

Robust Systems

From Today to the *N3XT 1,000X*


Subhasish Mitra



Department of EE & Department of CS
Stanford University

World Relies on Computing

Genomics



Smart Cities



Military




Intelligence Gathering


Security



Finance




Science



**Computational demands
exceed
Processing capability**

World Relies on Computing

Genomics



Smart Cities



Military



Health Care




Government




Security



Finance



Science



- Ensure robust operation
- Meet computation demands
- New application horizon

Research Topics

- Robust operation
 - Bugs, reliability, security
- Revolutionize nanosystems
 - 1,000X opportunity
- Program human brain
 - SNI Big Ideas in Neuroscience initiative

Outline

- Robust operation: silicon CMOS reliability
- Beyond silicon
- Conclusion

Silicon CMOS Reliability Challenges

- Radiation-induced soft errors
 - **Fatal** flip-flop errors
- Early-life failures (ELF)
 - Burn-in: difficult, expensive
- Variations: V_{dd} , thermal, circuit aging
 - Worst-case guardbands expensive

Definitions

- **Malfunction** (often referred to as **failure**)
 - Deviation from specified behavior
 - Underlying cause: **failure**
- **Error**: incorrect signal value
- **Fault model**
 - (Logic) representation of effect of failure

System Output Response to Failure

- **Error on output:** non-critical apps. (e.g., games ?)
- **Fault-secure:** correct outputs or error indication
 - Retry adequate (e.g., banks)
- **Fault masked:** correct outputs
 - Fault in specified class (e.g., spacecraft)
- **Fail safe:** correct or “safe” outputs

Definitions

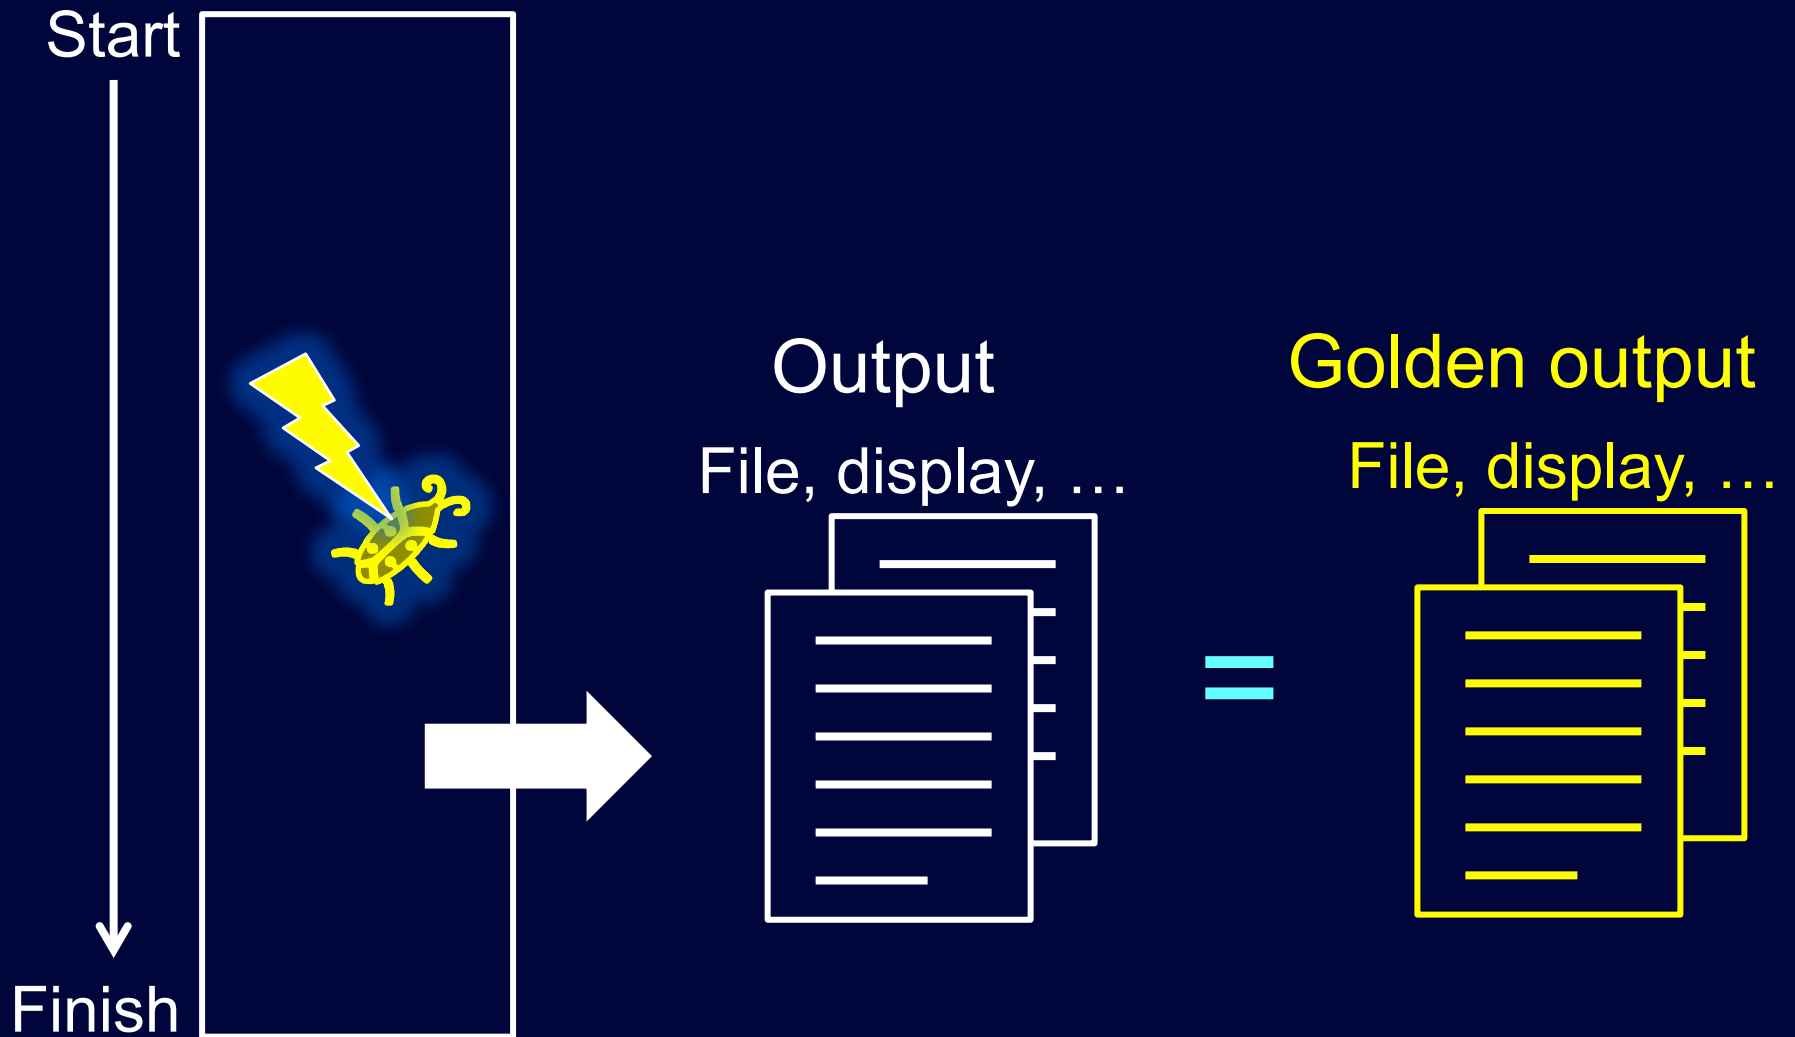
- Reliability: $R(t)$
 - Probability system works correctly up to time t
- Exponential model
 - $R(t) = e^{-\lambda t}$, λ = failure rate
- Mean Time to Failure (MTTF)

$$MTTF = \int_0^{\infty} t \times R(t) dt = \frac{1}{\lambda}$$

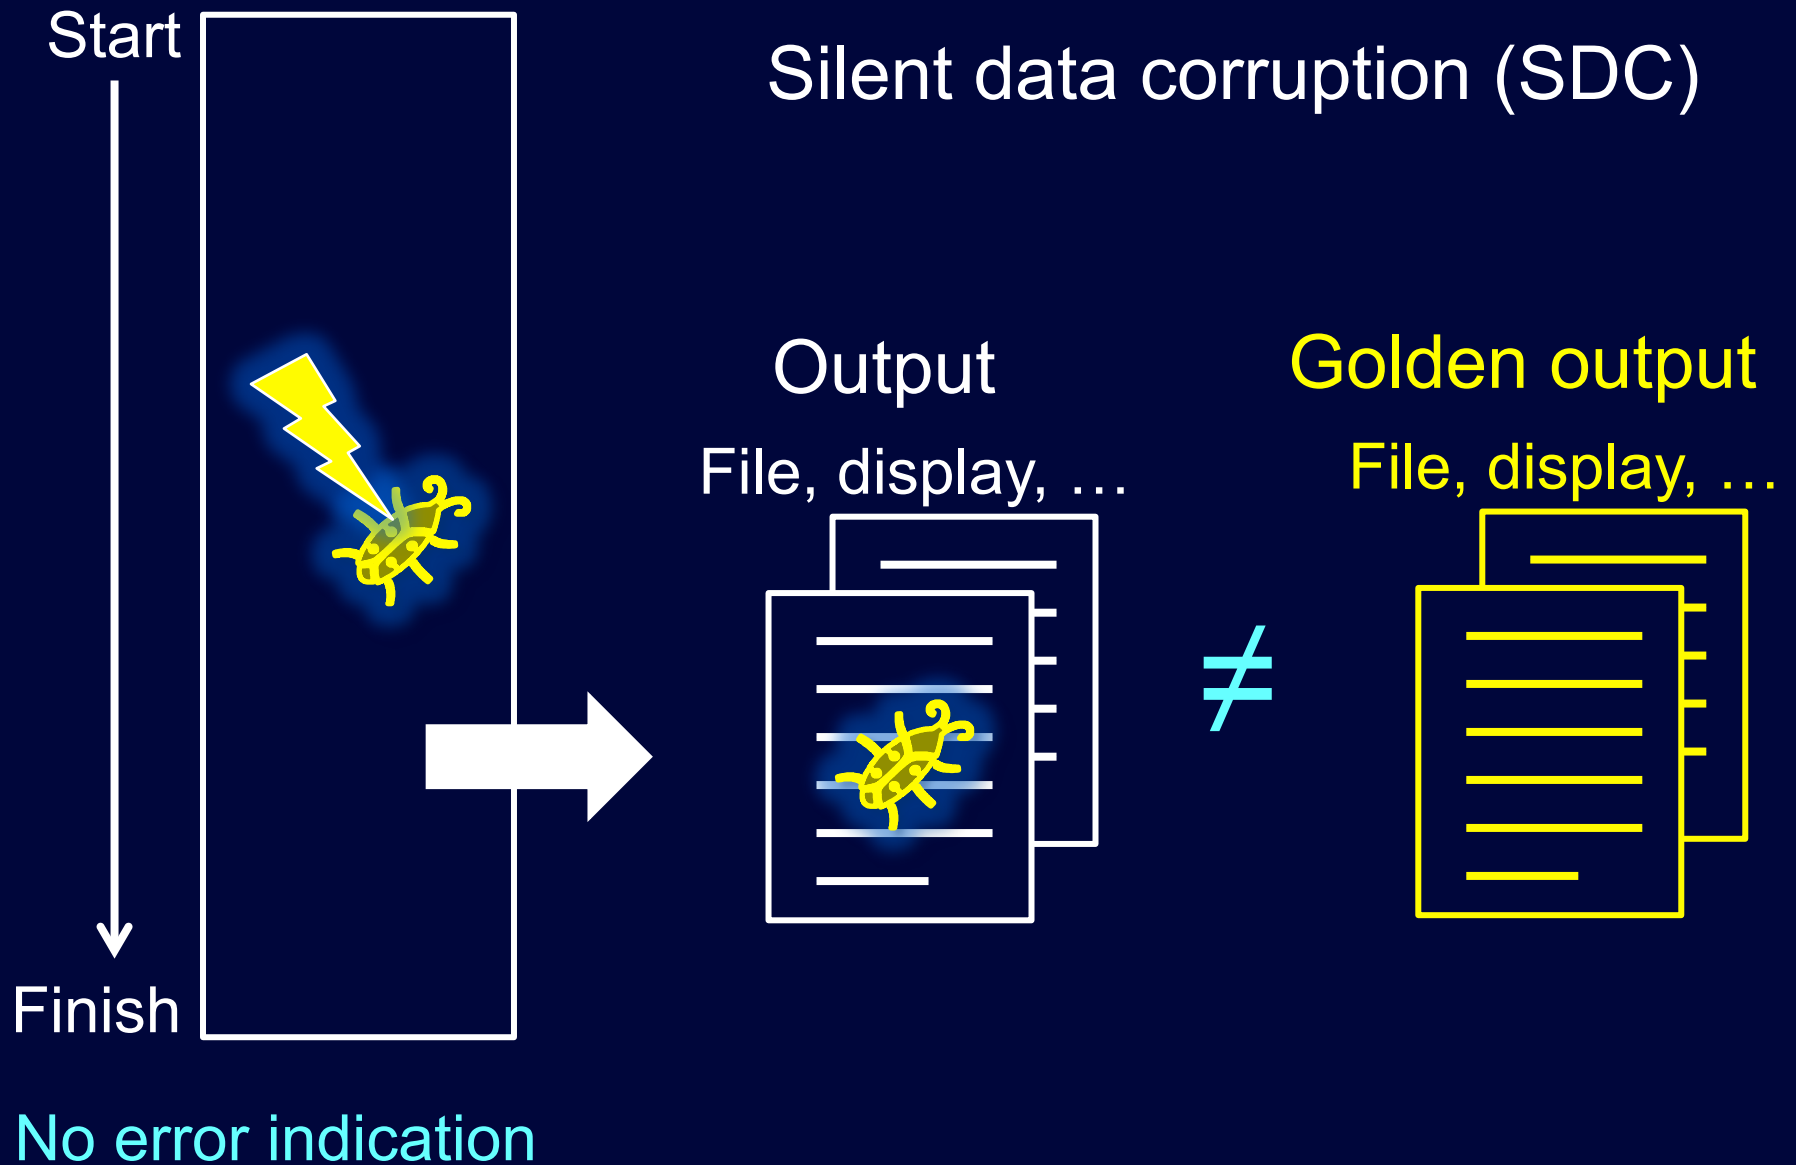
Definitions

- Availability: $A(t)$
 - Probability system works correctly **AT** time t
- Assume: system repaired after failure
 - Mean Time to Repair (**MTTR**)
- Steady-state availability:
$$\frac{MTTF}{MTTF + MTTR}$$
- How to improve availability ?

Error Effects: Vanished



Error Effects: Output Mismatch




Silent Data Corruption (SDC)

Error-free




Output Mismatch



- Output file incorrect
- No error indication

Error Effects: Unexpected Termination




e.g.,

- Divide-by-zero
- Memory access violation
- Application-detected errors
- ...

Hang


Start



$> 2 \times$ error-free execution time

Does not finish / terminate

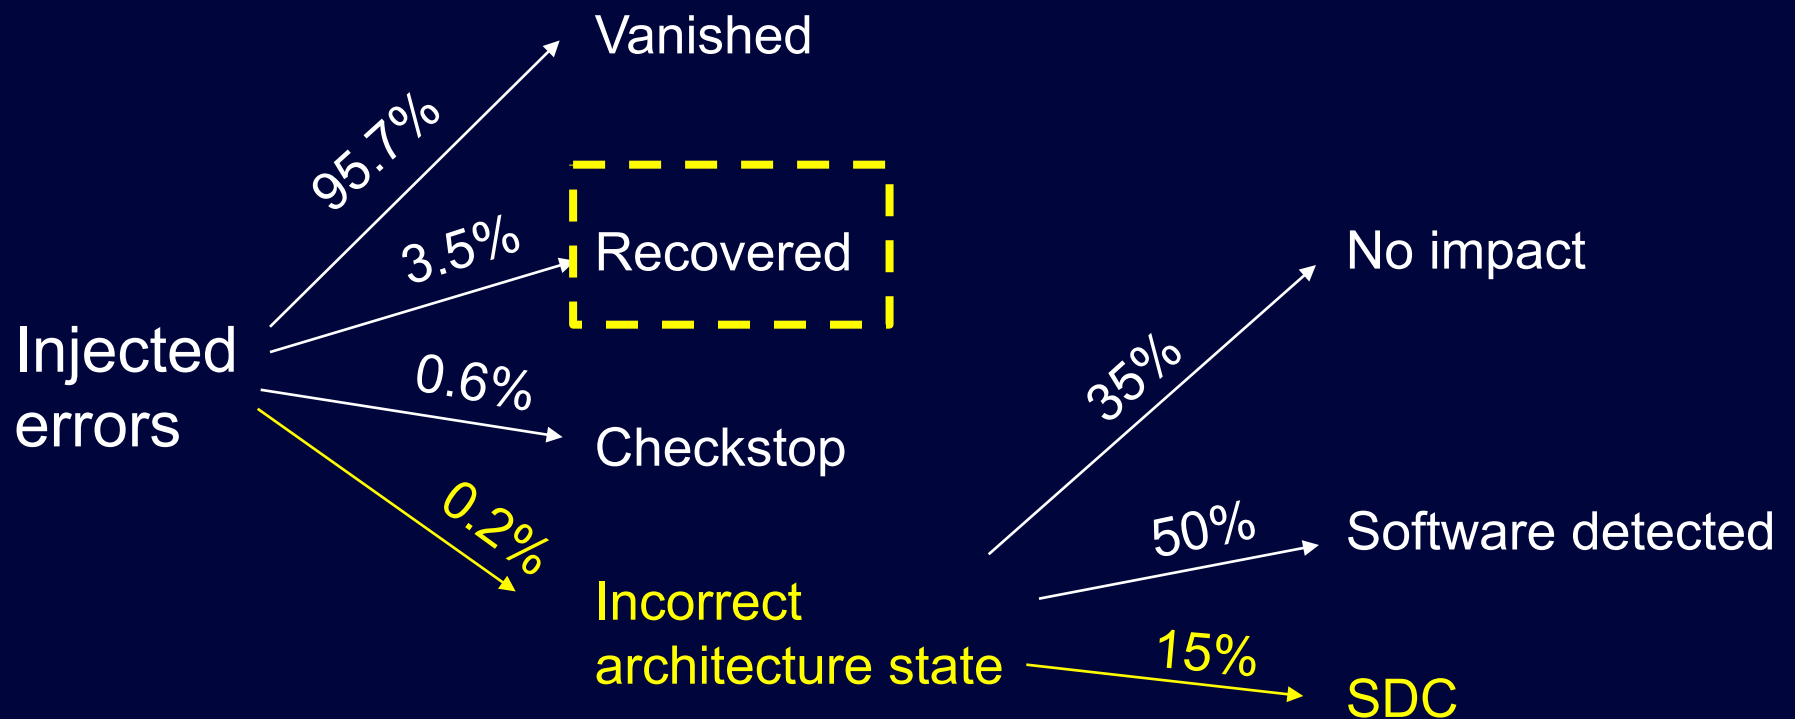
Soft Error Effects



All injected flip-flop errors

Detected but Uncorrected
Errors (DUE)

Soft Error Effects: BZip2 on IBM Power6



Fault-Tolerance: Rich Literature

Expensive

Early pioneers



Early systems



How Low Cost ?

Applications

Approach

No



Low-cost detection / correction

BISER, LEAP

Circuit failure prediction

Output errors



Yes

Light-weight correction

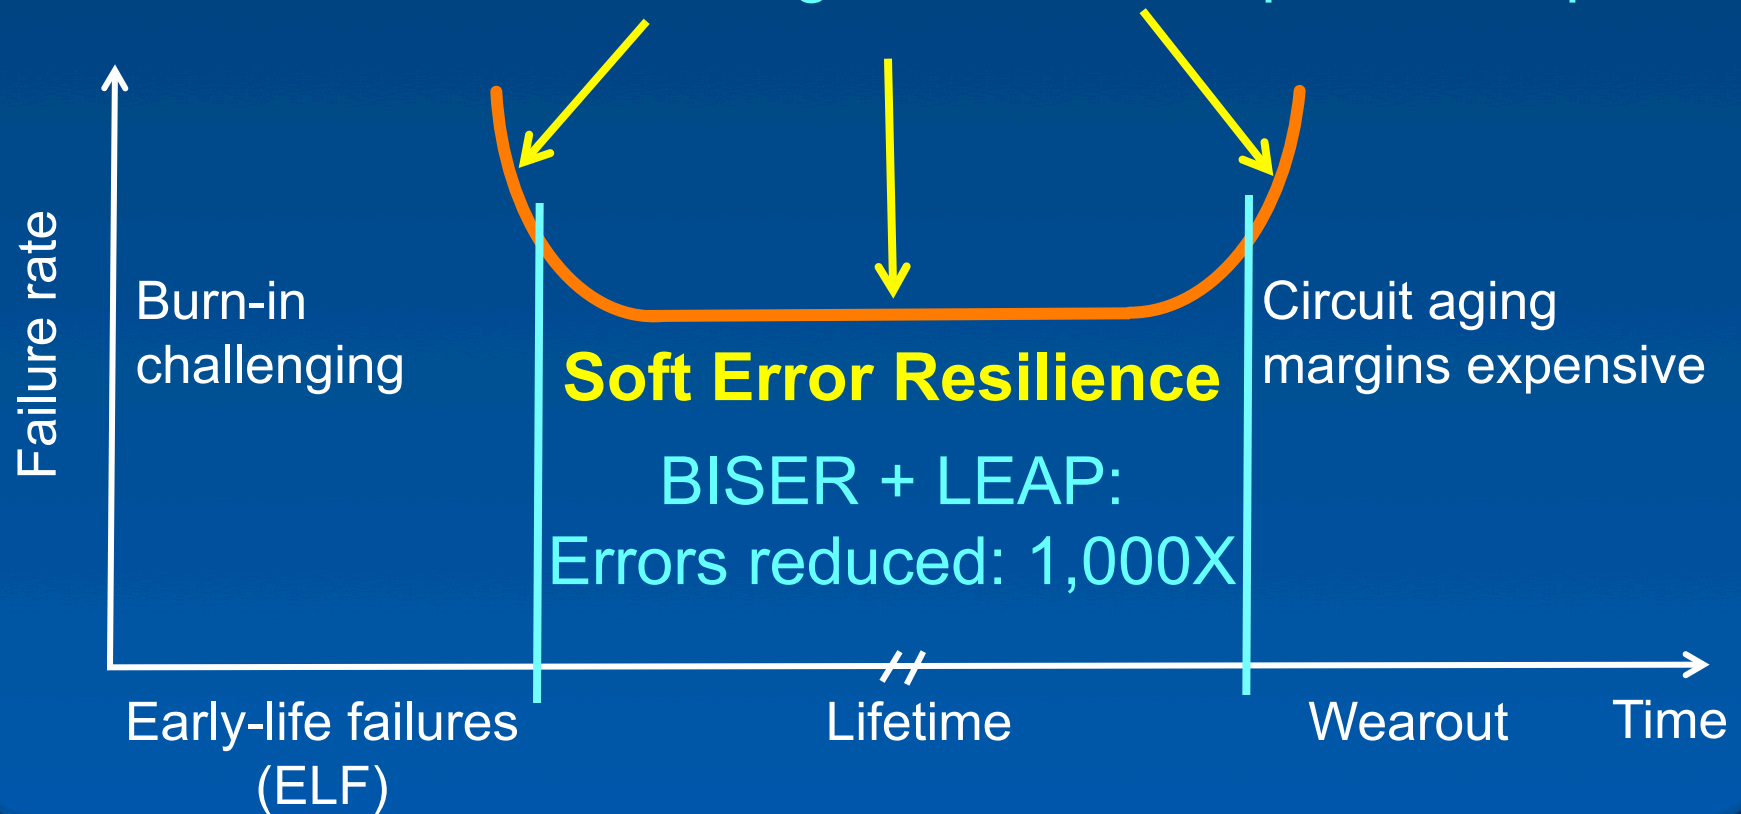
Error Resilient System Arch. (ERSA)

[Cho IEEE TCAD 12]

Low-Cost Techniques

Circuit Failure Prediction

CASP on-line self-test, diagnostics, self-repair & adaptation



How Low Cost?

Solution: cross-layer resilience?

“multiple error resilience techniques from different layers of the system stack cooperate to achieve cost-effective error resilience”

[DARPA, PERFECT BAA 12]

[Borkar Intel, IEEE Micro 05]

[Carter Intel, DATE 10]

[Pedram, NSF 12]

[Chandra ARM, DAC 14]

[Gupta IBM, IRPS 14]

[Henkel, DAC 14]

Existing Resilience Techniques

- Many point solutions

- Some cross-layer, some single-layer

- Missing

- End-to-end cross-layer resilience framework

ARGUS [Meixner, Micro 2007]

Parity checking [Spainhower, IBM Journal 1999]

DICE [Calin, Trans. Nuclear Science 1996]

BCDMR [Furuta, VLSI Circuits 2010]

EDDI [Oh, Trans. Reliability 2002]

Bose-Lin code [Bose, IEEE Trans. Computers 1985]

Constructive reliability [Chen, IEEE Trans. Computers 1982]

Reliability-driven transforms [Rehman, CADICS 2014]

Relyzer [Hari, DSN 2012]

Berger code [Berger, Information and Control 1961]

Fault screening [Racunas, HPCA 2007]

Delay [de Knijff, ISCA 2009]

Synclon [Szefran, IEEE Micro 2013]

Reliability [Ando, JSSC 2003]

DIVA [Austin, Micro 1999]

z990 [Meaney, T-DMR 2005]

Resilience Actuators [Kleeberger, IEEE Micro 2013]

HLS-driven [Campbell, DAC 2015]

Multi-Layer [Henkel, DAC 2014]

Scalable Hardware [Chippa, DAC 2010]

ABFT for matrix [Huang, IEEE Trans. Computers 1984]

LEAP [Lee, IRPS 2010]

ANT [Hedge, ISLPED 1999]

ABFT HPC [Bosilca, Journal Parallel and Distributed Computing 2008]

Shoestring [Feng, Comp. Arch. News 2010]

RSE [Nakka, DSN 2004]

Assertions [Sahoo, DSN, 2008]

Error detectors [Pattabiraman, DSC 2007]

RCC [Seifert, IRPS 2010]

SSNoC [Varatkar, ISSCC 2007]

ABFT arrays [Nair, IEEE Trans. Computers 1990]

Scalable cross-layer resilience [Schoen, ISCA 2012]

CFCSS [On, Trans. Reliability 2002]

Razor [Blaauw, ISSCC 2008]

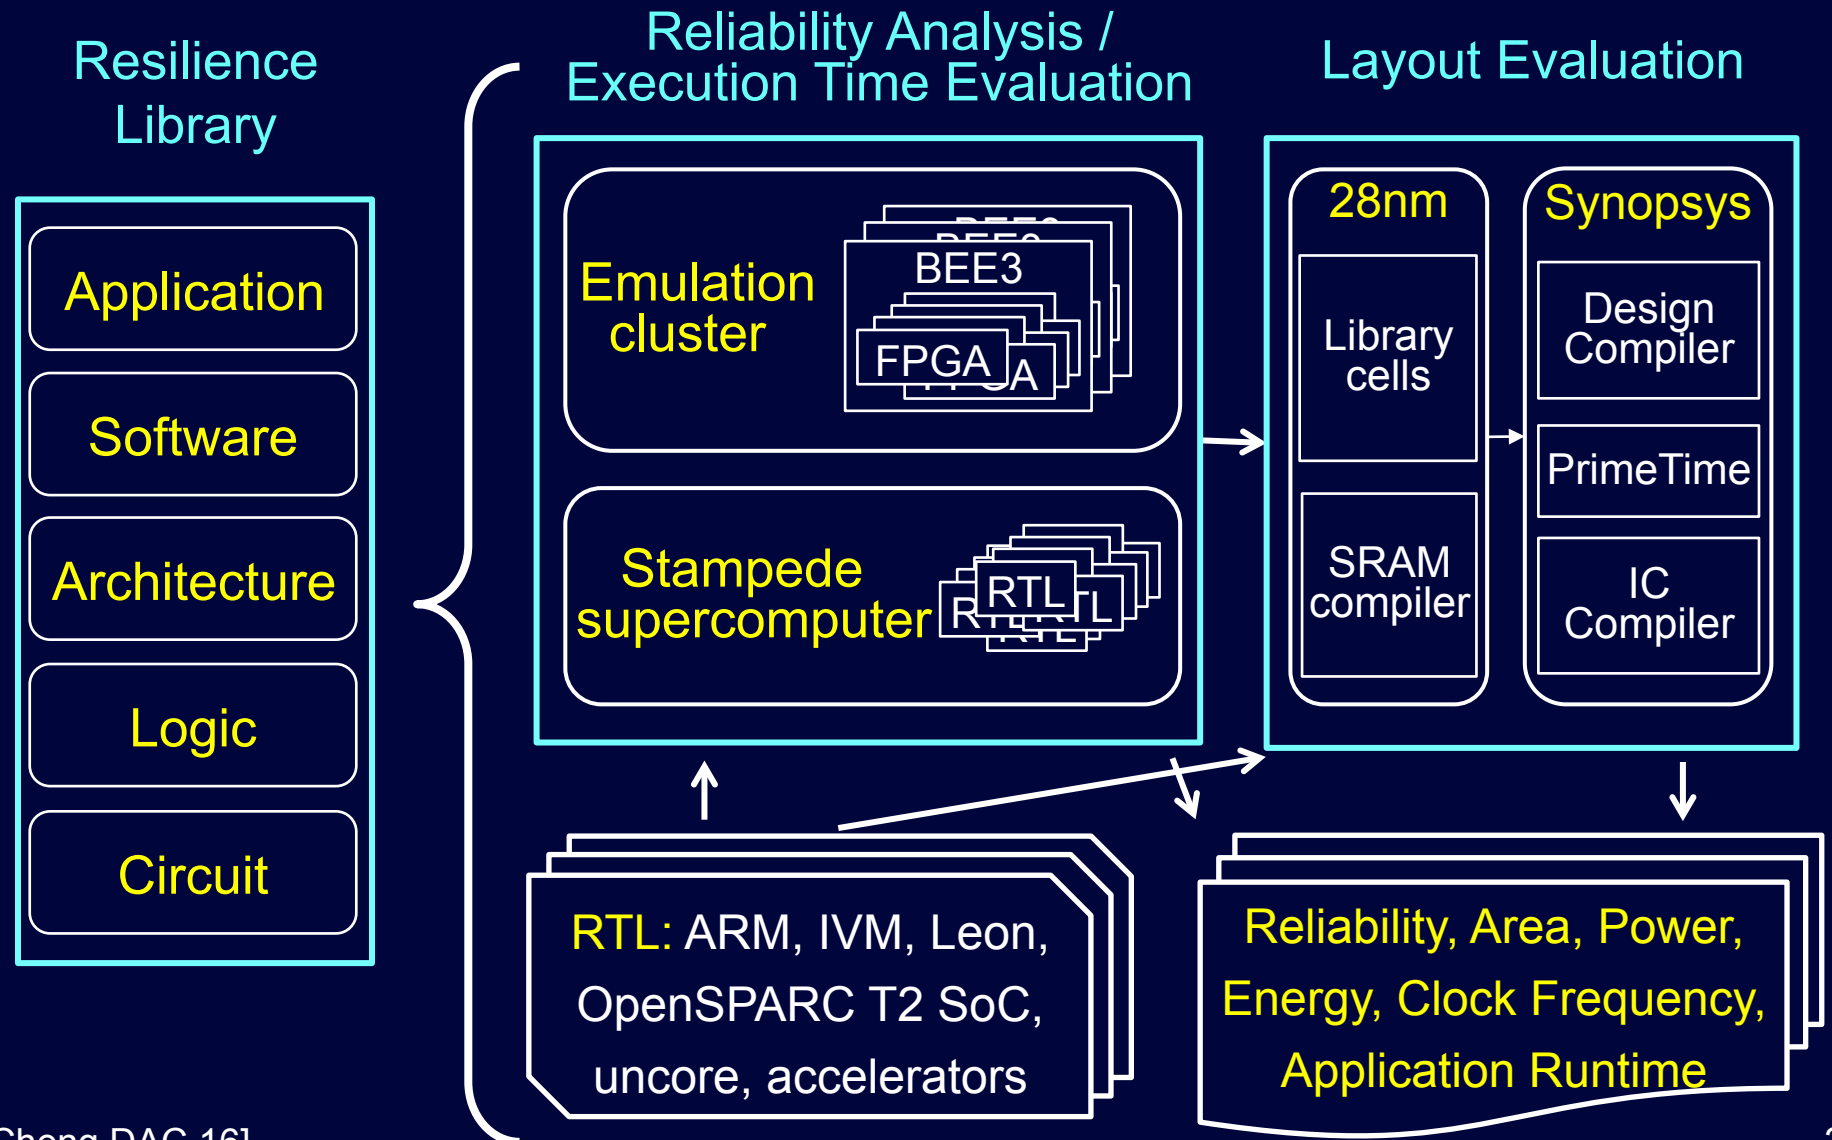
BISER [Mitra, IEEE Computer 2005]

SWIFT [Reis, ISCG 2005]

RMT [Mukherjee, ISCA 2002]

CLEAR

Cross-Layer Exploration for Architecting Resilience



Today's Focus

- Radiation-induced soft errors in flip-flops
 - Single-Event Upsets (SEUs)
 - Single-Event Multiple Upsets (SEMUs)
- Combinational logic soft errors not critical

CLEAR: Extensive Study

- Designs: wide variety
 - ARM, LEON3, Alpha, OpenSPARC multi-core SoC, accelerators
- Thorough flip-flop error injections
 - FPGA clusters, Stampede supercomputer (522,080 cores)
 - Full workloads (SPEC, PARSEC, PERFECT, proprietary)
- Detailed physical design
 - Wire routing, process / voltage / temperature corners


Many Cross-Layer Questions Answered

- Cross-layer always best?
- All cross-layer solutions equally good?
- Application constraints (e.g., soft real-time)?
- Benchmark dependence?
- Definitive guidelines for new resilience techniques

Key Message

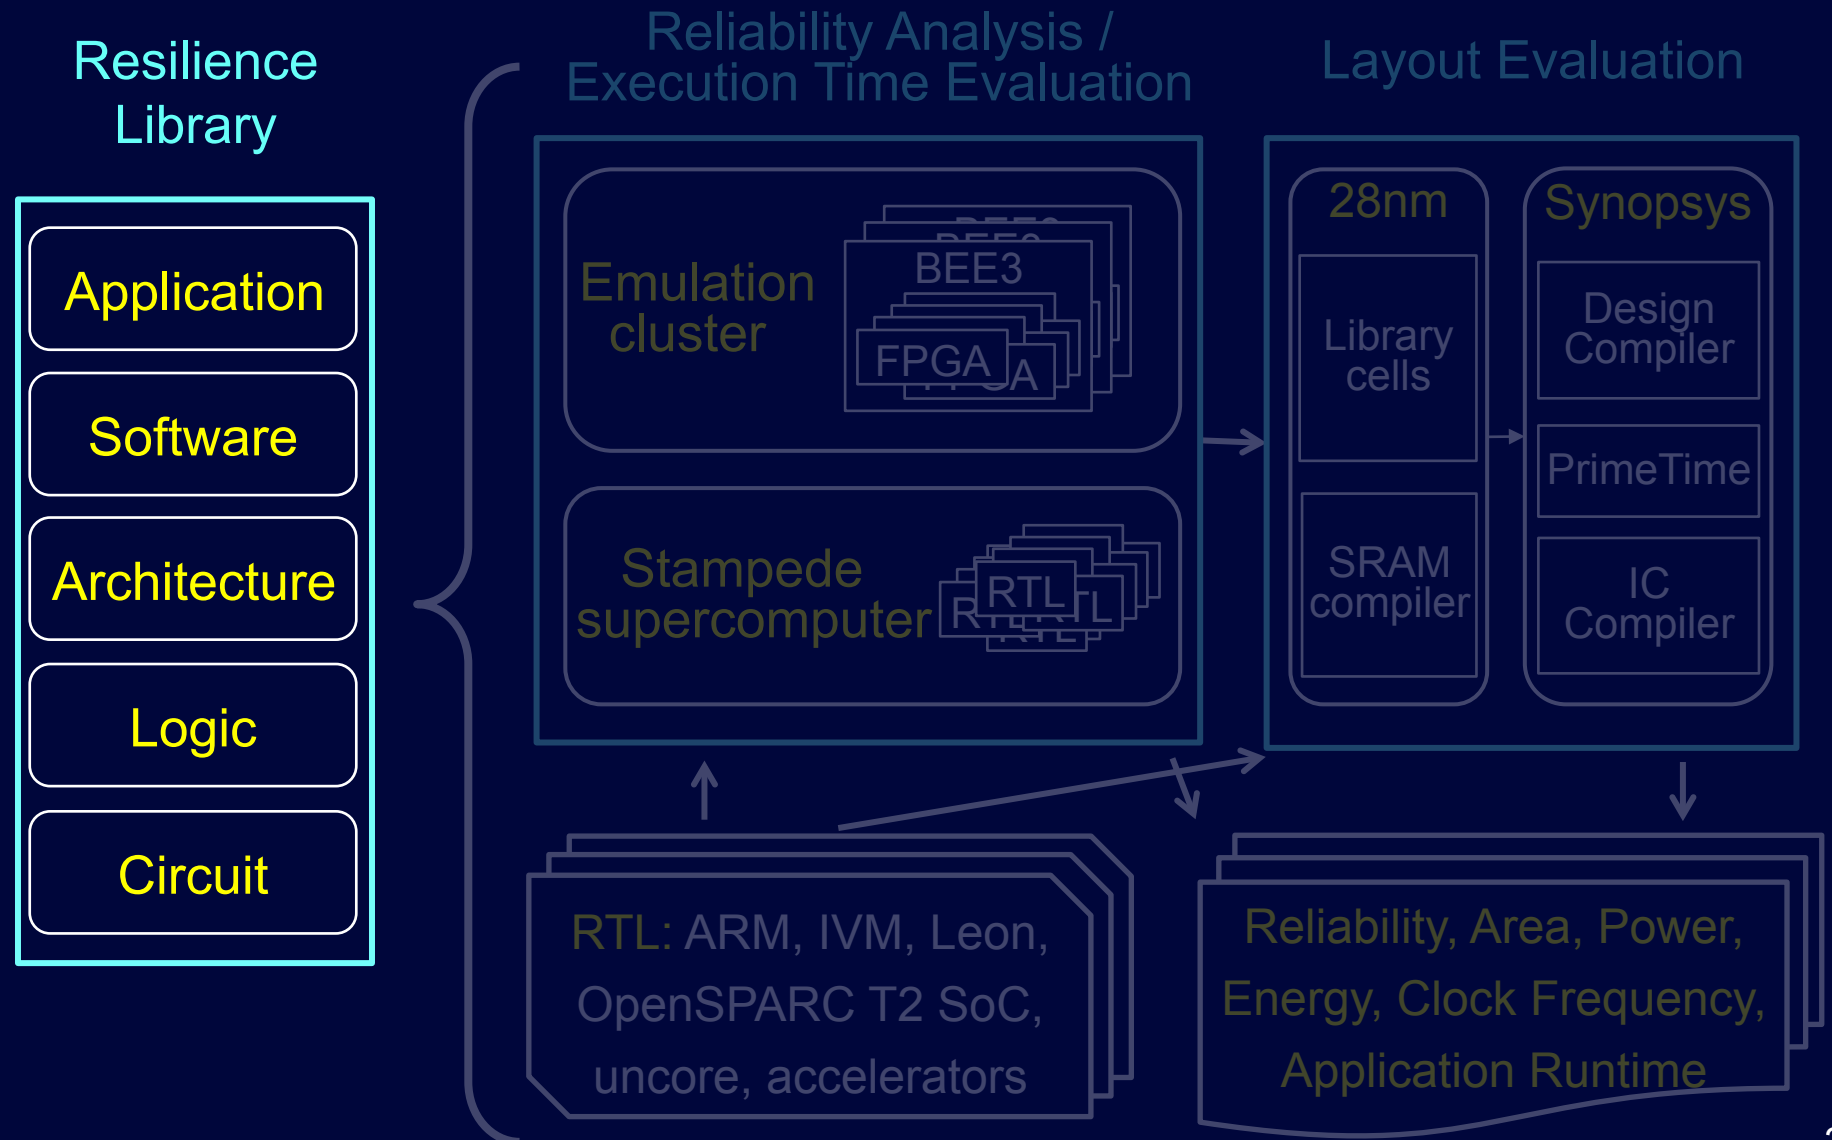
- From black art to science
- 5-50x resilience, 0.2-6% energy cost
 - Circuit + logic + micro-arch. recovery
- Circuit alone (application-guided)
 - ~1% extra energy vs. best cross-layer

Several layers,
Numerous combinations



CLEAR

Cross-Layer Exploration for Architecting Resilience



Representative Resilience Techniques

Algorithm

Algorithm Based Fault Tolerance (ABFT) Detection
Algorithmic Noise Tolerance (ANT)
Algorithm Based Fault Tolerance (ABFT) Correction

Software

Control Flow Checking by Software Signatures (CFCSS)
Error Detection by Duplicated Instructions (EDDI)
Shoestring Fault Screening
Assertions
Reliability Driven Fault Correction
Software Implemented Fault Tolerance (SWIFT)

Micro-architecture Recovery

Reorder Buffer (RoB) recovery
Extended Instruction Replay (EIR) recovery

Logic

Berger Code
Residue Code
Bose-Lin Code

Circuit

Razor
Built-In Soft Error Resilience (BISER)
Layout design through Error Aware transistor Positioning (LEAP)
Dual Interlocked Cell (DICE)
Bistable Cross-coupled Dual Modular Redundancy (BCDMR)
Error Detection Sequential (EDS)
Reinforced Charge Collection (RCC)

10 error detection /
correction techniques +
4 recovery techniques

798 combinations

1. Algorithm Based Fault Tolerance (ABFT) Correction

2. ABFT Detection

3. Software Assertions

4. Control Flow Checking by Software Signatures (CFCSS)

5. Error Detection by Duplicated Instructions (EDDI)

6. Data Flow Checking (DFC)

7. Monitor Cores

8. Logic Parity

Micro-arch. Recovery

1. Flush

2. Reorder Buffer (RoB)

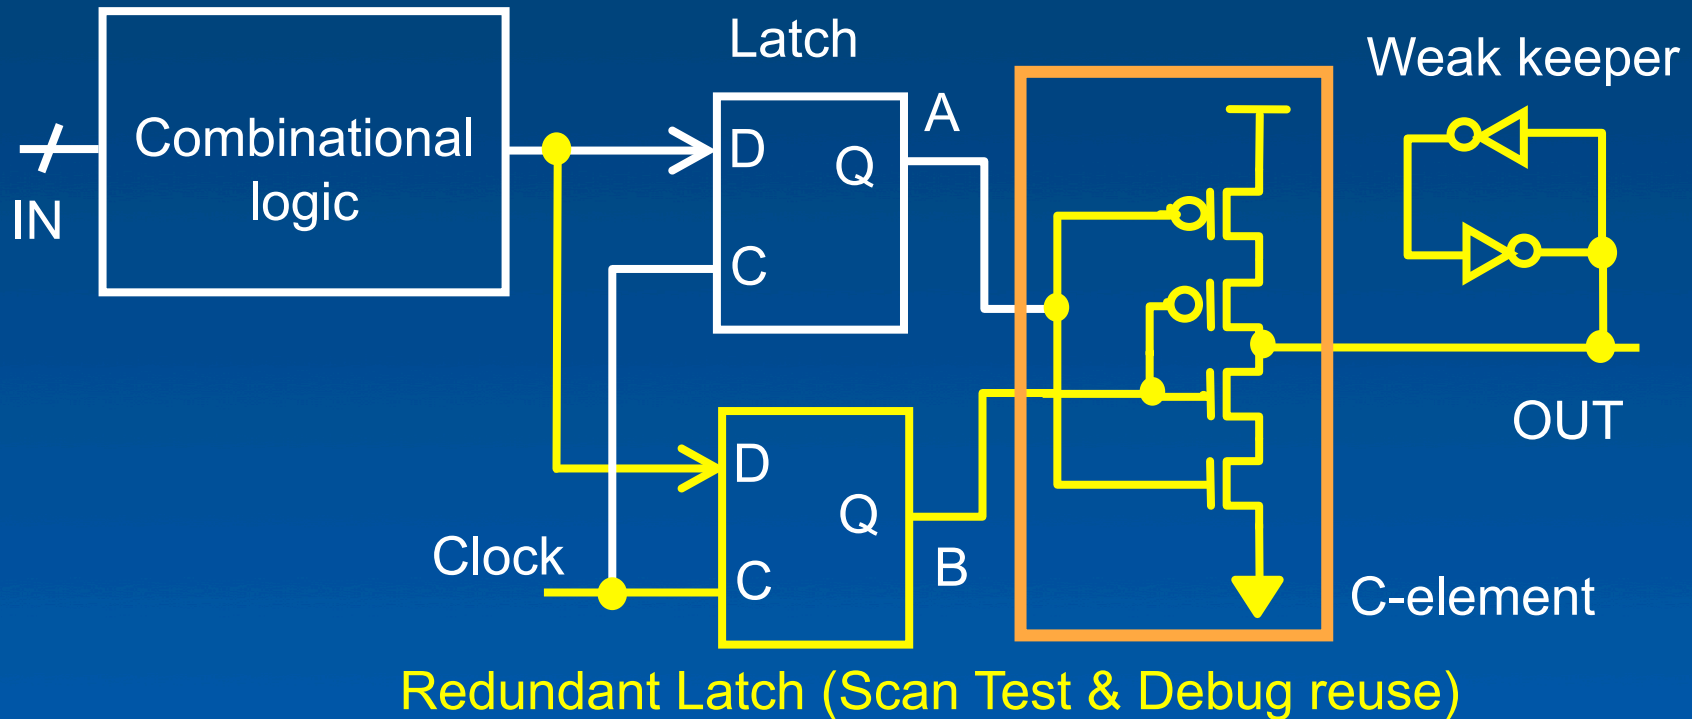
3. Instruction Replay (IR)

4. Extended IR (EIR)

9. Layout design through Error-Aware transistor Positioning (LEAP)

10. Error Detection Sequential (EDS)

BISER: Built-In Soft Error Resilience

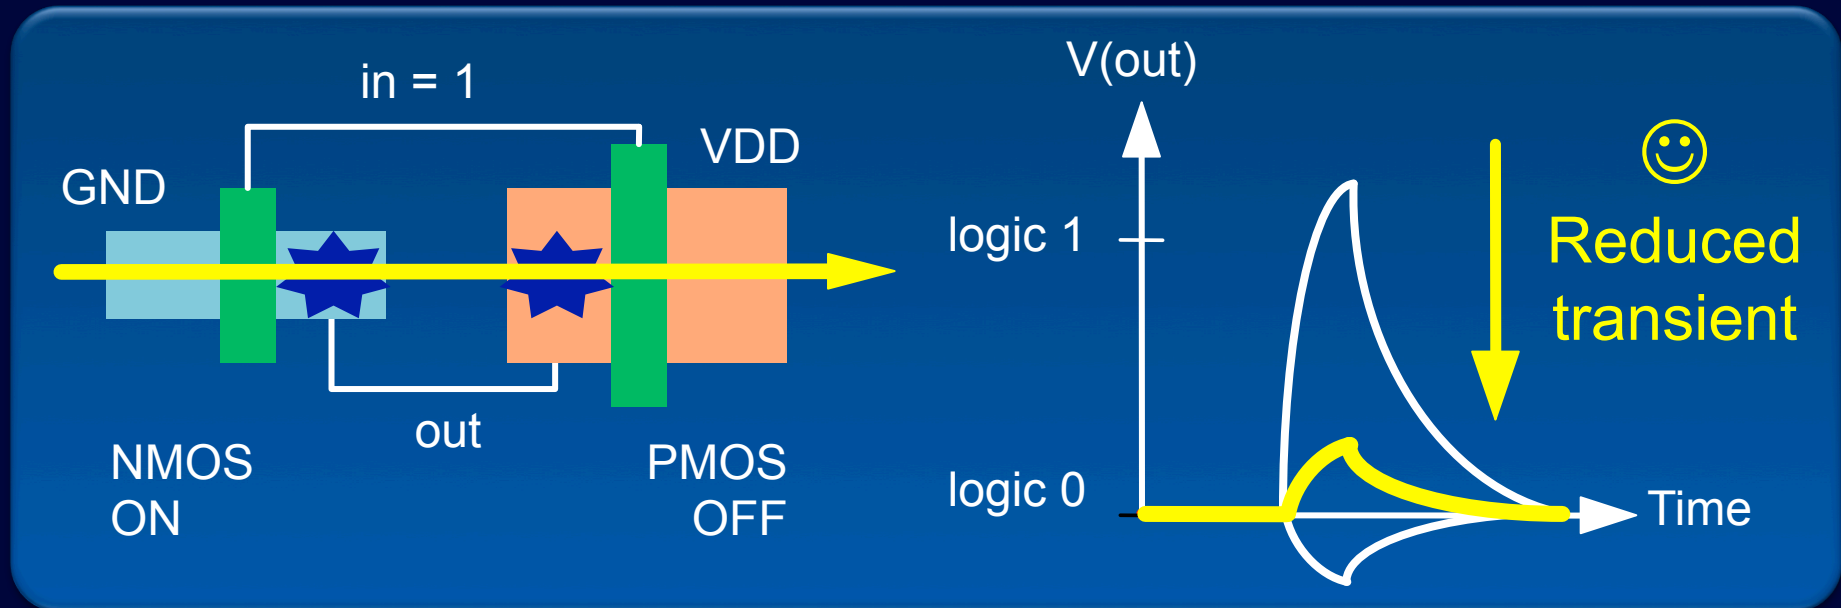


45nm: up to 1,000X benefits

Single Error Assumption Inadequate

- Single-event **multiple** upsets (SEMUs)

LEAP: Layout by Error Aware transistor Positioning



Errors **Corrected**: SEUs and SEMUs

Extensive LEAP Characterization

- Radiation beam experiments
 - 40nm, 28nm, 20nm, 14nm
 - Bulk, SOI
- VDD: nominal, near-threshold

Flip-flop	Soft Error Rate (SER)	Area	Power	Delay	Energy
Baseline	1	1	1	1	1
LEAP-DICE	2×10^{-4}	2	1.8	1	1.8


Memory ECC and SEMUs

- Don't implement multiple error correction blindly
- Multiple **physically adjacent** errors

Memory ECC and SEMUs

Option 1

- Memory interleaving
 - 2 physically adjacent errors
 - Single errors in 2 **separate** words
- ☹ Cost, difficult for smaller geometries



Option 2

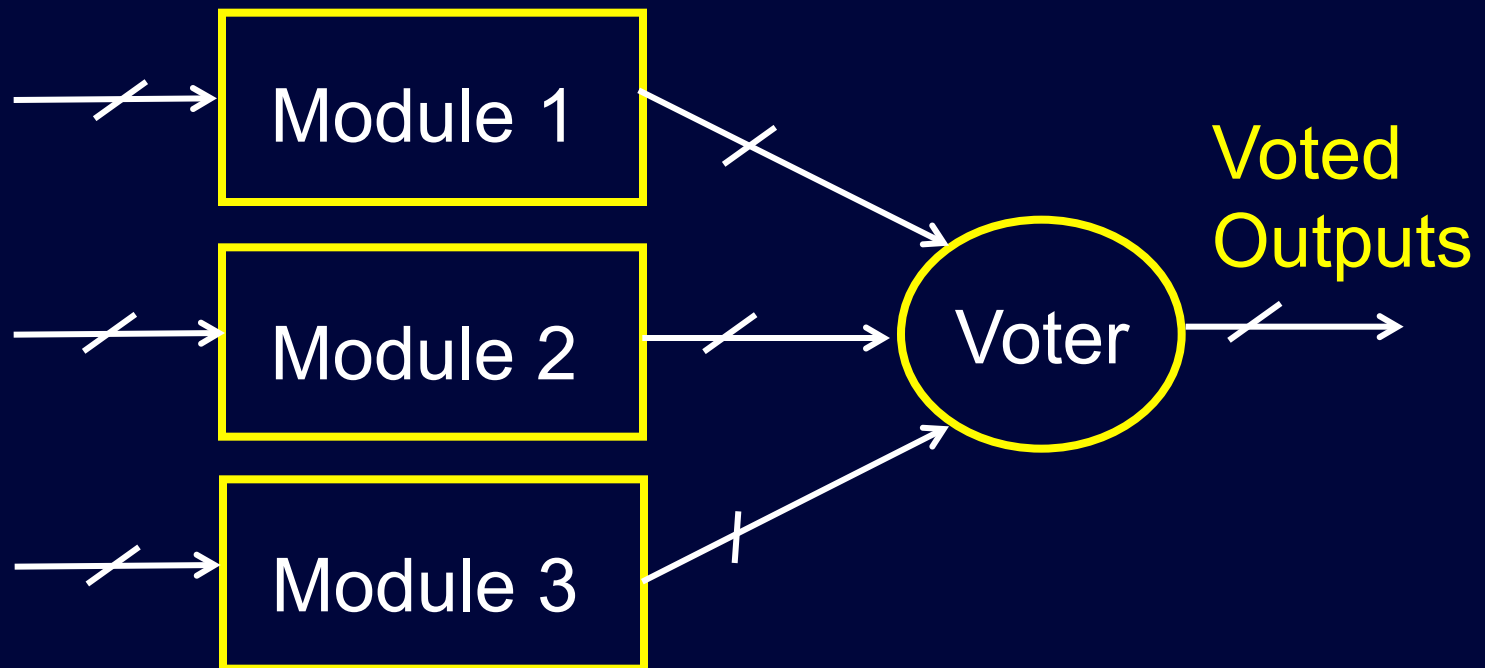
- **Adjacent bit** error correction [Dutta ITC 07]

Memory ECC Challenges

- Performance overhead
 - Pipelining
 - Additional latency, verification effort
 - Detection followed by correction
 - Variable latency, verification effort
- Small distributed memories

Error Masking

- No error on outputs
 - **T**riple **M**odular **R**edundancy (TMR)

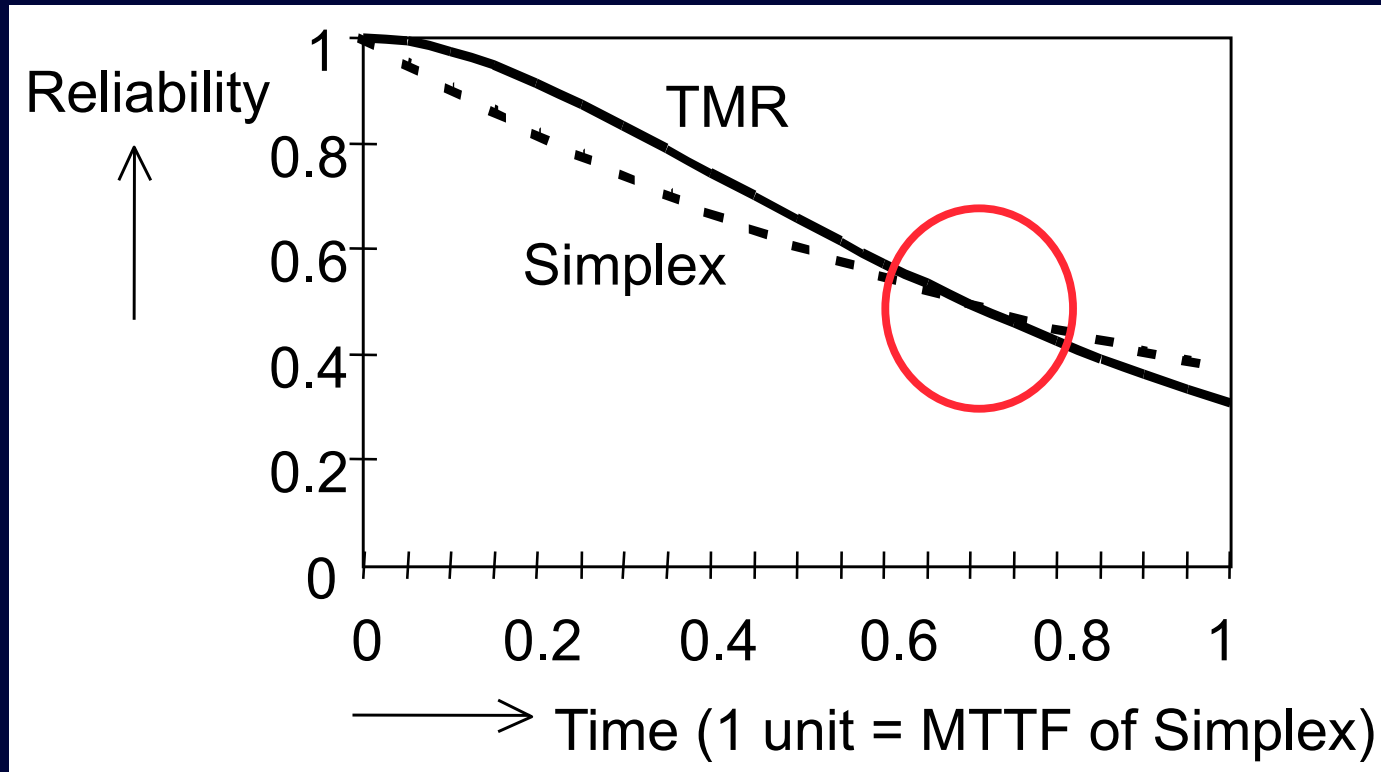


TMR Reliability

$$R_{TMR} = R_{voter} \times [R_m^3 + R_m^2 \times (1 - R_m)]$$

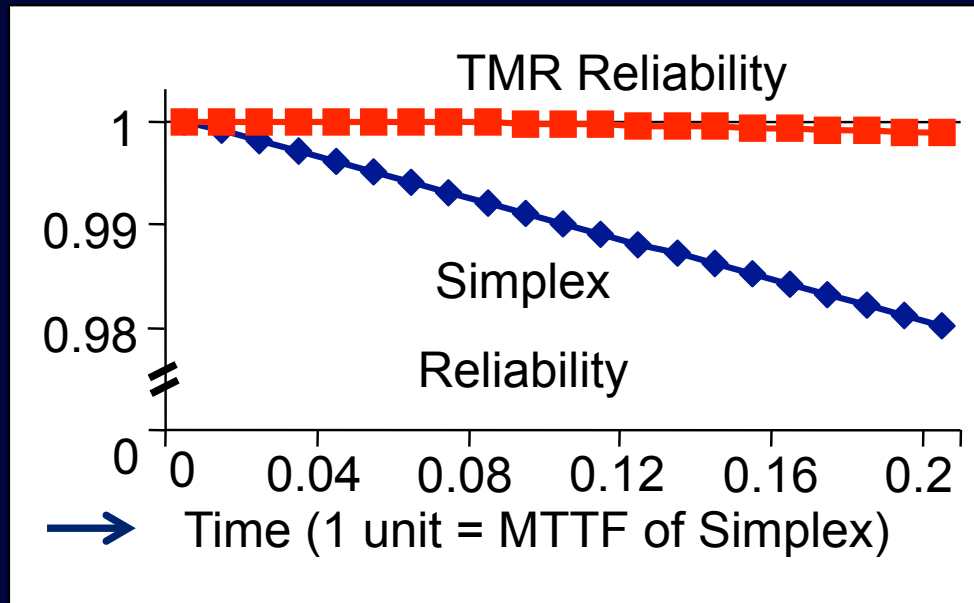
- R_m : individual module reliability
- Pessimistic: non-overlapping errors
- Optimistic: correlated / common-mode failures
- TMR MTTF < Simplex MTTF

TMR Reliability



- TMR reliability = simplex reliability
 - $\text{Time} = \log_e 2 \times \text{Simplex MTTF (perfect voter)}$

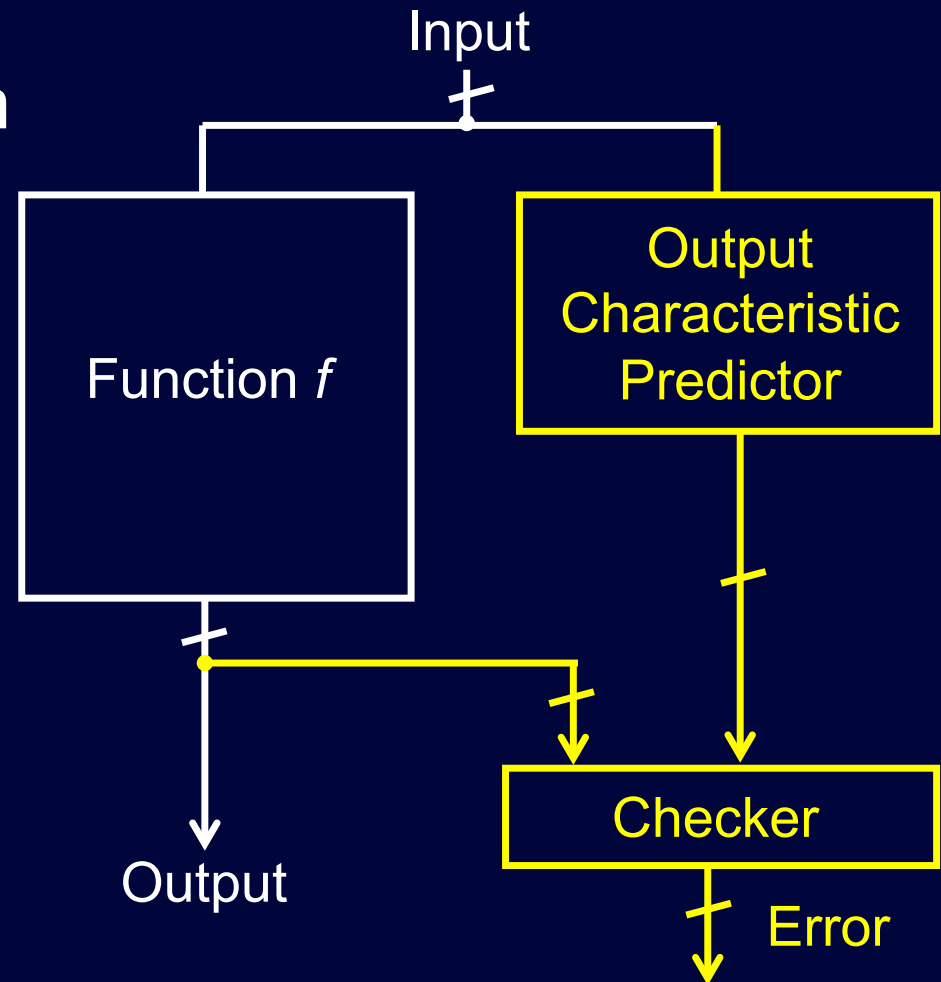
TMR Reliability vs. Mission Time



- TMR effective for “short” mission times
- Other options: TMR-Simplex, TMR + Duplex-Repair

Concurrent Error Detection (CED)

- Normal system operation
- Preserve **data integrity**
 - Correct outputs or
 - Error indicated
 - Incorrect outputs
 - aka **fault-secure**



Output “Characteristics”


- Output itself: duplication
 - Major challenges if not “fine-grained”
- Output parity
- Output residue
- 1s or 0s count in output word
- Many others (extensive literature)
- Self-checking checkers

Processor Duplication Challenges

- Synchronization !!
- False DUEs when out of sync
 - e.g., error correction event in one processor
 - Mismatch when output pins compared

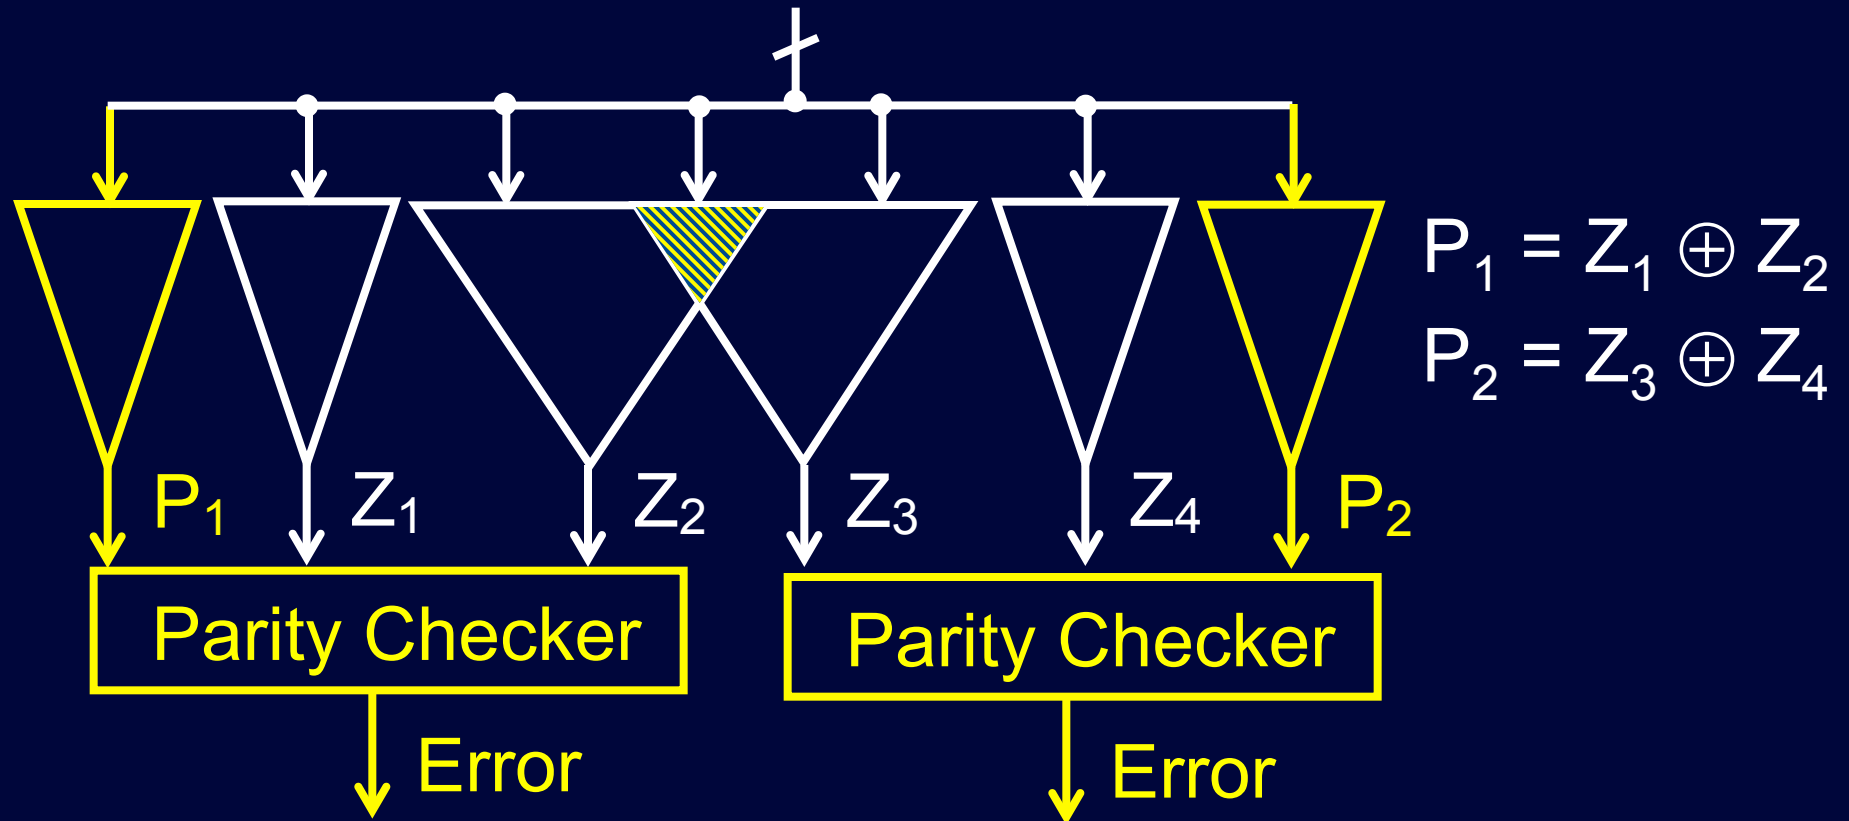
Single-Bit Logic Parity Prediction

- $P = Z_1 \oplus Z_2 \oplus \dots \oplus Z_m$
- Disjoint output logic (no logic sharing)
 - Only for combinational logic errors

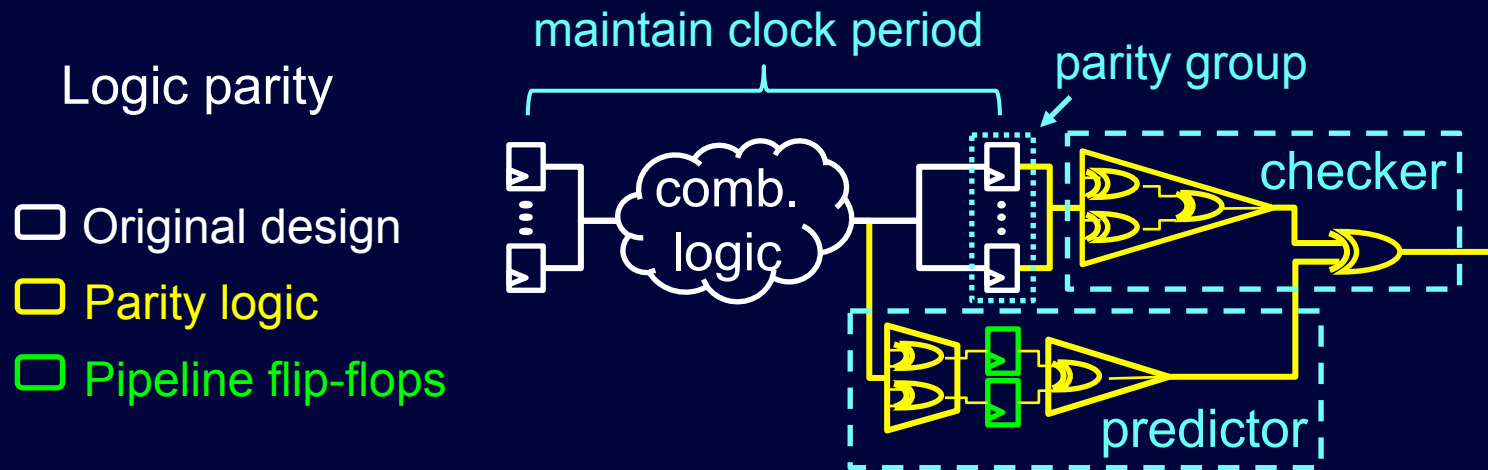


Multiple-Bit Logic Parity Prediction

- Main purpose: cost reduction (sharing, routing, logic)
 - Can be still expensive



Logic Parity Checking: Flip-Flop Errors



Logic parity: Naïve	200 MHz clock speed impact
Logic parity: Incorrect heuristic	80% additional energy impact
Logic parity: CLEAR heuristic	No clock speed impact Minimal energy impact

Parameters: parity size, flip-flop vulnerability, floorplan location, timing path slack, etc.

Parity Prediction for Datapath Circuits

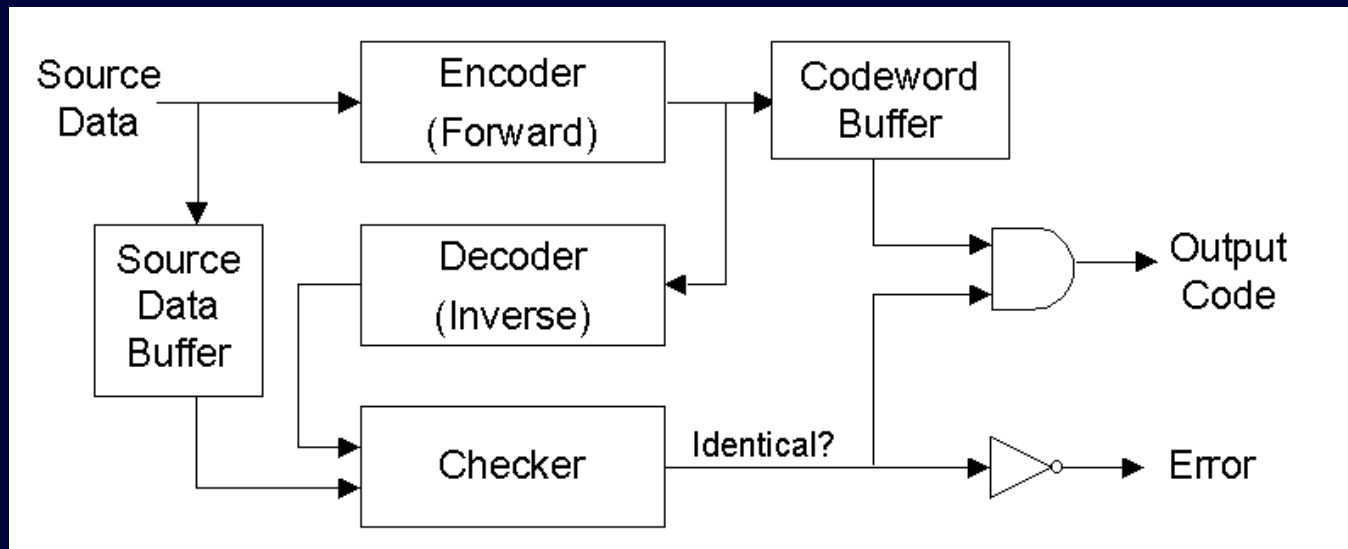
- $S = A + B$ (n -bit operation)
- $\text{Parity}(S) = S_1 \oplus S_2 \oplus S_3 \oplus \dots S_n$
 $= (A_1 \oplus B_1 \oplus C_1) \oplus (A_2 \oplus B_2 \oplus C_2) \dots (A_n \oplus B_n \oplus C_n)$
 $= (A_1 \oplus A_2 \dots A_n) \oplus (B_1 \oplus B_2 \dots B_n) \oplus (C_1 \oplus C_2 \dots C_n)$
 $= \text{Parity}(A) \oplus \text{Parity}(B) \oplus \text{Parity}(\text{internal carries})$
- Parity (internal carries) expensive
 - Several strategies for high-performance adders

Residue Codes for Datapath

- $y = x \bmod b$: y is residue of x (modulo b)
- $\text{Residue}(A + B) = \text{Residue}(A) + \text{Residue}(B)$
- $\text{Residue}(A \times B) = \text{Residue}(A) \times \text{Residue}(B)$
- Choice of b : Mersenne prime (form $2^m - 1$)
 - Coverage, checker complexity
- Issues: bit-wise logic, operand residue, checker cost
 - Often used for multipliers


Application-Specific CED

- LZ compression: loss-less, invertible
 - Compression: complex
 - Decompression: simple




9% area overhead, 0.5% delay overhead [Huang 00]

Error Detection Sequentials (EDS)




Errors in Processors



Control flow error

Incorrect instruction sequence



Computational error

Incorrect computation

Memory error

Incorrect value or address

Program Representation: Control Flow Graph

```
i=0; x=1; y=5;
```

```
While (i < 10) {
```

```
    z = x + y * i;
```

```
    i = i + 1;
```

```
}
```

**Basic Block
(BB)**

```
mov r1, 0  
mov r2, 1  
mov r3, 5
```

V1

L1:

```
inc r1  
bge r1, 10, L2
```

V2


```
mul r4, r3, r1  
add r5, r2, r4  
br L1
```

V3

L2:

```
...
```

V4



SIHFT

- Software Implemented Hardware Fault Tolerance
 - Automated by compiler
 - **EDDI** [Oh, IEEE Trans. Reliability 02]
 - **CFCSS** [Oh IEEE Trans. Reliability 02b]
 - **ED⁴I** [Oh IEEE Trans. Computers 02]
 - Lots of recent publications

EDDI

- Error Detection using Duplicated Instructions
- Duplicate instructions inside basic blocks
 - Different registers
- Duplicate data structures
- Comparison before memory stores
- Performance penalty 13% - 111%
 - Reduced by Instruction Level Parallelism (ILP)

EDDI Example

ADD R3, R1, R2	; $R3 \leftarrow R1 + R2$
MUL R4, R3, R5	; $R4 \leftarrow R3 * R5$
ST 0(SP), R4	; store R4 in location pointed by SP



ADD R3, R1, R2	; $R3 \leftarrow R1 + R2$ master
ADD R23, R21, R22	; $R23 \leftarrow R21 + R2$ shadow
MUL R4, R3, R5	; $R4 \leftarrow R3 * R5$ master
MUL R24, R23, R25	; $R24 \leftarrow R23 * R25$ shadow
BNE R4, R24, Error_Handler	; compare
ST 0(SP), R4	; store master result
ST offset(SP), R24	; store shadow result

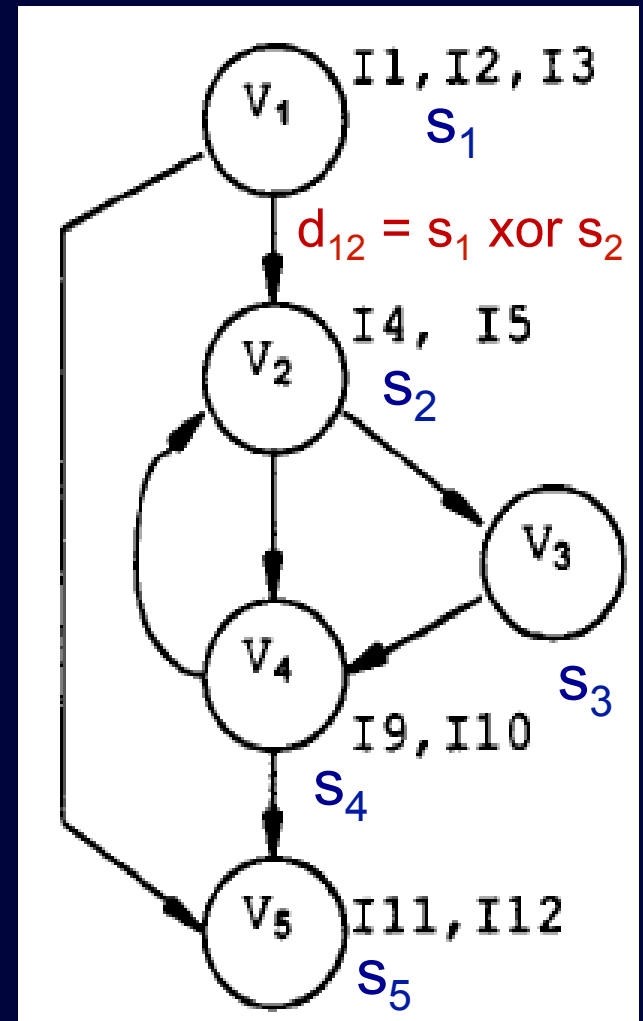
EDDI Design Choices

- Check after each instruction ?
- Storeless basic blocks (SBB) ?
 - No branch or store except final instruction
- Why SBB ?
 - Correctness defined by program output
 - Erroneous branches: stores skipped ?
 - Check at branches too

CFCSS

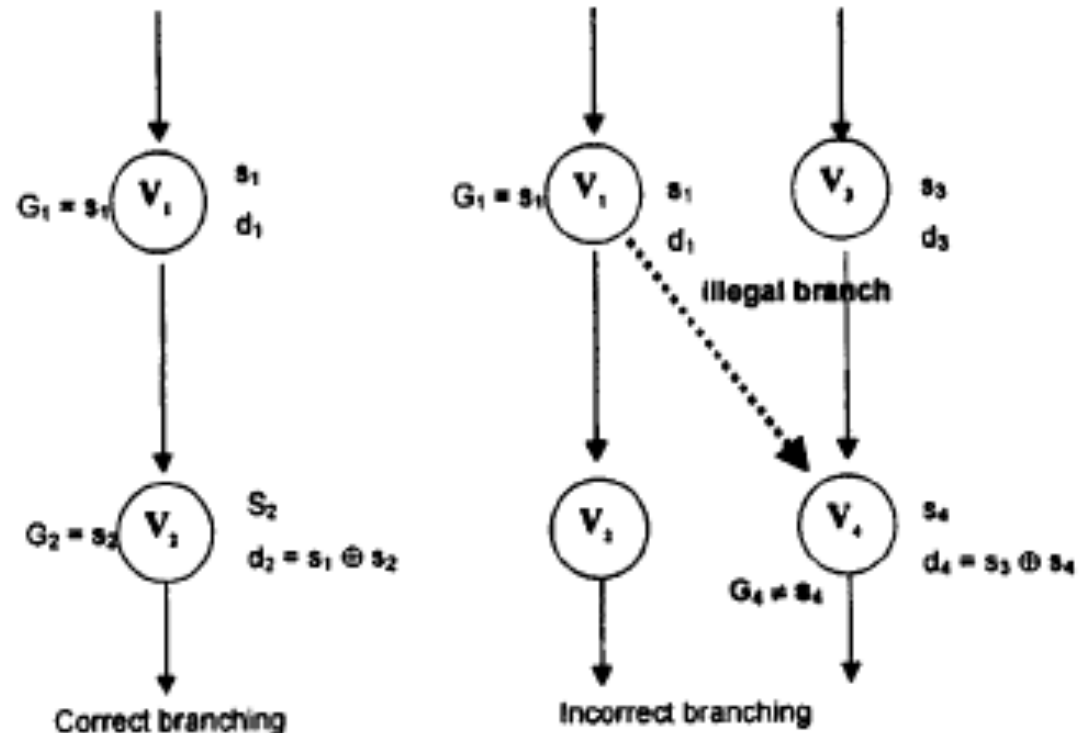
- Control Flow Checking using Software Signatures

- Each node
 - Unique signature
- Each edge
 - Transition between 2 signatures
 - Difference function: XOR



CFCSS


- Runtime signature G
- Basic block i to j
 - $G = s_i \text{ XOR } d_{i,j}$
 - Check $G = s_j$



G_n run-time signature at node V_n
 s_n signature assigned to node V_n
 d_n signature difference

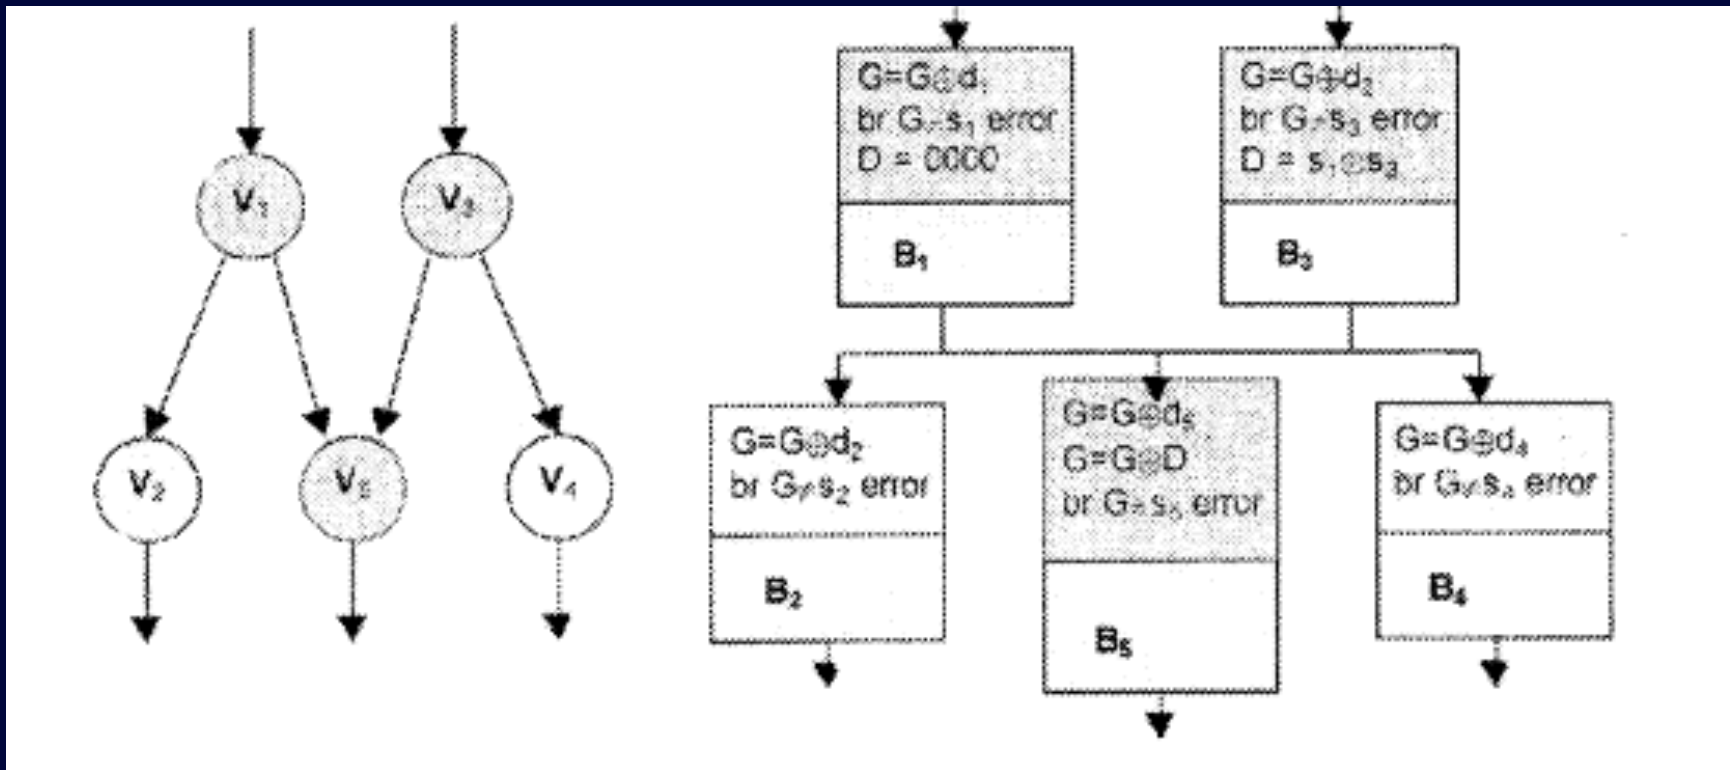
CFCSS Implementation

- Global variable G holds run-time signature
- Compute & check signature: start of each basic block



CFCSS: Branch Fan-in

- Basic block with multiple predecessors
- Run-time adjusting signature D differentiates fan-in




ED⁴I

- Error Detection using Diverse Data & Duplicated Instructions
- Duplicated instructions, **data diversity**
 - Expressions in shadows multiply by k (-1, -2, ...)
- $k = -2$: good choice
- Transient errors & most permanent faults detected
- Issues: floating point, pointers

SIHFT Results [Lovelette 02]

- COTS in space: no hardware redundancy
- ARGOS satellite experiment
 - Compare rad-hard processor vs. COTS
- Undetected errors in rad-hard processor
- COTS: 5.55 SEUs / Mbyte / day, 99.7% coverage
- 98.8% successful recovery: software ECC + restart
- COTS + SIHFT: faster than rad-hard

Multi-Threading for CED




- Same application computed by two threads
 - [Rotenberg 99, Saxena 00, Mukherjee 02]

What About Recovery?

Instruction Replay (IR)

Flush



Style 2:
“Flush”

Style 1:
“Instruction Replay”



- Cross-layer protected
- LEAP-DICE protected
- Recovery logic

Instruction Replay

Flush recovery

Overhead for recovery hardware

16% area,
21% energy

0.6% area,
0.9% energy

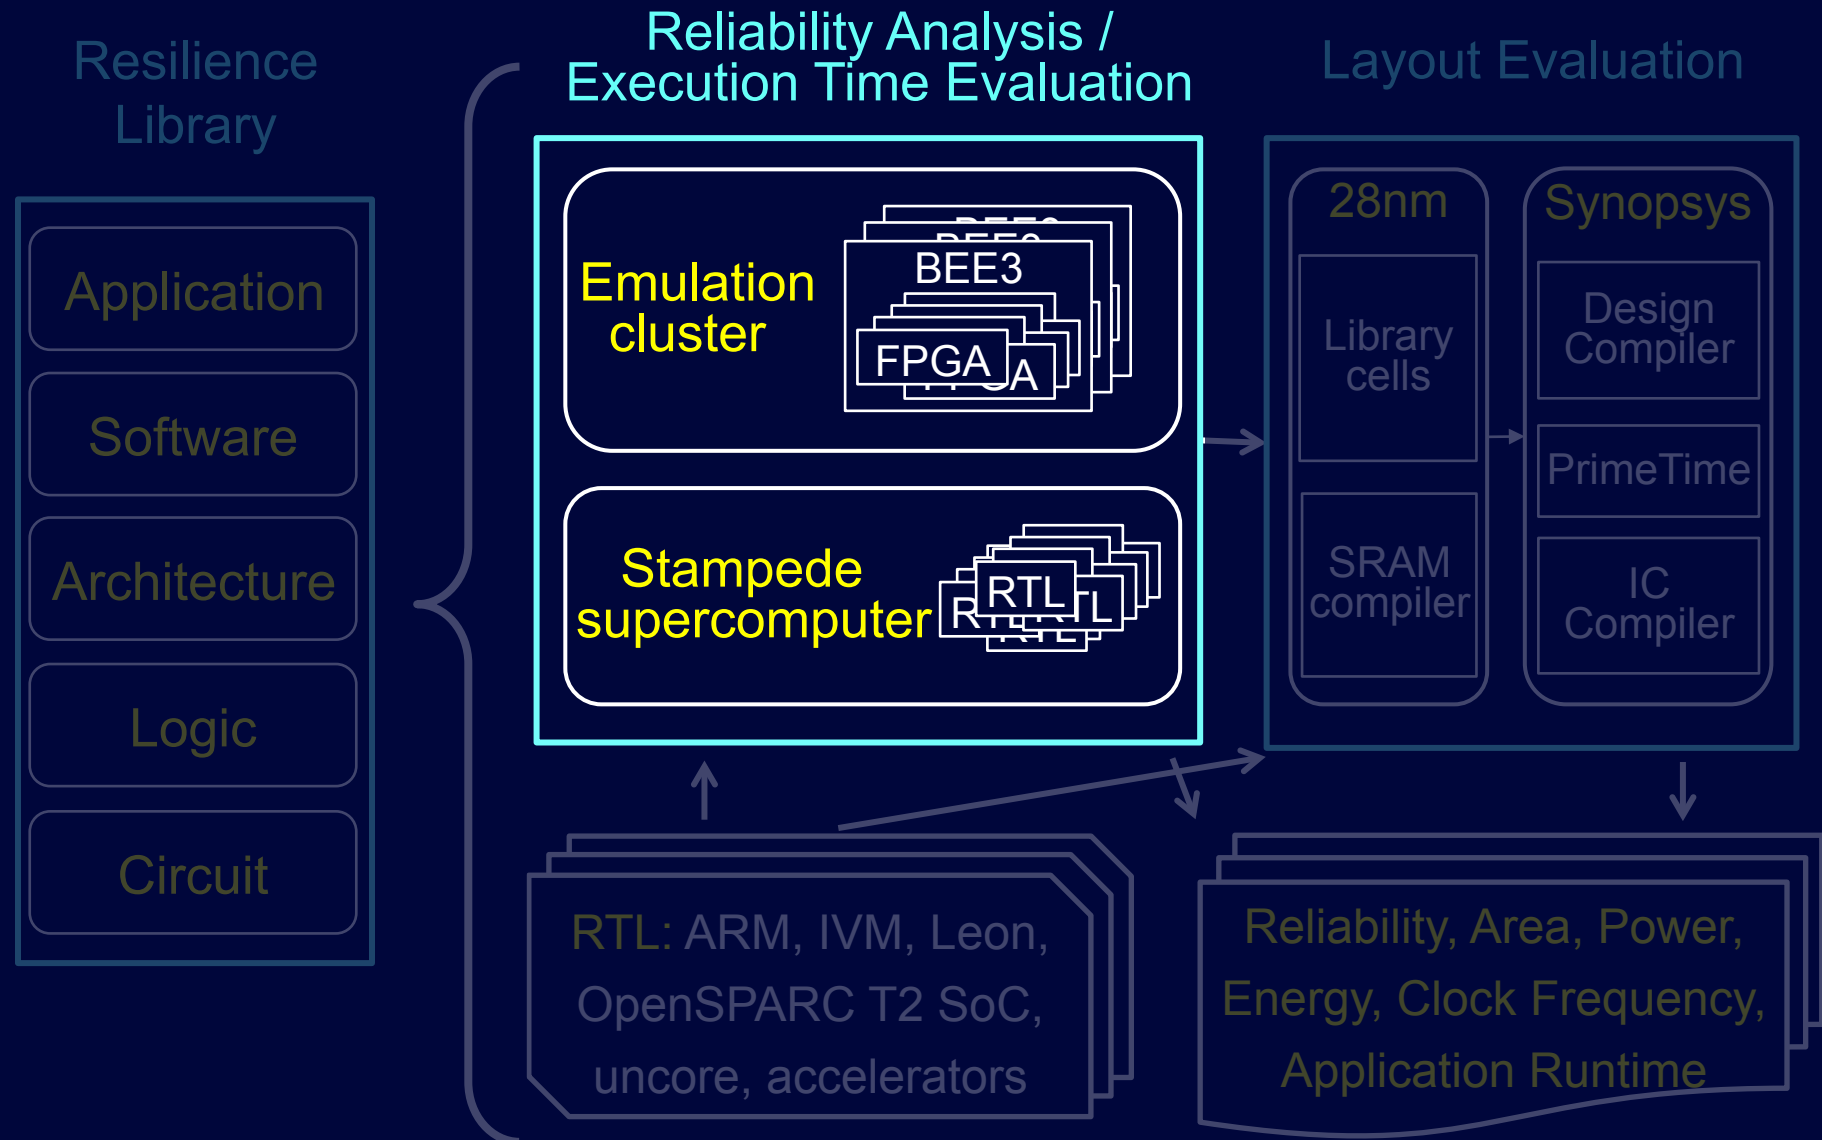
Recovery latency

47 cycles

7 cycles


CLEAR

Cross-Layer Exploration for Architecting Resilience




Radiation-Induced Soft Errors


Soft errors



Radiation
beam testing



Flip-flop
error injection




The Los Alamos
Neutron Science Center

Simulation /
Emulation

Soft Error Injection

Simulation speed

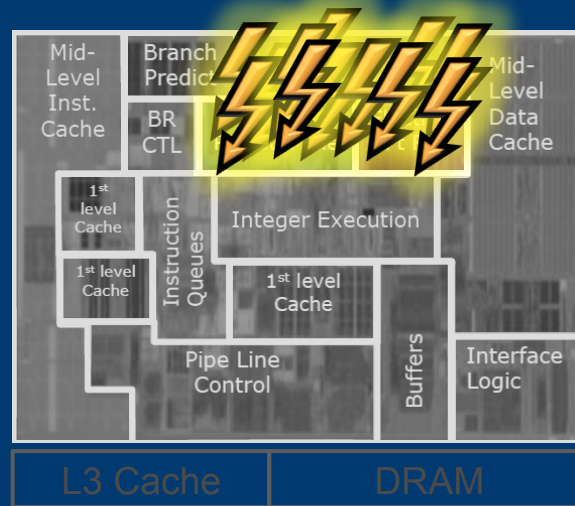


High-level error injection



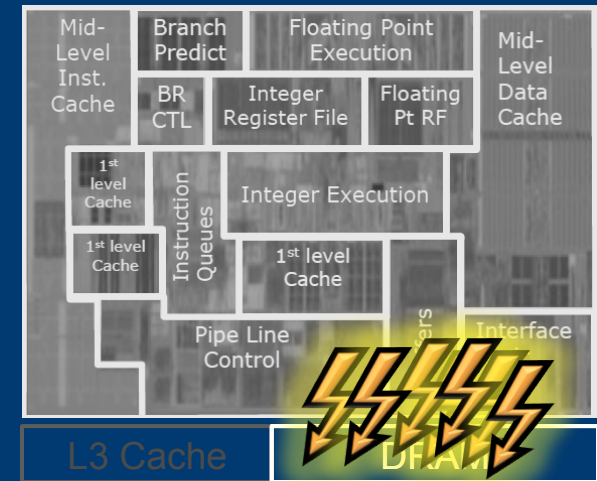
Flip-flop

10^2 cycles / sec



Architectural register

10^7 cycles / sec



Program variable

10^9 cycles / sec


Soft Error Injection

Accuracy




Ground truth

High-level error injection




Flip-flop

10^2 cycles / sec



Architecture register


10^7 cycles / sec



Program variable

10^9 cycles / sec

Silent Data Corruption (SDC)





- Output file incorrect
- No error indication

Detected but Uncorrected Error (DUE) also considered

Perils of Inaccurate Estimation

Overestimation

SDC





Overprotection



Underestimation

SDC



Unhappy customers



Little Prior Work

Error injection studies

[Sloan DSN 07] [Seward ITC 03] [Li HPCA 07]
[Lanigan DSN 07] [Chen PRDC 08] [Meixner Micro 07] [Reddy DSN 07]
[Wang Trans. Dependable and Secure Computing 06] [Ramachandran DSN 08]
[Hari MICRO 09] [Dimitrov PACT 10]
[Wang ISCA 07] [Sakata DSN 07] [Kalbarczyk Trans. Software Eng. 99]
[Feng ASPLOS 10] [Choi Trans. Reliability 90] [Lee DAC 01]
[Ejlali DSN 03] [Arlat TCOMP 03] [Ray MICRO 01]
[Hari ASPLOS 12] [Choi ICCAD 93] [Christmansson ISSRE 98]
[Racunas HPCA 07] [Nakka DSN 07] [Goswami FTCS 93]
[Kruijf ISCA 10] [Ferna TCAD 12] [Sahoo DSN 08]
[Choi Trans. Reliability 90] [Ignat DATE 06] [Miskov-Zivanov TCAD 10]
[Gschwind ICCD 11] [Kanawati AIAA 93] [Pellegrini DATE 12] [Mukherjee HPCA 05]
[Blome Workshop on Architectural Reliability 08] [Cheng DSN 07] [Rimen FTCS 94]
[Li HPCA 09] [Li ASPLOS 08] [Reis TACO 05] [Rebaudengo IOLTW 02]
[Yim DSN 10] [Reddy ICCD 06]
[Zhang PACT 10] [Maniatakos TCOMP 11] [Chen ASPDAC 06]
[Li DSN 05] [Pattabiraman Trans. Dependable and Secure Computing 11]
[Pandit DSN 09] [Avirmeni DSN 09] [Gracia DFT 01]
[Sterpone DDECS 11] [Constantinescu DSN 12]
[Michalak Trans. Device and Materials Reliability 12] [Goswami FTCS 93]
[Stott DSN 02] [Narayanasam DATE 07] [Baraza TVLSI 08]
[Vadlamani DATE 10] [Wang DSN 04] [Andres TVLSI 08] [Cheng TCAD 99]
[Alderighi Trans. Nucl. Sci. 08] [Romanescu PACT 08] [Thompto DAC 10]
[Saggese DSN 05] [Pattabiraman DSN 08] [Lima DAC 03]

Quantified comparison

[Rimen FTCS 94]

[Rebaudengo IOLTW 02]

What We Found

- Naïve high-level injections **highly** inaccurate
- How inaccurate?

What We Found

- Naïve high-level injections **highly** inaccurate
- How inaccurate?

Designs:

LEON3 (in-order, single-issue), ALPHA (out-of-order, superscalar)

Applications: SPEC 2000

Error injection samples: 6 million

Inaccuracy Question

Undetected
output error
rate

0.8%

1x

6x

13x

Flip-flop

Architectural
register

Program
variable

What We Found


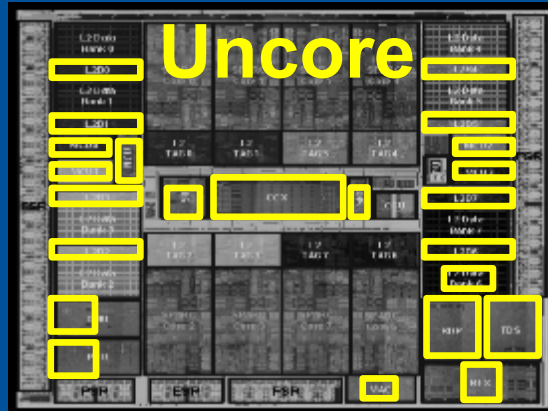
- Naïve high-level injections **highly** inaccurate
- How inaccurate?
 - Up to 45X
 - Neither optimistic nor pessimistic

What We Found


- Naïve high-level injections **highly** inaccurate
- How inaccurate?
 - Up to 45X
 - Neither optimistic nor pessimistic
- Why inaccurate?
 - Only 3% flip-flop error propagations modeled

Uncore Components

OpenSPARC T2 SoC



Intel i7 quad-core SoC



Power

Processor cores 60.2%	Uncore 39.8%
--------------------------	-----------------

[Gupta USENIX 12]

Existing Work

Errors in processor cores


[Sloan DSN 07] [Seward ITC 03] [Li HPCA 07]
[Lanigan DSN 07] [Chen PRDC 08] [Meixner Micro 07]
[Wang Trans. Dependable and Secure Computing 06] [Ramachandran DSN 08] [Reddy DSN 07]
[Hari MICRO 09] [Dimitrov PACT 10]
[Wang ISCA 07] [Sakata DSN 07] [Kalbarczyk Trans. Software Eng. 99]
[Feng ASPLOS 10] [Choi Trans. Reliability 90] [Lee DAC 01]
[Ejlali DSN 03] [Arlat TCOMP 03] [Ray MICRO 01]
[Hari ASPLOS 12] [Choi ICCAD 93] [Christmansson ISSRE 98]
[Racunas HPCA 07] [Nakka DSN 07] [Goswami FTCS 93]
[Kruijf ISCA 10] [Ferna TCAD 12] [Sahoo DSN 08]
[Choi Trans. Reliability 90] [Ignat DATE 06] [Miskov-Zivanov TCAD 10]
[Gschwind ICCD 11] [Kanawati AIAA 93] [Pellegrini DATE 12] [Mukherjee HPCA 05]
[Blome Workshop on Architectural Reliability 08] [Cheng DSN 07] [Rimen FTCS 94]
[Li HPCA 09] [Li ASPLOS 08] [Reis TACO 05] [Rebaudengo IOLTW 02]
[Yim DSN 10] [Reddy ICCD 06]
[Zhang PACT 10] [Maniatakos TCOMP 11] [Chen ASPDAC 06]
[Li DSN 05] [Pattabiraman Trans. Dependable and Secure Computing 11] [Avirmeni DSN 09]
[Pandit DSN 09] [Sterpone DDECS 11] [Constantinescu DSN 12] [Gracia DFT 01]
[Michalak Trans. Device and Materials Reliability 12] [Goswami FTCS 93]
[Stott DSN 02] [Narayanasam DATE 07] [Baraza TVLSI 08]
[Vadlamani DATE 10] [Wang DSN 04] [Andres TVLSI 08] [Cheng TCAD 99]
[Alderighi Trans. Nucl. Sci. 08] [Romanescu PACT 08] [Thompto DAC 10]
[Saggese DSN 05] [Pattabiraman DSN 08] [Lima DAC 03]

Errors in uncore?

HUNDREDS of publications


Uncore Soft Errors: First Extensive Study

- New error injection: fast & accurate
 - 20,000x speedup vs. RTL
- Reliability impact: uncore \approx processor cores
 - BUT, **long** error propagation latency




Lots of CLEAR Results

798 total combinations




High-Level Enough?

Silent Data Corruption (SDC) Rate



Cross-Layer Combinations

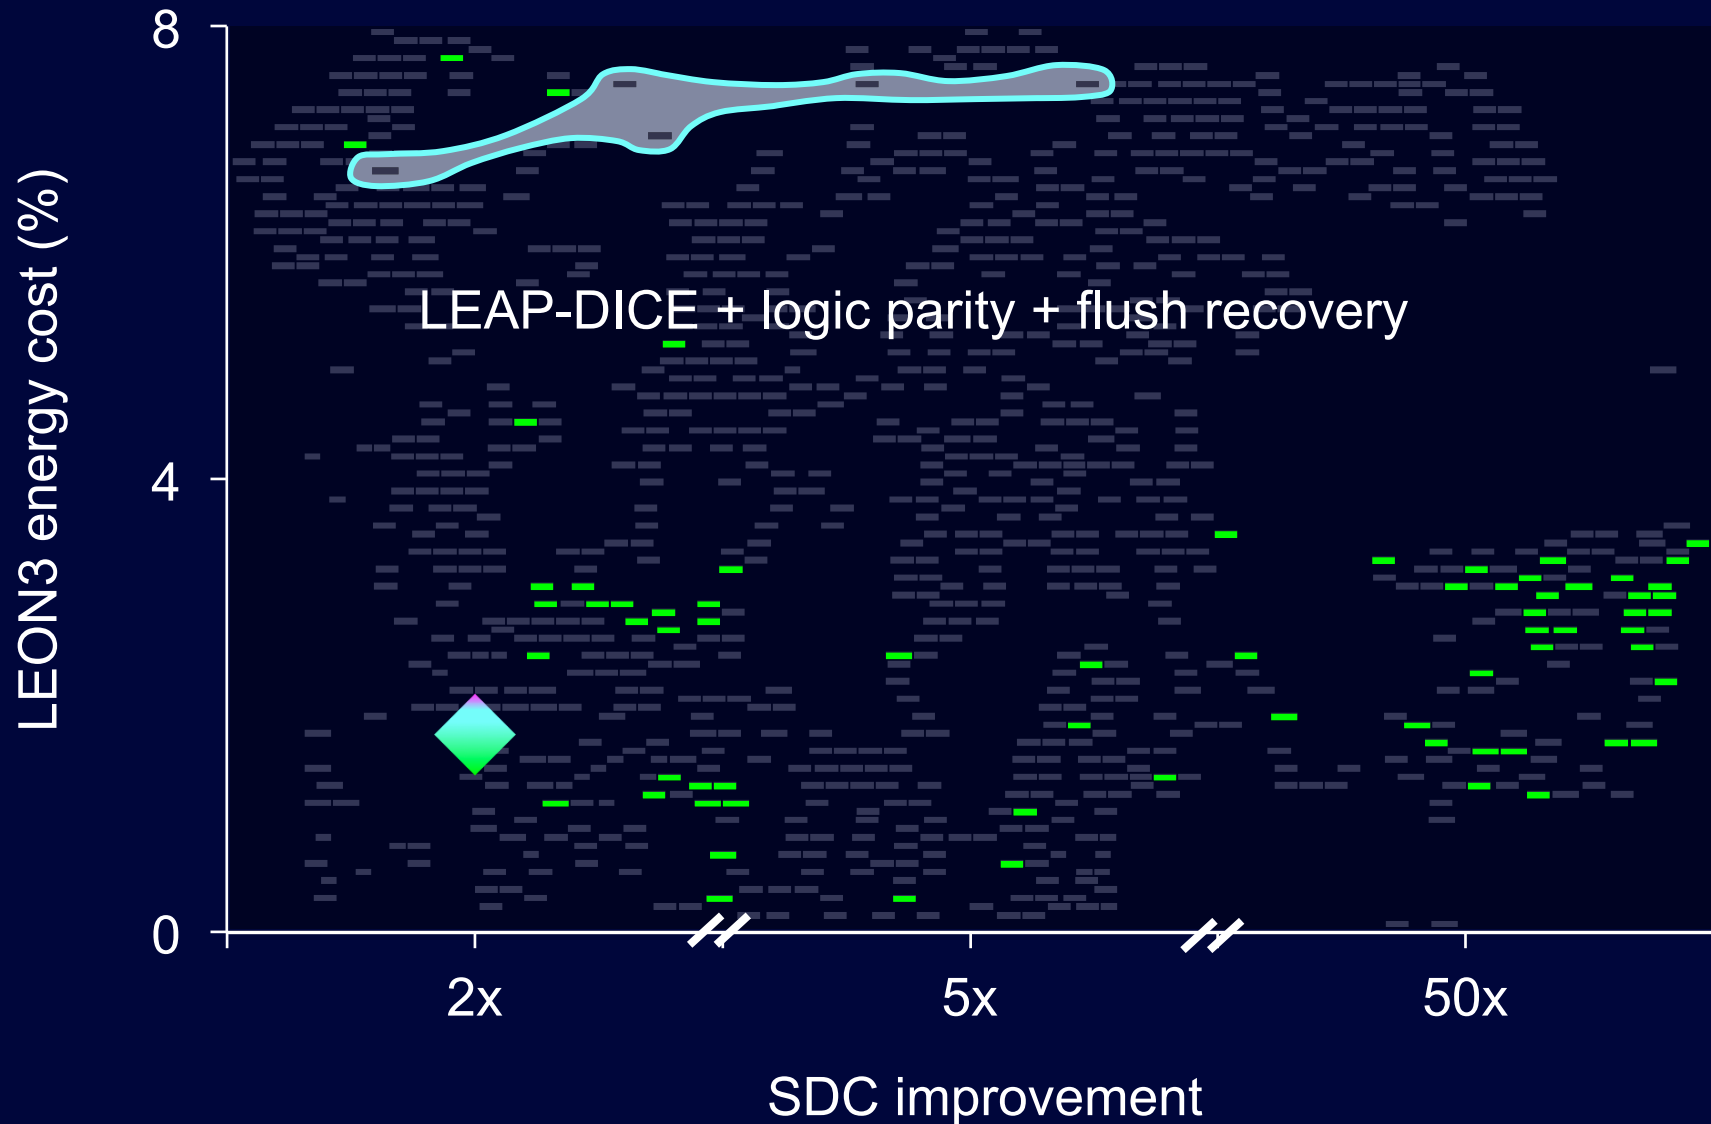
Algorithm
Software
Architecture
Logic
Circuit



Error bar: 0.5% energy cost (additive)

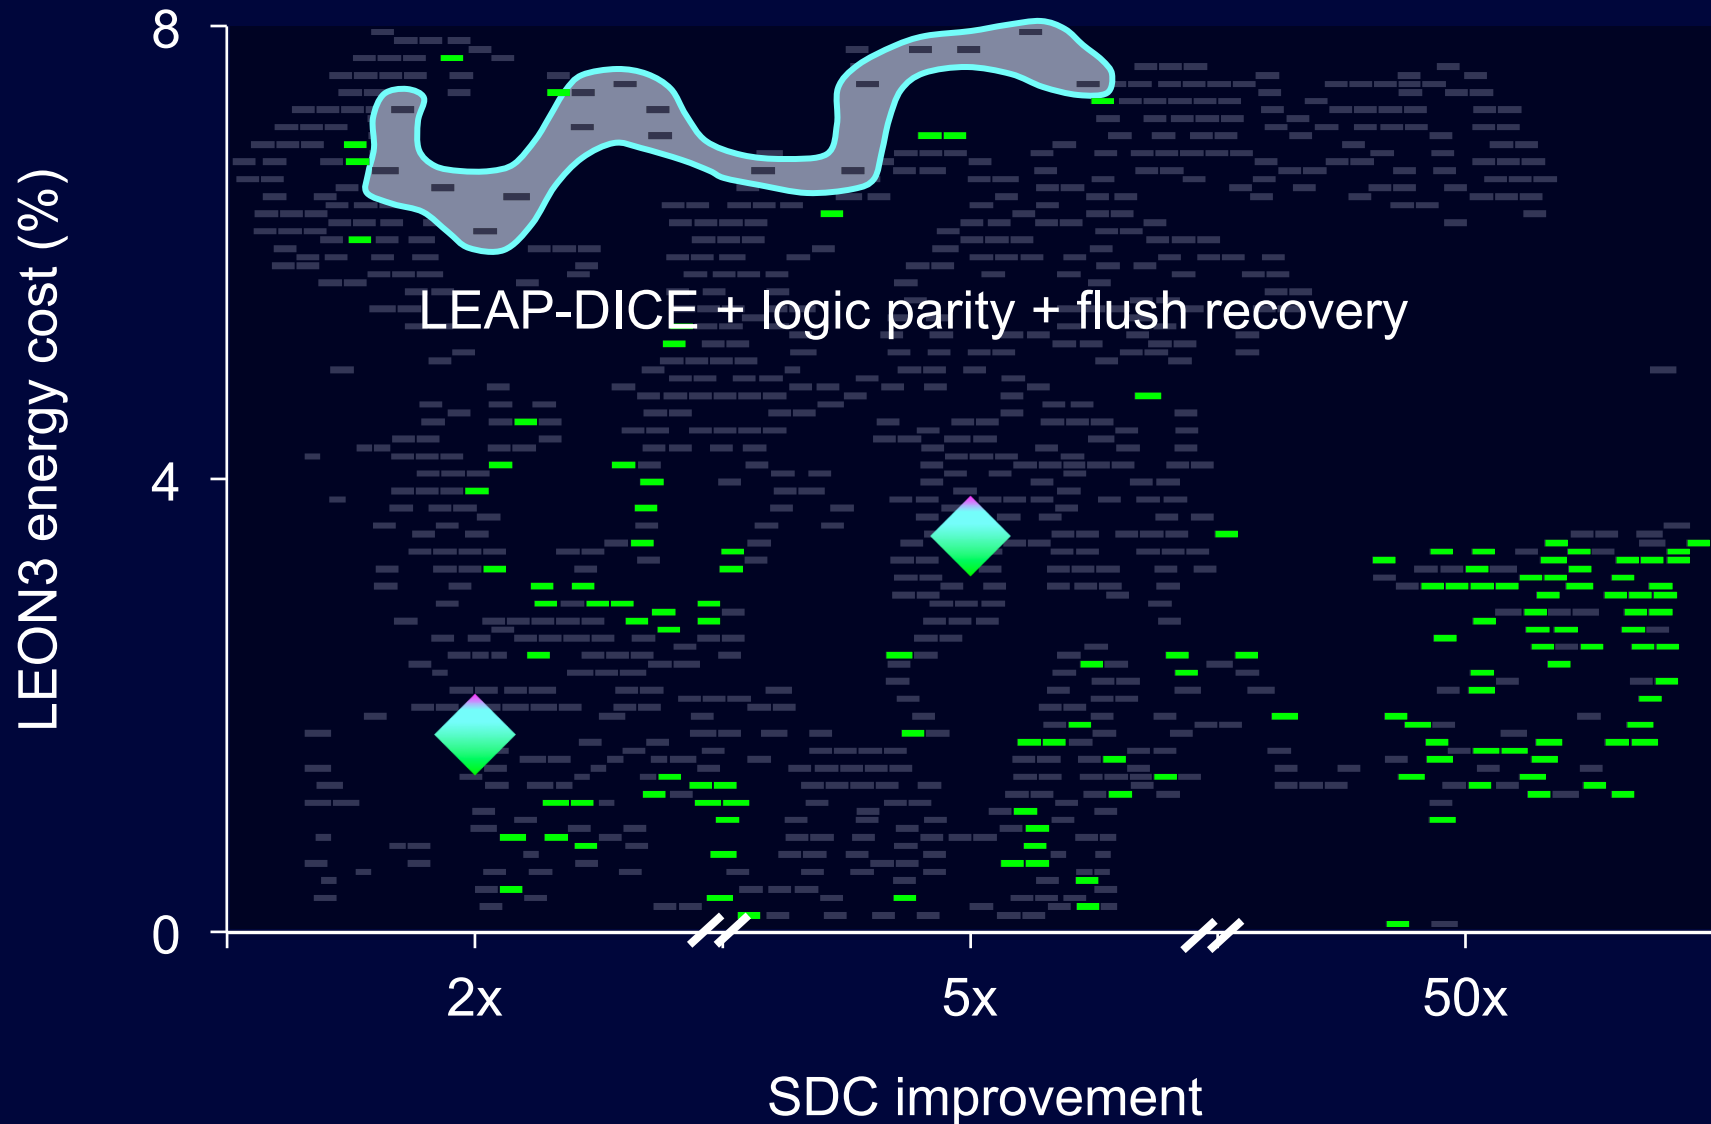
Cross-Layer Combinations

Algorithm
Software
Architecture
Logic
Circuit



Cross-Layer Combinations

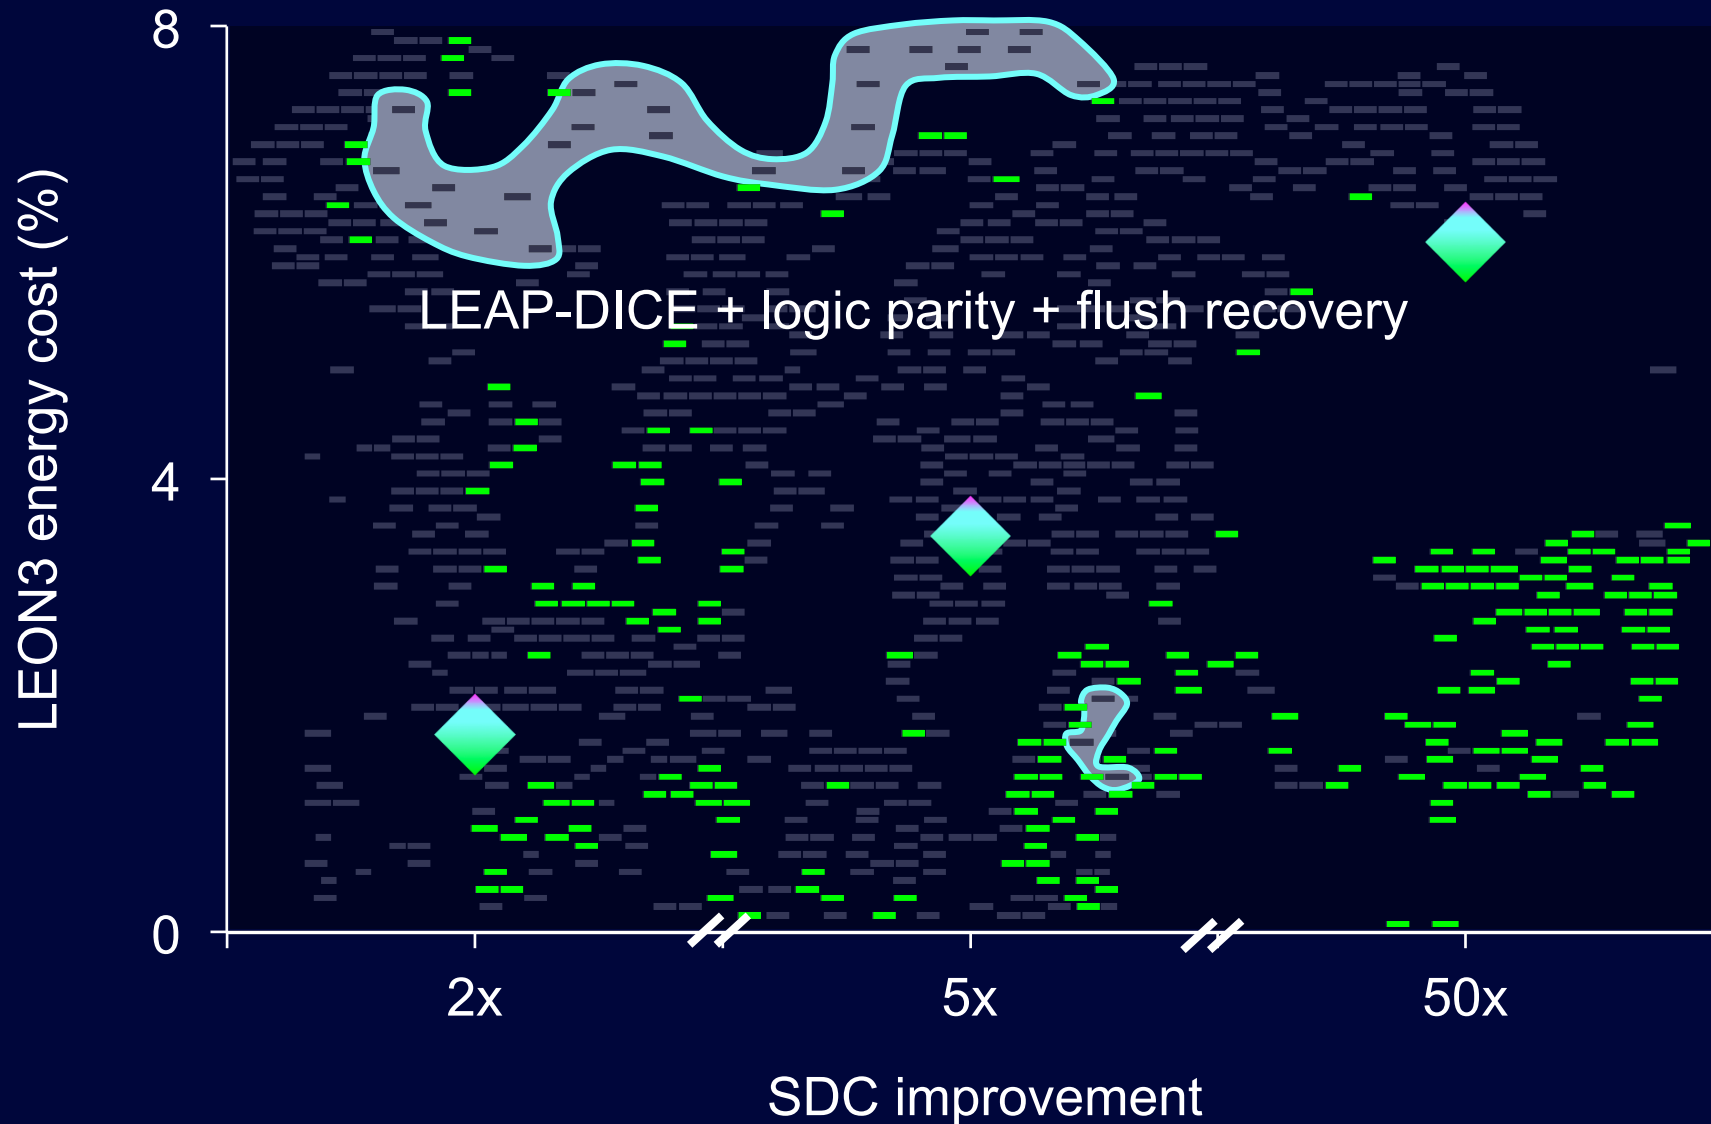
Algorithm
Software
Architecture
Logic
Circuit



Error bar: 0.5% energy cost (additive)

Cross-Layer Combinations

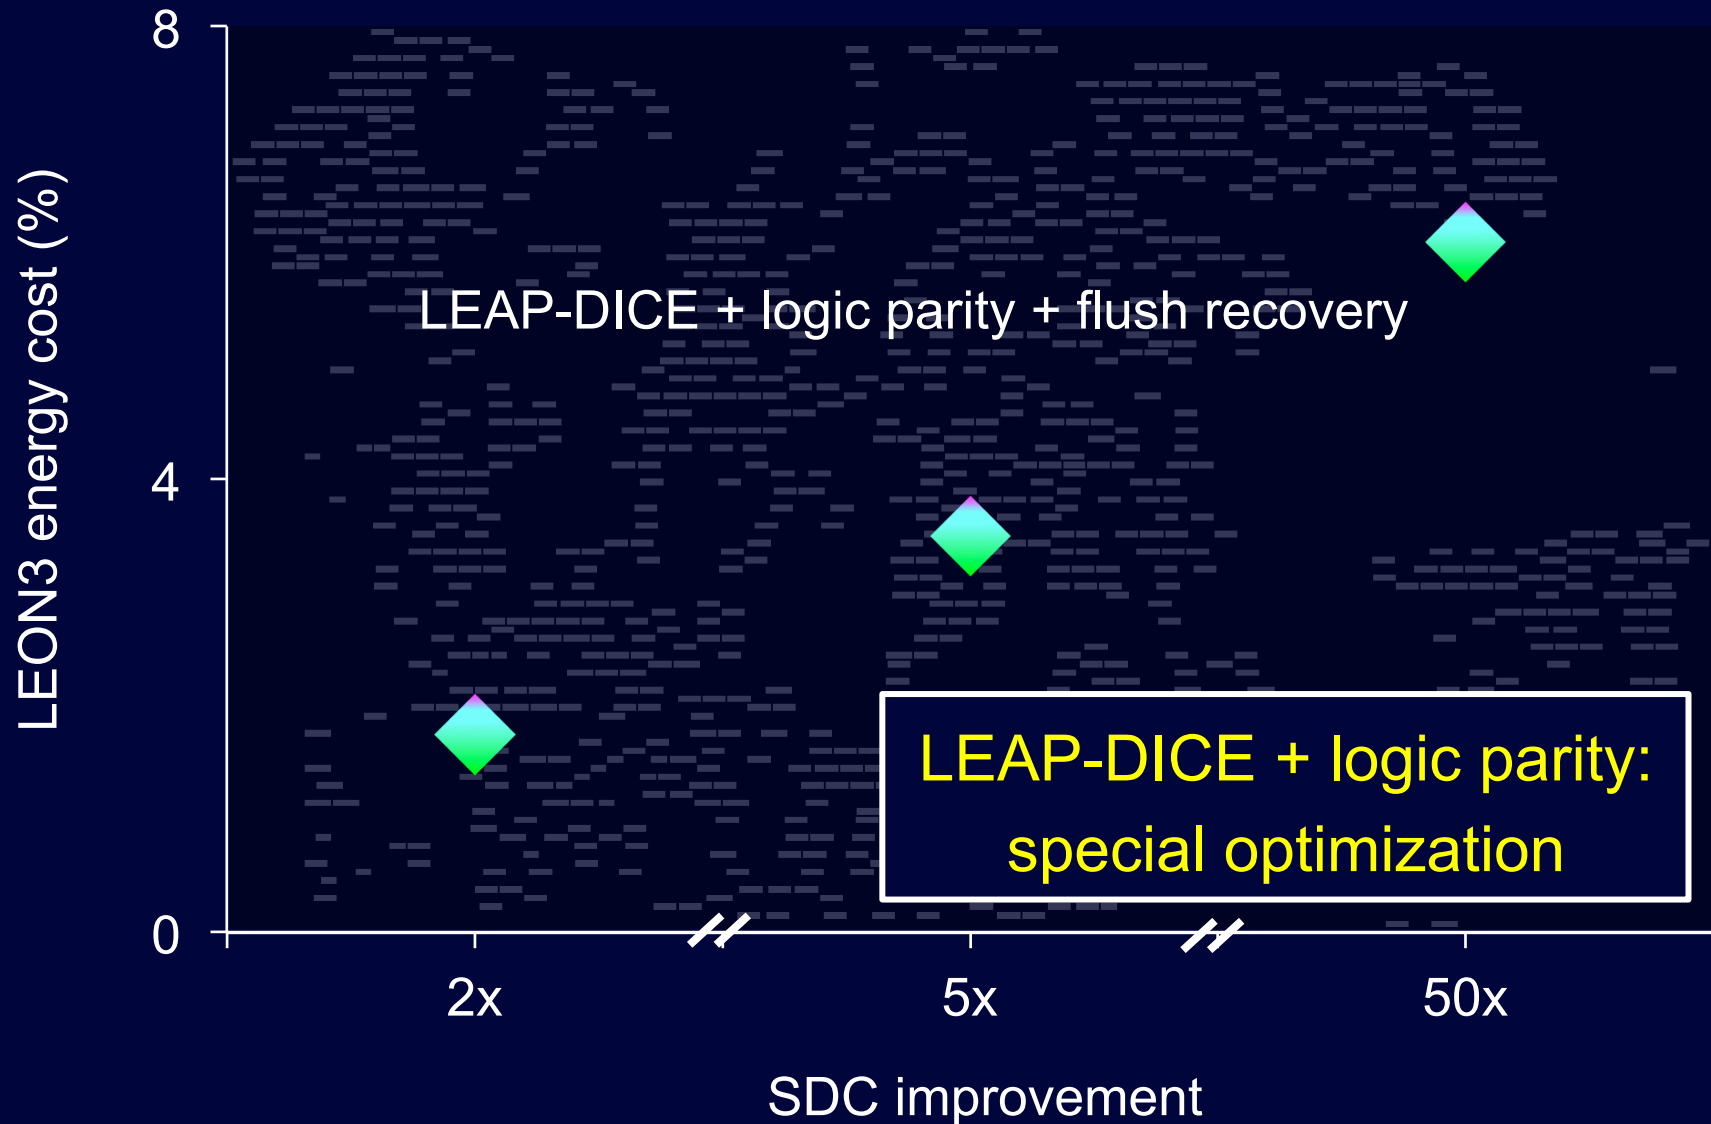
Algorithm
Software
Architecture
Logic
Circuit



Error bar: 0.5% energy cost (additive)



Cross-Layer Combinations

Algorithm
Software
Architecture
Logic
Circuit



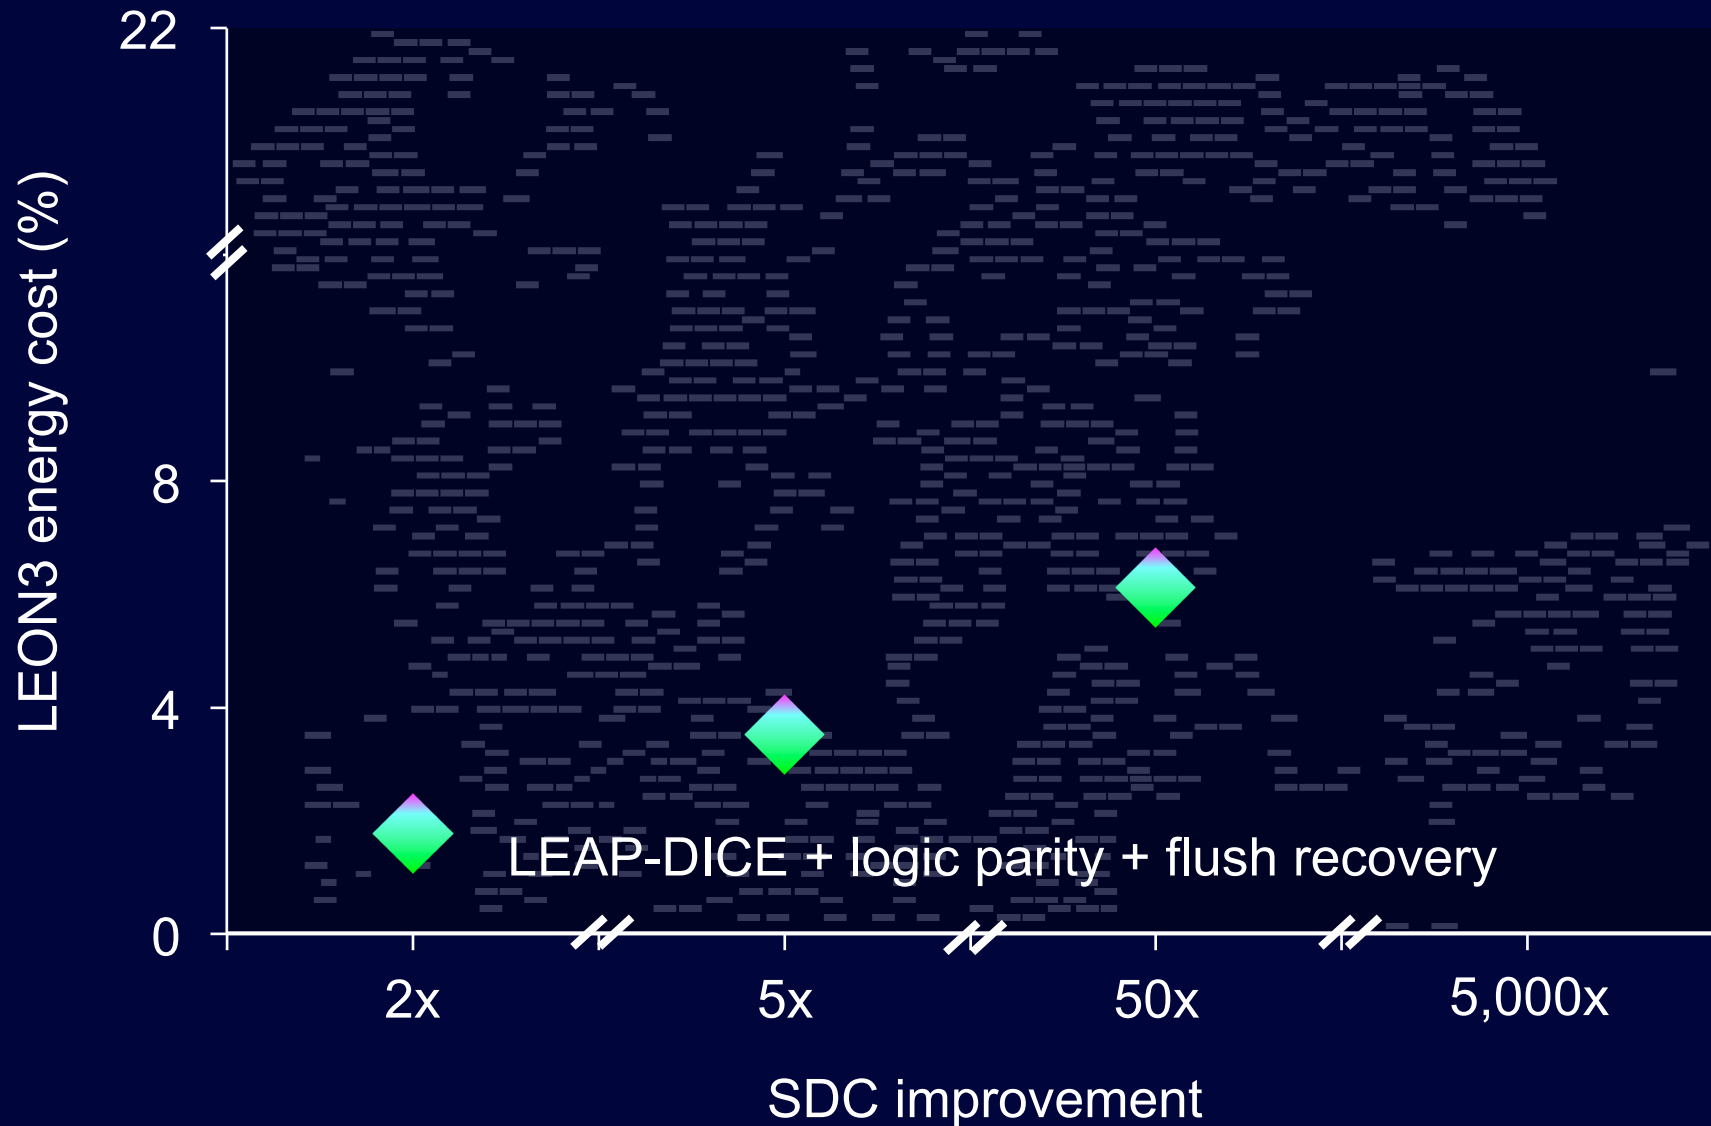
Error bar: 0.5% energy cost (additive)

Architecture & Software: Too Expensive



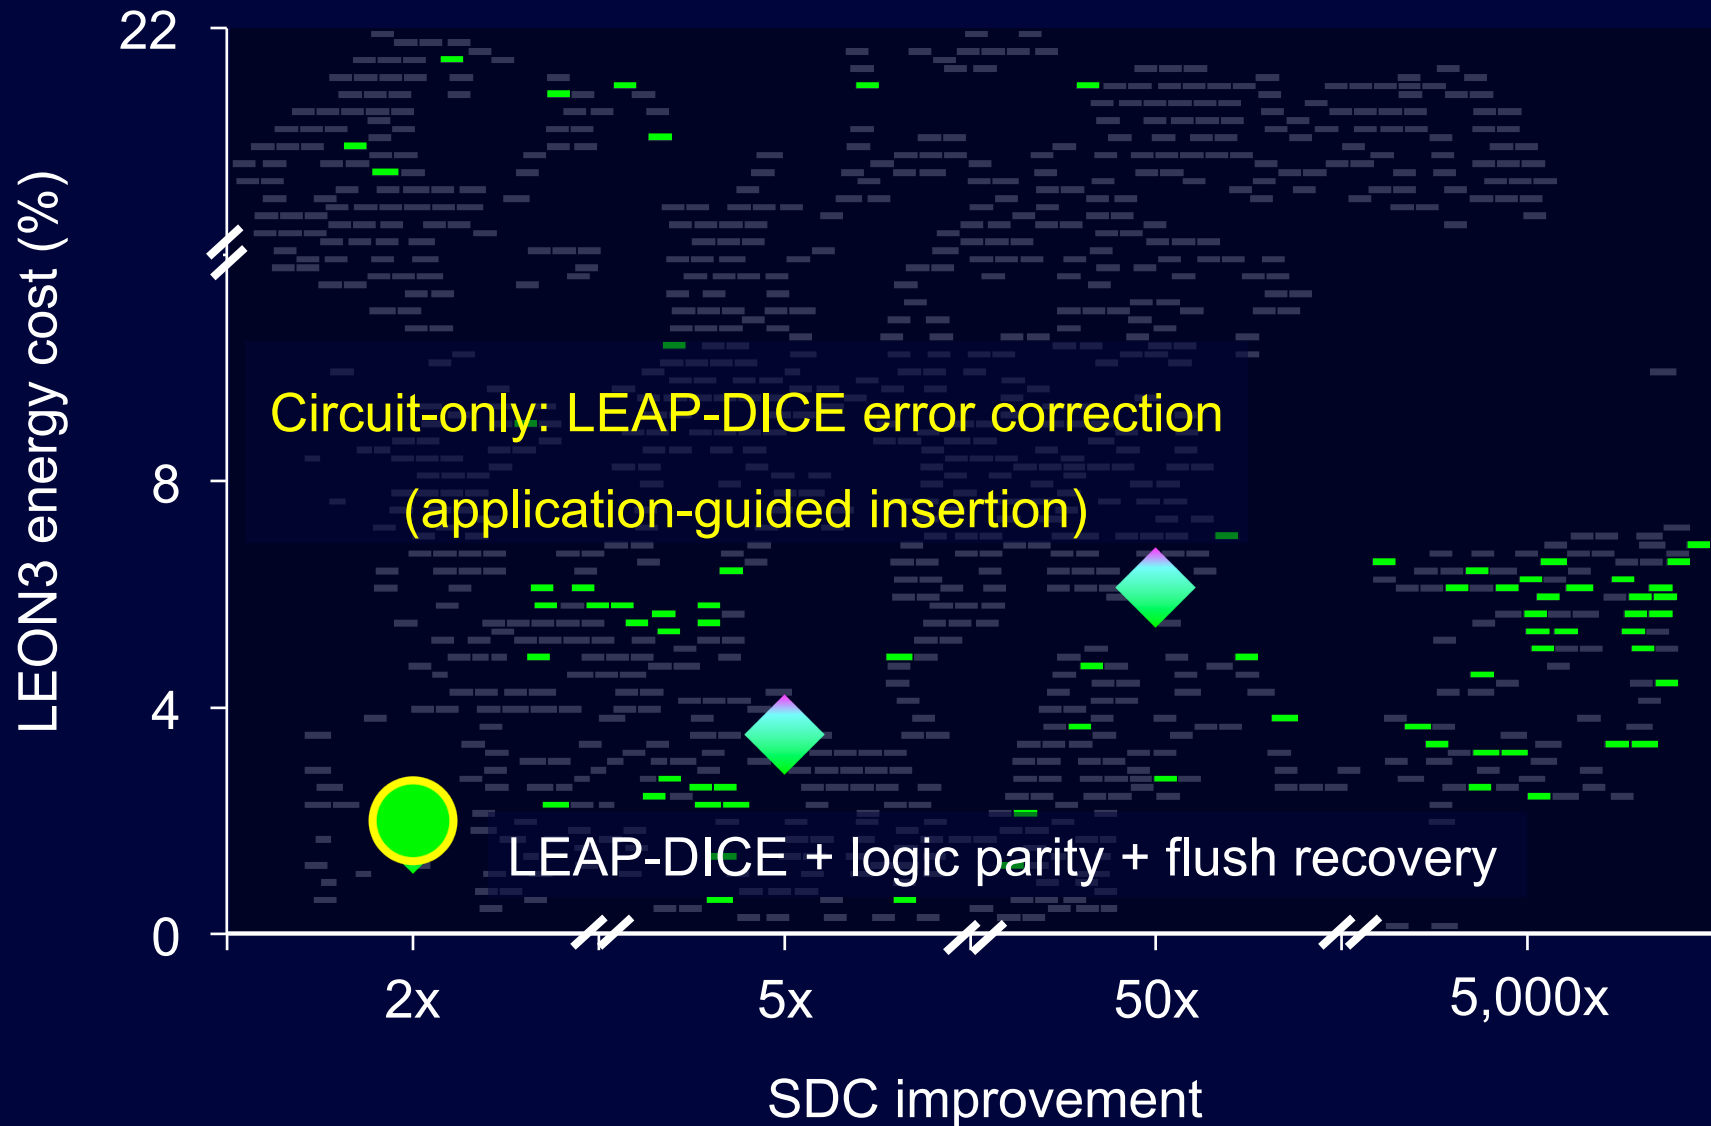
Circuit-only (Application-guided): Highly Effective

Algorithm
Software
Architecture
Logic
Circuit



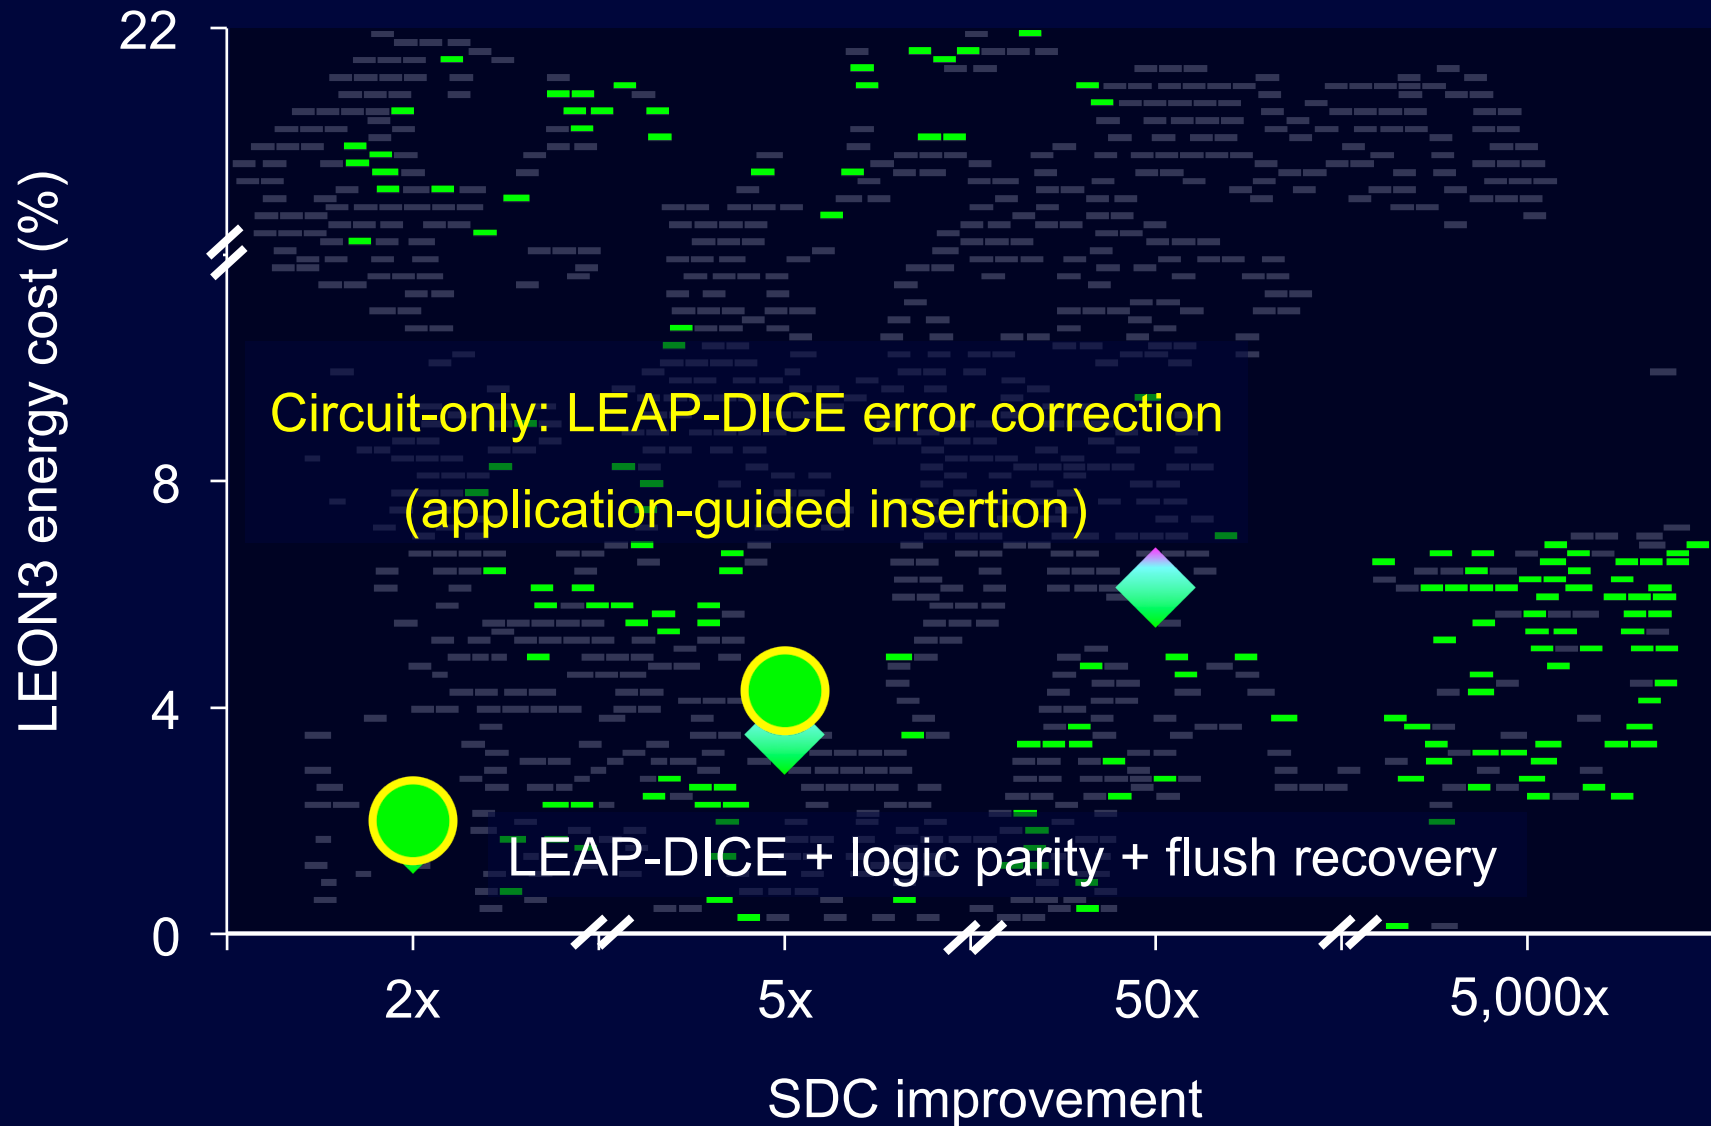
Circuit-only (Application-guided): Highly Effective

Algorithm
Software
Architecture
Logic
Circuit



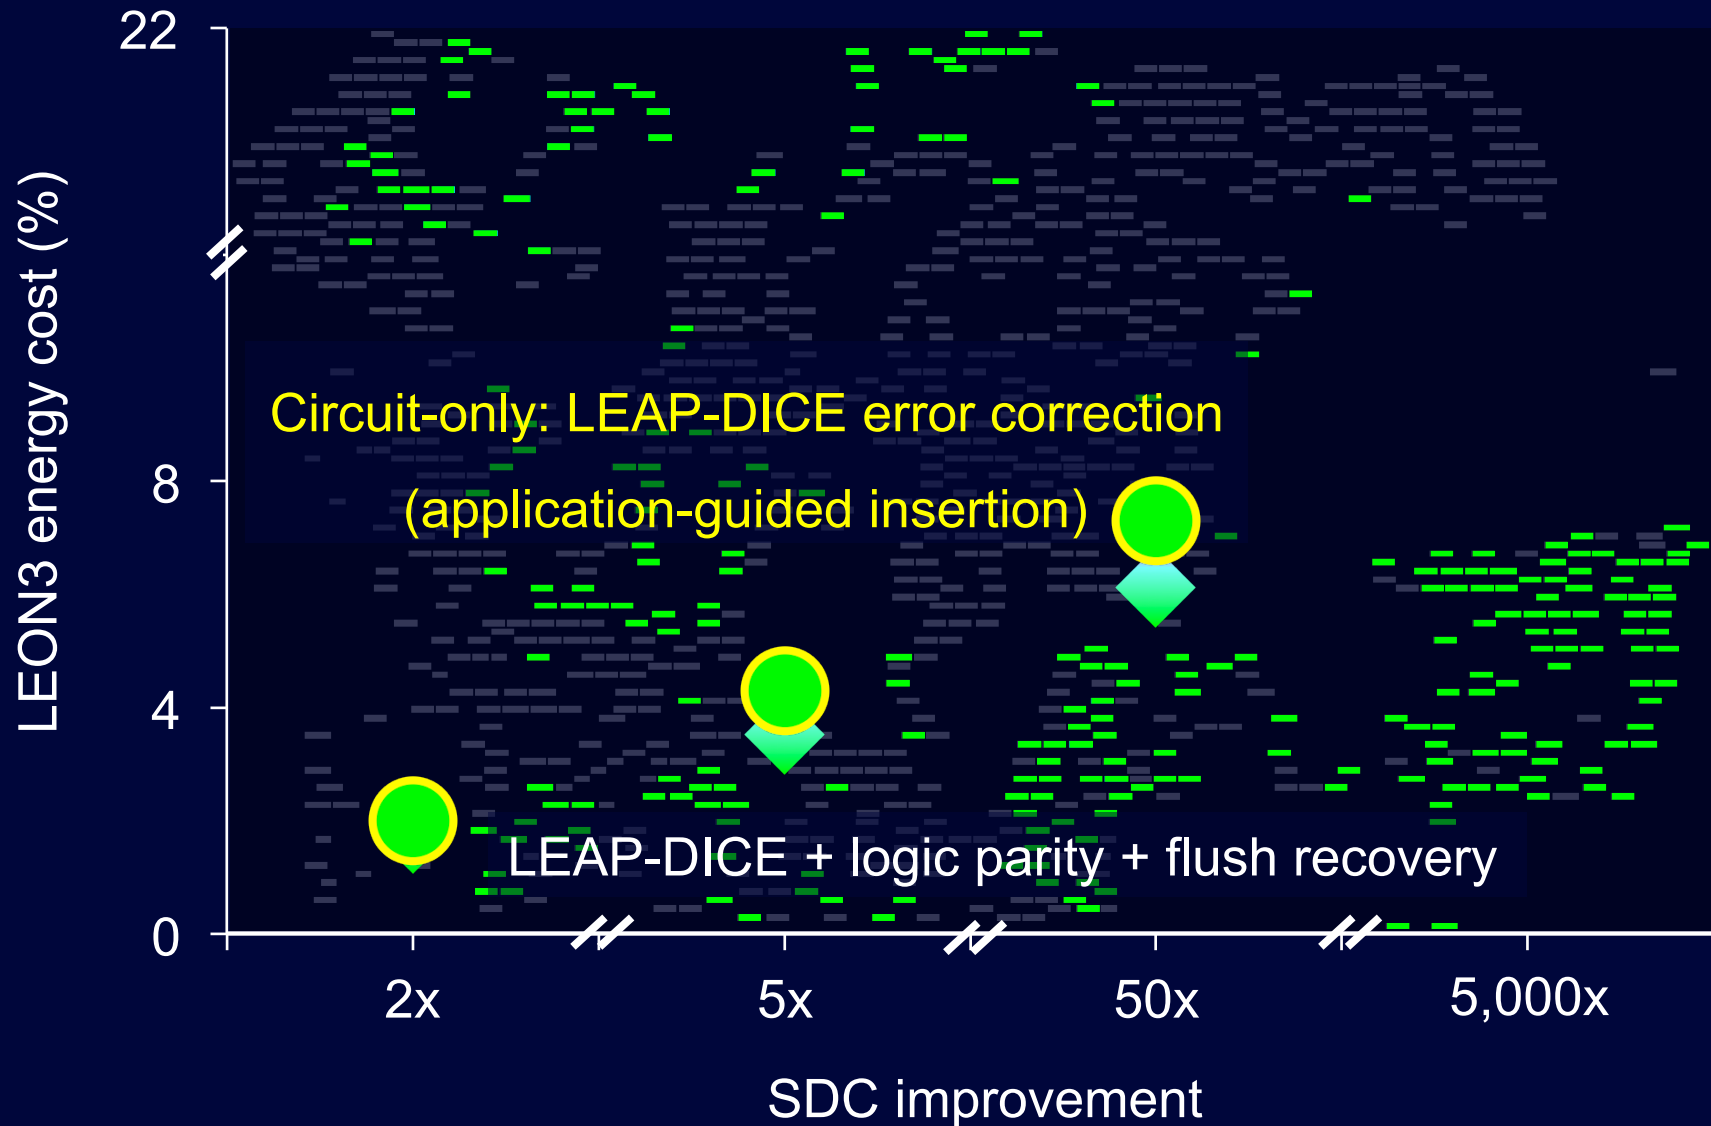
Circuit-only (Application-guided): Highly Effective

Algorithm
Software
Architecture
Logic
Circuit



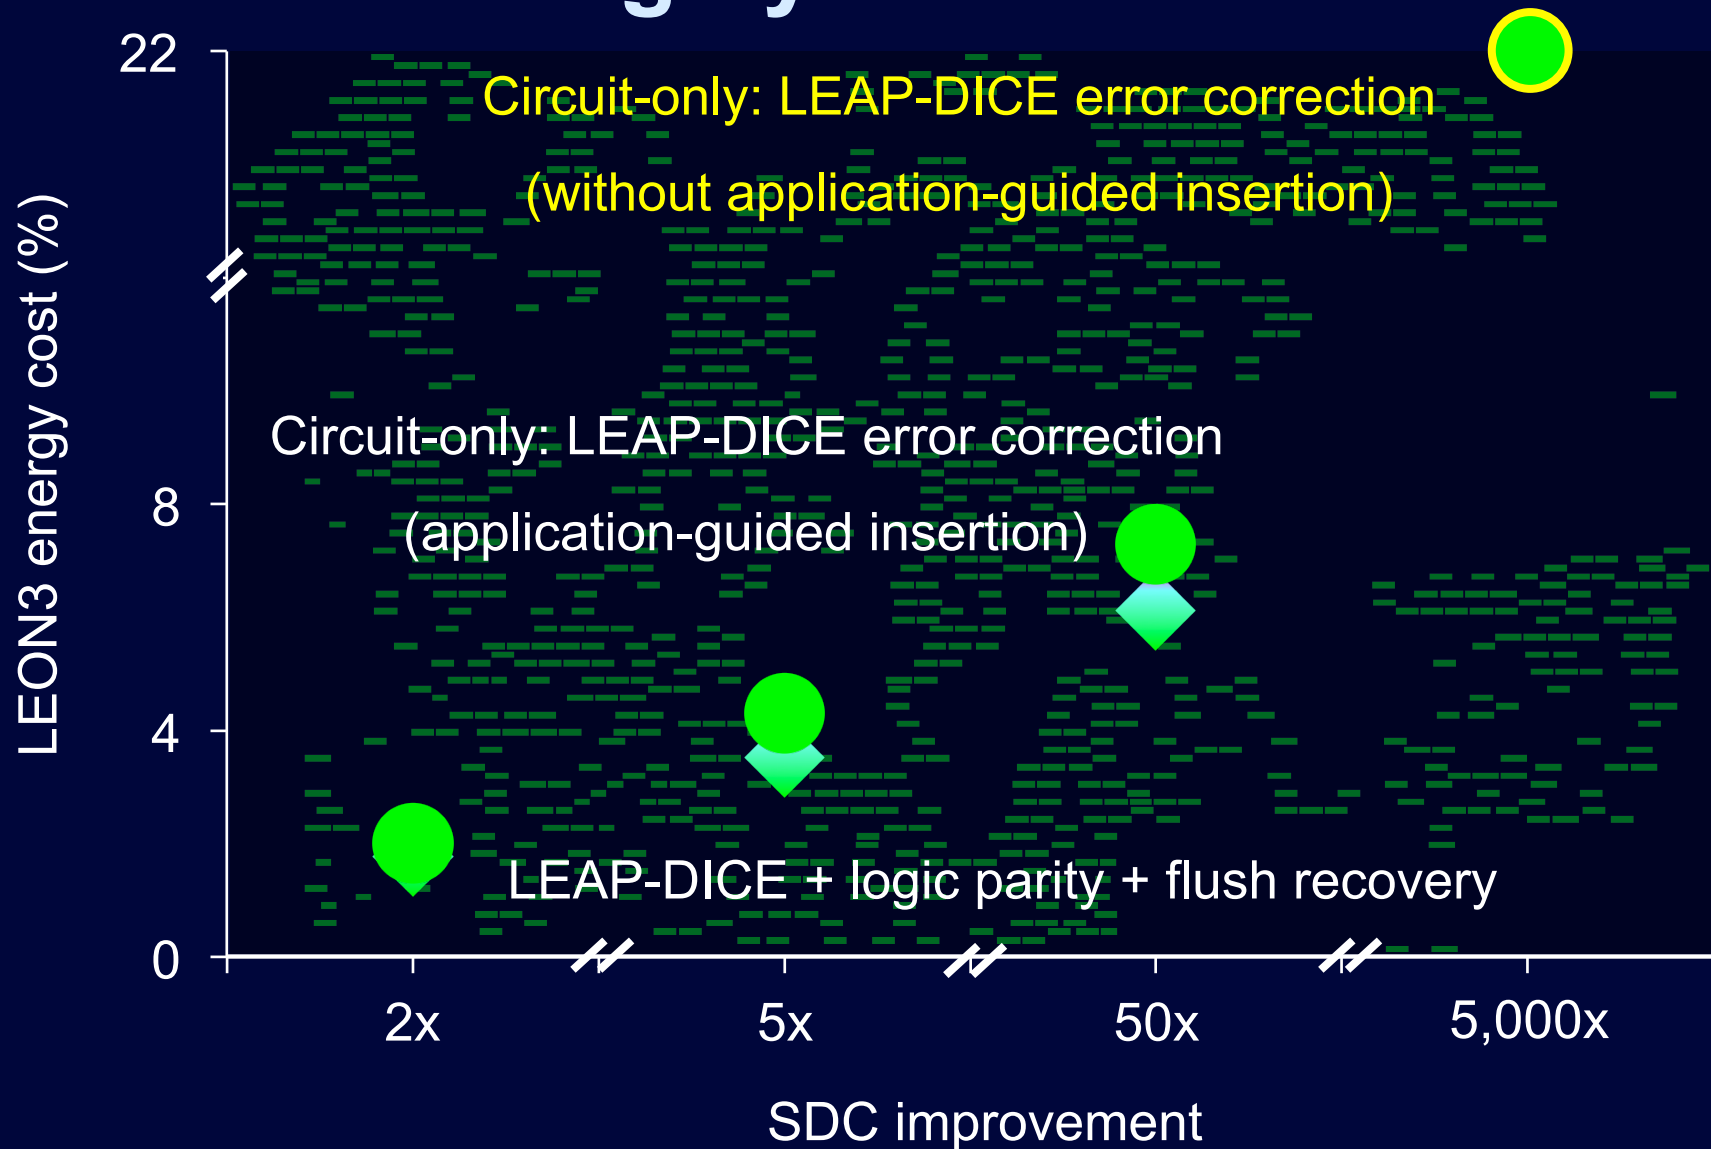
Circuit-only (Application-guided): Highly Effective

Algorithm
Software
Architecture
Logic
Circuit



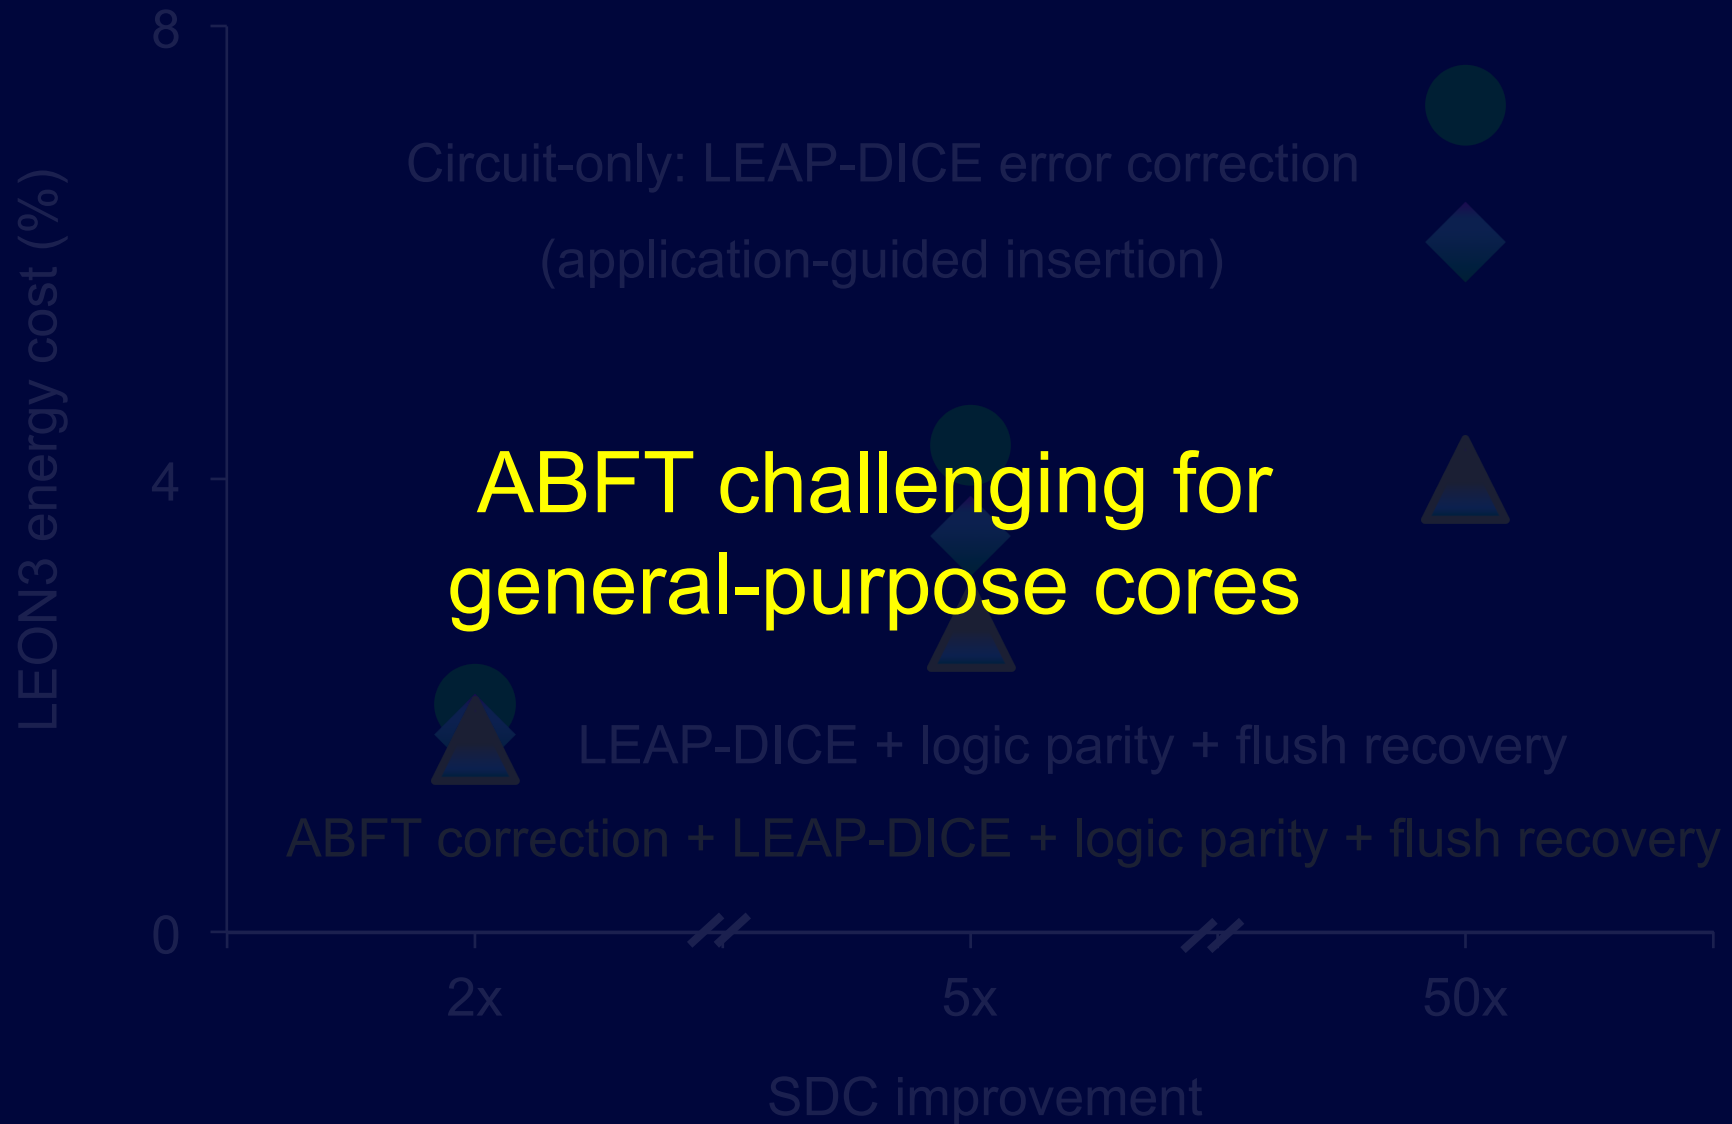
Circuit-only (Application-guided): Highly Effective

Algorithm
Software
Architecture
Logic
Circuit



What About ABFT?

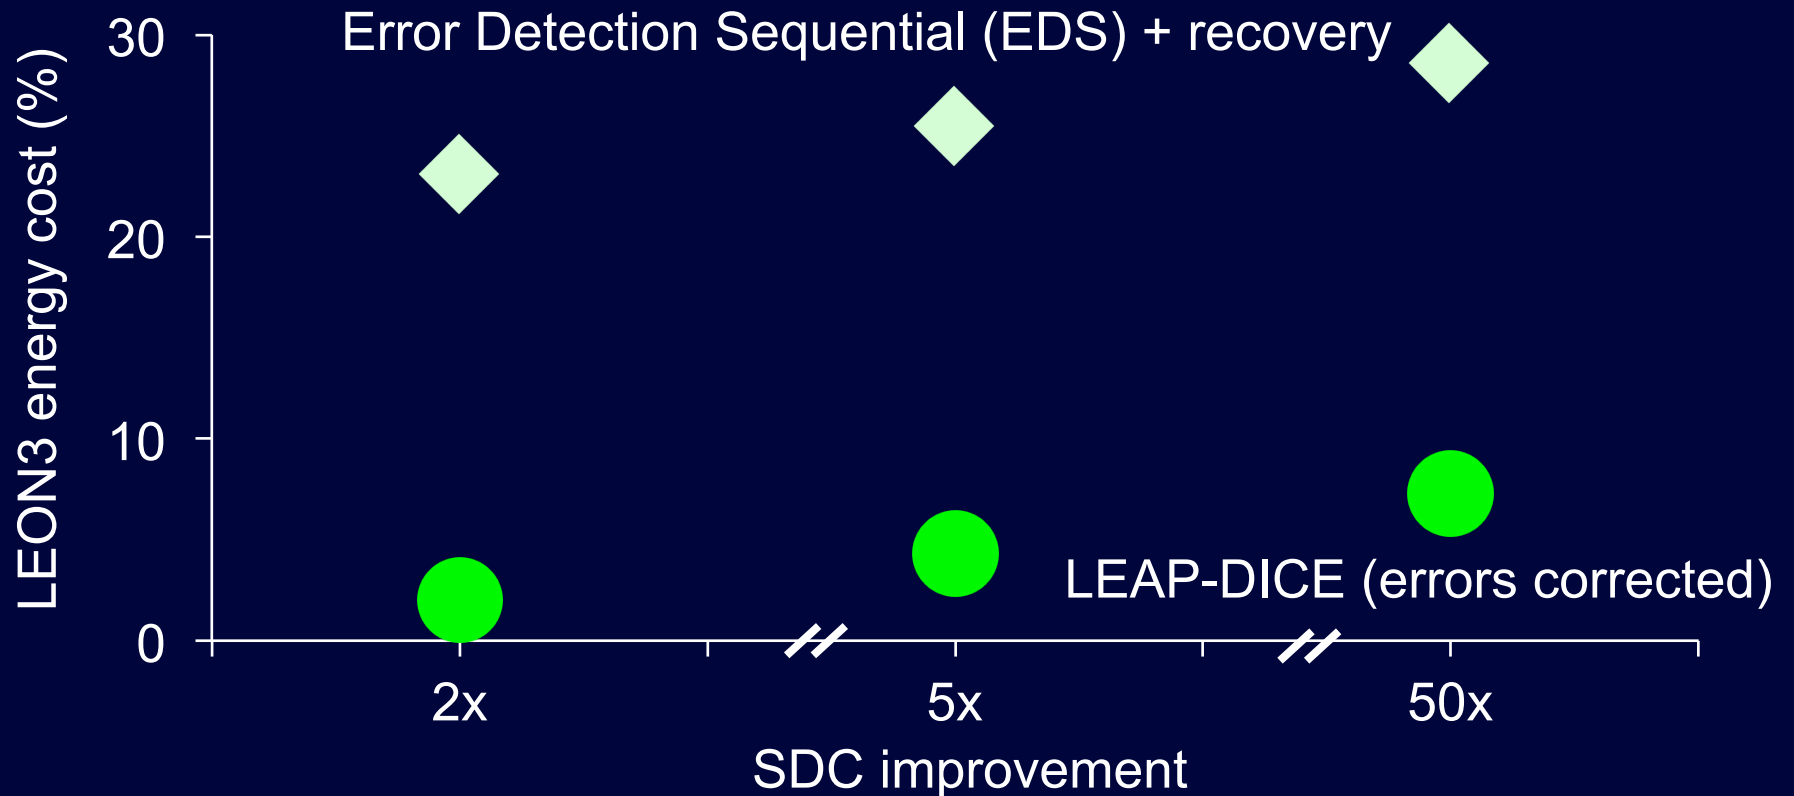
Algorithm
Software
Architecture
Logic
Circuit



CLEAR Insights

- Hidden costs & inefficiencies
- Implementation matters
- Inaccurate analysis

Example: “Hidden” Costs



Flip-flop	Area	Energy
Nominal	1	1
LEAP-DICE	2	1.8
EDS	1.5	1.4

Not just flip-flop overhead
Routing, recovery impact

Example: Inefficiencies

Few Flip-flops Protected

Data Flow Checking (DFC)
57%

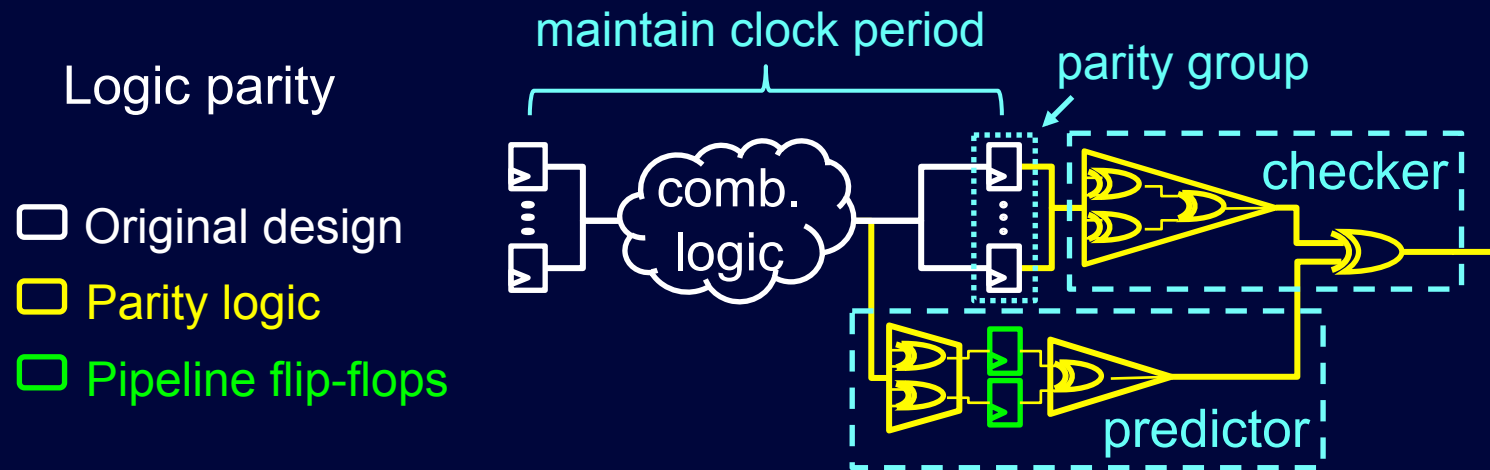
Low SDC Coverage per Flip-flop

Data Flow Checking (DFC)
30%

Result: Low SDC Improvement

Data Flow Checking (DFC)
1.2x

Example: Implementation Matters



Logic parity: Naïve	200 MHz clock speed impact
Logic parity: Incorrect heuristic	80% additional energy impact
Logic parity: CLEAR heuristic	No clock speed impact Minimal energy impact

Parameters: parity size, flip-flop vulnerability, floorplan location, timing path slack, etc.

Example: Inaccurate Analysis

- Software assertions: SDC improvement
 - Prior publications: 3.9x
 - Inaccurate error injection
 - Accurate analysis: 1.5x

	Flip-Flop error injection	Register Uniform error injection
SDC improvement	1.5x	4.8x

How About Benchmark Dependence?

- 50 *<training, evaluation>* pairs
 - Training: 4 SPEC, Evaluation: 7 SPEC

Trained SDC improvement	5x	50x	500x
Evaluated SDC improvement	4.8x	39x	433x




Add “lightweight” hardening
(e.g., LHL)

Extra energy cost (additive)	2%	1%	0.8%
Final SDC improvement	19x	152x	1,326x

Light-Hardened LEAP (LHL)

Flip-Flop	Soft Error Rate (SER)	Area	Power	Delay	Energy
Baseline	1	1	1	1	1
LHL	2.5×10^{-1}	1.2	1.1	1.2	1.3
LEAP-DICE	2×10^{-4}	2	1.8	1	1.8

Target for Future Resilience Techniques



Outline

- Robust operation: silicon CMOS reliability
- Beyond silicon
- Conclusion

T H E F U T U R E O F COMPUTING PERFORMANCE

Game Over or Next Level?

Samuel H. Fuller and Lynette I. Millett, *Editors*

Improve Computing Performance

System
integration




Device
performance



Option 1: Better Transistors

System
integration

- Few experimental demos
- Transistors \neq system






Device
performance






Option 2: Design Tricks

System
integration

Multi-cores



Power
management








- Limited “tricks”
- Complexity → design bugs

Device
performance

Improve Computing Performance

System
integration

Multi-cores



Device
performance


Target:
1,000× performance

New innovations required


Solution: NanoSystems

*Transform new nanotech
into new systems
enable new applications*


New devices




New fabrication



New sensors




New
Architectures




is?

abrication



Abundant-Data Explosion

“Swimming in sensors, drowning in data”




Wide variety & complexity



- Mine, search, analyze: near real-time
 - Data centers, mobile phones, robots

Today's System Bottlenecks

- **Separate** compute & memory chips
- Not enough on-chip memory
- Capacity & bandwidth critical





Abundant-Data Applications

Huge memory wall: processors, accelerators

Energy Measurements

Genomics classification

Natural language processing



■ ■ ■


■ Compute

■ Memory


Abundant-Data Applications

Huge memory wall: processors, accelerators

Deep Learning Accelerators



Nano-Engineered Computing Systems Technology



COVER FEATURE **REBOOTING COMPUTING**




Mohamed M. Sabry Aly, Mingyu Gao, Gage Hills, Chi-Shuen Lee, Greg Pittner, Max M. Shulaker, Tony F. Wu, and Mehdi Azabaghli, Stanford University
 Jeff Bolzer, University of California, Berkeley
 Franz Franchetti, Carnegie Mellon University
 Kenneth E. Goodson and Christos Kozmavakis, Stanford University
 Igor Markov, University of Michigan, Ann Arbor
 Kuntia Oskotun, Stanford University
 Larry Pileggi, Carnegie Mellon University
 Eric Pop, Stanford University
 Jan Rabney, University of California, Berkeley
 Christopher Ré, H.-S. Philip Wong, and Subhasish Mitra, Stanford University

Next-generation information technologies will process unprecedented amounts of loosely structured data that overwhelm existing computing systems. N3XT improves the energy efficiency of abundant-data applications 1,000-fold by using new logic and memory technologies, 3D integration with fine-grained connectivity, and new architectures for computation immersed in memory.

N3XT NanoSystems


Computation immersed in memory



N3XT NanoSystems

Computation immersed in memory

Increased functionality



Impossible with today's technologies

N3XT Computation Immersed in Memory

3D Resistive RAM
Massive storage


⋮

1D CNFET, 2D FET
Compute, RAM access

MRAM
Quick access

1D CNFET, 2D FET
Compute, RAM access

1D CNFET, 2D FET
Compute, Power, Clock




No TSV

Ultra-dense,
fine-grained
vias

Silicon
compatible

Carbon Nanotube FET (CNFET)

CNT: $d = 1.2\text{nm}$




Energy Delay Product

- $\sim 10\times$ benefit

Full-chip case studies

[IBM, IMEC, Stanford,
other commercial]


CNFET Inverter




Big Promise, Major Obstacles

- Process advances alone inadequate

Mis-positioned CNTs




Metallic CNTs




Imperfection-immune paradigm

CNT Growth *circa* 2005


- Highly mis-positioned




First Wafer-Scale Aligned CNT Growth




Quartz wafer
with catalyst




Aligned
CNT growth



Quartz wafer with CNTs



99.5% aligned CNTs




Stanford Nanofabrication Facility

Wafer-Scale CNT Transfer

High-temperature CNT growth


Low-temperature circuit fabrication



CNT transfer





120 °C



Before transfer




After transfer




Mis-Positioned CNT-Immune NAND

1. Grow CNTs



Mis-Positioned CNT-Immune NAND


1. Grow CNTs
2. Extended gate, contacts



Mis-Positioned CNT-Immune NAND


1. Grow CNTs
2. Extended gate, contacts
3. Etch gate & CNTs
4. Dope P & N regions

- Arbitrary logic functions
 - Graph algorithms




Imperfection-Immune VLSI

Mis-positioned CNTs



- Arbitrary logic functions

Metallic CNTs





- Scalable m-CNT Removal

Most Importantly

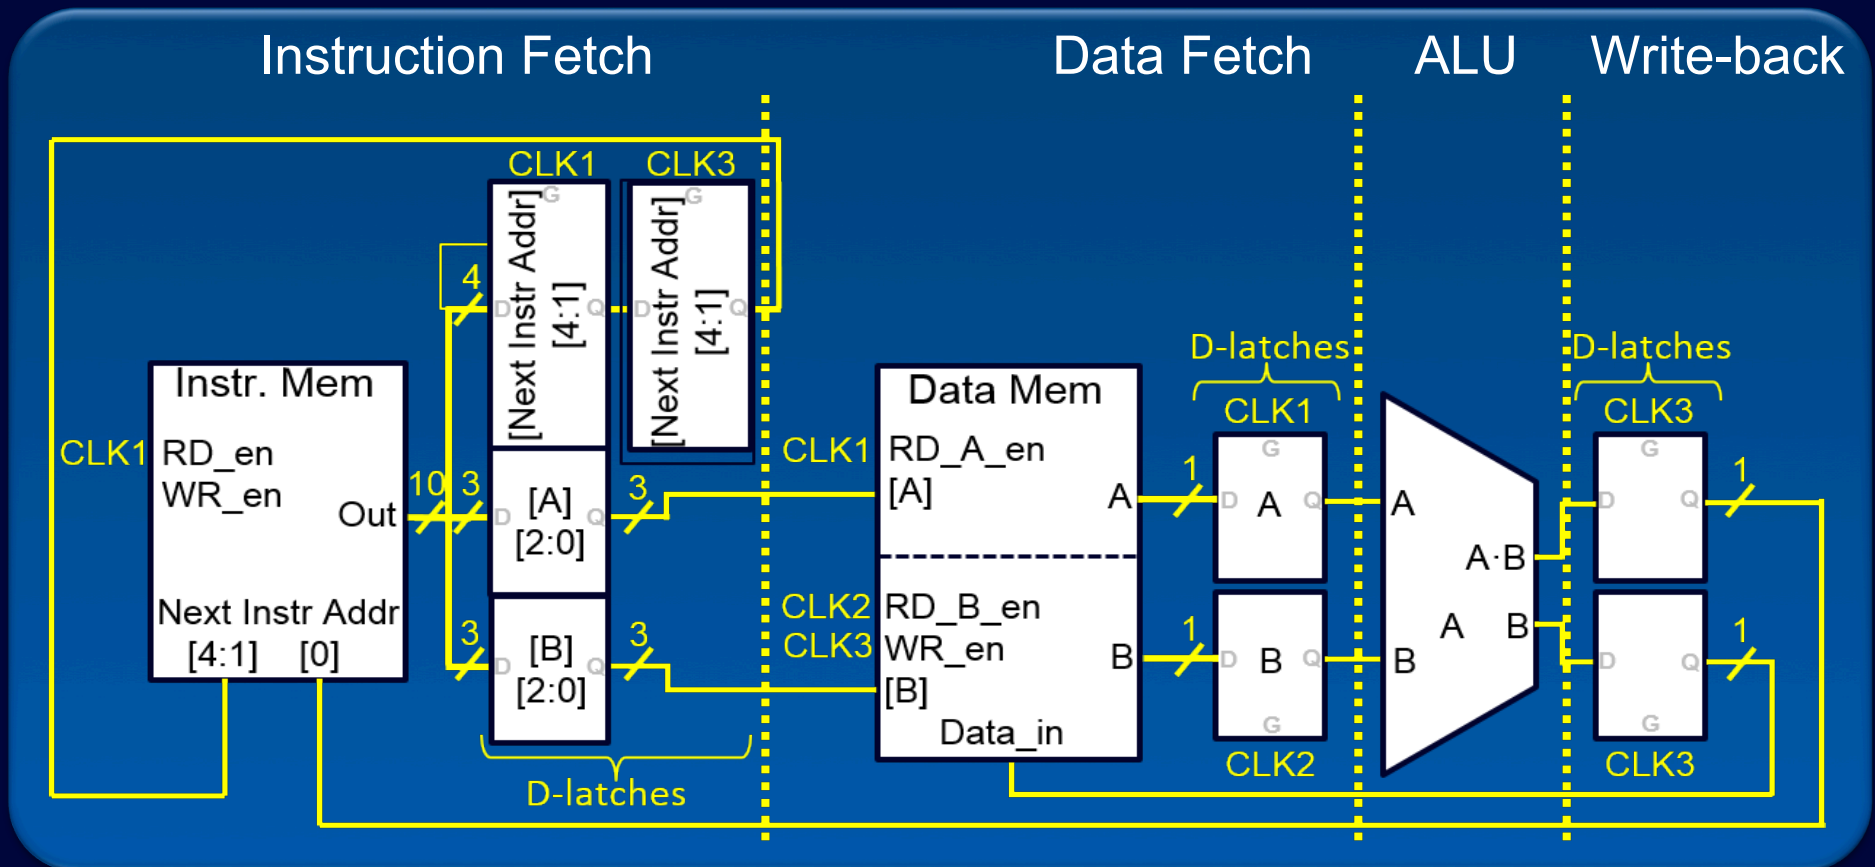
- VLSI processing
 - No per-unit customization
- VLSI design
 - Immune CNT library

CNT Computer



CNT Computer

- Turing-complete processor: entirely CNFETs



10× EDP, BUT...

How can we do better ?

N3XT Computation Immersed in Memory

3D Resistive RAM
Massive storage


⋮

1D CNFET, 2D FET
Compute, RAM access

MRAM
Quick access

1D CNFET, 2D FET
Compute, RAM access

1D CNFET, 2D FET
Compute, Power, Clock




No TSV

Ultra-dense,
fine-grained
vias

Silicon
compatible


Many Nano-scale Innovations

Memory & logic devices




3D Resistive RAM (RRAM)

MoS₂




$\downarrow \leq 1 \text{ nm}$
 \uparrow



2D FETs: large-area monolayer MoS₂

Embedded cooling



Phase change: hotspots suppressed




Vertical metal nanowire arrays

3D Integration


- Massive ILV density \gg TSV density

TSV (chip stacking)



Through silicon via
(TSV)

Dense, e.g., monolithic




Nano-scale
inter-layer vias (ILVs)

Realizing Monolithic 3D

- Low-temperature fabrication: **< 400 °C**


Device + Architecture Benefits

Naturally enabled




3D NanoSystem

Wafer-scale design + fabrication




3D NanoSystem

>2 Million CNFETs, 1 Mbit Resistive RAM



3D NanoSystem

- Interwoven compute + memory + sensing




3D NanoSystem

Abundant data: Terabytes / second



Millions of sensors


CNTs



Ultra-dense
vertical connections
x100,000

Memory


1 Megabit RRAM



CNT computing logic


classification accelerator

In-situ classification
Extensive, accurate



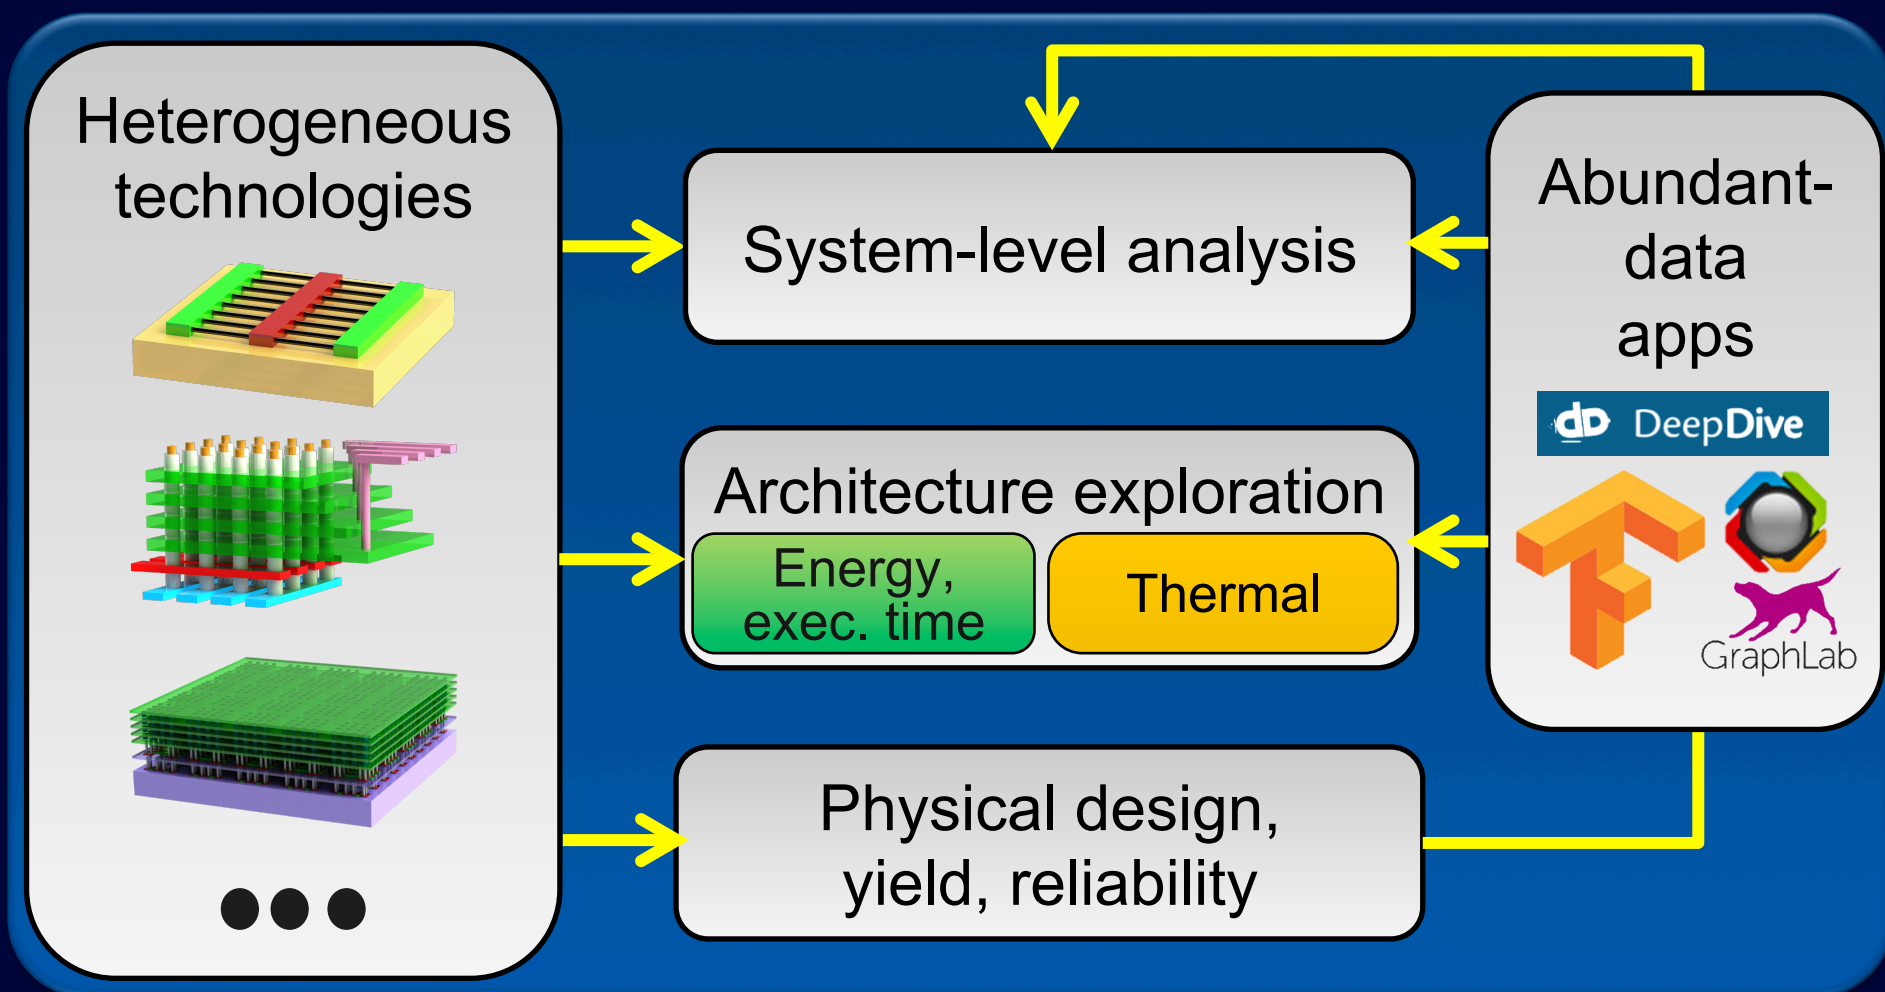
3D NanoSystem Results

Principle Component Analysis (PCA)



N3XT Simulation Framework

Joint technology, design & app. exploration



Massive Benefits: Deep Learning, Graph Analytics, ...



IBM graph analytics




DeepDive



2D

64 GB off-chip DRAM




64 processor cores

SRAM cache

Single-chip N3XT

64 GB on-chip 3D RRAM



“Simple”
interface

STTRAM
cache

64 processor cores

Massive Benefits: Deep Learning, Graph Analytics, ...




IBM graph analytics



DeepDive



~1,000× benefits, existing software



Massive Benefits: Deep Learning, Graph Analytics, ...



IBM graph analytics




DeepDive




851× benefits

Energy: 37×




Exec. Time: 23×




Massive Benefits: Deep Learning, Graph Analytics, ...

- 100× – 1,000× benefits (energy × execution time)
 - Deep learning accelerators



Chip stacking: 2 - 4× benefits

Complement with Software Solutions



**Co-optimized
s/w + h/w**



**Learning:
key
architectural
concept**

**Runtime
optimization**



**Cross-
Layer
Resilience**

**Yield,
reliability**

More Opportunities

Co-optimized hardware + software

Brain-inspired

Technology innovations



“Brain-Inspired Computing Exploiting Carbon Nanotube FETs and Resistive RAM: Hyperdimensional Computing Case Study,” ISSCC 2018.

Outline

- Robust operation: silicon CMOS reliability
- Beyond silicon
- Conclusion

Thanks to Research Group



Thanks to Sponsors & Collaborators



Conclusion

- Robust systems
 - New solutions: elegantly simple, effective

Silicon CMOS Reliability

BISER, LEAP

Failure prediction

CLEAR cross-layer

Beyond silicon

CNFET nanosystems

N3XT monolithic 3D

1,000X opportunity