Trading Off Lifetime, Fault-tolerance, and Power Consumption in Real-time MPSoC

Jacopo Panerati*, Samar Abdi†, and **Giovanni Beltrame***

**École Polytechnique de Montréal*, †*Concordia University*

MPSoC 2015 - Ventura Beach, CA, USA

Table of contents

Outline

Motivation

- Aerospace: high-frequency of Single Event Upsets
- Usually critical systems, requiring high availability
- Classical countermeasures:
  - Modular redundancy
  - Shielding
- Issues:
  - Cost
  - Extra hardware $\Longrightarrow$ more power $\Longrightarrow$ higher temperature $\Longrightarrow$ shorter lifetime
- What is a good trade-off?

Research Goal

- Reliability and fault-tolerance are essential for critical, autonomous systems
- We propose a methodology to quantify, and maximize, reliability in the presence of transient errors for MPSoC
- Fault-tolerance is traded-off with power consumption
- We target homogeneous multi-processor systems
  - Goal: keep a certain level of reliability/lifetime with varying fault rates

Outline

System Model

- Multiprocessor System-on-Chip (we're in the right place!)
- Identical processing elements (PEs) w/ private caches
- Voltage scaling: a set of operating points for each PE

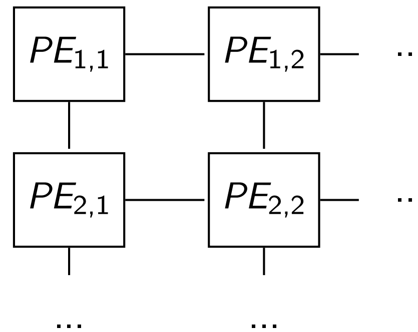

Fault models

- Transient faults (SEUs) w/ data scrubbing
- Permanent Faults
- Total Ionizing Does (TID) effects

Real-Time Application Model

- A set of tasks $\tau_1, \tau_2 ..., \tau_m$ is executed
- Each task has a WCET associaded with the slowes operating point of a PE
- The speedup is proportional to the frequency increase

$$WCET_{OP(f_i,-)} = WCET_{OP(f_0,-)} \cdot \frac{f_0}{f_i}$$

- Precedences via a Directed Acyclic Graph (DAG)

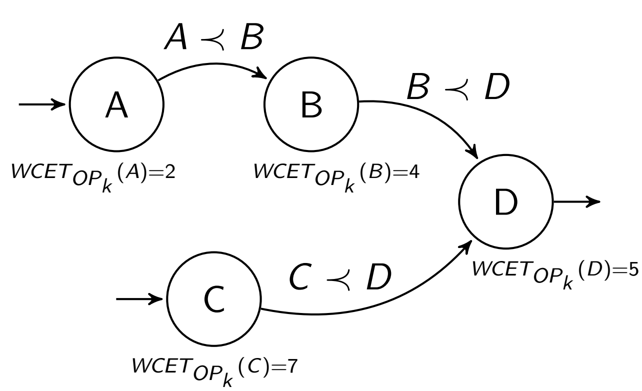

Single Event Upsets

We use probability theory to model the occurrence of faults. SEUs are caused by high-energy particles:

- Whose impacts are independent.
- Which happen at a constant average rate.
- The rate is mission phase-dependent.

The number of impacts in a scrubbing period of length T is a Poisson rand variable.

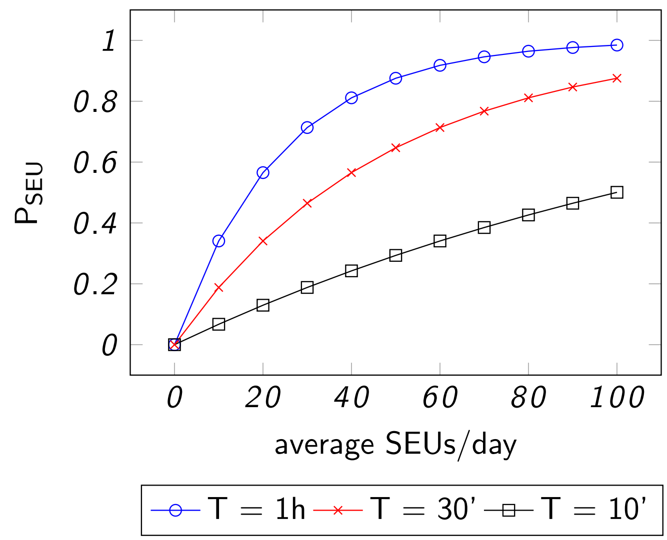

Permanent Faults

- We consider the most common wear-out phenomena: hot carriers, negative bias temperature instabiliti (NBTI), time dependent dielectric breakdown (TDDB), electromigration, and self-heating
- Hypothesize that Mean Time To Fail (MTTF) has an exponential relationship with PE load (utilization U)

$$MTTF_U \propto (MTTF_{100\%})^{U^{-1}}$$

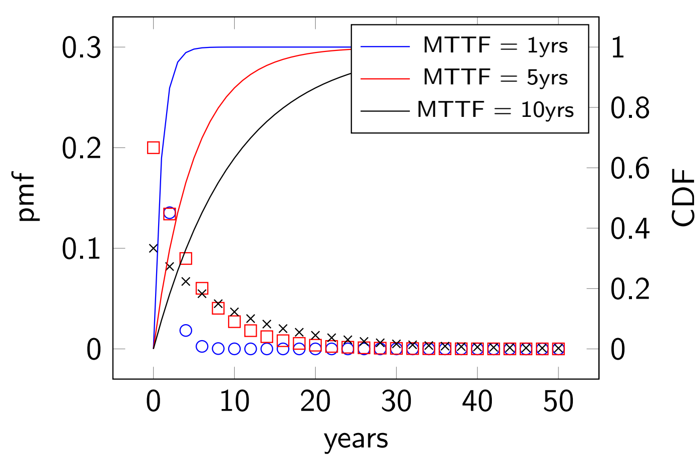

Power Model

- Total power = sum of each PE
- Standard model with capacitance, frequency, activation factor

$$P = \alpha \cdot C \cdot V^2 \cdot f$$

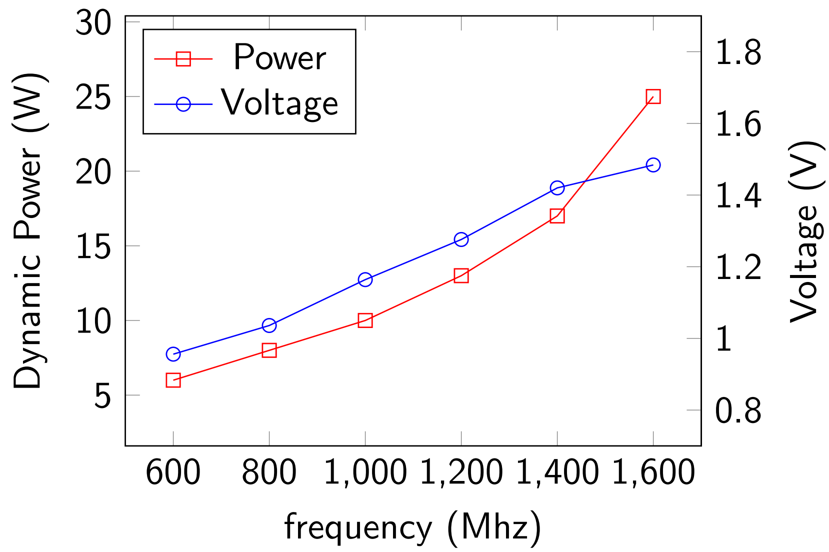

Outline

1. Introduction
2. System Model
3. Methodology
4. Case Study
5. Conclusions

Methodology

Task Mapping

- Enumerate all possible mappings
- Prune the design space according to WCET and slowest operating point
- Compute the utilization for each mapping

Power, Fault-tolerance, and Lifetime Optimization

- Compute the total energy according to utilization and operating points
- Utilizations reflect exponentially on the probability of system-wide error
  - Slack provides fault-tolerance
- We consider the effect of utilization on lifetime and the failure of multiple resources for lifetime optimization

Outline

Case Study (actually a toy example)

- Dual core, four tasks, each PE has four operating points
- Implementation on a Virtex 4 board
  - 16.5 faults/day in Low Earth Orbit (LEO)
  - 62 faults/day in Highly Elliptical Orbit (HEO)

		Operating Point		
		OP_1	OP_2	OP_3
		$f_1 = 600\text{MHz}$	$f_2 = 1.2\text{Ghz}$	$f_3 = 1.6\text{Ghz}$
Task	A	8.0	4.0	3.0
	B	4.0	2.0	1.5
	C	8.0	4.0	3.0
	D	12.0	6.0	4.5

Results

- Overall 29 acceptable points, 15 different points shown here
- Trade-offs for utilization (lifetime), power efficiency, or fault-tolerance

Average Utilization	Best Power Consumption	System Errors LEO	HEO
0.600	30.00W	**12**	**42**
0.650	27.70W	13	45
0.675	26.55W	14	47
0.700	25.40W	15	49
0.725	24.25W	15	50
0.800	**20.80W**	16	56
0.850	27.30W	17	59

Results

- Design space as an n-dimensional space of utilization levels, with reliability and power consumption design points

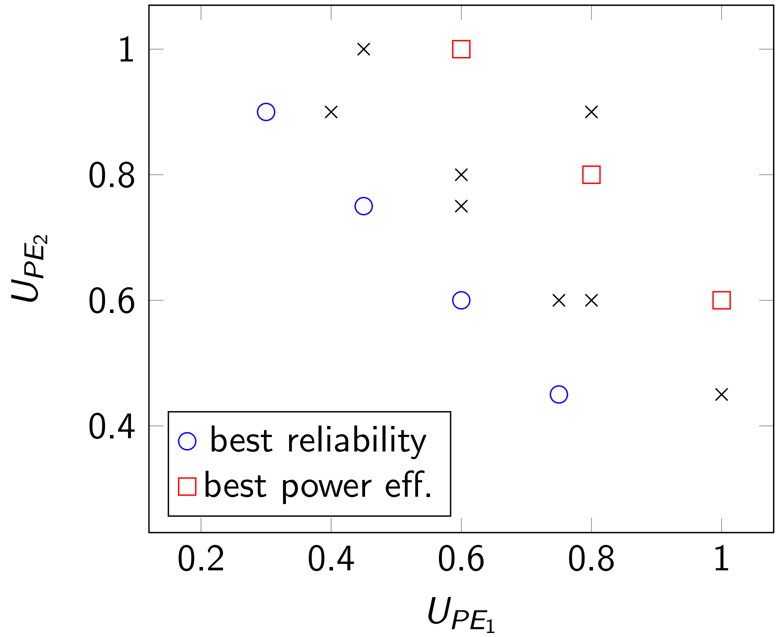

Outline

Conclusions

- Methodology for scheduling real-tiem tasks in homogeneous MPSoCs
- Energy, fault-tolerance, and lifetime-aware

Future Work

- Use a detailed temperature model instead of the utilization proxy
- Extend to the effects of interconnects
- More detailed modelling of permanent faults

POLYTECHNIQUE MONTRÉAL Introduction System Model Methodology Case Study Conclusions

The End

Questions?

`http://mistlab.ca`