Robust System Design with MPSoCs
Unique Opportunities

Subhasish Mitra

Robust Systems Group

Departments of Electrical Eng. & Computer Sc.

Stanford University

Email: subh@stanford.edu

Acknowledgment: Stanford Robust Systems Group Students & Collaborators

1

Robust Systems in Scaled CMOS

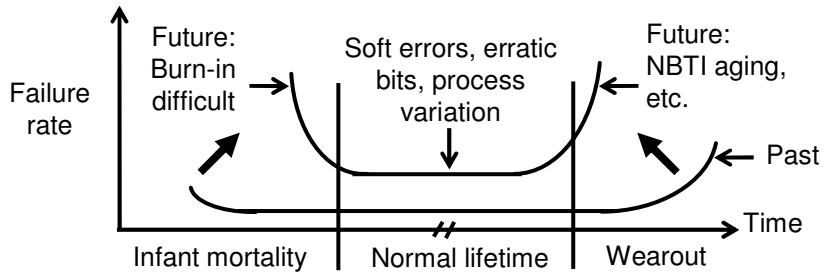

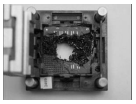

2001 Burn-in & Test Socket Workshop

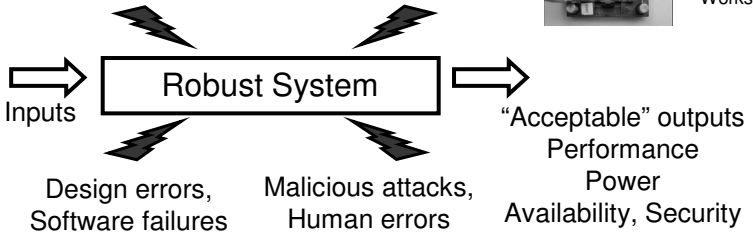

2

Globally Optimized Robust System Design

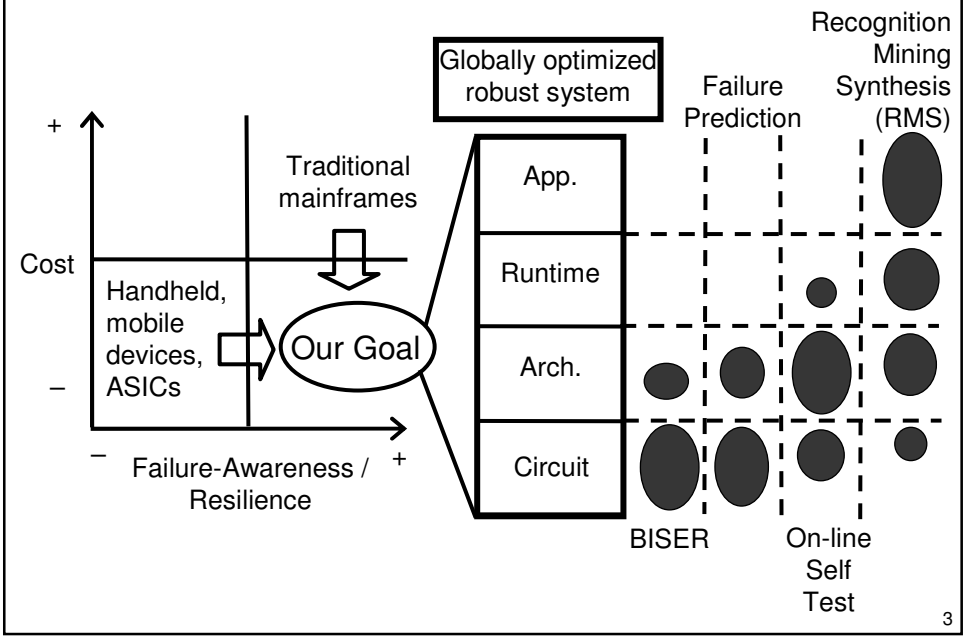

3

Multi-Core Robust System Opportunities

- Multi-core ≠ massive redundancy (e.g., TMR)
- Reconfigurable reliability – turn protection on / off
  - ❖ Power expensive, LOCAL gates inexpensive
- On-line self test – unique multi-core opportunities
  - ❖ Required for Circuit Failure Prediction
  - ❖ Essential for intelligent sparing
- Asymmetric core-level reliability
  - ❖ Globally optimized performance vs. resilience

4

Outline

- Introduction
- Built-In Soft Error Resilience (BISER)
- Circuit failure prediction
- Error Resilient System Architecture (ERSA)
- Conclusion

5

Key Takeaway

- Logic soft errors – extremely important
- Classical error detection + recovery
  - ❖ Expensive & complex
- Built-In Soft Error Resilience (BISER)
  - ❖ Best of all – Efficient & practical
  - ❖ Latch erratic errors also corrected

CORRECT Errors
DON'T Detect !!

6

Who Cares About Soft Errors ?

- Transient errors: α-particles, neutrons
- 20K processor server farm
  - ❖ 1 "major" error every 20 days
  - ❖ Silent data corruption (SDC)
    - 💣 $ 20K → $ 3,616 deposit
    - ☺ $ 20K → $ 52,768 deposit
  - ❖ Detected but uncorrected (DUE)
    - ➢ Downtime cost: $100K - $10M / hour

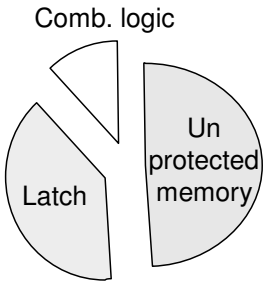

Soft error rate contributions

BISER for Latch Errors

A	B	C-element (A, B)
0	0	1
1	1	0
0	1	Previous value
1	0	Previous value

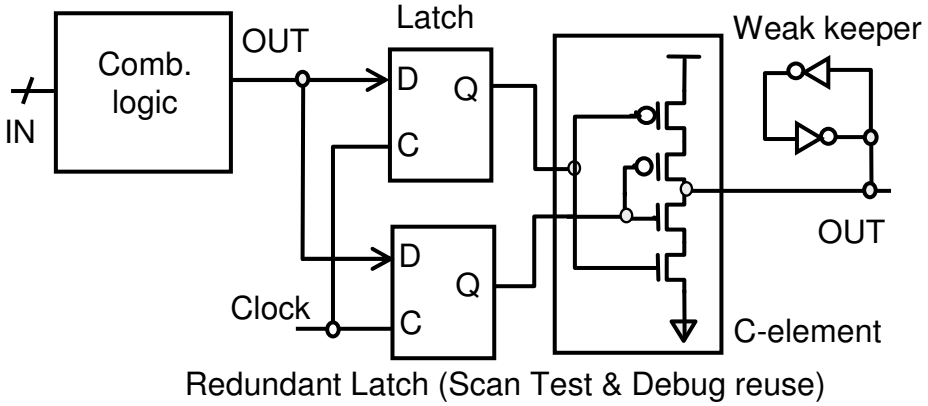

Redundant Latch (Scan Test & Debug reuse)

BISER Latch Error Correction Principle

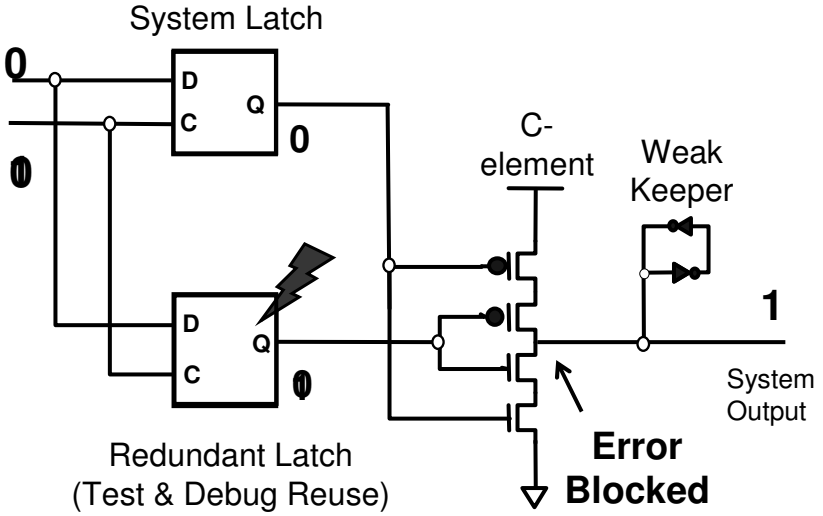

Key Observation: Latches vulnerable only in Opaque state (CLK = 0)

9

BISER Scan Flip-flop + Reconfigurable Protection – Design Reuse Enabled

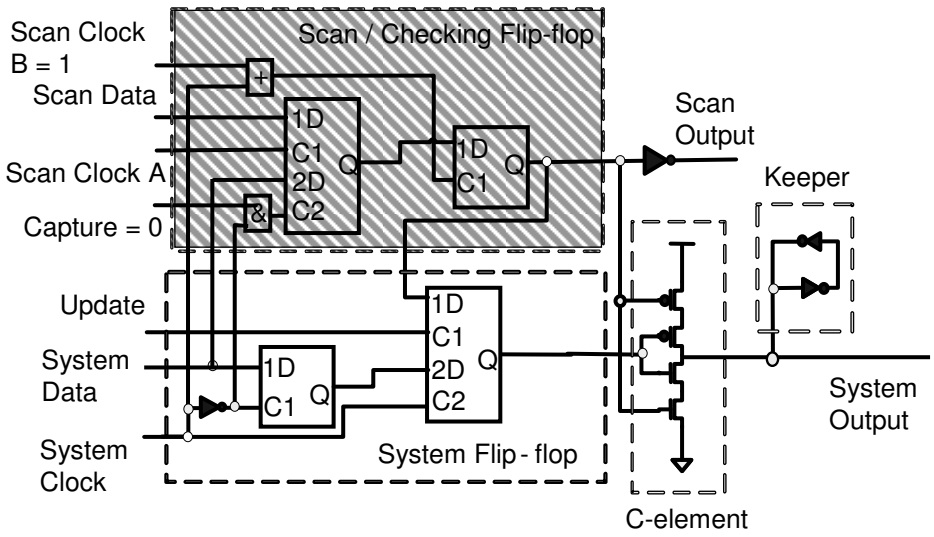

10

Architecture-Aware BISER Insertion

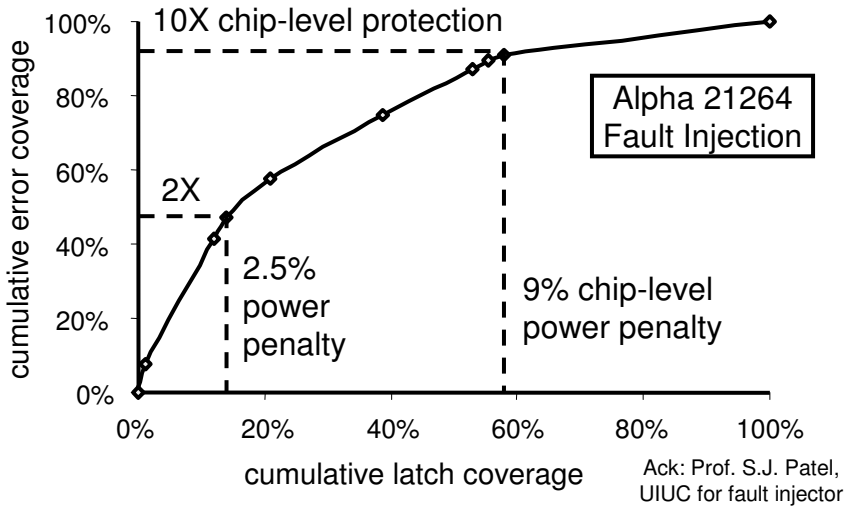

Ack: Prof. S.J. Patel, UIUC for fault injector

- Optimized BISER insertion [Seshia, Li & Mitra, DATE 07]
  - ❖ Maximize protection, minimize power penalty

11

BISER: Latch + Combinational Logic Errors

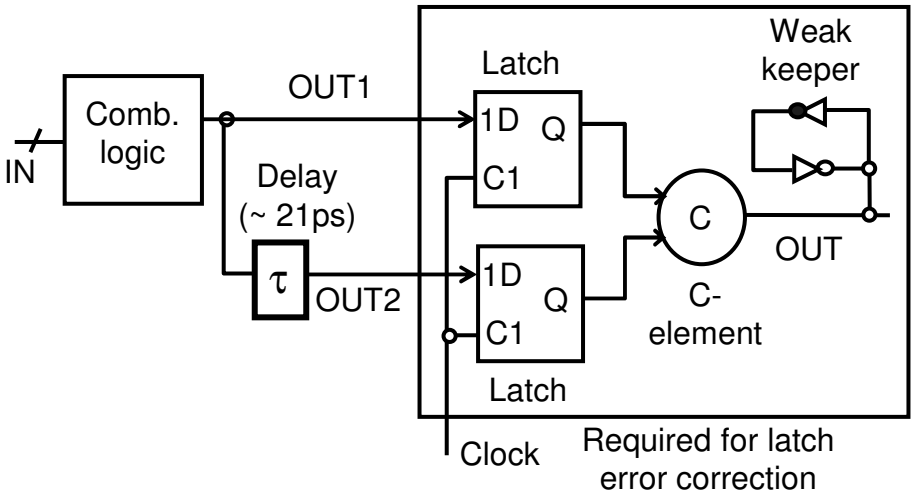

Incremental power cost over latch error correction: minimal

12

BISER vs. Traditional Techniques

	BISER	DICE	Duplicate	Multi-thread
Latch SER	> 25X ↓	25X ↓	Detected	Detected
Comb. logic SER	12 - 64X ↓	Ineffective	Detected	Detected
Chip energy	6 – 9%	> 15%	40-100%	~ 40 % (?)
Speed penalty	Latch: 1% Comb: 5%	1%	Minimal	~ 50%
Die size increase	~ None	~ None	Yes	Unclear
Downtime	None	None	Yes	Yes
Recovery	None	None	Complex	Complex
Configurability	Yes	No	Yes	Yes
Applicability	Unlimited	Unlimited	Unlimited	Processor

13

Outline

- Introduction
- Built-In Soft Error Resilience (BISER)
- Circuit failure prediction
- Error Resilient System Architecture (ERSA)
- Conclusion

14

Key Takeaway

Solution: Circuit Failure Prediction

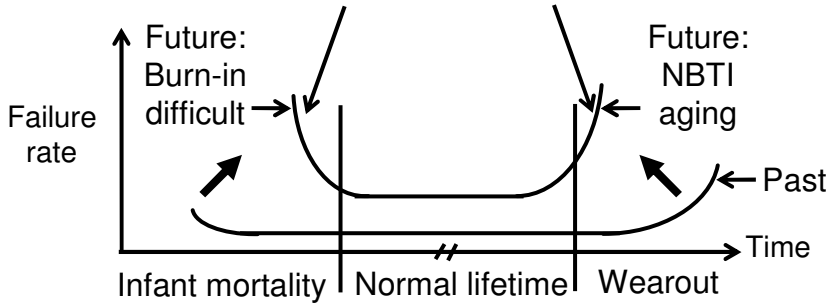

- Superior to error detection
- Practical – test chip prototype
- Effective – up to 4x aging guardband reduction
- MPSoC opportunities: On-line self-test & adaptation

15

Circuit Failure Prediction

- Predict failures – BEFORE errors appear
  - ➢ Collect data over time: on-chip sensors

☺ Applicability: Transistor aging, infant mortality

Failure Prediction	Error Detection
Before errors appear	After errors appear
+ No corrupt data & states	– Corrupt data & states
+ Self-diagnosis possible	– Limited diagnosis

"A little fire is quickly trodden out;
Which, being suffer'd, rivers cannot quench."

William Shakespeare
King Henry the Sixth,
Part III

16

Failure Prediction for Transistor Aging

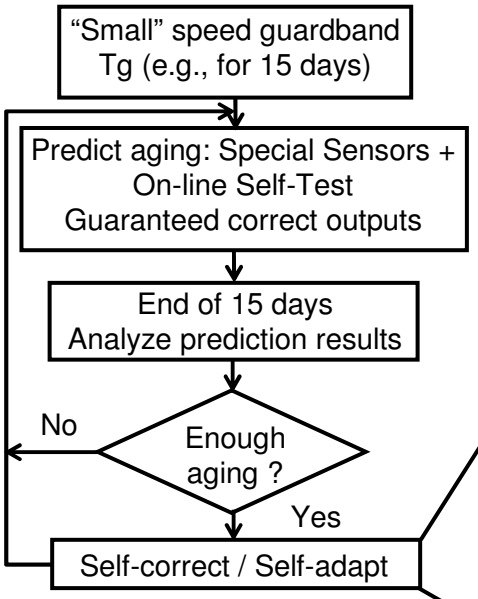

Special flip-flops with Aging Resistant Built-In Aging Sensor
OpenRISC core (65nm technology)
Ring osc. / Temp. sensors inadequate

Speed penalty	< 1%
Chip-level power penalty	< 0.4%
Die size increase	Not expected
Design flow impact	Minimal

Adjust V_t (Forward body bias ?)
Adjust V_{dd} (Runaway ?)
Adjust speed?
Core sparing

17

Aging Sensor Principle

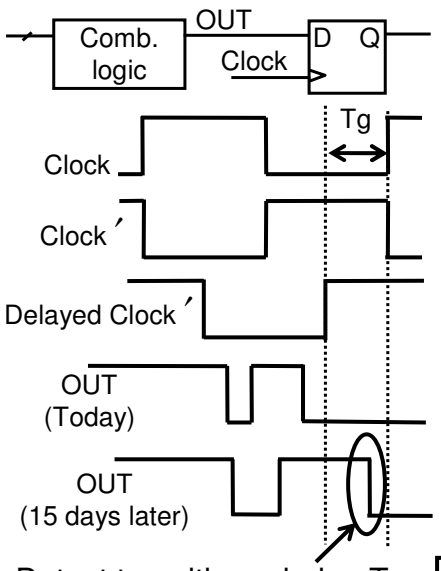

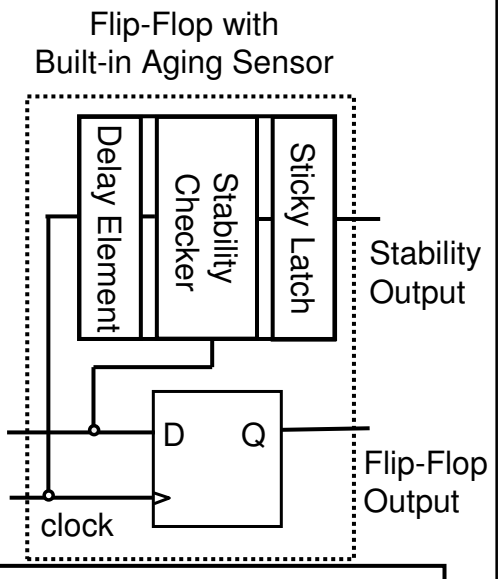

No additional hold time constraint

18

Circuit Failure Prediction Benefits

- Reduced guardband: close to best-case performance

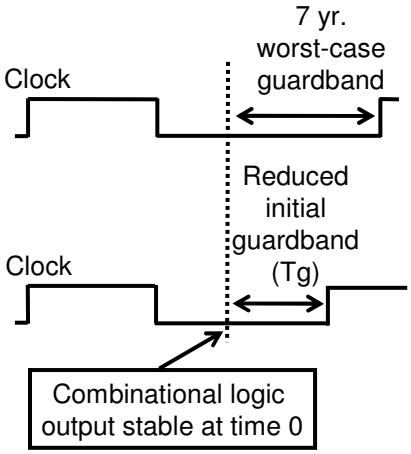

% Guardband Reduction

Clock speed	7-yr worst guardband (% clock period)	Initial guardband (Tg)	
		1 day	15 days
1 GHz	5	46	32
	10	66	52
	25	78	64
3 GHz	5	None	None
	10	26	12
	25	62	48

CORRECT OUTPUTS GUARANTEED

19

Built-In Aging Sensor: Test Chip Prototype

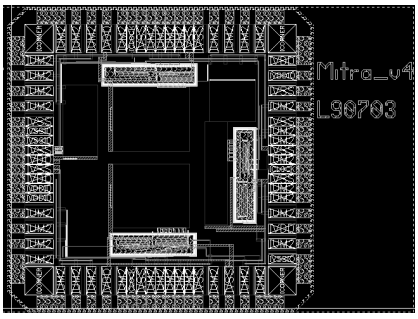

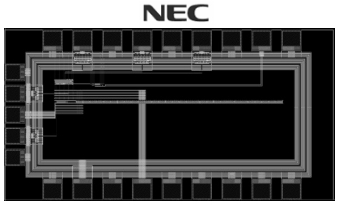

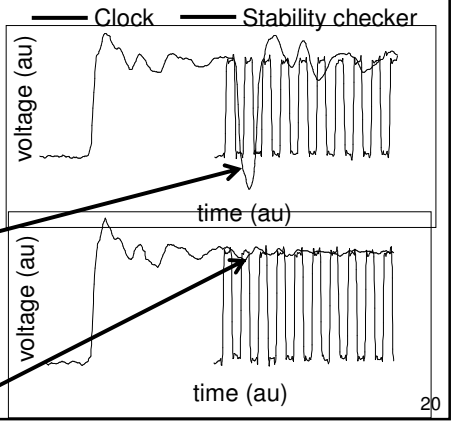

Test Chip Waveforms (Measured)

Signal transition outside Tg,
Stability checker output = 0

Signal transition inside Tg,
Stability checker output = 1

20

Outline

- Introduction
- Built-In Soft Error Resilience (BISER)
- Circuit failure prediction
- Error Resilient System Architecture (ERSA)
- Conclusion

21

Emerging Probabilistic Killer Applications

- Recognition, Mining, Synthesis (RMS)
  - ❖ Cognitive, computational genomics, vision, …
  - ❖ Large data sets
  - ❖ Highly parallel
- Core algorithms
  - ❖ Probabilistic belief propagation, K-means clustering, Bayesian networks

"Our society is creating massive amounts of complex data…There is a critical need …to **recognize**, **mine** and **synthesize** all of this digital data."

Pat Gelsinger
Technology @ Intel,
Feb. 2005

22

RMS Error Resilience Opportunities

- Algorithmic resilience – probabilistic & iterative
  - ❖ Low order bit-errors – minimal effects
    - ➢ Known for decades
- Cognitive resilience
  - ❖ "Acceptable" results OK

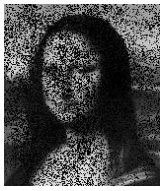

☹ BAD NEWS

RMS + unreliable H/W → Doesn't work

Control errors, High-order bit-errors

23

ERSA Hardware Prototype Results

No ERSA: RMS + unreliable H/W → Doesn't work

☹ Avg. GP + SP + BP errors to crash: 1.5 - 3.4

☹ FP errors: highly inaccurate results

RMS + ERSA at 10^{16} FITs (3,000 errors / sec.):

☺ No crashes, highly accurate results

☺ Useful throughput maintained

☺ Linear speedup

24

ERSA: Asymmetric Reliability is Key

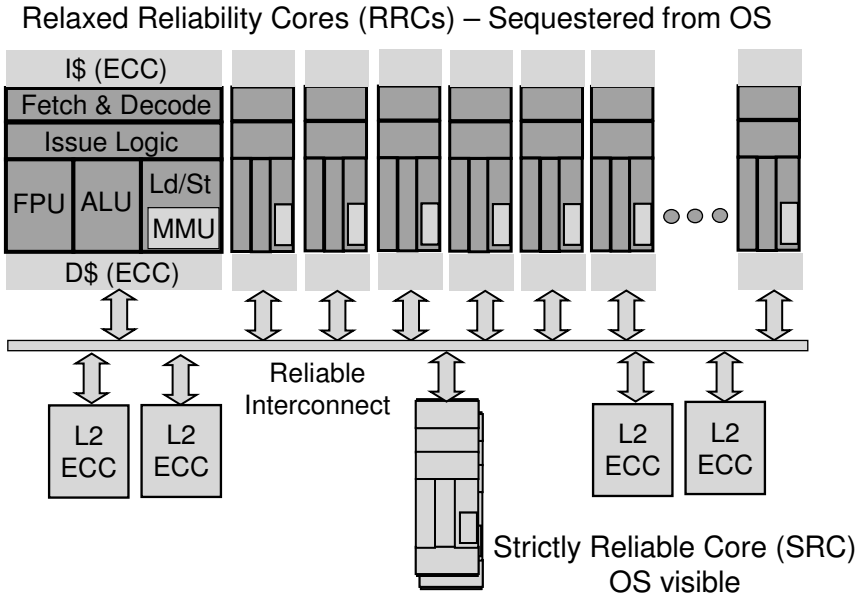

25

ERSA: Error Resilient System Architecture

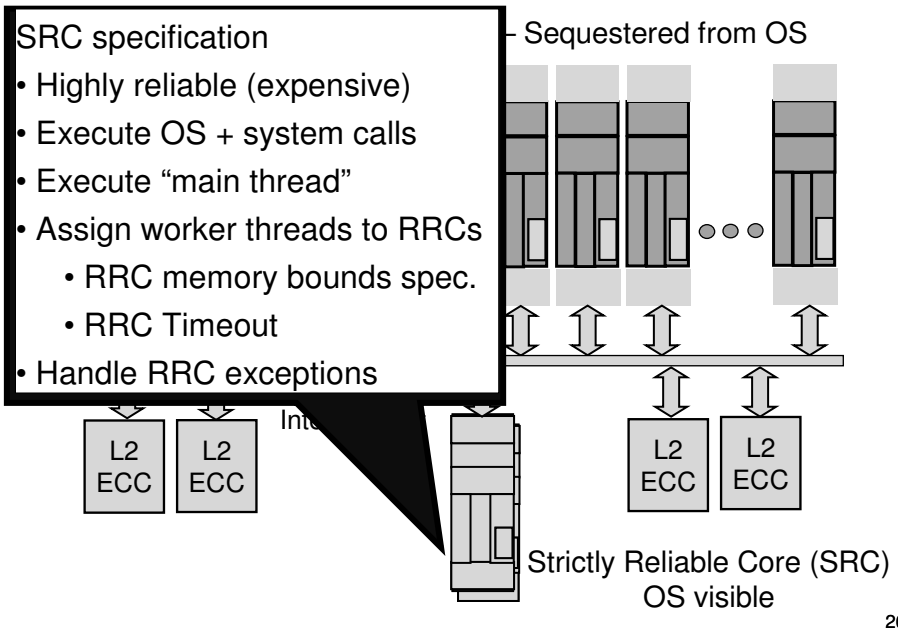

26

ERSA: Error Resilient System Architecture

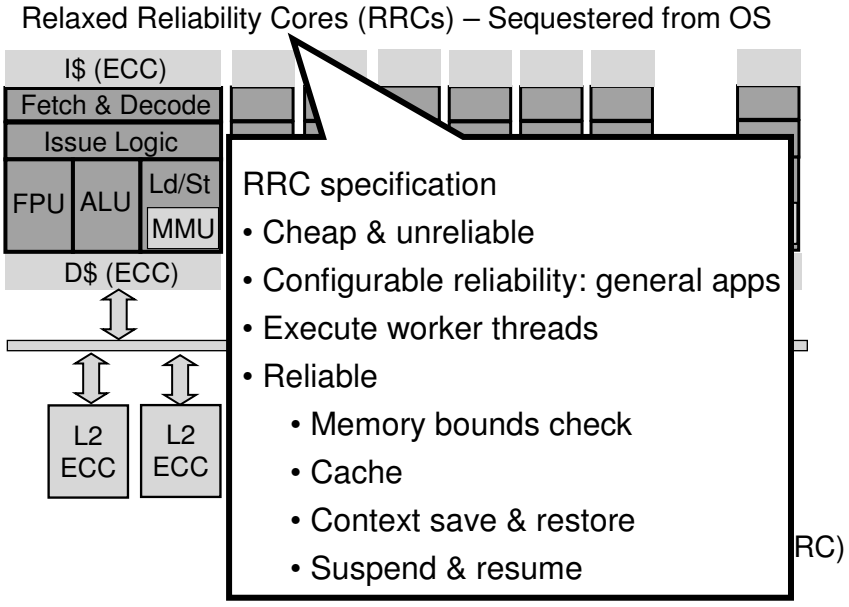

27

RMS on ERSA

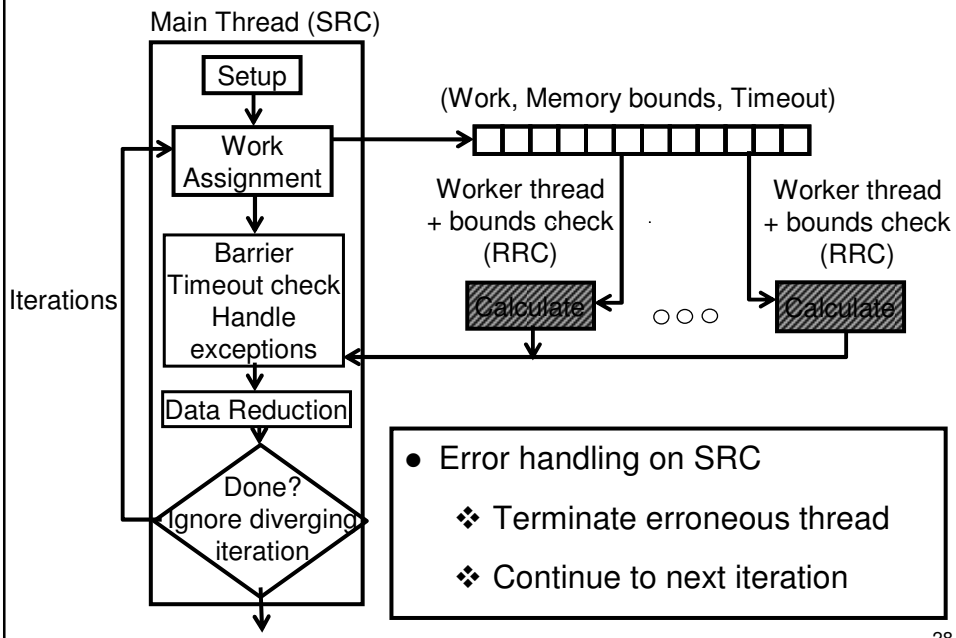

28

ERSA Prototype

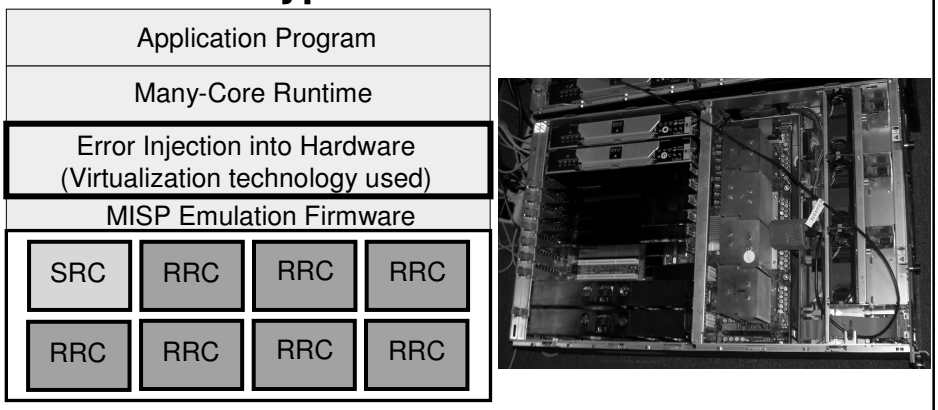

- 4-socket, 2.6 GHz, Dual-core IA-32 processors
- 2 GB main memory
- Multiple Instruction Stream Processor (MISP) infrastructure [Hankins ISCA 06]

29

ERSA Work Throughput Results

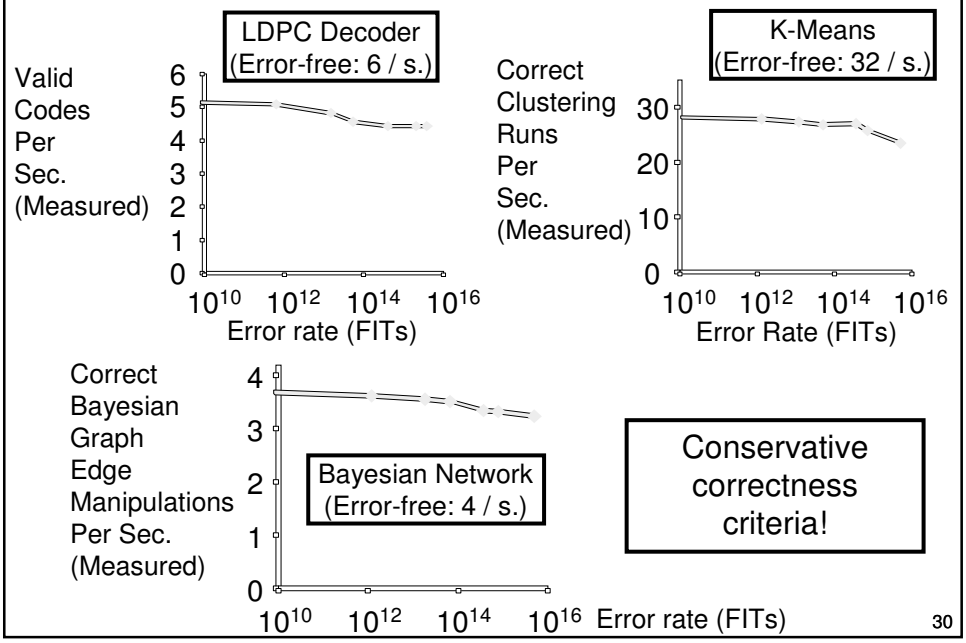

30

ERSA Results: High-order Bit Errors

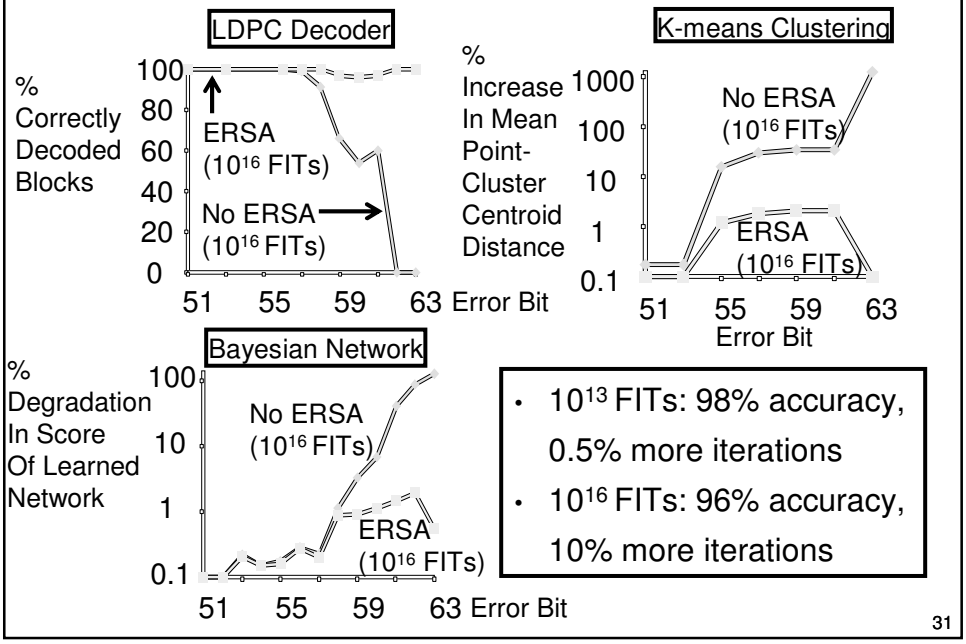

- 10^{13} FITs: 98% accuracy, 0.5% more iterations
- 10^{16} FITs: 96% accuracy, 10% more iterations

31

ERSA Speedup Results

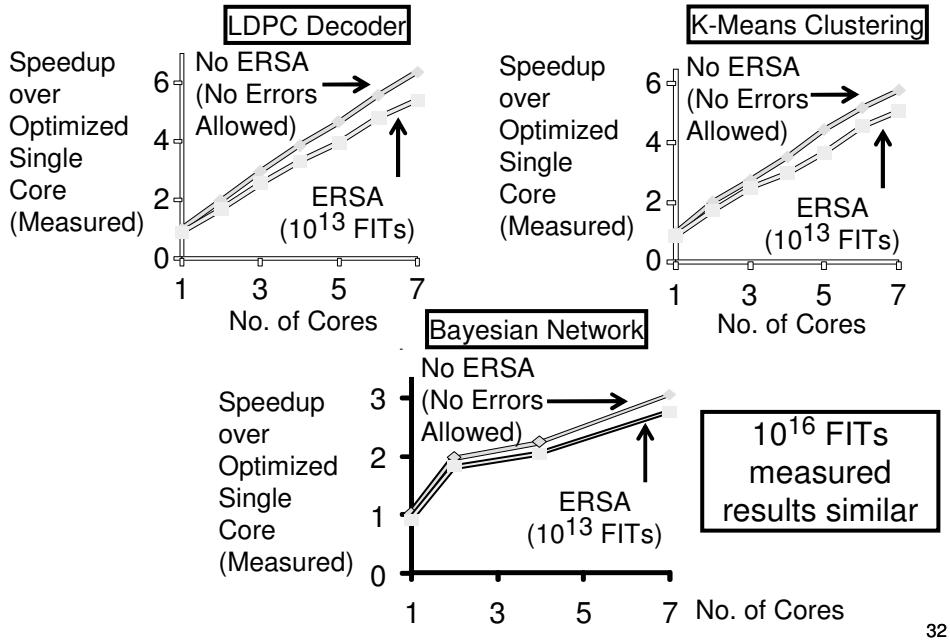

10^{16} FITs measured results similar

32

Outline

- Introduction
- Built-In Soft Error Resilience (BISER)
- Circuit failure prediction
- Error Resilient System Architecture (ERSA)
- Conclusion

33

Conclusion

- Robust system design → Global Optimization
  - ❖ Unique opportunities: New thinking required
- BISER – unique soft error properties
  - ❖ Cost-effective correction
- Circuit failure prediction – unique failure modes
  - ❖ Effective prediction possible
- ERSA – unique future killer apps
  - ❖ Resilient to extremely high error rates

34