

# QED & Symbolic QED


Subhasish Mitra



Department of EE & Department of CS  
Stanford University

Acknowledgment: Students & Collaborators

## Students, Sponsors, Collaborators



## Pre-Silicon Verification Inadequate



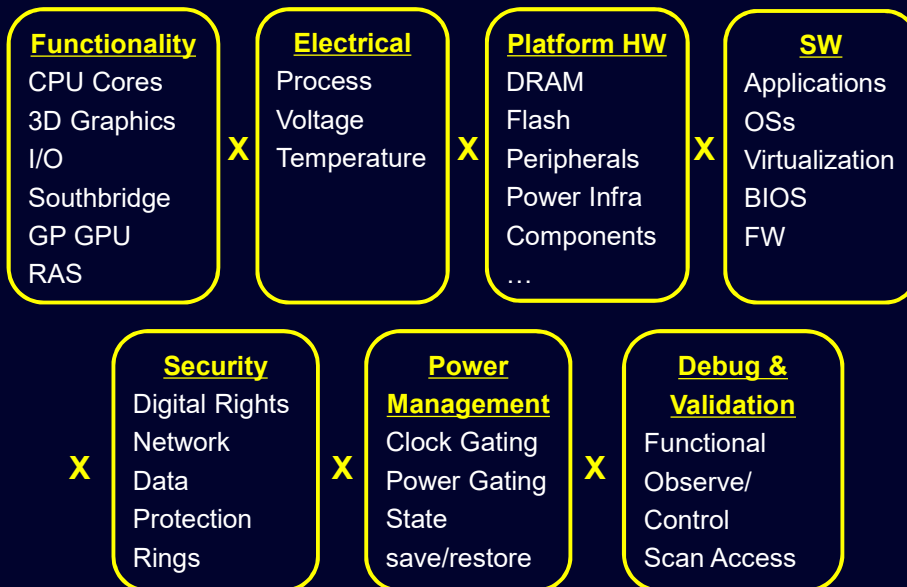
**It's only getting worse: custom hardware**



Source: Intel

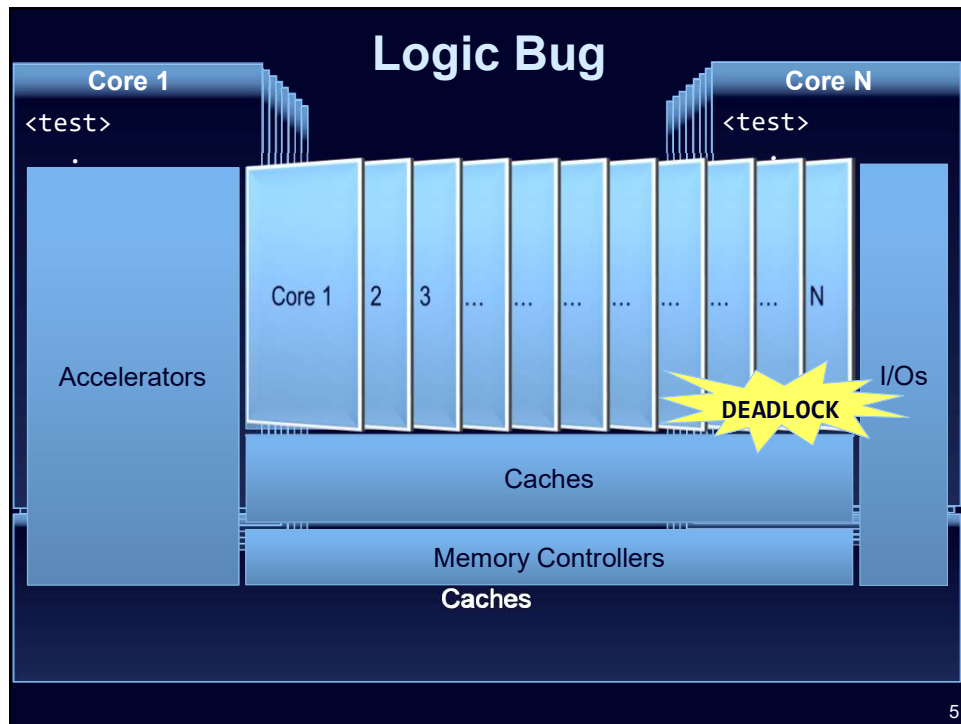
3

## Staggering Complexity



E. Rentschler, ex-AMD

4



## Electrical Bugs

- Speed-path, min-delay, noise, analog
  - Some P, V, T, F corners
  - “Heisenbug” characteristics

“Flakey” Shmoo

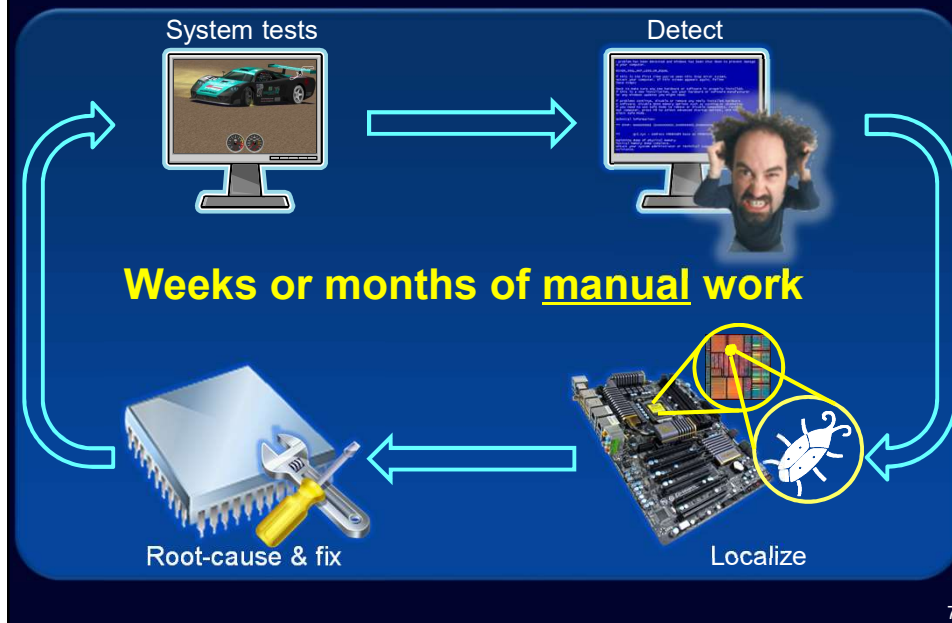
Clock period

Voltage droop

N. Hakim, Intel J. Stinson, ex-Intel

6

## Post-Silicon Validation Difficult



## Scalability Barriers

- System failure reproduction
  - Asynchronous I/Os, clock domains
- Full-system simulation for golden response
  - $10^6$  -  $10^9$ X slower than silicon



J. Stinson (ex-Intel)

Post-silicon costs rising faster than design cost

8

## Pre- vs. Post-Silicon



	Pre-Silicon	Post-Silicon
Debug & fix	Easier	Expensive
Internal signal access	Great	Limited
Electrical bugs	Model difficult	Already present
Platform interactions	Little (not real time)	Present
Speed	Slow simulation	Fast silicon

9

## Pre- vs. Post-Silicon vs. Test

		Pre-Silicon verification	Post-silicon validation	Manufacturing Test
TEST ATTRIBUTES	DELIVERY	Functional or pseudo-functional	TAPEOUT	PRODUCTION Scan-based
	VISIBILITY	White-box	Black-box	
	PASS/FAIL	Compare to Spec	Compare to Design	
	LOGIC	Functional (code, sequence)	Logic? Bug	Structural (stuck-at, bridge)
METRICS	TIMING	Timing analysis (STA, SSTA, ...)	Electrical Bug	Delay Fault (transition, path)

Prof. K.T. Cheng

10


## QED

- Post-silicon
  - Electrical bugs, logic bugs, system-level test



[Lin IEEE TCAD 14, Singh CAV 17]

11

## Error Detection Latency



12



## Existing Techniques Inadequate

- **Very long** error detection latencies

Deadlock detection

Design-specific assertions

Self-checking tests

Store readback tests

Checkpointing

etc...

15

## Quick Error Detection



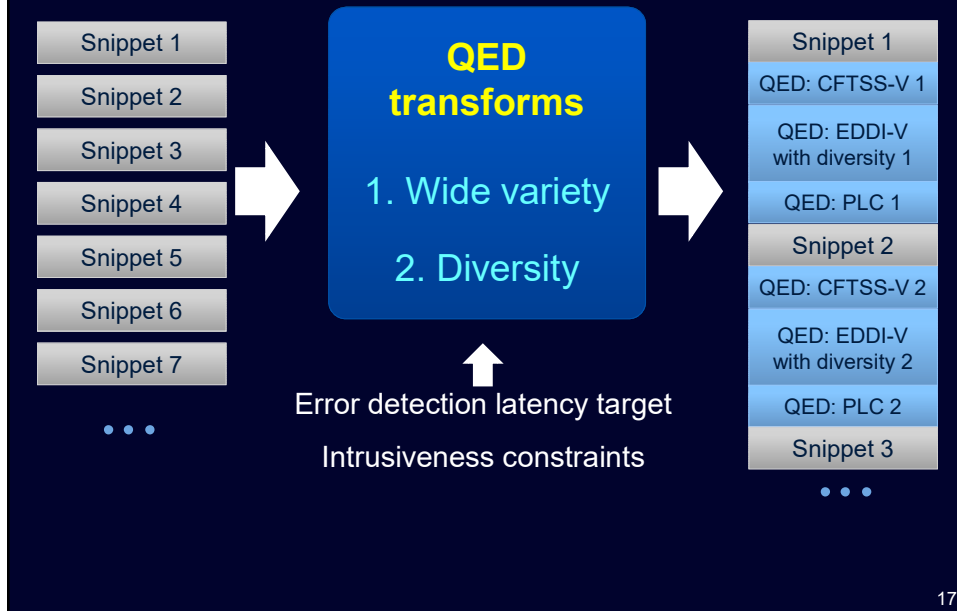
- Error detection latency: **guaranteed short**
- Coverage: **improved**
- Software-only: **readily applicable**

[Lin IEEE TCAD 14]

16

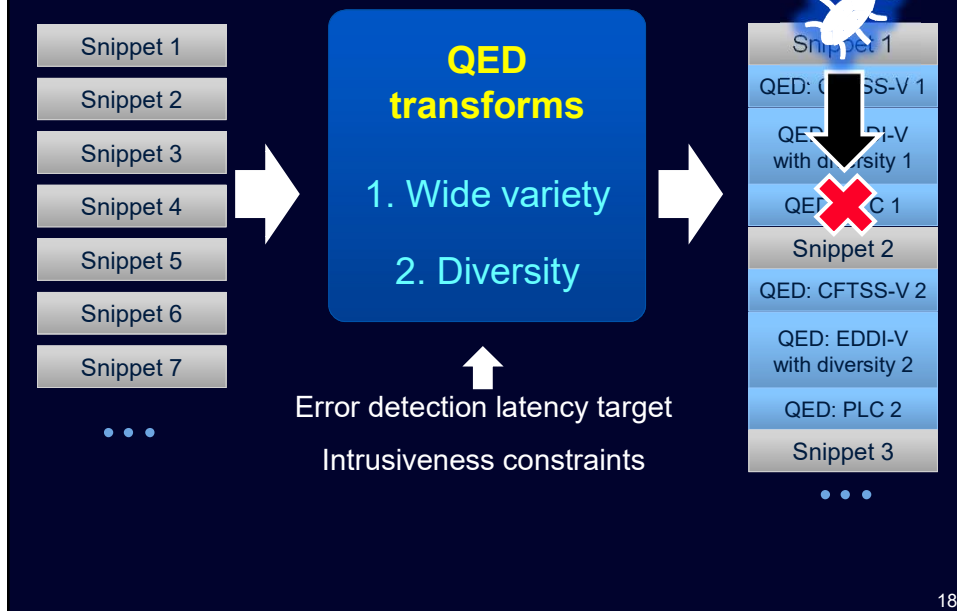


## Quick Error Detection



17

## Quick Error Detection



18



## QED Trace

21

## Validation program

## Trace

22

## QED Improves Coverage

## Validation program

[illegible]

## QED Trace

```

R1 ← R1 + 1
R18 ← R18 + 1
R1 == R18

```



```
R1 ← 0
```

• • •

Check R1 == 0 **masked!**

R1

R18

3 2	2
-----	---

## Error

## No QED: bug escape

QED: quick detection

23

## Diversity-Enhanced QED

## Validation program


[illegible]

## QED Trace

```
a ← 1; a' ← a; b ← 1, b' ← b
```

R1 ← a + b

$$R18 \leftarrow 5a' + 5b'$$

$5 \times R1 == R18$  



6 2

R1



14 :


R18

## Many diversity techniques

e.g., ED<sup>4</sup>I [Oh IEEE Trans. Comp. 02]

24

# QED Transforms



Processor cores

Uncore, accelerators

EDDI-V

PLC

CFTSS-V

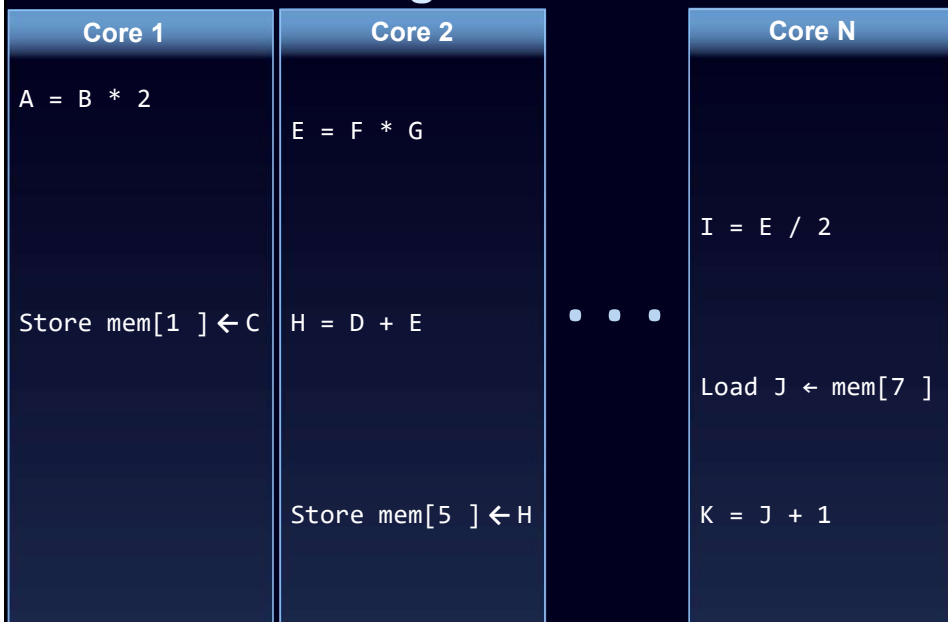
Fast QED

CFCSS-V

Hybrid QED

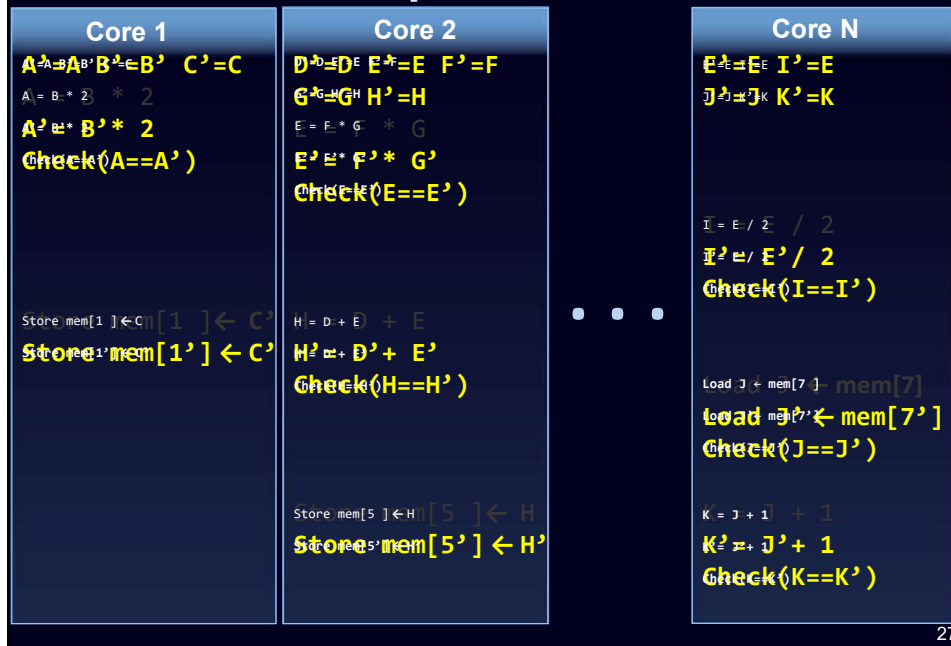
25

## Original Test

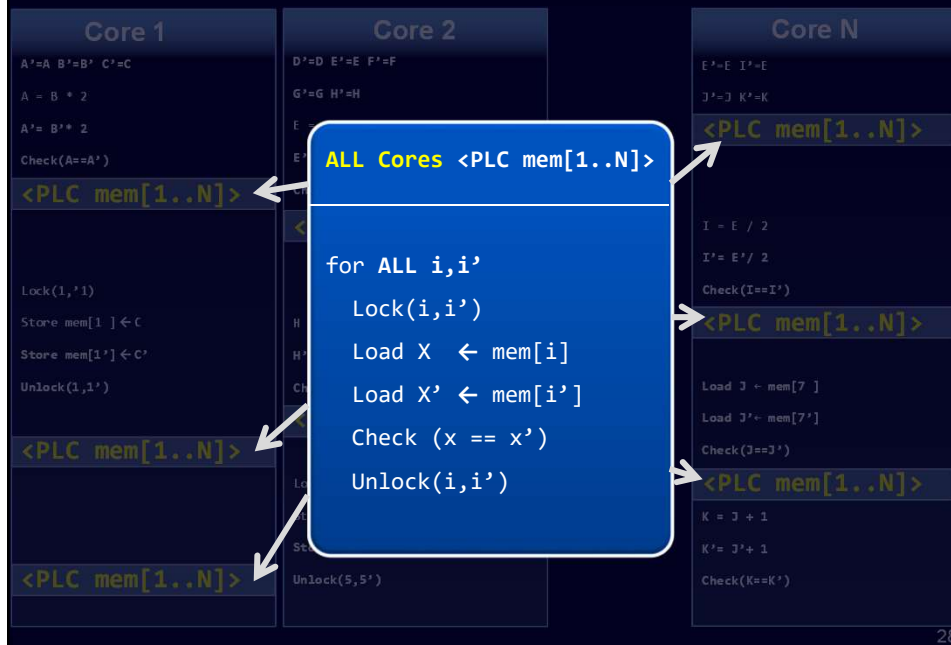


26

## QED: Duplicate & Check



## QED: Proactive Load & Check



## QED Coverage Considerations

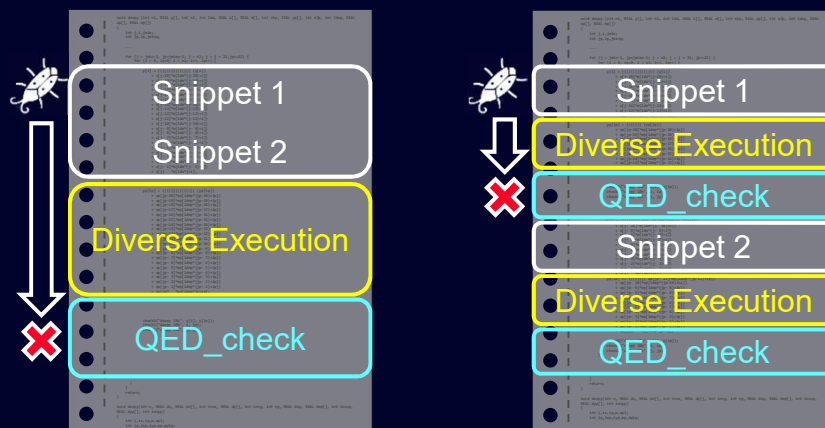
- Challenge
  - Intrusiveness: coverage impact ?
- Systematic solutions
  - QED family tests
  - Hardware-enhanced QED

29

## Error Detection Latency vs. Intrusiveness

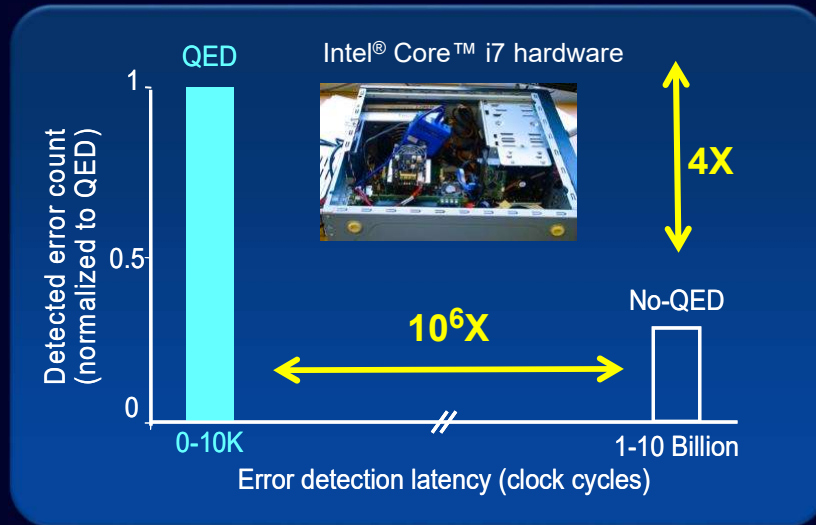
😊 Improved error detection latency

😊 Increased intrusiveness?



30


## QED Effective for Electrical Bugs



[Lin IEEE TCAD 14]

31

## Intel® Core™ i7 Hardware



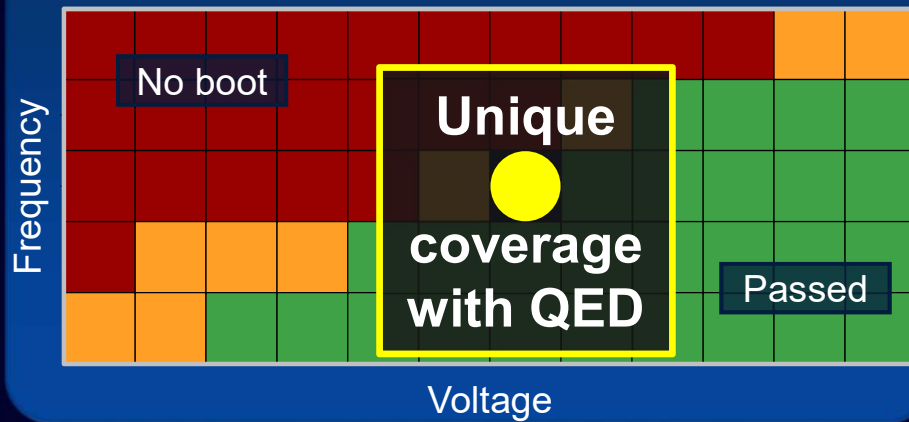
32



## Intel® Core™ i7 Hardware

■ Error detected (QED and no-QED)



● Error detected (QED only)



33

## Intel® SCC Hardware

48 processor cores  
0.9V, 800 MHz



● QED unique detect

● QED enhanced detect

● QED quick detect

34

## E-QED

Design RTL

```
// Memory Write Block
// Write Operation : When we_0 = 1,
cs_0 = 1
always @ (address_0 or cs_0 or
we_0 or data_0 or address_1 or
cs_1 or we_1 or data_1)
begin : MEM_WRITE
if (cs_0 && we_0) begin
mem[address_0] <= data_0;
end else if (cs_1 && we_1) begin
mem[address_1] <= data_1;
end end
```

Chip



No electrical bug or defect

Electrical bug / defect

Design

Recorded signals

Analyze inconsistencies

Flip-flop candidates (errors captured)

[Singh CAV 17]

35

## E-QED

Design phase: E-QED signature blocks

Test phase: Run QED tests

Localize error: formal methods

36

## E-QED Results

Flip-flop candidates	18 (out of 1 million) <b>50,000× localization</b>
Area impact	2.5% (actually 0%)
Effort	Automatic
Runtime	~ 9 hours
Failure reproduction	<b><u>None</u></b>
Full system simulation	<b><u>None</u></b>

OpenSPARC T2 SoC (500M transistors)

37

## QED Effective for Logic Bugs

### From Months to Seconds



Freescall  
multi-core SoC  
hardware

Error detection latency (cycles)


No QED	QED
<b>15 Billion</b>	<b>9</b>

[Lin IEEE TCAD 14]

38

## QED Effective for Logic Bugs

- “Difficult” logic bugs
  - Industrial bug databases
- Processor cores, accelerators, uncore, power management



[Lin IEEE TCAD 14]

39

## Activation Criteria × Effects


Activation Criteria
1. Two stores in X cycles
2. Two stores in X cycles to same cache line
3. Two stores in X cycles to adjacent cache line
4. Specific sequence of loads and stores
5. Two branches in X cycles
6. Forwarding data X to inst Y
7. Two icache misses in X cycles
8. Random

X

Bug effects
A. Cache coherence msg. dropped
B. Cache coherence msg. delayed
C. Store not allocated a cache line
D. Update delayed by Z cycles
E. Loaded value corrupted
F. Data in main memory corrupted
G. Data in L1\$ corrupted
H. Jump to wrong address
I. Error in instruction operand
J. Wrong instruction decoded

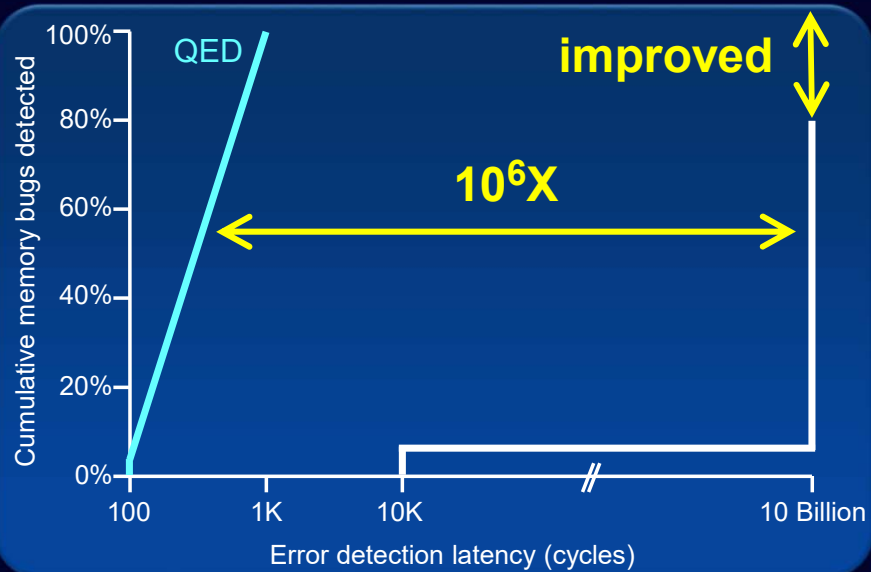
40

## 8-Core QED: Difficult Logic Bugs



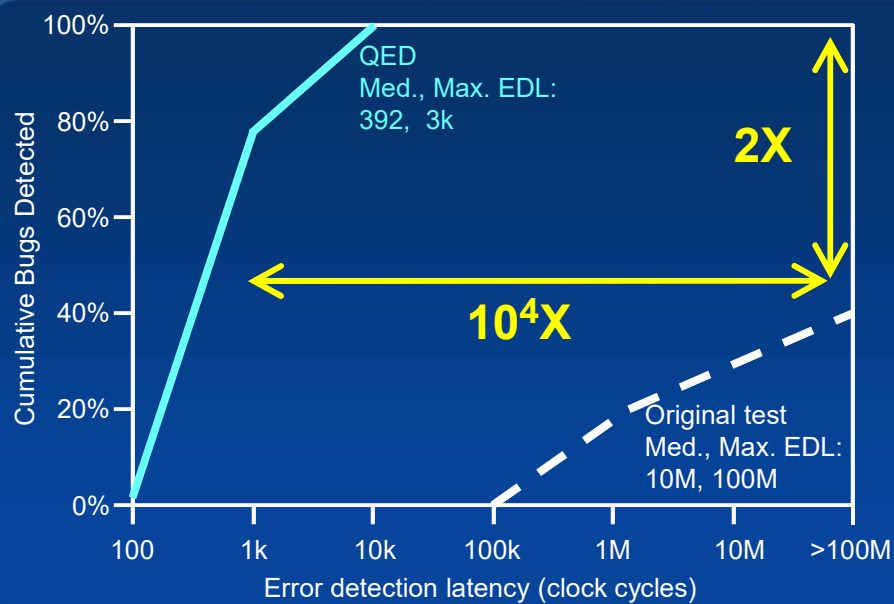
41

## 8-Core QED: Industrial Test



42

## 8-Core QED: Power Management Bugs



EDL = Error Detection Latency

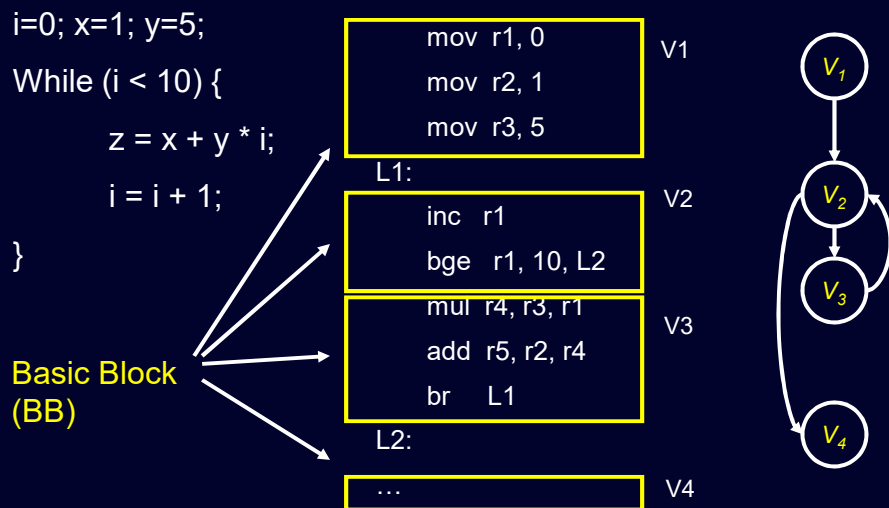
43

## QED Automation

- Application test, random test, ...
  - Source code, assembly, binary

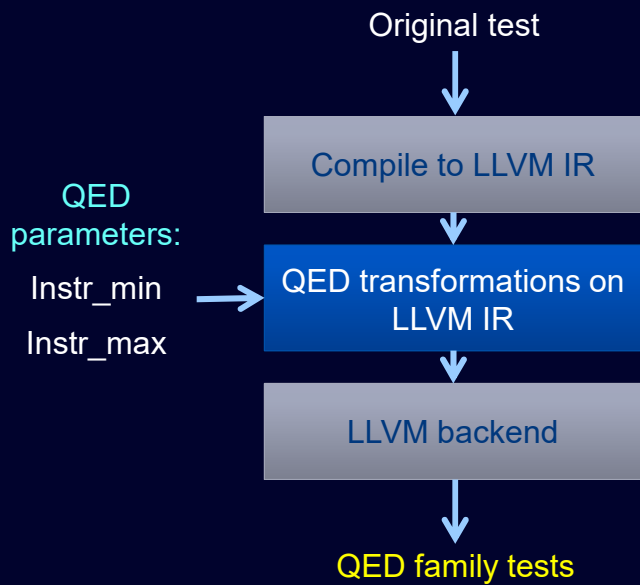
44

## Program Representation: Control Flow Graph



45

## QED Automation Example



IR = Intermediate Representation

46

## Symbolic QED

- Pre-silicon
  - Logic bugs: processors, accelerators, SoCs

[Lin ITC 15, Singh IEEE TCAD 18, Lonsing ICCAD 19]

47

## Traditional Bounded Model Checking

Property

Design



$\exists$  program of length  $\leq k$



Property violated ?

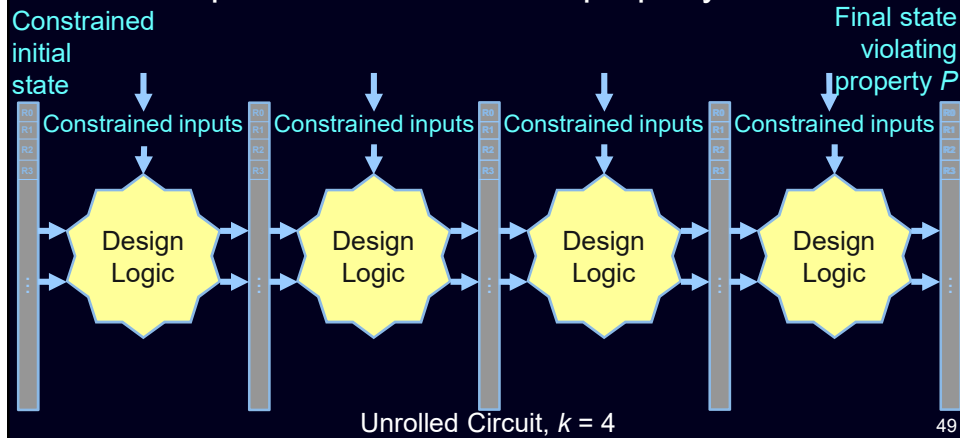
Bug found

48



## Bounded Model Checking (BMC)

- Input constraints, initial state constraints, property  $P$
- Exhaustively explore design for  $k$  clock cycles
- Return input values that violate property  $P$



## Traditional BMC vs. Symbolic QED

	Traditional BMC	Symbolic QED
Properties	Manual	<b>Automatic QED checks (Universal property)</b>
Design size	Small blocks	<b>Large SoCs</b>

## Traditional BMC Challenges

1. What property ?

2. Design size ?

51


## What Property ?

- Bugs not known *a priori*
- Manually written: time-consuming
- Automatic: challenging for difficult bugs
- Too many cycles to detect bugs

52

## Design Size ?

- Big designs, cannot load



OpenSPARC T2 SoC  
500 Million Transistors

53

## BMC Symbolic QED

Design



BMC Tool: *Bound = k*



$\exists$  QED program of length  $\leq k$  that fails ?



Bug found

54

## “Universal” Property: QED Check

**CMP Ra == Ra'**

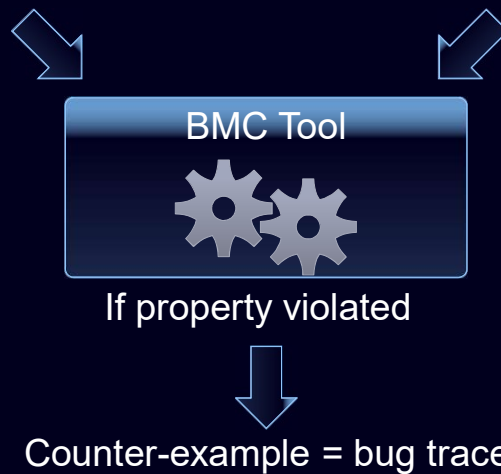
- Ra – original register
- Ra' – corresponding duplicated register
- $Ra \neq Ra'$  – error detected

55

## BMC Symbolic QED

“Universal” Property

Model



56

## Need Input Constraints

- If unconstrained, BMC may choose any inputs

$Ra \leftarrow 1$

$Ra' \leftarrow 2$

**CMP  $Ra == Ra'$**

- False fail – not a bug

57

## Input Constraints

- Original sequence then duplicated sequence
  - **CMP  $Ra == Ra'$  after both executed**
- Enforced by **QED module automatically**
  - **Only during BMC**
  - Not in fabricated design

58

## Solution: QED Module

Only during BMC, not in fabricated chip

Input  
(Chosen by BMC)

```
ST [0x10000], Ra
ST [0x10040], Rb
LD Rc, [0x10000]
```

**QED  
Module**

Output

```
ST [0x10000], Ra
ST [0x10040], Rb
LD Rc, [0x10000]
ST [0x20000], Ra'
ST [0x20040], Rb'
LD Rc', [0x20000]
<COMPARES>
```

59

## Solution: QED Module

Only during BMC, not in fabricated chip

**QED Module**

**Processor  
core  
fetch unit**

- Automatically duplicates instructions
- Determines when QED checks happen

60

## BMC Symbolic QED

“Universal” Property

Model + QED Module  
(no hardware overhead)



If property violated

Counter-example = bug trace

61

## Starting State Important

- If unconstrained, BMC may choose any initial state

$Ra = 1 \quad Ra' = 2 \quad // \text{initial state}$

$Ra = Ra + 1; Ra' = Ra' + 1; \text{CMP } Ra == Ra'$

- QED-consistent starting state

- Run “simple” QED test

- More sophisticated

- [Fadiheh DATE 18, Devarajegowda DATE 20]

62

## BMC Symbolic QED

“Universal” Property + QED-consistent starting state      Model + QED Module (no hardware overhead)



If property violated



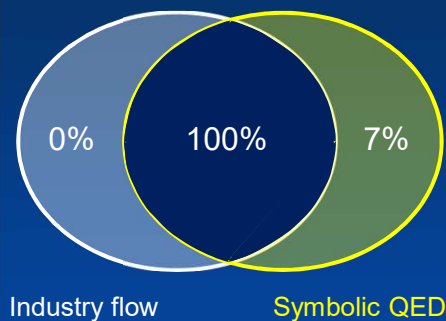
Counter-example = bug trace

63

## Symbolic QED: Infineon Study

- Several automotive microcontroller cores

All (known) bugs + more



60× productivity



[Singh DATE 19]

64



## Symbolic QED Study: A Major Company

- RISC V-based SoC (late-stage design)

### 1. New corner-case bugs found

**Previously unknown**

(despite countless simulations, emulation)

### 2. All “tricky” synthetic bugs correctly analyzed

**No spurious counterexamples**


### 3. Symbolic QED setup: 3 person days only

65

## Symbolic QED for Accelerators

- Stand-alone hardware accelerators for AI

**All (known) bugs + more**



**30× productivity**



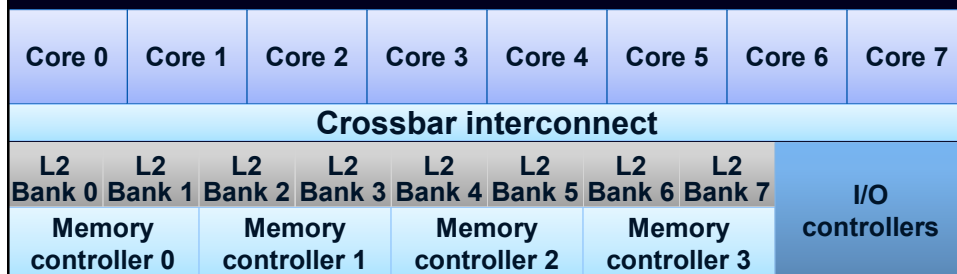
66

## BUT...

- Big designs ?
- **Solution: Partial instantiation**

67

## Solution: Partial Instantiation




68

## Solution: Partial Instantiation

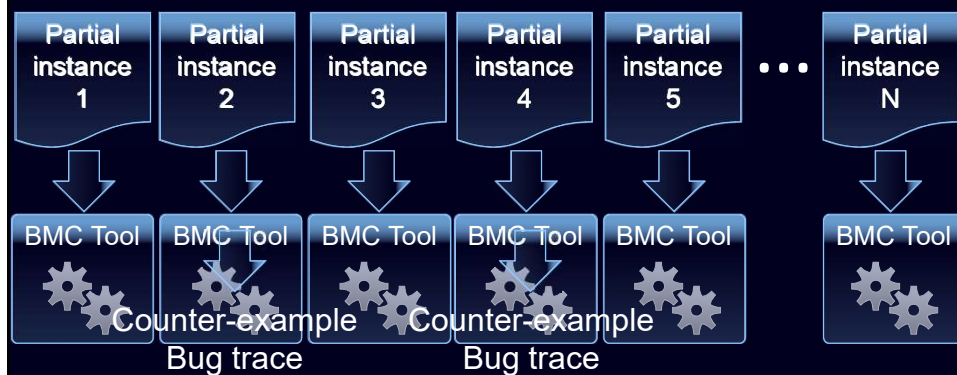
**Partial  
instances**

- At least 1 processor core per partial instance




69

## BMC on Partial Instances



70

## BMC on Partial Instances

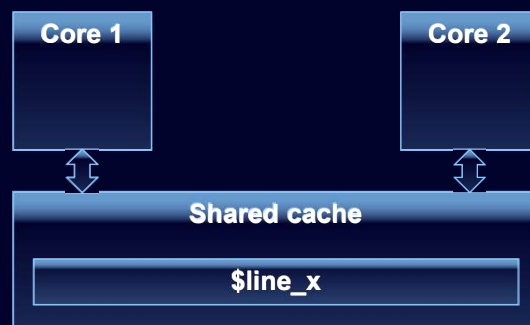


- Smallest partial instance → best localization

71

## Traditional Properties Not Compositional

- Not preserved across all partial instances




Core 1 writes \$line\_x ##1 Core 2 reads \$line\_x →  
\$line\_x in shared state

72

## Traditional Properties Not Compositional

- Not preserved across all partial instances



Core 1 writes \$line\_x ##1 **Core 2 reads \$line\_x** →  
\$line\_x in shared state

- Does not work if Core 2 missing

73


## QED Checks Compositional

- Preserved across partial instances
- Not design- / implementation-specific

74

## Case Study: OpenSPARC T2 SoC

- 92 difficult logic bugs
  - Commercial multi-core SoCs
  - Core, uncore, power management bugs



OpenSPARC T2 SoC  
500 Million Transistors


75

## 100X Reduction in Time

Traditional	<u>Automatic</u> Symbolic QED
Weeks to months (manual)	20 mins. to 7 hours

76

## Symbolic QED Runtime



77

## $10^6$ X Reduction In Bug Trace Length

Traditional	<u>Automatic</u> Symbolic QED
Millions of cycles	Less than 30 cycles

- Short bug trace
  - Easier to understand
  - Easier to fix bugs

78

## More Opportunities

- Hardware security
  - **Derive new attacks:** beyond Spectre, Meltdown
  - Trojans
- Firmware
- Large-scale systems


[Fadiheh DATE 19]

79

## Big Obstacles at Full System Level

1. Faulty chip localization

2. System-Level Testing (No Trouble Found)



Z. Conroy, Cisco

80



## QED & Symbolic QED

- Pre-silicon, post-silicon
  - Automatic, overnight, billion-transistor SoCs
- Widely applicable
  - Core, uncore, accelerator, bugs, defects
- More opportunities
  - Security, full systems, System-level Testing

81