CarnegieMellon

Faults and Uncertainty: Do We Need a Totally New Approach to Properly Address These Problems?

Radu Marculescu
Carnegie Mellon University
IWLS, Laguna Beach, May 29, 2003

Electrical & Computer ENGINEERING

CarnegieMellon

Q2: Re-thinking on-chip communication

- Interconnections become dominant in DSM era
  - Buses are not really scalable and consume too much power!
- Regularized, tile-based Network-on-Chip (NOC) architecture
  - Very good for predictable and affordable designs

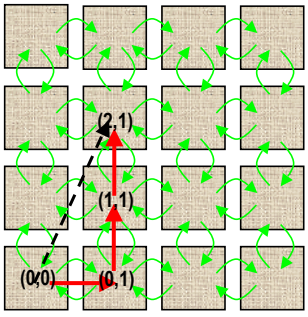

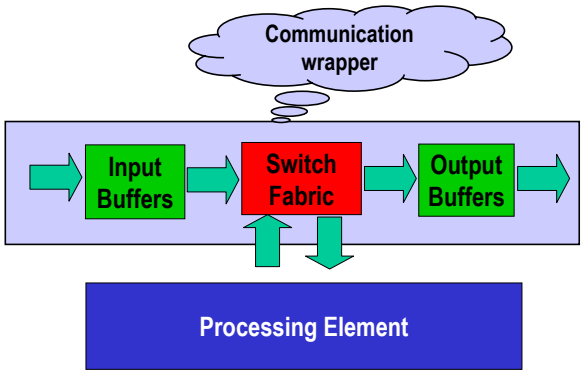

CarnegieMellon

An old story about bridges

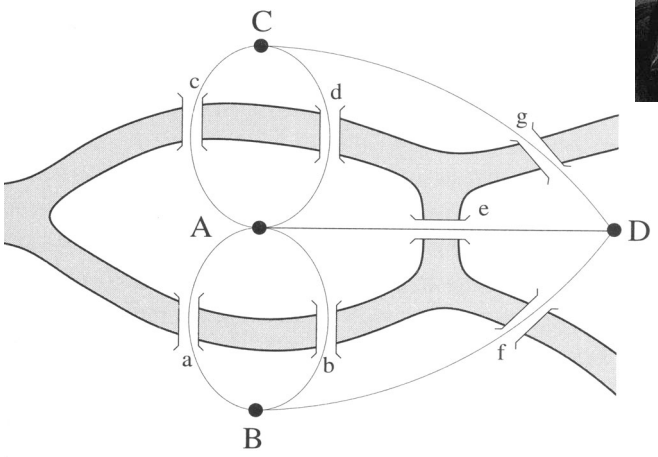

Big idea: Think Network = Think graphs

CarnegieMellon

…and a party example

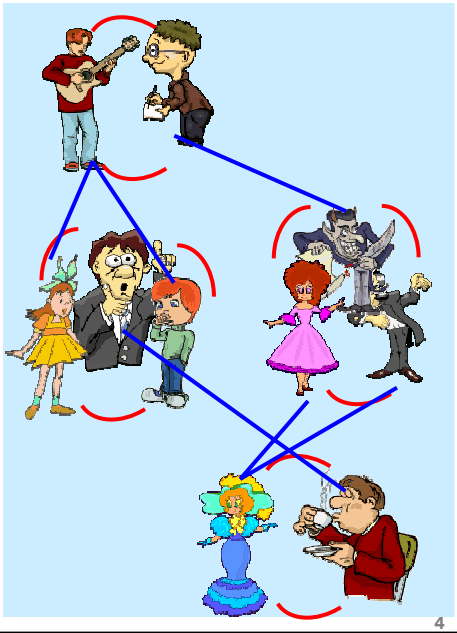

Big idea: Think random graphs

Q5: Inject Randomness into communication

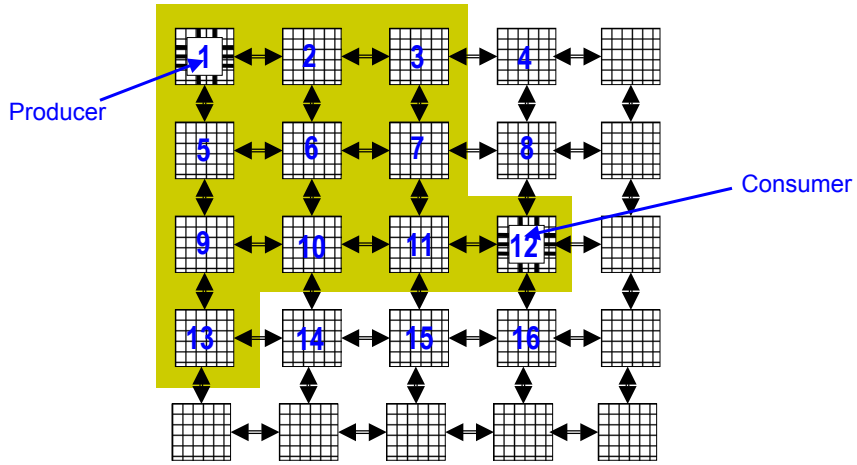

A random universe

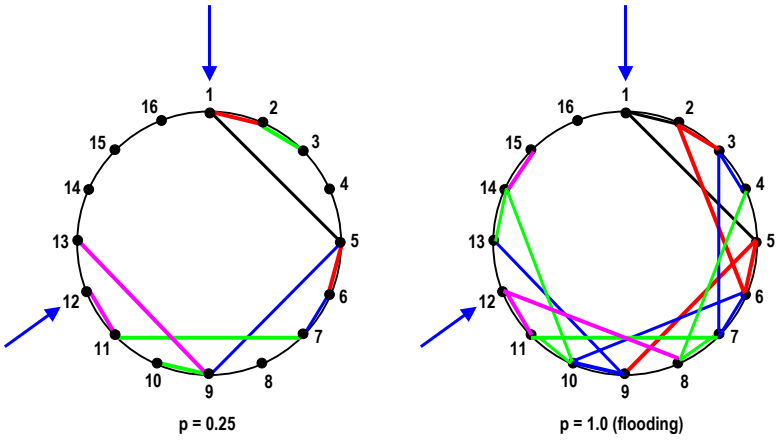

Big idea: Ambiguity about structure

Q6: Types of uncertainties we have to deal w/ Q3: Is this really a problem?

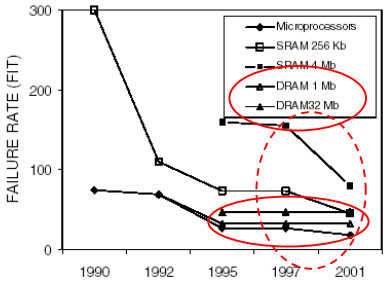

Fig. 1. Permanent failure rates for CMOS semiconductor devices (Source: Telcordia Technologies, Reliability Procedure for Electronic Equipment. FIT=failures in 10^9 hours)

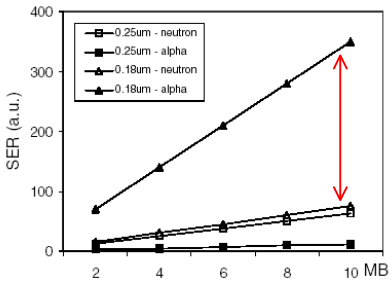

Fig. 2. SER for two CMOS SRAM technology generations

Big idea: System-level fault-tolerance is needed

Q4: How much Fault-Tolerance is needed?

Stochastic Communication – with faults

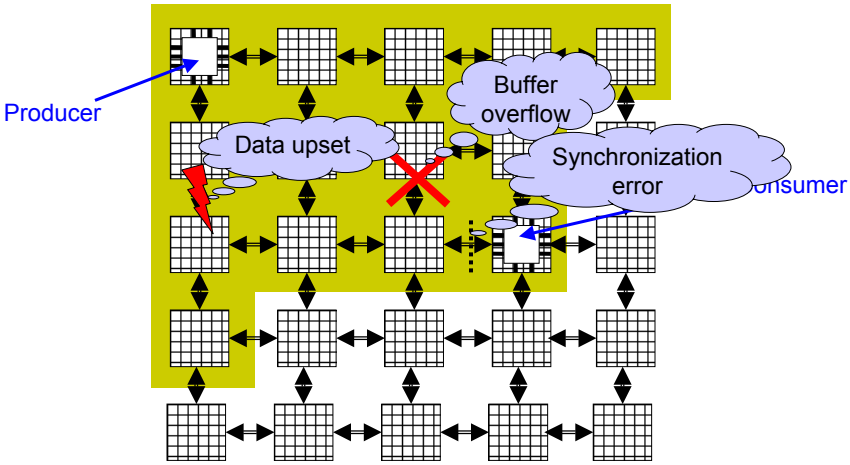

A simple MS example

$$\pi = \int_0^1 \frac{4}{1 + x^2}\,dx \approx \sum_{i=1}^{n} \frac{1}{n} \cdot \frac{4}{1 + \left(\left(i - \frac{1}{2}\right)\frac{1}{n}\right)^2}$$

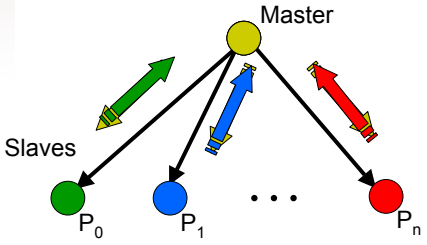

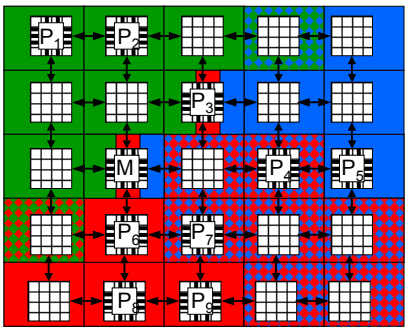

Fault-Tolerance

- **Tile failures vs. data upsets**
  - **Tile failures**
    - Little impact on latency
    - Will cause communication to fail
  - **Data upsets**
    - Bigger impact on latency
    - Will not cause communication to fail (except when ≈ 100%)

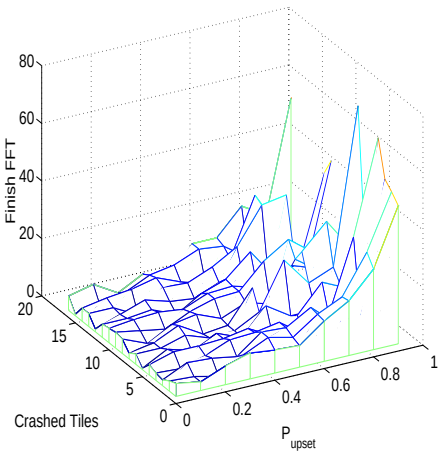

Summary

- *Yes, we need a paradigm shift to properly address faults and design uncertainty*
  - **Move from deterministic to probabilistic design**
  - **Route packets instead of wires. Exploit regularity**
  - **Fault-tolerance**
    - On-chip stochastic communication: a fundamentally new perspective on communication mechanisms for regular architectures
      - Affordable, scalable, and fault-tolerant communication scheme that can easily integrate synchronous and asynchronous domains
      - Low manufacturing, testing and design costs
      - Low latency (since it does not require costly retransmissions) and high flexibility
  - **Wide open area for research!**

Thank you!

More info: www.ece.cmu.edu/~sld