

E-QED: Electrical Bug Localization During Post-silicon Validation Enabled by Quick Error Detection and Formal Methods

Eshan Singh^(✉), Clark Barrett, and Subhasish Mitra

Stanford University, Stanford, USA
{esingh, clarkbarrett, subh}@stanford.edu



Abstract. During post-silicon validation, manufactured integrated circuits are extensively tested in actual system environments to detect design bugs. Bug localization involves identification of a bug trace (a sequence of inputs that activates and detects the bug) and a hardware design block where the bug is located. Existing bug localization practices during post-silicon validation are mostly manual and *ad hoc*, and, hence, extremely expensive and time consuming. This is particularly true for subtle electrical bugs caused by unexpected interactions between a design and its electrical state. We present E-QED, a new approach that automatically localizes electrical bugs during post-silicon validation. Our results on the OpenSPARC T2, an open-source 500-million-transistor multicore chip design, demonstrate the effectiveness and practicality of E-QED: starting with a failed post-silicon test, in a few hours (9 h on average) we can automatically narrow the location of the bug to (the fan-in logic cone of) a handful of candidate flip-flops (18 flip-flops on average for a design with ~ 1 Million flip-flops) and also obtain the corresponding bug trace. The area impact of E-QED is $\sim 2.5\%$. In contrast, determining this same information might take weeks (or even months) of mostly manual work using traditional approaches.

1 Introduction

For complex integrated circuits (ICs), difficult design flaws (bugs) increasingly escape pre-silicon design verification and are only detected during post-silicon validation when manufactured ICs are extensively tested in actual system environments (Foster 2015). Design bugs can be broadly classified into *logic bugs* that are caused by (logic) design errors and *electrical bugs* that are caused by unexpected interactions between a design and its electrical state. Examples include errors introduced by crosstalk, power-supply noise, thermal effects or process variations. Traditional pre-silicon verification is slow; but, more importantly, it is generally incapable of detecting electrical bugs that appear only after ICs are manufactured. **This paper focuses on electrical bugs.**

Typical post-silicon validation involves: 1. bug detection by applying a variety of test stimuli (e.g., random instruction tests, end-user applications) at various voltage,

This work is supported in part by DARPA and the Semiconductor Research Corporation (SRC).

temperature, and clock frequency corners; 2. *bug localization* which identifies a *bug trace* (a sequence of inputs that activates and detects the bug) and a hardware design block where the bug is located; and, 3. bug fixing using techniques such as software patches, clock frequency/operating voltage selection, or silicon respin. Existing post-silicon validation and debug practices are mostly manual and *ad hoc*, and, hence, very expensive (Mishra et al. 2017). The effort to localize bugs from observed system failures (e.g., crashes, output errors) dominates the overall cost (Dusanapudi et al. 2015; Friedler et al. 2014; Nahir 2014). For example, it might take weeks (or even months) of (manual) work to localize a single bug (Dusanapudi et al. 2015; Reick 2012; Vermeulen and Goossens 2014).

Post-silicon bug localization is difficult because of long error detection latencies. *Error detection latency* is the time elapsed between when a test activates a bug and creates an error and when that error manifests as an observable failure (e.g., system crash). Error detection latencies of difficult bugs can exceed several millions or even billions of clock cycles (Hong et al. 2010; Lin et al. 2014). It is extremely difficult to trace that far back into the history of system operation for complex ICs. In addition, IC design size and complexity pose major challenges. Full-chip simulation to obtain expected responses (for various internal states, not just software-visible states) is several orders of magnitude slower than actual silicon and may be impractical. Formal analysis and Boolean Satisfiability techniques can be severely limited by design size. System-level failure reproduction, which involves returning the system to an error-free state and re-running the system (perhaps with some modifications) to reproduce the “exact” failure, is difficult (due to non-deterministic behaviors such as asynchronous I/Os and multiple-clock domains). In order to limit the number of cycles that must be traced and analyzed during bug localization, techniques such as Quick Error Detection (QED (Lin et al. 2014)) that ensure short error detection latencies are crucial. An overview of existing bug localization approaches is presented in Sect. 4.

New techniques are essential to overcome post-silicon bug localization challenges. There have been some recent publications that address detection and localization of logic bugs (e.g. (Lin et al. 2015)) during pre-silicon and post-silicon validation. Here, we present E-QED, a new technique to automatically localize electrical bugs during post-silicon validation and debug. Key features of E-QED are: 1. It is broadly applicable to most digital designs. 2. It can localize electrical bugs inside processor cores as well as in uncore components (interconnection networks, cache controllers, memory controllers) that occupy large portions of System-on-Chip (SoC) designs. 3. It does not require manual intervention during design or during post-silicon validation and debug. 4. It does not rely on full system-level simulation. 5. It scales to large designs.

We demonstrate the effectiveness and practicality of E-QED using OpenSPARC T2, a 500-million-transistor open-source SoC. Our results (details in Sect. 3) show that: 1. E-QED correctly and automatically localizes electrical bugs in a few hours (between 7–13 h). Such bugs would generally take weeks (or even months) of manual work to localize using traditional approaches. 2. E-QED achieves very fine-grained electrical bug localization. For each localized electrical bug, using formal analysis, E-QED automatically generates a small list of candidate flip-flops (FFs) that might have captured error(s) caused by the bug. For the OpenSPARC T2 SoC with ~ 1 Million FFs, E-QED automatically localizes electrical bugs to only 18 candidate FFs on

average. Thus, E-QED achieves a *localization factor* (total number of FFs in a design divided by the total number of candidate FFs that an electrical bug is localized to) of over 50,000. 3. For each localized electrical bug, E-QED automatically generates a short bug trace using formal analysis. 4. E-QED incurs only a small area overhead ($\sim 2.5\%$ for OpenSPARC T2) and has practically no clock-speed impact. 5. E-QED enables flexible trade-offs between area overhead, electrical bug localization granularity, and bug localization runtime. For example, with 1.5% area overhead (vs. 2.5%) the average FF candidate count increases only by a factor of 3 (i.e., E-QED still achieves a localization factor of four orders of magnitude).

E-QED uses the following three steps that work together in a coordinated fashion to overcome electrical bug localization challenges: 1. Low-cost hardware structures called E-QED signature blocks are automatically inserted during the design phase (Sect. 2.1). 2. QED tests that achieve short error detection latencies are run during post-silicon validation.¹ 3. Formal techniques are used to reason about the signatures collected by the E-QED signature blocks, automatically localizing bugs to a handful of candidate FFs (Sects. 2.2–2.4).

Motivating Example

Consider the following electrical bug example for the OpenSPARC T2 SoC (OpenSPARC), shown in Fig. 1. Suppose that a bug occurs during the following sequence of events: (1) Processor core 0 writes the value 0 to the crossbar (in order to store the value 0 in address [C]); (2) the entry corresponding to address [C] in L2 cache bank 0 is updated with value 0 (from the crossbar); (3) in response to a request to load

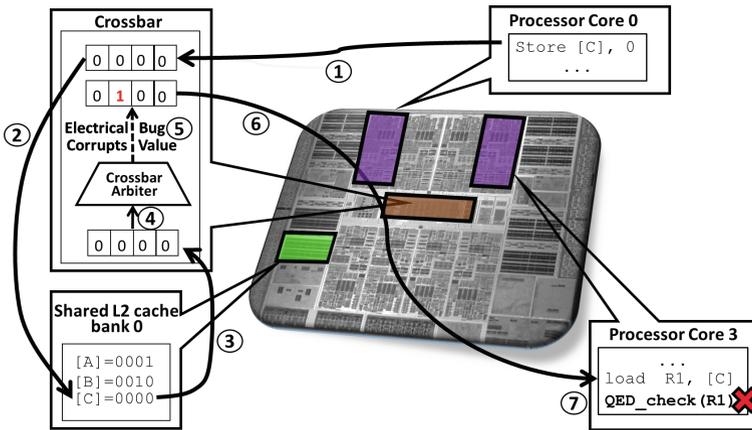


Fig. 1. An example of an electrical bug corrupting the value (corresponding to address [C]) stored by processor core 0 to L2 cache bank 0 (steps 1–2) as it passes through the arbitration logic of the crossbar (steps 3–5) while being loaded by processor core 3 (step 6). The bug is detected by a QED test (step 7).

¹ We use Quick Error Detection (QED) tests that typically achieve error detection latencies of 1,000 clock cycles or fewer (Lin et al. 2014); however, our approach can work with other tests that achieve similar error detection latencies.

from address [C] by processor core 3, the cached value corresponding to address [C] is written into the crossbar; (4) when this value passes through the crossbar arbitration logic, (5) an electrical bug causes a single-bit error to be captured in the output register of the crossbar arbiter; (6) the corrupted value is loaded into core 3; and then, (7) the corrupted value is detected by a QED test.

Note that without QED, the error may not be detected by the post-silicon validation test or the error detection latency can be millions of clock cycles; e.g., the error may be detected during an end-result check (which checks for expected output values upon test completion). Using a QED test, the error detection latency improves (i.e., reduces) to only a few hundred cycles. This drastically reduces the amount of data that needs to be analyzed to localize the bug. However, upon error detection by the QED test in processor core 3, it is impossible to directly determine where the error actually occurred and consequently how to localize the bug to a specific design block. An error in the datapath during any of steps 1–6 (in processor core 0, the crossbar, L2 cache bank 0, or processor core 3) would have been detected in the same way by the QED test running on processor core 3 (using the QED check in Fig. 1, step 7).

E-QED automatically localizes this bug not only to the crossbar (containing 40,000 of the nearly 1 million FFs in the design), but to a subset of 8 candidate FFs within the crossbar that could have captured this single-bit error. The 8 candidates include the FF in the output register that did capture the actual error. E-QED also provides a 544-Kbit bug trace, the sequence of inputs to the crossbar that triggered the bug during post-silicon validation. The total runtime required to obtain this result by running E-QED is 488 min, and the hardware area overhead is 2.5%.

In contrast, traditional post-silicon bug localization approaches would require significant manual effort, additional hardware (e.g., trace buffers, details in Sect. 4) with significantly higher area overhead, or both. For example, even if the bug was detected quickly using a QED test, saving a full trace of all inputs and outputs of the crossbar alone (instead of using E-QED signature blocks, details in Sect. 2.1) for just 1,000 cycles would require over 34 Mbits of data.

The rest of the paper is organized as follows. Section 2 presents our new E-QED technique. Results are then presented in Sect. 3, followed by related work in Sect. 4. We conclude in Sect. 5. An extended version of this paper appears in (Singh et al. 2017).

2 Electrical Bug Localization Using E-QED

E-QED relies on several steps that are summarized in Fig. 2. During the chip design phase, E-QED signature blocks are automatically inserted (details in Sect. 2.1). These E-QED signature blocks are used during post-silicon validation to capture and compress the logic values of selected signals (*signatures*). During post-silicon validation, a suite of tests is run—it is crucial to run tests with short error-detection latencies (e.g., QED tests (Lin et al. 2014)). When an error is detected by such a test, the test is immediately halted, and all the captured signatures are scanned out (using on-chip scan chains (Abramovici et al. 1990)). In the last phase, formal analysis, the collected signatures are analyzed by a Bounded Model Checking (*BMC*) tool to first identify which design block produced the error(s), and then to find the FFs in that block that

could have captured the error(s), and finally to narrow this list even further by checking for consistency with signatures captured by neighboring design blocks. We explain each of these steps in more detail below.

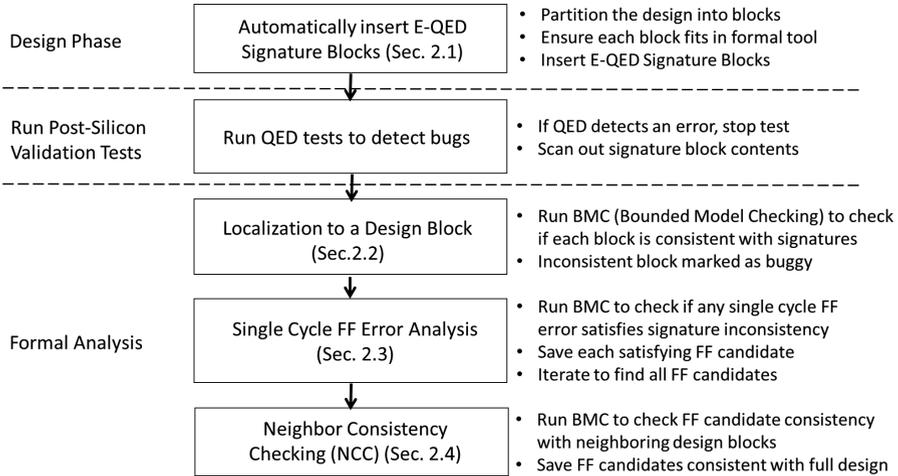


Fig. 2. An overview of E-QED.

2.1 Automatic Insertion of E-QED Signature Blocks

E-QED relies on being able to use a formal tool to perform bug localization after error detection by a post-silicon validation test. If the entire design can be loaded into the formal tool and analyzed, then this requirement can be satisfied by externally tracing the design's inputs and outputs during post-silicon testing (i.e., by saving all the signals at the external design interface). In this case, no additional internal hardware is required. However, most designs (particularly large SoCs) cannot be analyzed by existing formal tools without being partitioned into smaller blocks. In order to be able to analyze these smaller design blocks, we insert additional hardware (*E-QED Signature Blocks*) to capture logic values of signals at the boundaries of these smaller design blocks.

Design Partitioning. As mentioned above, unless it is possible to fit the full design into a formal tool, the design must be divided into smaller blocks. We use a simple algorithm that builds a list of design blocks by recursively descending through the design hierarchy. At each step, the current design block is tested to see if it can be loaded into the formal tool. If so, the recursion terminates and the block is added to the list. The result is a partition of the design into design blocks, each of which can be analyzed by the formal tool. Next, input and output signals for each block are grouped into *interfaces* as follows. For each design block A, the set of output signals driven by A and captured by the same design block B (or the same set of design blocks in the case that the signals fan out to more than one design block) are grouped into a single interface. In addition, if a design block C receives inputs from a design block in a

different clock domain, then all such input signals for design block C are also grouped into a separate interface. Each interface gets a single E-QED Signature Block (explained next) as illustrated in Fig. 3. As explained above, the primary inputs and outputs are traced externally.

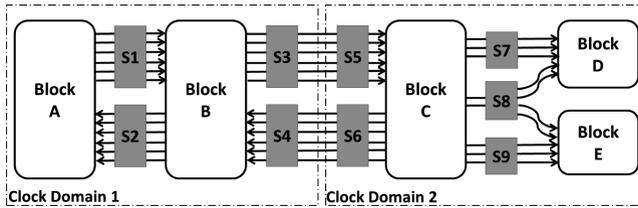


Fig. 3. Insertion of E-QED Signature Blocks. A and B use the same clock, so they can share signature blocks S1 and S2 at their interfaces. Since C is in a separate clock domain, the interfaces between B and C require separate signature blocks within each clock domain. Each signal from C captured in signature block S8 fans out to both D and E.

In our evaluation in Sect. 3, we explored one further enhancement beyond this algorithm. We observed that uninitialized memory arrays contributed too many degrees of freedom during formal analysis. To overcome this issue, we inserted additional signature blocks. Specifically, we added signature blocks to the signals between each cache bank and the cache controller logic and to the signals between the instruction cache and the instruction fetch unit for each thread on each processor core. This provided additional information on the data values being stored and loaded in the caches and the instructions being executed by the processors. As shown in Fig. 12 (in Sect. 3.2), this significantly improves the precision of E-QED.

E-QED Signature Block Design. E-QED Signature Blocks store logic values called signatures that represent a (lossy) compression of the sequence of logic values for a set of signals over time. We use multiple-input signature registers or *MISR*s (which have been extensively used for circuit test response compaction (Saxena and McCluskey 1997)) for this purpose. A MISR is a shift-register in which certain bits are XORed and fed back into the first bit, and, at the same time, a set of input signals are XORed with the values being shifted. An example of a 6-bit MISR is shown in Fig. 4.

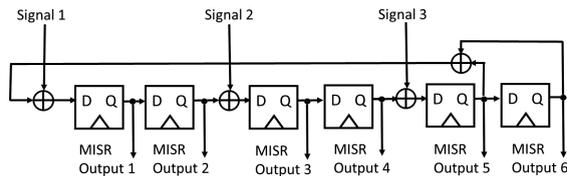


Fig. 4. An example of a 6-bit Multiple Input Signature Register (MISR) with 3 input signals and feedback generated from bits 5 and 6.

If a MISR has been operating for N cycles since reset, we refer to the number N of captured cycles as the *capture window* of the MISR. For post-silicon bug localization, it is crucial that the capture window is long enough to include the point when a bug gets activated and an error is created. However, the capture window must also not be too long, since the design behavior over the entire window is analyzed using the formal tool (see Sects. 2.2–2.4). If N is too large, the formal tool will fail as the unrolled design grows too large to analyze. This is why tests with short error detection latencies are necessary, and QED in particular enables our E-QED approach. For most cases, QED ensures that capturing the last 1,000 cycles prior to error detection is sufficient to also capture the point of bug activation. In our experiments (Sect. 3), we set N to 1,024. Since post-silicon validation tests run longer than 1,024 cycles, each MISR will be reset to a known state periodically to maintain the capture window length (1,024 cycles in this case).

In choosing the MISR size, we select the number of bits K to be at least equal to the number of signal bits being captured (M). K should also be large enough such that the MISR has more states than the capture window length (i.e., K must be greater than $\log_2(N)$). For M signal bits over N cycles, a complete trace would require $M*N$ bits. Thus, if K is $M*b$ (can be viewed as b MISR bits per input signal), the parameter b decides the compression ratio N/b .

Based on empirical analysis (discussed in Sect. 3.2), we used $b = 8$ for all the blocks in the OpenSPARC T2 design except for inside the processor core, where we used $b = 4$. After selecting the MISR size, the choice of which bits to use as feedback can be made based on extensive existing work which optimizes for various good properties of MISRs (e.g. (Bardell et al. 1987)).

As mentioned above, we periodically reset each MISR to a known state (to mark the start of a capture window). However, this raises the possibility that the error might be captured right after the MISR is reset, in which case the number of cycles captured will be much less than our target N (and the point when the bug got activated may not be included in the capture window as a result). To avoid this situation, in our *E-QED Signature Block* (shown in Fig. 5), we use two MISRs (operating in parallel) for each signature block. The two MISRs are paired with a counter, and they are reset (the MISRs to a defined starting

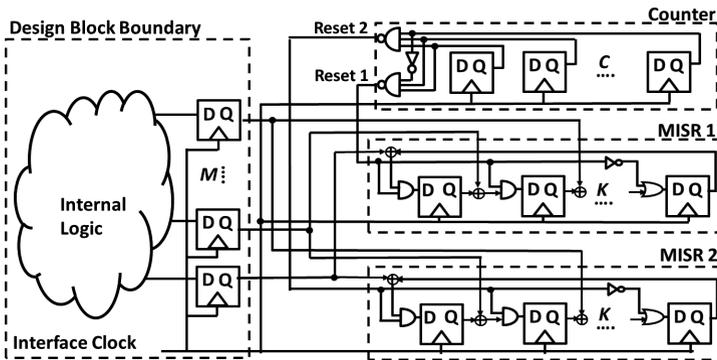


Fig. 5. The complete E-QED Signature Block design, with M interface signals generating a K -bit signature and the reset logic shown with the FFs of a C -bit counter (detailed counter implementation and MISR feedback not shown).

state and the counter to 0) when the design is powered on with a global reset. After that, the counter repeatedly counts to $2N$, alternately resetting one of the MISRs (to its starting state) every N cycles (waveforms in Fig. 6). This ensures that at the instant an error is detected, at least one of the MISRs has captured a signature covering at least N (and no more than $2N$) cycles. By setting N to a value greater than the expected error detection latency (e.g., 1,000 cycles for typical QED tests), we can ensure that with high likelihood, at least one signature covers the entire period between bug activation and error detection. The counter is sized to C , the ceiling of $\log_2(2N)$. The logic values in the MISRs and the counter can be scanned out when an error is detected by the QED test.

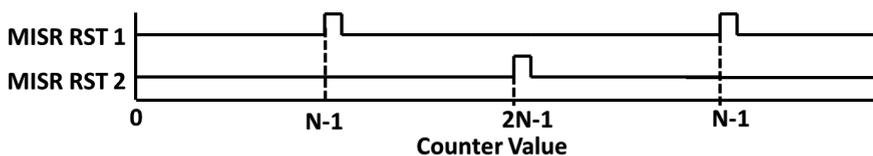


Fig. 6. Reset of MISR 1 and MISR 2, each resetting every $2N$ cycles. Note that at any instant, one MISR has captured data covering at least N cycles.

2.2 Bug Localization to a Design Block

We now explain how our design partitioning approach in Sect. 2.1 supports the first step in the electrical bug localization process: isolating an error detected by a post-silicon validation test to a specific design block. Following the detection of an error by QED tests, the logic values captured in the counters and MISRs are scanned out from each E-QED signature block. Formal analysis of the signature values is then performed using BMC (Clarke et al. 2001). BMC works by taking a model of the system (e.g., Verilog RTL), unrolling it some fixed number of time steps, and employing iterative time frame expansion to model sequential operation over this window of time. Then, automated tools such as Boolean Satisfiability solvers are used to check whether a set of constraints on the unrolled design is consistent. Constraints can be placed on any part of the unrolled design, commonly including the initial state, the inputs, and the final state.

Recall that when an error is detected, at least one of the two MISRs in each E-QED signature block has captured at least N cycles. Let T be the larger of the two capture windows (obtained from the counter in the E-QED signature block). For each design block B to be analyzed, we set up a BMC problem over T unrolled cycles (time frames), where T is the larger of the two windows captured by the E-QED Signature Blocks for the interfaces associated with B (i.e., signals that are either inputs or outputs for B).² In the BMC problem, we constrain the appropriate MISR (the one capturing T cycles) in each signature block to be equal to its scanned-out value in the final state following error detection, and also to be equal to its reset value at the beginning of the capture

² We assume that all interfaces for a design block are in the same clock domain and thus all have the same value of T . This assumption can easily be satisfied by modifying the design partitioning algorithm to continue its recursive descent if the current design block uses more than one clock domain.

window (T cycles before the final state). This corresponds to whether the design can satisfy the following property:

```
MISR reset = 1
##<T>
Input MISRs = <Input Signatures> && Output MISRs = <Output Signatures>
```

If the BMC tool is able to solve these constraints (note that, the initial state of the design block is not specified and so can be chosen arbitrarily by the BMC tool), we conclude that no bug occurred in this particular design block during the captured cycles. The solution produced by BMC also provides one possible error free T -cycle trace for the design block consistent with the signatures. On the other hand, if BMC is unable to find a solution, this means that the actual signals observed by the E-QED signature blocks are not consistent with the design logic, indicating that an electrical bug must have occurred within the block.³ Because this process identifies the buggy design block, it can already produce a significant degree of bug localization. Note that, it is theoretically possible for all design blocks to pass the BMC test, either because (i) the bug triggered (and produced error) outside the capture window; or (ii) some valid trace shares the same set of signatures as the actual captured buggy trace (because of lossy compression by E-QED signature blocks); or (iii) the symbolic initial states (corresponding to internal flip-flops in the design block where E-QED signature blocks are not inserted) introduce too many degrees of freedom. However, in our simulations based on injected errors (see Sect. 3), using QED tests and $N = 1024$, we did not encounter any such scenario: we always found exactly one block that failed the BMC test. These points are discussed in more detail in (Singh et al. 2017).

We illustrate the analysis above through a simple example, shown in Fig. 7. The inputs and outputs are both compressed using the 6-bit MISR shown in Fig. 4. For the Input MISR, its inputs, signals 1, 2, and 3 (corresponding to Fig. 4), are B, A, and 0 respectively. For the Output MISR, signals 1, 2, and 3 are Z, Y, and X respectively. During cycle 1, the MISRs are at their initial state (000001) and the circuit inputs A and B are 1 and 0. The initial outputs at cycle 1 are the contents of F6, F7 and F8 (011), and they are compressed into the output MISR on the next clock edge. A transient electrical bug occurring somewhere in the circuit results in the output X having the incorrect value 1 at cycle 5. At cycle 6, the error is captured in the output MISR signature, making it 110010 (instead of the correct output MISR signature 110000). A BMC run for this design over the 6 cycles (based on the captured output MISR signature 110010 and the captured input MISR signature 111010) reports unsatisfiable, indicating that an error has occurred in the circuit⁴.

³ As an optimization, in this case, we also run the BMC analysis with the smaller value of T for the block. If this analysis is also inconsistent, it provides a much shorter window in which the bug occurred.

⁴ E-QED uses symbolic initial values for FFs. To keep this example simple, we set the initial values of the FFs as shown in Fig. 7.

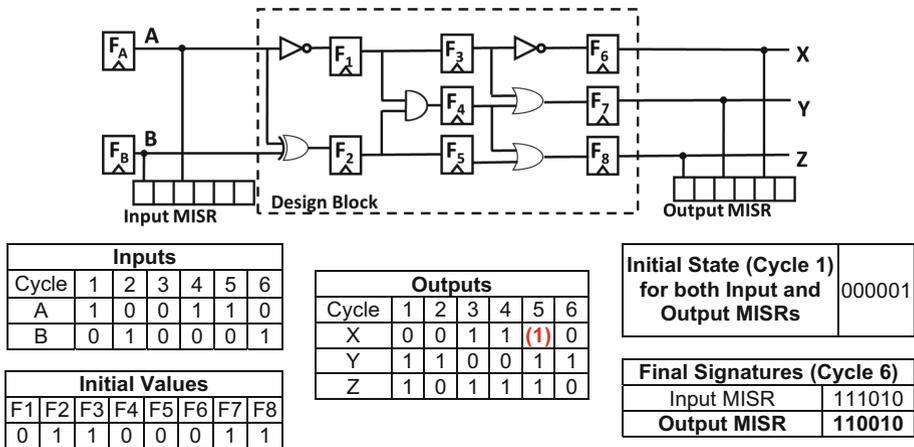


Fig. 7. A simple design block, a sequence of inputs and outputs, the initial values in the FFs and MISRs, the final signature captured by the Input MISR and the expected and captured (corrupted) signature in the Output MISR.

2.3 Single-Cycle FF Error Analysis with Bounded Model Checking

After the electrical bug has been localized to a single block (recall that in Sect. 2.1, we choose block sizes based on what fits within the BMC tool), E-QED performs further localization within the block by using an error model. In this paper, we use a *single-cycle FF error model*, which assumes that the electrical bug causes a transient error that affects one (arbitrary) flip-flop in the design during a single (arbitrary) cycle. Such an error model has been commonly used in the literature (e.g. (McLaughlin et al. 2009)) for electrical bugs. Our E-QED approach can be extended to include other models as discussed in (Singh et al. 2017). To demonstrate how error analysis can find a list of candidates, one of which corresponds to the FF and cycle that captured the single-cycle error due to an electrical bug, consider the earlier example of Fig. 7 presented again in Fig. 8.

For simplicity, let us first assume that we can directly observe all of the inputs and outputs on every cycle (we will consider the case where we have only the MISR signatures next). Recall that the expected output at cycle 5 is XYZ = 011, but we observe: XYZ = 111; i.e., the value of X is incorrect. In this simple example, there are only three FFs along the path to the X output; F₁, F₃ and F₆. According to our error model, an electrical bug could have caused an error in any one of them, and all three could have potentially caused the observed error at X; so there are three possible error scenarios. In the first, the error causes F₆ to capture a 0 instead of a 1 in cycle 4 (due to an electrical bug along the path from F₃ to F₆). In the second, the error causes F₃ to capture a 0 instead of a 1 in cycle 3 (due to an electrical bug along the path from F₁ to F₃). This error then propagates to F₆ and then to X. In the third, the error causes F₁ to capture a 0 instead of a 1 in cycle 2 (due to an electrical bug along the path from F_A to F₁). But notice that in this case, the input to F₄ in cycle 3 would then change to 0, and crucially, the input to F₇ in cycle 4 would also change to 0, and therefore output Y would be 0 in cycle 5. Since the error did not affect Y, this third scenario is not

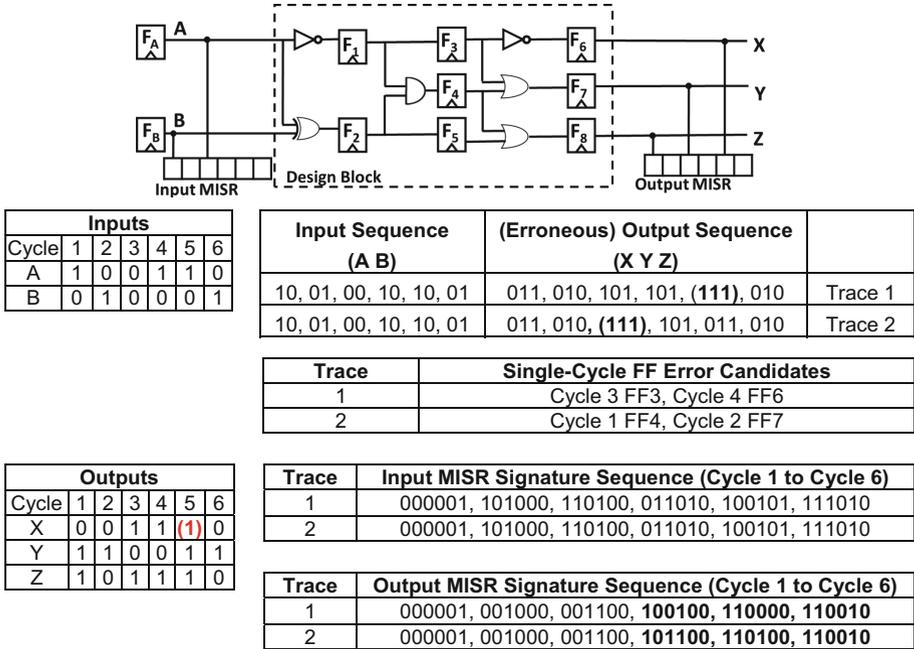


Fig. 8. The results of single-cycle FF error analysis for the example from Fig. 7. Each candidate (cycle and FF pair) has the property that if an error occurs at the input of the candidate FF at the candidate cycle (captured in the following cycle), it would result in the captured Output MISR signature at cycle 6. Note that, the candidates that generate Trace 1 match the actual sequence but the candidates that generate Trace 2 are aliased candidates in that they correspond to a different output sequence but result in the same final output MISR signature.

consistent with the observed outputs, leaving only errors in FFs F_3 and F_6 as viable candidates. Importantly, observe that, given our error model, this is an exhaustive list of the potential sources of the error observed at X in cycle 5.

Now, consider the case where we do not have direct access to the input or output signals, but only to the compressed 6-bit MISR signatures. Figure 8 once again lists the captured Output MISR signatures in sequence ending at cycle 6. The two candidate FFs that were identified as possible sources of the error above both generate the captured erroneous output signature by recreating the exact output sequence (corresponding to Trace 1) that was captured due to the error. However, there are also other traces that can be created with different candidates that end with the same final signatures in the MISRs. We call these *aliased* FF candidates and traces. An example is shown as Trace 2 in Fig. 8. Aliasing occurs due to lossy compression in the MISRs: errors introduced into the output MISR at different times can result in the same final signature. Observe from the output sequences in Fig. 8 that for Trace 2, the erroneous candidate FFs cause an output error at Y in cycle 3. This error shifts through the output MISR (as shown in the MISR sequence for Trace 2) to still match the final captured output MISR signature. Since in practice we do not have the actual output signals, but only work with the MISR signatures, aliasing can increase the number of candidate FFs and traces returned

by the formal analysis, reducing the resolution of the bug diagnosis. We discuss a technique for significantly improving (i.e., reducing and potentially eliminating) these aliased candidates in the next section.

We use BMC to systematically compute all candidate FFs consistent with the signature values for the interfaces of the buggy design block. To do this, we first modify the design (this modification is only to support the formal analysis at this step; the actual manufactured design is **not** affected). In the modified design, a two-input multiplexer is added before the data input of every FF in the design (Fig. 9(a)). The inputs to the multiplexer are the original data input to the FF and an inverted data input value to help inject an error as needed. By setting the select line of the multiplexer, an error is captured in the FF, matching the model for electrical bugs described earlier.

The select lines are all controlled by a decoder, allowing an error to be injected into a single FF on exactly one cycle as shown in Fig. 9(b). The decoder inputs are left unconstrained and provided to the BMC tool to control. The least significant output bit is not used, so an input combination of all 0s to the decoder generates an output of all 0s (on decoder output bits $1 \dots R$), meaning no error is injected into the design on that cycle. An additional FF (Fig. 9(b)) is added to ensure that only one single FF in the design block is chosen to have an error for only one cycle during a single BMC run. This additional FF (in Fig. 9(b)) is initially reset to 0 (reset is asserted at the start of BMC) and is then set to 1 (and remains 1) once the decoder produces a non-zero value on output bits $1 \dots R$. Once the value of this FF is 1, it forces the “Error Select” signals to all be 0 for all remaining cycles.

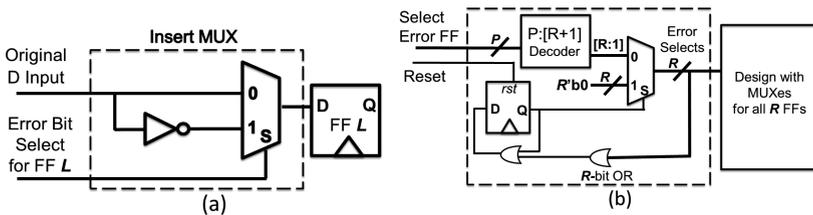


Fig. 9. (a) The inserted multiplexer allows an error to be injected into FF L . (b) control logic to inject an error during one cycle into the design when the BMC provides a FF index as input “Select Error FF” to the decoder. No error is injected when P is all 0s. The first time a non-zero input is provided, exactly one of the bits of the R “Error Select Lines” is set, causing the corresponding FF to capture an error during the next cycle.

Thus, the BMC, in its attempt to satisfy the error trace, can only inject an error into a single FF in the design, and only in one cycle.

As before, a BMC problem is set up with an unroll limit of T cycles, and constraints are added corresponding to the reset and final signature values. This time, however, the BMC tool can additionally control the decoder so that it can attempt to satisfy the input and output constraints by injecting a single error into a FF in the design block sometime during the T cycles. If BMC finds a solution, it provides a candidate FF which could (potentially) be the source of the erroneous bit flip (bug) that caused the error, together with a trace of length T of the input and output signals that is consistent with the

MISR signatures. Of course, it is possible that there is more than one way to achieve this. Therefore, the BMC tool is run repeatedly (each time with an added constraint to rule out all previously found FF candidates) until no more candidate FFs can be found.

2.4 Neighbor Consistency Checking

As explained in the example in Sect. 2.3 above, a consequence of using compressed signatures to constrain BMC is that aliasing can occur: different candidates for a single FF error can lead to the same output signature. Aliasing can be reduced by increasing the number of bits in the MISRs. However, in this section we introduce an alternative technique for reducing aliasing that does not require longer MISRs. This technique, which we call *Neighbor Consistency Checking (NCC)*, checks whether the trace generated by BMC corresponding to a particular candidate FF is consistent with the signatures captured in the MISRs at the interfaces of **other** blocks in the rest of the design. Any inconsistency can rule out the candidate FF as the one which captured the error, thereby improving diagnosis resolution.

The NCC strategy is as follows. Recall that for each FF candidate in the buggy block, the BMC tool also returns a full sequence of input and output values for T cycles (where T is the number of time frames analyzed). These values are applied as input or output constraints, as appropriate, to each of the neighboring blocks. The BMC tool is then used to check each neighboring block to see whether these newly-added constraints are consistent with the block's logic and the signatures on its interfaces. Note that, because the initial state of a neighboring block (initial states of FFs in the neighboring block) is not known, these checks cannot be done using simulation.

If a neighboring block is consistent, then the BMC tool returns a trace for inputs and outputs of that neighboring block. This sequence can then be recursively tested against additional neighbors, continuing across the chip. The limit would be reaching a clock domain boundary, where the interface is not shared directly with another block (as discussed in Sect. 2.1). If the BMC tool reports an inconsistency while analyzing the **immediate** neighbors of a buggy block, the trace under consideration (and the corresponding candidate FFs) can be eliminated. If the analysis has progressed further, then we backtrack to the previously analyzed neighbor and check to see if a different trace can be found that is consistent with its signatures. We continue the analysis in this way until either a fully consistent (across the whole design) set of traces is found or no such set of traces can be found in which case the candidate trace is eliminated.

Let us return to our running example. This time, we add a small neighboring block whose three inputs are driven by the three outputs (X, Y, Z) as shown in Fig. 10. The outputs of this block are captured using another 6-bit MISR (identical to Output MISR 1, capturing U, V, W instead of X, Y, Z respectively). Recall from Sect. 2.3 (Fig. 8) that there are 4 FF candidates and 2 traces consistent with the signature from Output MISR 1 at cycle 6. If we now use those two traces as the input constraints on the neighboring block, we can test if they are consistent with Output MISR 2. As seen in Fig. 10, the output sequence for aliased Trace 2 (in Fig. 8) results in a different final signature (at cycle 6) in Output MISR 2. This means it can be ruled out as a candidate, narrowing the candidate FFs to those from Trace 1 (in Fig. 8).

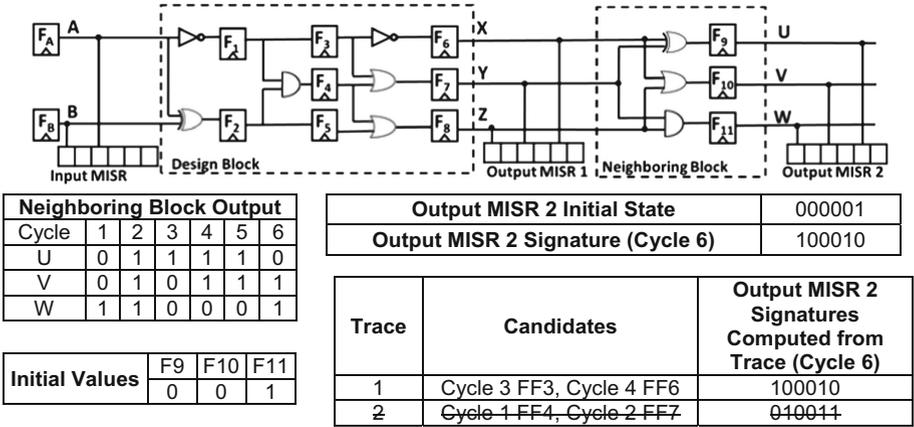


Fig. 10. Example circuit (continuing from Fig. 8) demonstrating NCC.

3 Results

In this section, we present simulation results demonstrating the effectiveness and practicality of E-QED. For our simulations, we use the OpenSPARC T2 SoC (OpenSPARC), an open-source version of the UltraSPARC T2 SoC (a 500-million-transistor SoC with 8 processor cores with private L1 caches supporting 64 hardware threads, 8 banks of shared L2 cache memory using a directory-based cache coherence protocol, 4 on-chip memory controllers, and a crossbar-based interconnect). We used the following design parameters for the E-QED Signature Blocks (Sect. 2.1): target capture window $N = 1,024$ cycles, counter size $C = 11$ bits; MISR size $K = M * b$, where $b = 8$ for all MISRs except those inside the processor core, where $b = 4$. A total of 118 E-QED Signature Blocks were used, adding 185,057 FFs, with an area impact of 2.5%.

In addition to an overview of results (in Sect. 3.1) with the parameters listed above, we also present an analysis of various trade-offs (in Sect. 3.2) that influence total area impact (by varying the parameter b in various signature blocks and also the signature block locations) and the granularity of electrical bug localization (in terms of candidate FFs as well as the number of candidate traces). The results in this section are for the single-cycle FF error model of electrical bugs (Sect. 2.3). We present an expanded set of results in (Singh et al. 2017).

3.1 Overview of Results

We randomly injected single-cycle FF errors in the OpenSPARC T2 design (with E-QED signature blocks inserted using our algorithm in Sect. 2.1) through RTL simulations (using Synopsys VCS). We ran 1,000,000 simulation runs, half with the SPLASH-2 Fast Fourier Transform (FFT) benchmark (Woo et al. 1995) and the other half with an in-house version of the Matrix Multiply (MMULT) benchmark. We ran 64-threaded versions of both benchmarks to ensure all threads on all processor cores of

OpenSPARC were active. The single-cycle FF errors were injected into: the processor core (all instances), crossbar, L2 cache controller (all instances) and the memory control unit (all instances). For each simulation run, we ran both the original benchmark and the benchmark transformed using QED (EDDI-V and PLC transforms with *Inst_min* and *Inst_max* set to 100, see (Lin et al. 2014)). During each simulation run, the injected error was allowed to propagate until it either vanished (i.e., was masked before impacting any program outputs) or resulted in observable failure (e.g., crash, incorrect outputs, exception or error detection by QED versions of the benchmarks). As soon as a QED test detected an error, we performed electrical bug localization using E-QED (Fig. 2). E-QED results for errors that cause hangs and deadlocks are discussed in (Singh et al. 2017) (they require QED tests with CFTSS-V transform (Lin et al. 2014)). For BMC, we used the Questa Formal tool (version 10.5) from Mentor Graphics on an Intel Xeon E5-2640 with 128 GB of RAM.

Table 1 summarizes the results. The bug trace size is the number of signals for all input and output interfaces captured in the signature blocks for that design block, multiplied by the length of the trace generated by the BMC tool. Recall that this length is determined by the counter value (T in Sect. 2.2) used to set up BMC analysis.

In Table 1, the “Number of Detected Errors” is the number of injected errors detected by the corresponding validation test (deadlocks and hangs caused by either test were not included). Note that, the vast majority of injected errors vanished (which is consistent with published literature (Cho et al. 2015; Sanda et al. 2008)). QED tests detected many more errors than the original tests ($\sim 4x$ more than original tests due to fine-grained QED checks, which is also consistent with published literature (Hong et al. 2010; Lin et al. 2014)).

Table 1. Electrical bug localization results for 1,000,000 single-cycle FF error simulations in the OpenSPARC T2 SoC.

	Original tests (no QED)	E-QED (after running QED versions of the original tests)
Number of detected errors	2,832	12,555
Error detection latency (cycles) [min, avg, max]	[2K, 976K, 8.8M]	[19, 168, 763]
Number of FF candidates [min, avg, max]	$9.3 * 10^5$ (\sim all flip-flops in the design)	[5, 18, 36]
Number of candidate bug traces [min, avg, max]	N/A	[1, 1, 1]
Number of bits in a bug trace [min, avg, max]	N/A	[109k, 483k, 1.74M]
BMC runtime (minutes)	N/A	[420, 526, 772]
Total E-QED signature block area impact	N/A	2.5%
Clock speed impact	N/A	$\sim 0\%$
Localization factor [min, avg, max]	1x	[25,837x, 51,674x, 186,025x]

Observation 1: E-QED automatically localized every error detected by the EDDI-V and PLC checks in the QED tests, and produced a list of 18 candidate FFs on average (out of a total of close to a million FF candidates) – a localization factor of over four orders of magnitude. These results were achieved with BMC runtime of 7–13 h. Very importantly, in every case, E-QED was able to report a unique trace with an average trace size of 483 kbits.

Table 2 shows the breakdown of results with respect to various design blocks of the OpenSPARC T2 SoC and various techniques presented in Sects. 2.2, 2.3 and 2.4.

Table 2. Electrical bug localization results (similar to Table 1) with a detailed breakdown.

		Design Block localization (Sect. 2.2)	Single-cycle FF error localization (Sect. 2.3)	NCC localization (Sect. 2.4)
Processor core (Core)	Number of traces [min, avg, max]	N/A	[1, 3, 8]	[1, 1, 1]
	Number of FF candidates [min, avg, max]	44,084	[11, 54, 91]	[8, 16, 29]
	Avg. localization factor	22x	17,224x	58,133x
	BMC runtime (minutes) [min, avg, max]	[62, 74, 77]	[303, 416, 478]	[49, 85, 114]
L2 cache (L2C)	Number of traces [min, avg, max]	N/A	[2, 14, 19]	[1, 1, 1]
	Number of FF candidates [min, avg, max]	31,675	[28, 77, 154]	[7, 19, 33]
	Avg. localization factor	29x	12,080x	48,954x
	BMC runtime (minutes) [min, avg, max]	[51, 58, 64]	[468,504,535]	[42,187,242]
Crossbar (CCX)	Number of traces [min, avg, max]	N/A	[2, 12, 17]	[1, 1, 1]
	Number of FF candidates [min, avg, max]	41,521	[31, 74, 130]	[5, 23, 36]
	Avg. localization factor	22x	12,569x	40,440x
	BMC runtime (minutes) [min, avg, max]	[53, 61, 72]	[379,421,450]	[54,142,205]
Memory control unit (MCU)	Number of traces [min, avg, max]	N/A	[3, 19, 24]	[1, 1, 1]
	Number of FF candidates [min, avg, max]	18,068	[21, 67, 143]	[5, 11, 18]
	Avg. localization factor	51x	13,882x	84,557x
	BMC runtime (minutes) [min, avg, max]	[35, 37, 41]	[315,387,428]	[78,163,251]
Total E-QED signature block area impact		2.5%		
Clock speed impact		~0%		

Observation 2: Signature-based electrical bug localization at the design block level (Sect. 2.2) achieves a localization factor of 22–51x. Single-Cycle FF Error Localization improves the localization factor by another 200–1,000x. Neighbor Consistency Checking (NCC) further improves localization factor by another 5–10x (for an average overall localization factor of 50,000x) and produces just a single candidate trace.

3.2 Design Trade-Offs

For the results in Tables 1 and 2, the total area impact of E-QED signature blocks is 2.5% (obtained through synthesis of the OpenSPARC T2 SoC using the Synopsys Design Compiler with the Synopsys EDK 32 nm library for standard cells and memories). In Fig. 11, we analyze the area impact vs. the granularity of electrical bug localization. We vary the MISR parameter b (recall from Sect. 2.1 that the MISR size is $b \cdot M$, where M is number of signals captured by the signature) from 2 to 14.⁵ For the entire design, this varies the total number of FFs in the signature blocks from 48,436 to 338,930. We use the number of candidate traces here for comparison since the primary aim of NCC is to reduce the number of candidate traces found (ideally to 1). Reducing the number of candidate traces also reduces the number of FF candidates.

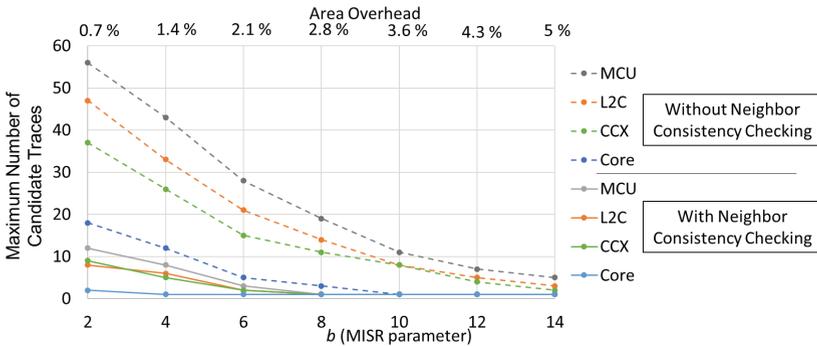


Fig. 11. The maximum number of unique candidate traces observed, before and after Neighbor Consistency Checking (NCC), for different area overheads, varied by changing the size of the MISRs in the signature blocks (MISR size is $b \cdot M$, where M is number of signals captured by the signature).

Observation 3: Neighbor Consistency Checking is able to reduce the number of candidate traces to just 1 when the area impact is 2.8% (this corresponds to a design where MISRs are constructed with $b = 8$ for the entire design). The area impact can be further reduced to 2.5% (while still ensuring only a single candidate trace is found) by using a hybrid set of parameters (the parameters used in Sect. 3.1), namely $b = 4$ within the processor cores and $b = 8$ everywhere else.

⁵ Note that in our simulation experiments, above (in Sect. 3.1), we used different fixed values of b for different parts of the design. In this section, we allow b to vary but all MISRs use the same value of b .

As discussed in Sect. 2.1, we made two enhancements to our approach to signature block insertion: signature blocks were added for signals around the L1 and L2 cache memory array banks (adding 16,640 FFs for all L2 banks and 3,840 FFs for all L1 banks), and signature blocks were also added at the instruction fetch units of the processors (adding 8,192 FFs across all 8 cores). As shown in Fig. 10, NCC analysis with these additional E-QED signature blocks significantly reduces the number of candidate traces for each value of b .

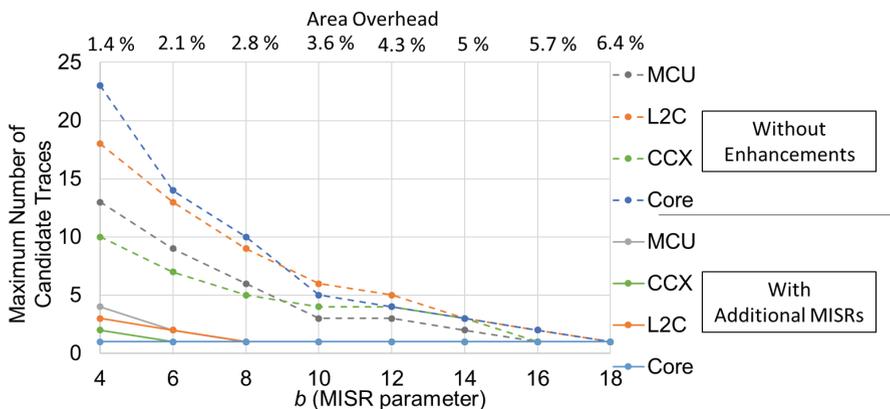


Fig. 12. The maximum number of unique candidate traces observed after performing NCC, with and without the addition of extra signature blocks around the L2 and L1 caches and processor instruction fetch unit, for different MISR parameters b (From Sect. 2.1, MISR size is $b \cdot M$, where M is number of signals captured).

4 Related Work

Most existing publications on post-silicon bug localization (especially for electrical bugs) focus on ways to improve observability as well as techniques for localizing bugs inside the design (and combinations thereof). For example, many existing techniques capture system behavior (logic values of various signals during post-silicon validation runs) using trace buffers: small on-chip memories that record logic values of a selected set of internal signals (Abramovici 2006; Anis and Nicolici 2007; Ma et al. 2015, and many others). However, with typical long error detection latencies (millions or billions of clock cycles) trace buffers can quickly become ineffective. E-QED leverages the QED technique (Lin et al. 2014) to create post-silicon validation tests with sufficiently short error detection latencies. (Other techniques that create tests with short error detection latencies can also be used with E-QED). Even when tests with short error detection latencies are available, inserting trace buffers for fine-grained electrical bug localization still imposes unacceptably high area overheads. For example, as discussed in Sect. 1, saving a full trace of all inputs and outputs of the crossbar block alone of the OpenSPARC T2 SoC for just 1,000 cycles would require over 34 Mbits of data.

Recorders that store microarchitectural information in processor cores (e.g. (Park et al. 2009, 2010)) require much less area. But they are applicable for processor cores only (while uncore components and accelerators occupy large portions of complex SoCs). Moreover, approaches in (Park et al. 2009, 2010) require manual analysis. Techniques for compressing trace buffer contents generally provide limited benefits. Trace compression techniques such as state restoration (Ko and Nicolici 2009)) (or related approaches that use RTL simulation to restore signals from partial traces) cannot be used for electrical bugs (because electrical bugs are not present in the RTL description). (Vali and Nicolici 2016) presents a trace signal selection approach that directly tries to improve the ability to diagnose FF errors caused by electrical bugs. However, the reported improvements are modest.

Signature blocks used by E-QED overcome these limitations. Although MISRs have been extensively used for manufacturing testing (e.g. (Saxena and McCluskey 1997 and numerous others)), E-QED does not require the generation of fault-free signatures by simulating post-silicon validation tests (unlike manufacturing testing). Instead, it uses formal techniques to reason about consistency between signatures captured by E-QED signature blocks at the inputs and outputs of a design block, as well as the logic inside the design block itself. While it may be possible to use hardware-implemented assertions (e.g. (Bayazit and Malik 2005)) to enhance E-QED, finding the “right” set of (hardware implementation-friendly) assertions remains a major challenge for automatic assertion generation. Other observability techniques (e.g. (DeOrio et al. 2011; Li et al. 2010)) can result in false positives.

The problem of electrical bug localization is similar to the classical fault diagnosis problem in manufacturing testing (Abramovici et al. 1990). The use of Scan Design for Testability makes the problem somewhat simpler for manufacturing testing vs. post-silicon validation. Many publications have used formal techniques for these purposes (e.g. (Larrabee 1992; Mangassarian et al. 2007, and others)). Similarly (Zhu et al. 2011, 2014) use formal techniques, aided by backbones and a sliding window approach, for electrical bug localization. Techniques such as BackSpace and its derivatives (De Paula et al. 2008, 2011, 2012; Le et al. 2013; Sengupta et al. 2012) also use formal methods for bug localization purposes. However, as explained in Sect. 2, the biggest challenge is to create an automatic approach that scales for large SoCs without incurring large area impact. Some of the above techniques attempt to overcome the scalability challenge by “consistently” reproducing the buggy behavior over multiple runs; this can be very difficult for complex designs, as explained in (De Paula et al. 2011). Our E-QED approach overcomes these challenges: E-QED signature blocks enable scalability for large SoC designs with small area impact; short error detection latencies of QED tests enable us to apply BMC in conjunction with the E-QED signature blocks; and, E-QED analysis techniques in Sect. 2 enable us to perform bug localization despite lossy compression by E-QED signature blocks (with minimal reliance on consistent reproduction of buggy behaviors). Note that, the E-QED Neighbor Consistency Checking technique in Sect. 2.4 is different from consistency checking techniques presented in (Jones et al. 1996; Park et al. 2009, 2010, and others).

5 Conclusion

E-QED overcomes electrical bug localization challenges during post-silicon validation and debug. It automatically localizes electrical bugs and provides a comprehensive list of components that may contain the bugs (together with corresponding bug traces). It is an automatic approach which is highly effective and practical for large designs, as demonstrated on the OpenSPARC T2 SoC: automatic electrical bug localization in a few hours (9 h on average) that can narrow an electrical bug to a handful of candidate flip-flops (18 flip-flops on average for a design with ~ 1 Million flip-flops) and identify a single candidate bug trace. The area impact of E-QED is $\sim 2.5\%$. In contrast, it might take weeks (or even months) of mostly manual work (per bug) using traditional approaches. E-QED is made possible through a unique combination of Quick Error Detection techniques for bug detection, E-QED signature blocks that are automatically inserted during design, and new consistency checking techniques enabled by formal methods.

There are several future directions. E-QED can be extended to: 1. leverage already-existing Scan Design for Testability techniques to further enhance bug localization; 2. localize bugs in analog and mixed-signal components of SoCs; 3. understand the interplay between scalability of BMC tools for bug localization vs. error detection latencies of QED tests vs. design of more sophisticated E-QED signature blocks; 4. diagnose defects that are detected using system-level testing during manufacture; 5. enable full system-level (consisting of many ICs chips) bug localization; and, 6. correct/fix bugs after manufacture.

References

- Abramovici, M., Breuer, M.A., Friedman, A.D.: *Digital Systems Testing and Testable Design*. Computer Science Press, New York (1990)
- Abramovici, M.: A reconfigurable design-for-debug infrastructure for SoCs. In: *Proceedings of IEEE/ACM Design Automation Conference*, pp. 7–12 (2006)
- Anis, E., Nicolici, N.: On using lossless compression of debug data in embedded logic analysis. In: *Proceedings of 2007 IEEE International Test Conference (ITC)* (2007)
- Bardell, P.H., McAnney, W.H., Savir, J.: *Built-in test for VLSI: Pseudorandom Techniques*. Wiley, New York (1987)
- Bayazit, A.A., Malik, S.: Complementary use of runtime validation and model checking. In: *Proceedings of ICCAD-2005, IEEE/ACM International Conference on Computer-Aided Design*, pp. 1052–1059 (2005)
- Cho, H., et al.: Understanding soft errors in uncore components. In: *Proceedings of 2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6 (2015)
- Clarke, E., Biere, A., Raimi, R., Zhu, Y.: Bounded model checking using satisfiability solving. *Formal Methods Syst. Des.* **19**(1), 7–34 (2001)
- DeOrio, A., Khudia, D.S., Bertacco, V.: Post-silicon bug diagnosis with inconsistent executions. In: *Proceedings of 2011 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, San Jose, CA, pp. 755–761 (2011)

- De Paula, F.M., et al.: BackSpace: formal analysis for post-silicon debug. In: Proceedings of International Conference on Formal Methods in Computer-Aided Design, pp. 1–10 (2008)
- De Paula, F.M., et al.: TAB-BackSpace: unlimited-length trace buffers with zero additional on-chip overhead. In: Proceedings of IEEE/ACM Design Automation Conference (2011)
- De Paula, F.M., Hu, A.J., Nahir, A.: nuTAB-BackSpace: rewriting to normalize non-determinism in post-silicon debug traces. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 513–531. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31424-7_37](https://doi.org/10.1007/978-3-642-31424-7_37)
- Dusanapudi, M., et al.: Debugging post-silicon fails in the IBM POWER8 bring-up lab. IBM J. Res. Dev. **59**(1), 1–10 (2015)
- Foster, H.D.: Trends in functional verification: a 2014 industry study. In: Proceedings of IEEE/ACM Design Automation Conference, pp. 48–52 (2015)
- Friedler, O., et al.: Effective post-silicon failure localization using dynamic program slicing. In: Proceedings of IEEE/ACM Design Automation Test in Europe, pp. 1–6 (2014)
- Hong, T., et al.: QED: quick error detection tests for effective post-silicon validation. In: Proceedings of IEEE International, Test Conference, pp. 1–10 (2010)
- Jones, R.B., Seger, C.-J.H., Dill, D.L.: Self-consistency checking. In: Srivas, M., Camilleri, A. (eds.) FMCAD 1996. LNCS, vol. 1166, pp. 159–171. Springer, Heidelberg (1996). doi:[10.1007/BFb0031806](https://doi.org/10.1007/BFb0031806)
- Ko, H.F., Nicolici, N.: Algorithms for state restoration and trace-signal selection for data acquisition in silicon debug. IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst. **28**(2), 285–297 (2009)
- Larrabee, T.: Test pattern generation using Boolean satisfiability. IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst. **11**(1), 4–15 (1992)
- Le, B., Sengupta, D., Veneris, A., Poulos, Z.: Accelerating post silicon debug of deep electrical faults. In: Proceedings of 2013 IEEE 19th International On-Line Testing Symposium (IOLTS), Chania, pp. 61–66 (2013)
- Li, W., Forin, A., Seshia, S.A.: Scalable specification mining for verification and diagnosis. In: Proceedings of Design Automation Conference (DAC), pp. 755–760 (2010)
- Lin, D., et al.: Effective post-silicon validation of system-on-chips using quick error detection. IEEE Trans. Comput. Aided Des. Integr. Circ. Syst. **33**(10), 1573–1590 (2014)
- Lin, D., et al.: A structured approach to post-silicon validation and debug using symbolic quick error detection. In: Proceedings of 2015 IEEE International Test Conference (ITC), October 2015
- Ma, S., et al.: Can't see the forest for the trees: state restoration's limitations in post-silicon trace signal selection. In: Proceedings of 2015 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 1–8 (2015)
- Mangassarian, H., et al.: A performance-driven QBF-based iterative logic array representation with applications to verification, debug and test. In: Proceedings of International Conference on Computer-Aided Design (ICCAD) (2007)
- McLaughlin, R., Venkataraman, S., Lim, C.: Automated debug of speed path failures using functional tests. In: Proceedings of 2009 IEEE VLSI Test Symposium, pp. 91–96 (2009)
- Mishra, P., Morad, R., Ziv, A., Ray, S.: Post-silicon validation in the SoC era: a tutorial introduction. In: IEEE Design & Test, April 2017
- Nahir, A., et al.: Post-silicon validation of the IBM POWER8 processor. In: Proceedings of IEEE/ACM Design Automation Conference, pp. 1–6 (2014)
- OpenSPARC: World's First Free 64-bit Microprocessor. <http://www.opensparc.net>
- Park, S.-B., Hong, T., Mitra, S.: Post-silicon bug localization in processors using instruction footprint recording and analysis (IFRA). IEEE Trans. Comput. Aided Des. Integr. Circ. Syst. **28**(10), 1545–1558 (2009)

- Park, S.-B., et al.: BLoG: post-silicon bug localization in processors using bug localization graph. In: Proceedings of IEEE/ACM Design Automation Conference, pp. 368–373 (2010)
- Reick, K.: Post-silicon debug – DAC workshop on post-silicon debug: technologies, methodologies, and best-practices. In: Proceedings of IEEE/ACM Design Automation Conference (2012)
- Sanda, P.N., et al.: Soft-error resilience of the IBM POWER6 processor. *IBM J. Res. Dev.* **52**(3), 275–284 (2008)
- Saxena, N.R., McCluskey, E.J.: Parallel signature analysis design with bounds on aliasing. *IEEE Trans. Comput.* **46**(4), 425–438 (1997)
- Sengupta, D., et al.: Lazy suspect-set computation: fault diagnosis for deep electrical bugs. In: Proceedings of the Great Lakes Symposium on VLSI. ACM (2012)
- Singh, E., Barrett, C., Mitra, S.: E-QED: electrical bug localization during post-silicon validation enabled by quick error detection and formal methods, [arXiv:1705.0125](https://arxiv.org/abs/1705.0125) [cs.OH] (2017)
- Vali, A., Nicolici, N.: Bit-flip detection-driven selection of trace signals. In: Proceedings of 2016 21th IEEE European Test Symposium (ETS), Amsterdam, pp. 1–6 (2016)
- Vermeulen, B., Goossens, K.: Debugging Systems-on-Chip: Communication-Centric and Abstraction Based Techniques. Springer, Heidelberg (2014)
- Woo, S.C., et al.: The SPLASH-2 programs: characterization and methodological considerations. In: Proceedings of International Symposium on Computer Architecture (1995)
- Zhu, C.S., Weissenbacher, G., Malik, S.: Post-silicon fault localisation using maximum satisfiability and backbones. In: Proceedings of IEEE/ACM Formal Methods Computer-Aided Design, pp. 63–66 (2011)
- Zhu, C.S., Weissenbacher, G., Malik, S.: Silicon fault diagnosis using sequence interpolation with backbones. In: Proceedings of 2014 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), San Jose, CA (2014)