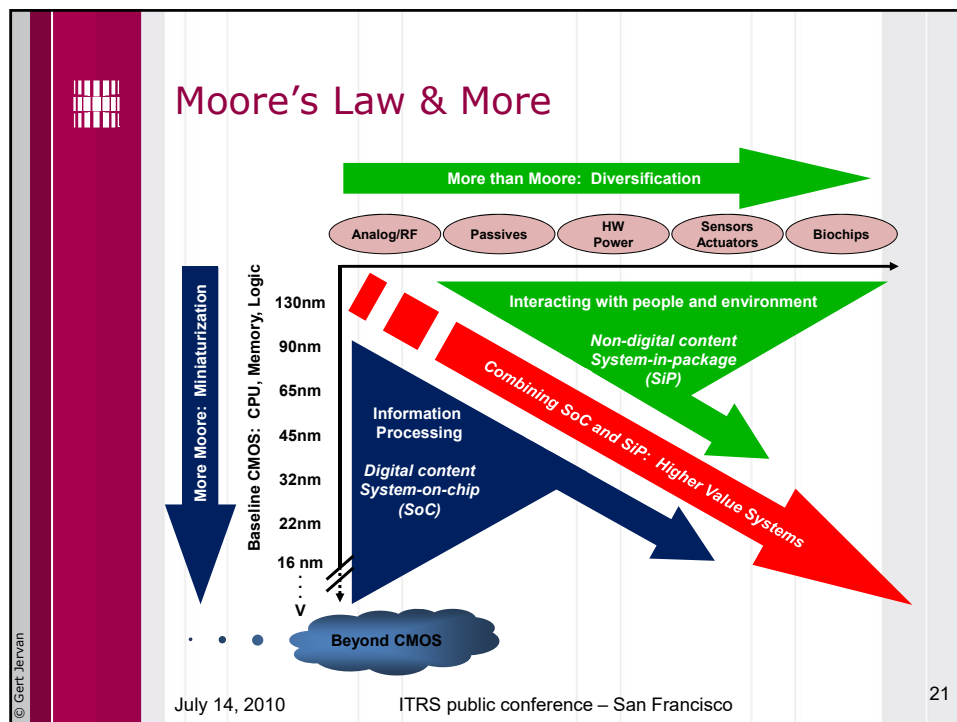
2016: Physical gate length 9 nm = 30x30x30 atoms (22 nm technology) 50 Si atoms in a channel

2025 (?): 4 nm technology 10 Si atoms in a channel

Courtesy: A. Asenov

18






Hardware - Background

- Chip designers, device engineers and the high-reliability community recognize that reliability concerns ultimately limit the scalability of any generation of microelectronics technology
- Statistical methods and reliability physics provide the foundation for better understanding the next generation of scaled microelectronics
 - Microelectronics device physics
 - Reliability analysis and modeling
 - Experimentation
 - Accelerated testing
 - Failure analysis
- The design, fabrication and implementation of highly aggressive advanced microelectronics requires expert controls, modern reliability approaches and novel qualification strategies

© Gert Jervan

22




Scaling Trends & Reliability Considerations

- Dramatic increase in processing steps with each new generation
 - approx. 50 more steps per generation and a new metal level every 2 generations
- Rush to market - Less time to characterize new materials than in the past
 - e.g. reliability issues with new materials not fully understood and potential new failure modes
- Manufacturers' trends to provide 'just enough' lifetime, reliability, and environmental specs for commercial & industrial applications
 - e.g. 3-5 yr product lifetimes, trading off 'excess' reliability margins for performance

© Gert Jervan

23





Scaling Trends & Reliability Considerations

- Significant rise in the amount of proprietary technology and data developed by manufacturers, reluctance to share information with hi-relevance customers
 - e.g. process recipes, process controls, process flows, design margins, MTTF
- Next generation microelectronics focus on the performance needs of the commercial customer, with little or no emphasis on the extreme needs
 - e.g. extended life, extreme environments, high reliability
- Increasingly difficult testability challenges due to device complexity

© Gert Jervan

24






Product Technical Trends

	1990	2000	2010
Operating temperature, °C	-55 to 125	-40 to +85	0 to 70
Supply voltage	5v	1.5v	0.6v
Max. power (high perf.)	5	100	170
No. of package types	<10	<60	??
Design support life	>10 yrs.	1-5 yrs.	<1yr.
Production life	>10 yrs.	3-5 yrs.	<3yrs.
<u>Service life</u>	<u>>20 yrs.</u>	<u>5-10 yrs.</u>	<u><5yrs.</u>

© Gert Jervan

*MRQW-2002, Bernstein

26



Growing Internet Traffic

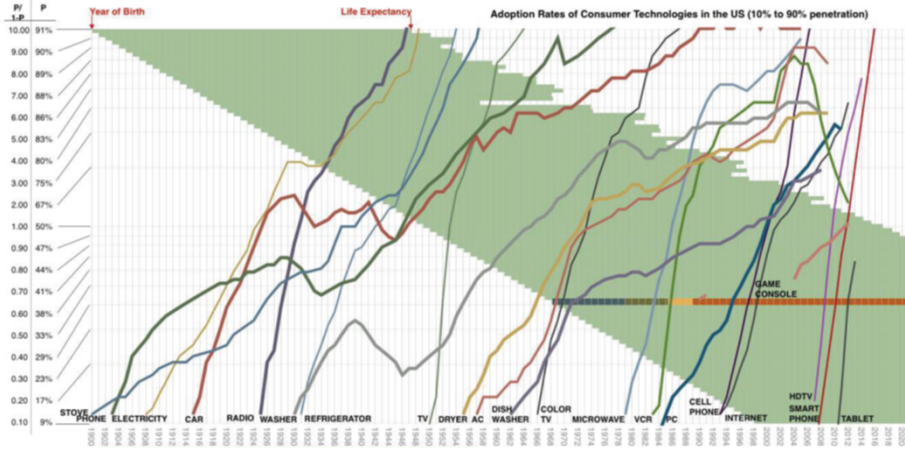
Year	Global Internet Traffic
1992	100 GB/Day
1997	100 GB/Hour
2002	100 GB/Sec
2007	2 000 GB/Sec
2012	12 000 GB/Sec
2017	35 000 GB/Sec


Cisco VNI, 2013

© Gert Jervan

27

Ubiquitous Computing



 School of Computer and Communication Sciences

End of the World

9

© Gert Jervan

James Larus, EPFL

28

Software complexity is a challenge

Aviation:

- Boeing 747 → 0.4 M LOC
- Boeing 777 → 4 M LOC
- Technology Review 2002

Software:

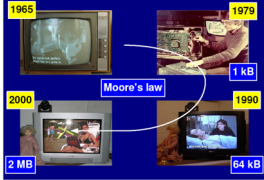
- Exponential increase in software complexity
- In some areas code size is doubling every 9 months [ST Microelectronics, Medea Workshop, Fall 2003]
- ... > 70% of the development cost for complex systems such as automotive electronics and communication systems are due to software development [A. Sangiovanni-Vincentelli, 1999]

Automotive:


- ✓ 2010 Premium → 100 M LOC
- ✓ 1995 – 2000 → 52%/Year
- ✓ 2001 – 2010 → 35%/Year

Tony Scott, GM CIO

- ✓ 2011 – BMW is the first manufacturer to break the 1Gb barrier




Rob van Ommering, COPA Tutorial, as cited by: Gerrit Müller: Opportunities and challenges in embedded systems, Eindhoven Embedded Systems Institute, 2004





Linux Complexity



31




Big Data

- An increasingly sensor-enabled and instrumented business environment generates HUGE volumes of data with MACHINE SPEED characteristics



- 1 Billion lines of code
- EACH engine (A380 has 4 of them) generating 10 TB every 30 minutes!

32




Course Overview

- To get an insight into the broad area of system safety
- We cover techniques for high availability, fault tolerance, monitoring, detection, diagnosis, and confinement of failure, ways to improve availability through fast recovery and graceful service degradation, and techniques for using redundancy and replication.
- We also discuss the utopia of flawless software, the impact of scale on availability, ways to cope with human operator error, and metrics for evaluating dependability.

© Gert Jervan

33




Contents

- Fault tolerance
- System reliability
- Hardware redundancy
- Error detection techniques
- Coding techniques
- Processor-level detection and recovery
- Disk arrays
- Checkpointing and recovery
- Software fault tolerance
- Testing distributed real-time systems
- ...

© Gert Jervan


34



Lecture Outline

- ✓ **Historical perspective and famous incidents/accidents**
- **Basic terminology**

35




Murphy's Law

- "If something can go wrong, it will go wrong"
Major Edward A. Murphy, Jr.
US Air Force, 1949
- "Every component than can be installed backward, eventually will be"


© Gert Jervan

36




Genesis Space Capsule

- \$260 million Genesis capsule was collecting samples of the solar wind over 3 years period
- Crashed in Sept 2004 due to the failure of the parachutes
- Reason:
 - the deceleration sensors — the accelerometers — were all installed backwards. The craft's autopilot never got a clue that it had hit an atmosphere and that hard ground was just ahead.



© Gert Jervan

37



Mars Orbiter

- One of the Mars Orbiter probes crashed into the planet in 1999.
- It did turn out that engineers who built the Mars Climate Orbiter had provided a data table in "pound-force" rather than newtons, the metric measure of force.
- NASA flight controllers at the Jet Propulsion Laboratory in Pasadena, Calif., had used the faulty table for their navigation calculations during the long coast from Earth to Mars.

© Gert Jervan

38



Lockheed Martin Titan 4

- In 1998, a LockMart Titan 4 booster carrying a \$1 billion LockMart Vortex-class spy satellite pitched sideways and exploded 40 seconds after liftoff from Cape Canaveral, Fla.
- Reason: frayed wiring that apparently had not been inspected. The guidance systems were without power for a fraction of a second.



© Gert Jervan

39



Therac-25

- Therac-25:
 - the most serious computer-related accidents to date (at least nonmilitary and admitted)
 - machine for radiation therapy (treating cancer)
 - between June 1985 and January 1987 (at least) six patients received severe overdoses (two died shortly afterward, two might have died but died because of cancer, the other two had permanent disabilities)
 - scanning magnets are used to spread the beam and vary the beam energy
 - dual-mode: electron beams for surface tumors, X-ray for deep tumors

© Gert Jervan

40



Denver Airport

- Denver International Airport, Colorado: intelligent luggage transportation system with 4000 "Telecars", 35km rails, controlled by a network of 100 computers with 5000 sensors, 400 radio antennas, and 56 barcode readers. Price: \$186 million (BAE Automated Systems).
- Due to SW problems about one year delay which costs \$1.1 million per day (1993).
- Abandoned in 2005 to save \$1 million per month on maintenance
- Today we have the on-going story with the new Berlin Brandenburg Airport
 - Scheduled to open in 2011, the new estimate is 2014

42



Boeing 787 Dreamliner

- Program launched in 2003, roll-out in 2007, first delivery in 2011. 114 delivered so far.
- Grounded on January 16, 2013 due to the problems with electrical circuitry
 - Leading to thermal runaway of Li-ion batteries and causing several fires in the battery compartment (several emergency landings, one aircraft (ET) was heavily damaged on ground)
 - Comprehensive review of the 787's critical systems, including the design, manufacture and assembly.
 - Japanese ANA alone lost 1.1 M USD per day (17 aircrafts)
- Grounding lifted on April 26, 2013




43



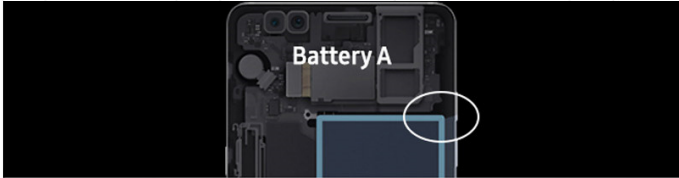
LAX airport ATC software failure

- 2,4 billion USD system (developed by Lockheed Martin) crashed on April 30, 2014.
 - Reason: U-2 spy plane that was Flying „too high“
 - Result: The system attempted to calculate all possible flight paths and run out of memory
- The “new \$40 billion air traffic control system, known as NextGen, which encompasses ERAM, including its reliance on Global Positioning System data that could be faked” is “very over-budget and behind schedule,” Moss (founder of Def Con) told Reuters. It “doesn't surprise me that it's got some bugs - it's the way it presented itself' that's alarming.” You can expect at least two upcoming Def Con talks to delve into exploiting weaknesses in the system.

44



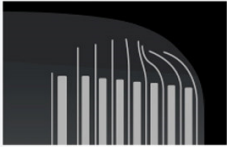
Samsung Galaxy Note 7



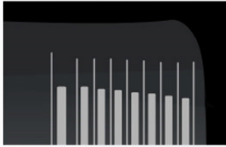
Abnormal

Normal

Main Cause




The negative electrode was deflected in the upper-right corner of the battery




The negative electrode is not deflected

© Gert Jervan

45



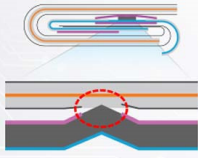
Samsung Galaxy Note 7




Abnormal

Normal

Main Cause



High welding burrs on the positive electrode resulted in the penetration of the insulation tape and separator which then caused direct contact between the positive tab with the negative electrode



The positive tab is appropriately attached to the positive electrode

Negative Electrode
Separator
Insulation Tape
Positive Tab
Positive Electrode

© Gert Jervan

46