Arvutitehnika instituut
ati.ttu.ee

Reliability and technology scaling: friends or foes?

Gert Jervan

Department of Computer Engineering
Tallinn University of Technology
ati.ttu.ee

Today: Hot Chips – but not only...

Market Shares

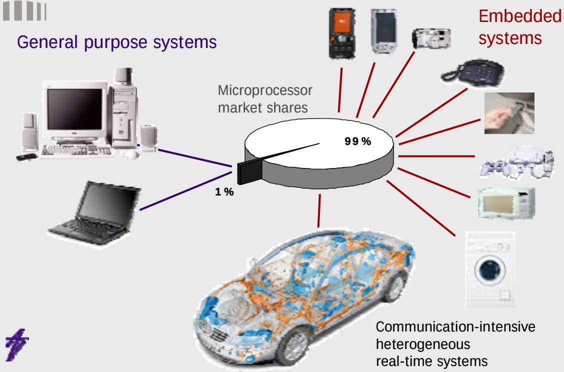

General purpose systems

Embedded systems

Microprocessor market shares

99 %

1 %

Communication-intensive
heterogeneous
real-time systems

Req'd Performance for Multi-Media Processing

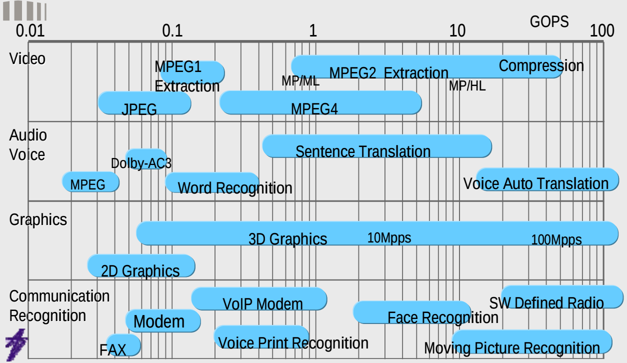

GOPS: Giga Operations Per Second

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

---

Hardware and Environment Failures

- ✓ Moving parts, high speed, low tolerance, high complexity: disks, tape drives/libraries
- ✓ Lowest MTBF found in fans and power supplies
- ✓ Often fans fail gradually → subtle, sporadic failures in CPU, memory, backplane
- ✓ Environment: power, cooling, dehumidifying, cables, fire, collapsing racks, ventilation, earthquakes, ...

A Few Layers of Computer Systems

- Algorithm
- Programming
- Architecture
- Organization
- Logic
- Integrated Circuit
- Device

What is Technology Scaling

Drawn to the same scale

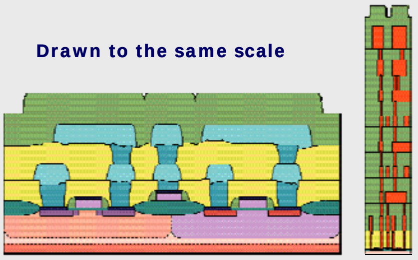

1.0 µm
Mid 1980s
Speed: 10 MHz

0.1 µm
Early 2000's
Speed: 3 GHz

Today: 65 nm, going down to 22 nm by 2016

Benefits of Technology Scaling

- ✓ Benefits of scaling the dimensions by 30%:
  - Reduce gate delay by 30% (increase operating frequency by 43%)
  - Double transistor density
  - Reduce energy per transition by 65% (50% power savings @43% increase in frequency)

Technology Scaling

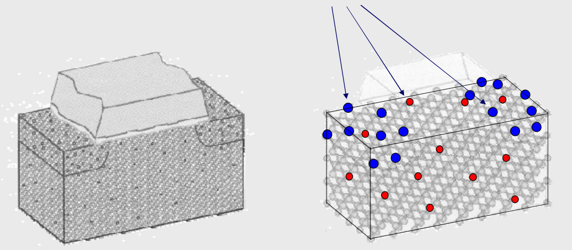

20 nm MOSFET (2015?)
50 Si atoms along the channel

4 nm MOSFET (2020?)
10 Si atoms along the channel

Technology Roadmap

International Technology Roadmap for Semiconductors 2002

Year	2001	2003	2005	2007	2010	2013	2016
DRAM ½ pitch [nm]	130	100	80	65	45	32	22
MPU transistors/chip	97M	153M	243M	386M	773M	1.55G	3.09G
Wiring levels	8	8	10	10	10	11	11
High-perf. phys. gate [nm]	65	45	32	25	18	13	9
High-perf. VDD [V]	1.2	1.0	0.9	0.7	0.6	0.5	0.4
Local clock [GHz]	1.7	3.1	5.2	6.7	11.5	19.3	28.8
High-perf. power [W]	130	150	170	190	218	251	288
Low-power phys. gate [nm]	90	65	45	32	22	16	11
Low-power VDD [V]	1.2	1.1	1.0	0.9	0.8	0.7	0.6
Low-power power [W]	2.4	2.8	3.2	3.5	3.0	3.0	3.0

Node years: 2007/65 nm, 2010/45 nm, 2013/32 nm, 2016/22 nm

Moore's Law in Microprocessors

Transistors on lead microprocessors double every 2 years

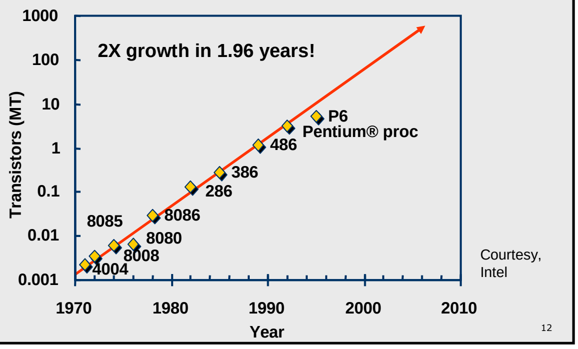

Courtesy, Intel

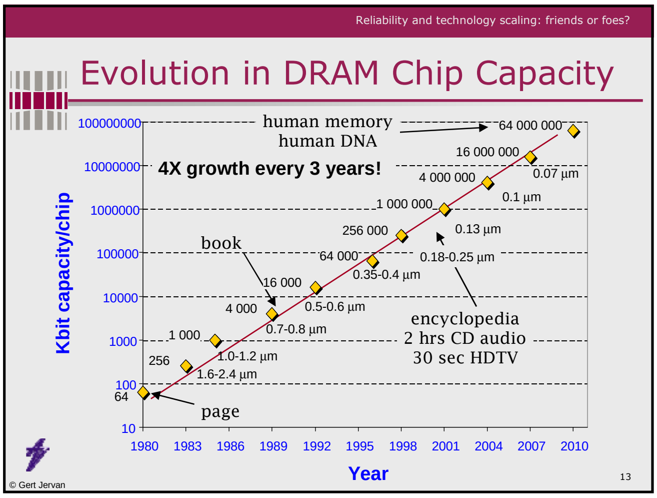

Reliability and technology scaling: friends or foes?

Evolution in DRAM Chip Capacity

© Gert Jervan

13

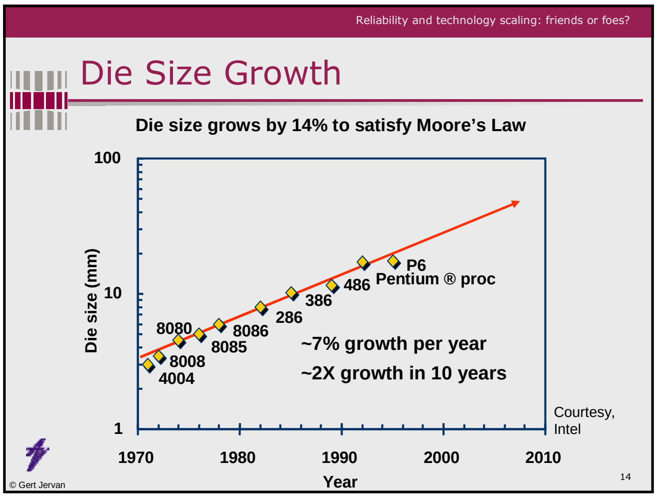

Reliability and technology scaling: friends or foes?

Die Size Growth

Die size grows by 14% to satisfy Moore's Law

Courtesy, Intel

© Gert Jervan

14

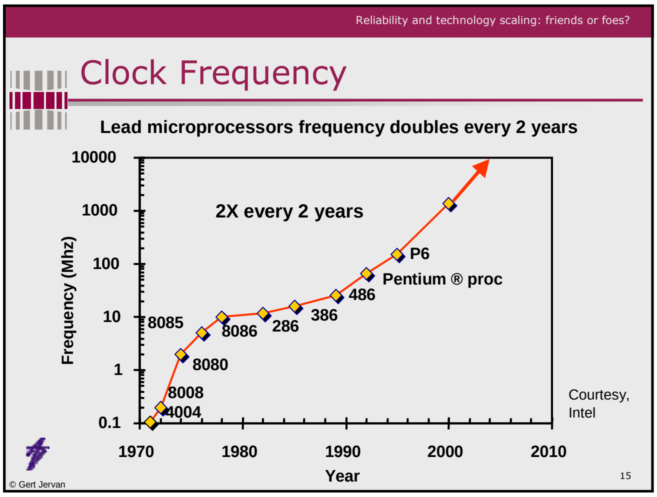

Reliability and technology scaling: friends or foes?

Clock Frequency

Lead microprocessors frequency doubles every 2 years

Courtesy, Intel

© Gert Jervan

15

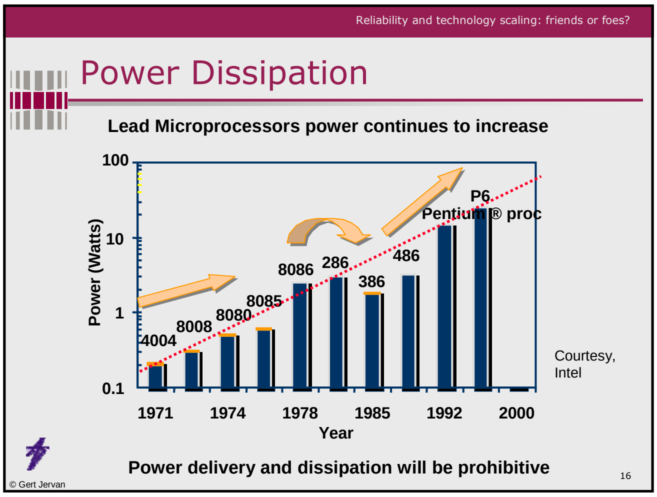

Reliability and technology scaling: friends or foes?

Power Dissipation

Lead Microprocessors power continues to increase

Courtesy, Intel

Power delivery and dissipation will be prohibitive

© Gert Jervan

16

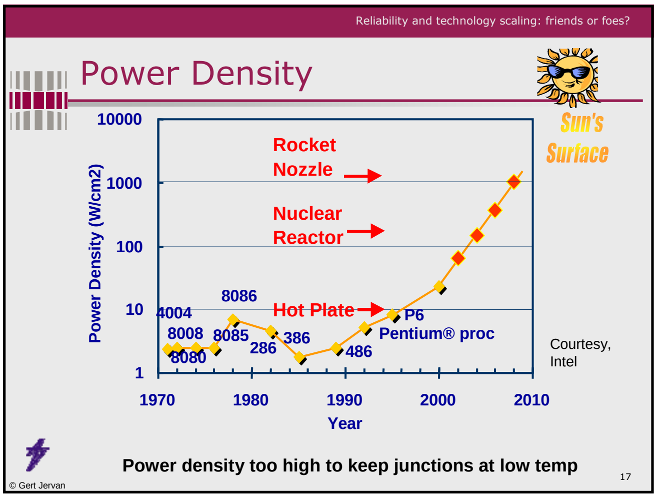

Reliability and technology scaling: friends or foes?

Power Density

Courtesy, Intel

Power density too high to keep junctions at low temp

© Gert Jervan

17

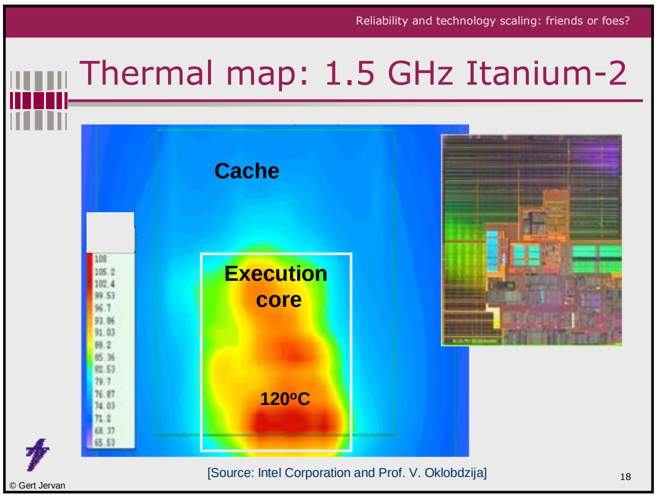

Reliability and technology scaling: friends or foes?

Thermal map: 1.5 GHz Itanium-2

[Source: Intel Corporation and Prof. V. Oklobdzija]

© Gert Jervan

18

Industry Scaling Trends & Reliability Considerations

- ✓ Reduced gate oxide thicknesses
- ✓ Increased thermal/power densities
- ✓ Reduced interconnect dimensions
- ✓ Higher device operating temperatures
- ✓ Increased sensitivity to defects and statistical process variations
- ✓ Introduction of new materials with each new generation, replacing proven materials
  - e.g. Cu and low K inter-level dielectrics for Al and SiO2

Industry 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

---

Industry Scaling Trends & Reliability Considerations

- ✓ Significant rise in the amount of proprietary technology and data developed by manufacturers, reluctance to share information with hi-rel. 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

---

Implications to Design

- ✓ Design fabric will be Regular
- ✓ Will look like Sea-of-transistors interconnected with regular interconnect fabric
- ✓ Shift in the design efficiency metric
  - From Transistor Density to Balanced Design

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.
Service life	***>20 yrs.***	***5-10 yrs.***	***<5yrs.***

*MRQW-2002, Bernstein

Commercial Chip Reliability Estimation

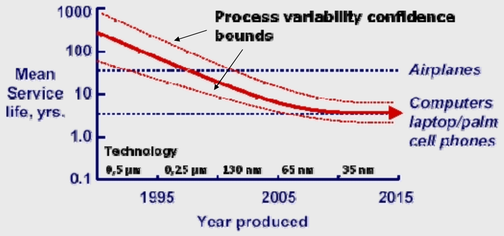

*Extrapolated from ITRS roadmap, MRQW-2002, Bernstein

Device Reliability Trends

As technology progresses, wearout failures become statistically indistinguishable from infant mortality failures with the same wearout drivers.

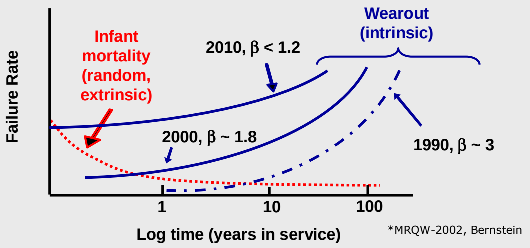

*MRQW-2002, Bernstein

Transient faults

- Happen for a short time
- **Corruptions of data, miscalculation in logic**
- Do not cause a permanent damage of circuits
- Causes are outside system boundaries

Radiation

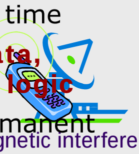

Electromagnetic interference (EMI)

Lightning storms

Intermittent faults

- Manifest similar as transient faults
- Happen repeatedly
- Causes are inside system boundaries

Internal EMI

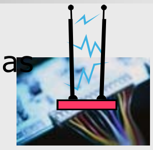

Crosstalk

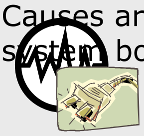

Power supply fluctuations

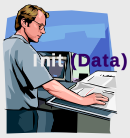

Software errors (Heisenbugs)

Soft Errors

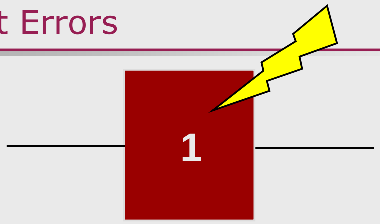

- ✓ Transient bit-flip (soft memory error)
  - Random event
  - Corrupts the value but not the cell
  - Can be corrected (in contrast to hard errors caused by faults in the hardware itself)
  - Happen continuously during system lifetime (*i.e.*, can not be screened by burn-in tests)

Sources

- ✓ First traced to alpha particlce emissions from chip packaging materials
  - Most sources removed (pure materials, different designs, shielding)
- ✓ Today’s main problem: cosmic radiation
  - Cosmic particles from deep space (actually 5th- or 6th-hand collision particles)
    - At ground level ca 95% neutrons, 5% protons
  - Radioactive material in manufacturing process

Sources (cont.)

- ✓ Four main sources:
  - Low-energy alpha particles
  - High-energy cosmic particles
  - Thermal neutrons
  - Poor system design

SER type	Source	Mechanism	Trend
Alpha	Thorium and uranium contamination in-mold compound, silicon, or lead bumps	2- to 9-MeV alpha particle creating electron-hole funnel traveling 25 microns in silicon	Exponential increase with scaling
Cosmic	Intergalactic sources modulated by solar flares	High-energy neutrons/protons (10 MeV to 1 GeV) colliding with silicon nuclei	Decrease in failures in time per megabit
Thermal neutron	Boron present in BPSG25-meV neutrons	Collision with B10 in BPSG	Highest, always dominates if present

Soft Errors

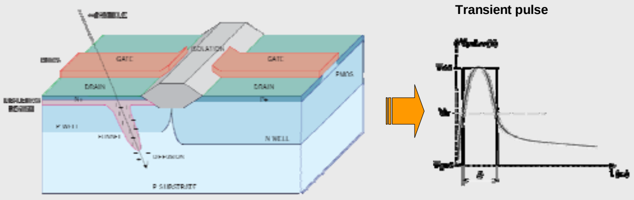

The electric field in the depletion region directly generates electron-hole pairs in its wake, causing the charges to drift so that the transistor sees a current disturbance

---

The Altitude Factor

- ✓ A person on an airplane over the Atlantic at 35,000 ft working on a laptop with 256 Mbytes (2 Gbits) of memory. At this altitude, the SER of 600 FITs per megabit becomes 100,000 FITs per megabit, resulting in a potential error every five hours.
- ✓ 1 FIT (failures in time), is the number of failures in 1 billion device-operation hours. A measurement of 1000 FITs corresponds to a MTTF (mean time to failure) of approximately 114 years.

Evidence of Cosmic Ray Strikes

- ✓ Documented strikes in large servers found in error logs
  - ▪ Normand, “Single Event Upset at Ground Level,” IEEE Transactions on Nuclear Science, Vol. 43, No. 6, December 1996.
- ✓ Sun Microsystems, 2000 (R. Baumann, Workshop talk)
  - ▪ Cosmic ray strikes on L2 cache with defective error protection
    - • caused Sun’s flagship servers to suddenly and mysteriously crash!
  - ▪ Companies affected
    - • Baby Bell (Atlanta), America Online, Ebay, & dozens of other corporations
    - • Verisign moved to IBM Unix servers (for the most part)

Current Situation

- ✓ Soft errors induced the highest failure rate of all other reliability mechanisms combined

Rober Baumann, TI

Physical Solutions are hard

- ✓ Shielding?
  - ▪ No practical absorbent (e.g., approximately > 1 m of concrete)
  - ▪ unlike Alpha particles
- ✓ Technology solution: SOI?
  - ▪ Partially-depleted SOI of some help, effect on logic unclear
  - ▪ Fully-depleted SOI may help, hard to manufacture in high volumes
- ✓ Radiation-hardened cells?
  - ▪ 10x improvement possible with significant penalty in performance, area, cost
  - ▪ 2-4x improvement may be possible with less penalty
- ✓ Some of these techniques will help alleviate the impact of Soft Errors, but not completely remove it

---

Arvutitehnika instituut
ati.ttu.ee

So?

Should we stop to develop?

Arvutitehnika instituut
ati.ttu.ee

Definitely not!

But,

we have to have different assumptions

System Design & Evaluation Top-Level View

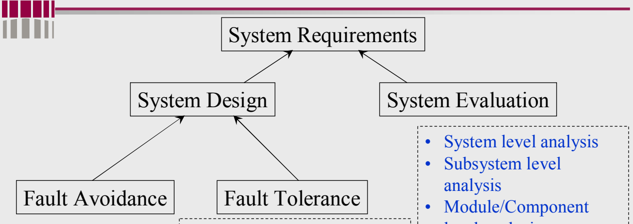

System Requirements

System Design

System Evaluation

Fault Avoidance

Fault Tolerance

Possible techniques
- Parts selection
- Design reviews
- Quality control
- Design Methodology
- Documentation

Possible techniques
- Redundancy (Hardware, Software, Information, Time)
- Fault detection
- Fault masking
- Fault containment
- Reconfiguration

- System level analysis
- Subsystem level analysis
- Module/Component level analysis

Possible Techniques
- FMEA
- FTA
- RBD
- Markov
- Petri net

Failure Rate

- ✓ Assuming a unit works correctly in [0,t], the conditional probability λ(t) that a unit fails in [t, t +Δt]
- ✓ Typically the failure λ rate depends on
  - Temperature
  - Time (burn-in and aging)
  - Environmental exposure
    - Soft errors, EMI
- ✓ Often the component failure rate is assumed to be constant for simplicity

---

Failure Rate: The Bathtub Curve

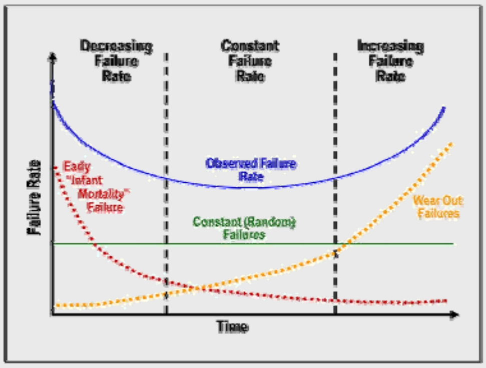

Reliability

- ✓ The probability function R(t) that a system works correctly in [0, t] without repairs
- ✓ Reliability is a function of time
  - If the system consist of a single component with constant failure rate λ, then
    - R(t) = exp (– λt)
  - The mean time to failure is MTTF = 1/ λ
- ✓ In general, the MTTF is E[t] = ∫ R(t)dt

Dependability Concepts

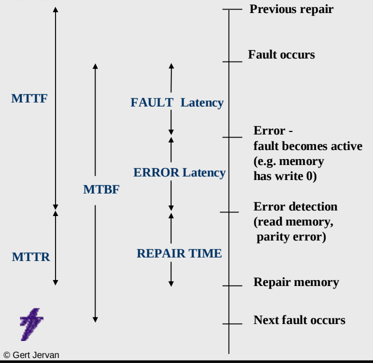

Reliability:
a measure of the continuous delivery of service;
R(t) is the probability that the system survives (does not fail) throughout [0, t];
expected value: ***MTTF(Mean Time To Failure)***

Maintainability:
a measure of the service interruption
M(t) is the probability that the system will be repaired within a time less than t;
expected value: ***MTTR (Mean Time To Repair)***

Availability:
a measure of the service delivery with respect to the alternation of the delivery and interruptions
A(t) is the probability that the system delivers a proper (conforming to specification)service at a given time t.
expected value: ***EA = MTTF / (MTTF + MTTR)***

Safety:
a measure of the time to catastrophic failure
S(t) is the probability that no catastrophic failures occur during [0, t];
expected value:
MTTCF(Mean Time To Catastrophic Failure)

Reliability is a System Issue

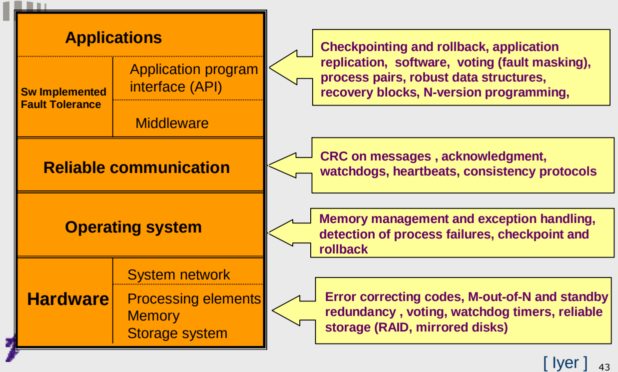

Checkpointing and rollback, application replication, software, voting (fault masking), process pairs, robust data structures, recovery blocks, N-version programming,

CRC on messages , acknowledgment, watchdogs, heartbeats, consistency protocols

Memory management and exception handling, detection of process failures, checkpoint and rollback

Error correcting codes, M-out-of-N and standby redundancy , voting, watchdog timers, reliable storage (RAID, mirrored disks)

[Iyer]

Classical Methods

- ✓ Hardware redundancy (space redundancy)
- ✓ Software redundancy
- ✓ Time redundancy
- ✓ Information redundancy

Hardware redundancy

Tripple-Modular Redundancy Technique

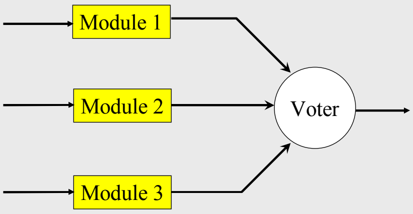

Tolerates N/2 faults

Fault-Tolerance Capability

- Assuming perfect voter, how many module faults can the TMR technique tolerate?
- What if 2 modules fail the same way?
- Does TMR technique provide fault detection capability?
- How about imperfect voter?
- Performance impacts from the voter in the TMR technique

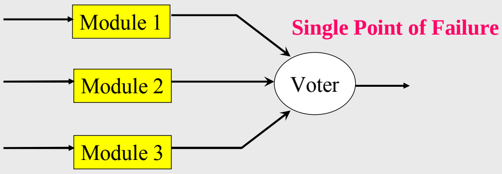

Single Point of Failure

Reliability of a TMR System

$$R_{TMR} = R_1R_2R_3 + (1-R_1)R_2R_3 + R_1(1-R_2)R_3 + R_1R_2(1-R_3)$$

$$R_1 = R_2 = R_3 = R$$

$$R_{TMR} = 3R^2 - 2R^3$$

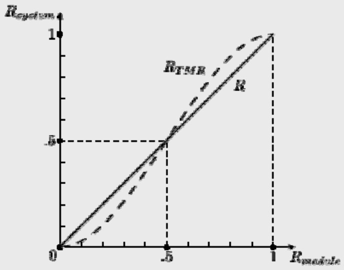

Fault Tolerance at System Level

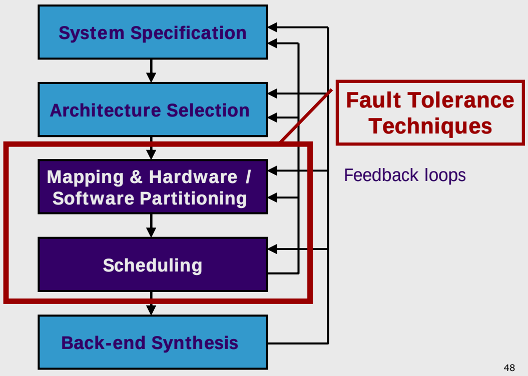

Fault Tolerance Techniques

Feedback loops

Fault Tolerance Techniques

Re-execution

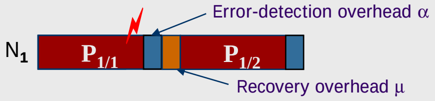

Rollback recovery with checkpointing

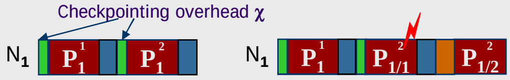

Active replication

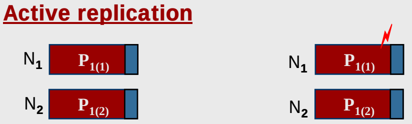

New Architectures

- ✓ Massively parallel architectures (Von Neuman is dead...) based on hundreds (millions) of (non-)reliable components
  - Multiple Input stream, Multiple Data stream machines
  - Wide use of network infrastructures (Networks-on-Chip)
  - Built-In Self-Repair will become a widespread technology
    - Dynamic reconfiguration

Arvutitehnika instituut
ati.ttu.ee

The problem to be solved:

How to design reliable system out of non-reliable hardware?

Arvutitehnika instituut
ati.ttu.ee

Thank You!

Gert Jervan

Department of Computer Engineering
Tallinn University of Technology
ati.ttu.ee