

TIUP: Effective Processor Verification with Tautology-Induced Universal Properties

1st Yufeng Li

Institute of Software, Chinese Academy of Sciences
University of Chinese Academy of Sciences
Beijing, China
e-mail: yufeng@nfs.iscas.ac.cn

2nd Yiwei Ci

Institute of Software, Chinese Academy of Sciences
University of Chinese Academy of Sciences
Beijing, China
e-mail: yiwei@iscas.ac.cn

3rd Qiusong Yang*

Institute of Software, Chinese Academy of Sciences
University of Chinese Academy of Sciences
Beijing, China
e-mail: qiusong@iscas.ac.cn

Abstract—Design verification is a complex and costly task, especially for large and intricate processor projects. Formal verification techniques provide advantages by thoroughly examining design behaviors, but they require extensive labor and expertise in property formulation. Recent research focuses on verifying designs using the self-consistency universal property, reducing verification difficulty as it is design-independent. However, the single self-consistency property faces false positives and scalability issues due to exponential state space growth. To tackle these challenges, this paper introduces *TIUP*, a technique using tautologies as universal properties. We show how *TIUP* effectively uses tautologies as abstract specifications, covering processor data and control paths. *TIUP* simplifies and streamlines verification for engineers, enabling efficient formal processor verification.

Index Terms—Formal verification, Universal property, Tautology, Processor

I. INTRODUCTION

Verification is crucial in processor design and manufacturing [1], with modern processors having optimization components like branch predictors, prefetchers, and out-of-order execution units that make verification expensive. Simulation-based verification involves many test cases and an oracle to compare instruction execution results. In contrast, formal verification, like model checking [2], [3], constructs a mathematical representation of the system and proves it adheres to prescribed properties. It excels at discovering corner cases [4], but specifying the properties of the design under verification (DUV) is challenging, the quality of verification heavily depends on expertise in property formulation [5], [6].

To simplify property formulation, ISA-Formal [4] proposed an end-to-end approach for formal processor verification. It automates architectural specification language (ASL) into properties for each instruction. This sparked research on ISA-level verification for RISC-V processors (RISCV-Formal [7]) and ILA [8]–[10] for accelerators. Defining properties for individual instructions is a non-trivial task [11]. Moreover, instructions are usually verified independently, and the modeling of inter-instruction interactions is insufficient. *QED* (Quick Error Detection) proposes a design-independent universal property called self-consistency, based on the idea that a processor executes two identical operations, the result

should be consistent [12]. *SQED* [1], [13], [14] and *S²QED* [15] are its symbolic execution versions. However, relying solely on the self-consistency property has issues. Instruction enumeration becomes complex and time-consuming, especially for bugs with long activation sequences [16]. Additionally, self-consistency does not address single-instruction bugs. *C-S²QED* [17] extends *S²QED* to target single-instruction bugs by generating a comprehensive set of properties using metamodeling techniques [18]. Unfortunately, *C-S²QED* properties are no longer universal because of the requirement to formally specify the semantics of individual instruction.

In this paper, we propose an end-to-end formal verification approach for processors based on tautology-induced universal properties (*TIUP*). Instead of relying on a single universal property, *TIUP* explores a set of universal properties present in functional units. The fundamental properties, referred to as seeds, focus on data paths and simple control logic. Proposition tautology templates are replaced with seed formulas to generate tautologies covering complex control paths. By symbolically executing instruction sequences related to properties through model-checking engines, we can then verify whether the results of the instruction execution conform to tautology specifications. Unlike *SQED* and variants, *TIUP* verifies multiple properties to eliminate missed detection. For instance, consider the associative law $x - y - z = x - (y + z)$ in first-order logic (FOL) with arithmetic. It should hold in any processor that implements the addition and subtract operations, serving as one of the essential seed properties. These seed properties play a pivotal role in focusing on the data paths and simple control logic relevant to processor verification. By replacing elements P or Q in the proposition tautology templates $(P \& Q) \rightarrow P$ with the aforementioned associative law, we can generate a new tautology that covers jump and logic computation. After correctly executing instruction sequences from tautology, the value in the result register should be 1 (true). In $x - y - z = x - (y + z)$, two distinct sub-instruction sequences are generated from $x - y - z$ and $x - (y + z)$, respectively. The target result register serves as an indicator of whether they yield identical outcomes. In comparison to *SQED* and variants, our approach overcomes the inherent vulnerability in repetitive comparisons, which are susceptible to identical flaws.

*Corresponding author

In summary, the contributions of this paper are the following:

- We propose the use of universal properties derived from FOL tautologies that are independent of microarchitectural details for processor verification;
- We design and implement *TIUP* that can formally verify the implementation of a processor based on a set of universal properties;
- We evaluate the effectiveness and performance of *TIUP* on two real open-source processors (in-order and out-of-order) based on the RISC-V architecture. *TIUP* can effectively detect anomalies related to both the data path and control path of processors.

The rest of this paper is structured as follows. Section II provides additional details regarding the motivation behind this study, demonstrating the necessity of extending the universal property. Section III elaborates on *TIUP* design. We first introduce an overview and formal model of *TIUP* and then proceed to illustrate each component of the design. Section IV presents the results of our experimental evaluation of the proposed approach. Finally, in Section V, we conclude this paper.

II. MOTIVATION

The universal property is inherently design-independent, meaning that no manual property formulation is required for verification purposes [1]. *QED* and variants leverage the idea of instruction execution self-consistency to formulate a single universal property. The self-consistency property specifies that when both the original and duplicate instruction sequences are executed from a *QED-consistent* state (i.e., registers and memory locations that are partitioned separately for the original and duplicate instruction sequences hold identical values), the resulting state must also be *QED-consistent*.

A. False Positives with Self-consistency

As mentioned earlier, both the original and duplicate sequences can be affected by identical flaws, resulting in false positives during self-consistency checks. For instance, consider Listing 1, where both the original and duplicate instruction sequences contain multiple *ADD* instructions that are mistakenly interpreted as *SUB* instructions. Despite this discrepancy, the self-consistency property remains intact, making it impossible for the solver to identify a counterexample.

A more intricate case arises in the detection of control flow bugs. If both the original and duplicate instruction sequences have incorrect branch directions, the self-consistency check alone cannot identify the bug. The enhanced EDDI-V approach [19] is designed to detect incorrect branch directions. As demonstrated in Listing 2, it captures the target address of each control flow instruction (branch or jump), such as “2” and “5”. If the original sequence jumps correctly (“2”), and the duplicate sequence jumps erroneously (“5”), a non-*QED* instruction leads to self-consistency being violated. However, there is also a scenario where the original and duplicate sequences mistakenly jump to “5”, resulting in false positives.

To eliminate missed detections caused by false positives, researchers have supplemented self-consistency by specifying

```
// Original instruction sequence
R4 = R1 + R2    =>    R4 = R1 - R2
R5 = R4 + R3    =>    R5 = R4 - R3

// Duplicate instruction sequence
R20 = R17 + R18 =>    R20 = R17 - R18
R21 = R20 + R19 =>    R21 = R20 - R19

BNE R4, R20 ERROR_DETECTED
```

Listing 1. False positive for self-consistency

Original sequence	Duplicate sequence
R0 = 0, R4 = 1, R5 = 2	R16 = 0, R20 = 1, R21 = 2
0: R3 = R5 - R4	0: R19 = R21 - R20
1: BEQ R3, R0 #5	1: BEQ R19, R16 #5
// R3 != 0	// R19 != 0
// Branch not taken	// Error: Branch taken
2: R2 = R4 + R5	// Ignore QED instruction
	(2: R18 = R20 + R21)
// R2 = 3	
	// Use Non-QED instruction
	5: lui R18, 0
	// R18 = 0

Listing 2. Enhanced EDDI-V: erroneous branch direction

the semantics of all individual instruction [1], [17]. However, these measures compromise the advantages of design-independent universal properties. For *TIUP*, the determination of correctness is based on the requirement that the execution of the processor must meet all tautologies, thus false positives will not appear when checking all universal properties. For instance, during the verification of the property $x - y - z = x - (y + z)$, if “+” is mistakenly decoded as “-” in the implementation, the solver can promptly give a counterexample as $x - y - z \neq x - (y - z)$.

B. Under-constrained State Space with Self-consistency

Additionally, to verify the self-consistency property, it is necessary to make a copy of each original instruction sequence and enumerate various combinations of instructions. As the instruction sequence increases, the solver will time out in the face of an exponentially expanding state space. For BMC, deep unrolling may be prohibitive [16]. Although *S²QED* [15] reduces the instruction sequence by instantiating two processors, it has limitations in terms of false positives and larger design sizes due to two processors.

TIUP verifies multiple universal properties, which is equivalent to a more detailed decomposition of the self-consistency property, thus reducing the state space that needs to be proven for the system.

III. EFFECTIVE PROCESSOR VERIFICATION WITH TAUTOLOGY-INDUCED UNIVERSAL PROPERTIES (*TIUP*)

TIUP is an end-to-end formal verification approach that employs tautologies as abstract specifications. Fig. 1 illustrates the workflow of *TIUP*.

A. Overview

Specifications define the anticipated behaviors of the processor. In end-to-end verification, instructions serve as the interface between the specifications and the implementations




Fig. 1. Workflow of *TIUP*

of the processor.

Abstract specification. To verify the correctness of target hardware, one option is to specify its properties and identify at which level of abstraction these properties can be evaluated for truth or falsity [3]. Tautology is a formalism that can specify the validity of the processor. An abstract specification is expressed as follows:

$$\models f \quad (1)$$

$\models f$ [20] refers to f as a tautology which is always true.

Formal model. A processor can be conceptualized as a state transition system: $\langle S, S_0, r, i \rangle$, where S represents a vector of state variables, and S' denotes primed variables in the next state; S_0 is a predicate that denotes initial states; r denotes the transition relation, and $r(S, S', i)$ indicates that the system transitions from S to S' after the instruction i is executed. For a safety property SP , the verification process is to check if the transition system will ever reach a violating state S_v in which SP does not hold. Inv refers to an invariant. The verification model is presented as follows:

$$S_0(S) \rightarrow Inv(S) \quad (2)$$

$$Inv(S) \&\& r(S, S', i) \rightarrow Inv(S') \quad (3)$$

$$Inv(S) \rightarrow SP \quad (4)$$

(2) indicates that $Inv(S)$ holds in the initial state; (3) indicates that the invariant should hold in the next state if it holds in the previous one; (4) indicates that $Inv(S)$ implies SP , reflecting the fact that SP holds in the whole system. For an abstract specification $\models f$, it can map to an instruction sequence I_f (refer to III-C1). The meaning behind a tautology violation in a transition system can be formally defined as follows:

$$S_0(S) \&\& I_f = C(f) \&\& r(S_0(S), S_v, I_f) \&\& !SP \rightarrow \perp \quad (5)$$

$I_f = C(f)$ represents that mapping of tautology f to instruction sequence I_f . (5) asserts that if the transition system moves from a safety state $S_0(S)$ to a violating state S_v while following the instruction sequence specified by the tautology f , the system fails to meet the requirements of the abstract specification $\models f$.

Establishing abstract specifications. The objective of *TIUP* is to perform formal verification to determine if a processor violates a given set of abstract specifications. Therefore, as a prerequisite for the verification process, it is essential to establish a well-defined set of abstract specifications. In *TIUP*, a synthesis methodology is employed to establish this set of abstract specifications.

Mapping abstract specifications. The specifications determine the instructions under verification (IUVs). *TIUP* generates a symbolic instruction sequence using a specific tautology, with the aim of finding a mapping that leads the transition system from the initial state $S_0(S)$ to the violating state S_v through the instruction sequence $I_f = \{i_0, i_1, \dots, i_m\}$. This process involves lexical and syntax analysis of the tautology, transforming it into an ISA-independent Intermediate Representation (IR). The IR is then further transformed into a symbolic instruction sequence, which is stored in a scheduler module. The scheduler module is integrated into the processor's fetch stage for verification purposes and is not included in the final manufactured IC.

Interfacing RTL design and model checking. To perform formal verification of the processor's RTL using formal tools, the properties to be verified need to be expressed as assertions and integrated into the RTL code. Subsequently, the entire RTL is transformed into a file format compatible with model-checking tools.

B. Establishing Abstract Specifications

An abstract specification is created by several seeds and a template. Seeds are first-order tautologies that encompass universal properties pertaining to processors' basic functions. Instructions in processors can be classified into several categories, such as arithmetic and logic, memory access, and control. For arithmetic and logic, the seeds include associative law ($x - y - z = x - (y + z)$), De Morgan's theorem for Boolean algebra ($x \oplus y = \sim ((x \& y) | (\sim x \& \sim y))$), etc. For memory access, the tautology $ld(st(mem, j, v), j) = v$ is used as an abstract specification. The function st takes memory mem , index j , and value v as inputs, returning a memory variable. The function ld receives a memory and an index as its parameters and returns the value fetched from memory. For control, *TIUP* utilizes tautologies that contain implication connectives that can express conditional jumps.

To cover the complex control logic of processors, abstract specifications need to describe complex control flow. *TIUP* utilizes template synthesis to construct such specifications. Templates are propositional tautologies, where the value of propositional is true or false, for example, implication introduction $((P \rightarrow Q) \rightarrow (!P | Q))$. Non-logical connective elements (e.g. P, Q) of a template can be replaced by seeds through a process, called instantiation. During the replacement process, multiple occurrences of a non-logical connective are replaced with the same seed, thus each instantiated template is still a tautology.

Algorithm 1 adopts a depth-first search strategy to instantiate templates. It takes a set of tautology templates, denoted as F , and a set of tautology seeds, denoted as E , as its inputs. A template is selected from the set and instantiated in sequence (line 2). V keeps non-logical connectives of a

Algorithm 1 SynthesizeTautology

Input: F is a tautology template set, E is a tautology seed set
Output: G is a collection of instantiated tautology templates

```

1:  $G \leftarrow \emptyset$ ;
2: for  $f \in F$  do
3:    $V \leftarrow \emptyset$ ; ▷  $V$  keeps non-logical connectives
4:    $N \leftarrow 0$ ;
5:   for  $o \in f$  do
6:     if  $o$  is non-logical connective &&  $o \notin V$  then
7:        $V \leftarrow V \cup \{o\}$ ;
8:        $N \leftarrow N + 1$ ;
9:     end if
10:  end for
11:   $stack \leftarrow \emptyset$ ;
12:   $stack.push((0, f))$ ;
13:  while  $!stack.empty()$  do
14:     $(m, rf) = stack.top()$ ;  $stack.pop()$ ;
15:    if  $m == N$  then ▷ all non-logical connectives have been
16:       $G \leftarrow G \cup \{rf\}$ ; ▷  $rf$  is an instance of  $f$ 
17:      continue;
18:    end if
19:    for  $e \in E$  do
20:      replace element  $V[m]$  in  $rf$  with  $e$ ;
21:       $stack.push((m + 1, rf))$ ;
22:    end for
23:  end while
24: end for

```

template formula (line 3). Seeds are selected sequentially from the seed set (E) to replace elements in V . Once the replacement is completed, an instance of f is obtained, followed by backtracking to the previous element in V (replace P before replacing Q). Finally, the output of Algorithm 1 is a collection of instantiated tautologies.

Depending on the complexity of the processor or computational resource constraints, these universal properties can be verified separately, which can alleviate the state explosion problem encountered during search in comparison with a single self-consistency property.

C. Mapping Abstract Specifications

The instantiated tautologies are leveraged to map into instructions under verification ($I_f = C(f)$), which can be utilized to determine whether the tautologies hold. Firstly, an IR can be produced via lexical and syntax analysis of the instantiated tautologies. Secondly, the IR is transformed into a specific instruction sequence according to the ISA model. These symbolic instructions are then stored in a scheduler connected to the pipeline's fetch stage.

1) *Intermediate representation*: *TIUP* compiles an instantiated tautology into an IR that is ISA-independent according to the results of lexical and syntax analysis. IR comprises a sequence of directives that describe computations and data flow transformations within the tautology. In this representation, constant symbols and variables are represented as operands, whereas predicates, functions, and connectives are translated into opcodes. The IR for $\models (x + y > 0) \&\& (y + z < 0) \rightarrow (x + y) \times (y + z) < 0$ is shown in Listing 3. The logical implication describes a control-related specification. The specification $\models (x + y > 0) \&\& (y + z < 0) \rightarrow (x + y) \times (y + z) < 0$ describes that when the execution result of antecedent $((x + y > 0) \&\& (y + z < 0))$ is true, the control flow jumps to the instruction block of consequent $((x + y) \times (y + z) < 0)$ and the execution result of consequent must be true. The

%tmp1 = add i32 %x, %y	+
%tmp2 = sgt i32 %tmp1, 0	>
%tmp3 = add i32 %y, %z	+
%tmp4 = slt i32 %tmp3, 0	<
%tmp5 = logic_and i32 %tmp2, %tmp4	&&
beq %tmp5 1, label %if, label %else	→
if:	
%tmp6 = mul i32 %tmp1, %tmp3	×
%tmp7 = slt i32 %tmp6, 0	<
%result_reg = and i32 %tmp7, %result_reg	&
else:	
other instruction blocks	

Listing 3. IR of $\models (x + y > 0) \&\& (y + z < 0) \rightarrow (x + y) \times (y + z) < 0$

result of each tautology is accumulated through the *and* operation and written back the updated result to a special register ($\%result_reg = \text{and } i32 \ \%tmp7, \ \%result_reg$). A well-defined IR enables the ease of employing *TIUP* for the verification of processors with different ISAs. IR is transformed into specific instruction sequences according to the ISA model, for example, under the RV32I model, the directive $\%tmp1 = \text{add } i32 \ \%x, \ \%y$ is transformed into a 32-bit addition instruction, where the opcode is 0110011 (*add*), two source registers are $\%x$ and $\%y$, destination register is $\%tmp1$.

2) *Scheduler*: The IUVs mapped from tautologies are stored in a scheduler module, which is attached to the instruction fetch stage of the pipeline. These IUVs are dispatched to the pipeline according to the fetch instruction logic. The scheduler stores three types of instructions: IUVs, $Finish_Reg \leftarrow 1$, $Result_Reg \leftarrow 0$. $Finish_Reg$ and $Result_Reg$ are two special registers. $Result_Reg$ is initially set to 1 and used to preserve the final result. After all the instruction sequences corresponding to a tautology are committed, the satisfaction of that tautology is written back to $Result_Reg$. Once the processor violates a tautology, the value of $Result_Reg$ will be set to 0. The instruction $Finish_Reg \leftarrow 1$ flags that all IUVs have been committed.

TIUP is able to verify control flow. As depicted in Fig. 2, a part of IUVs stored in the scheduler derived from tautology $\models P \rightarrow Q$. If P is false, the branch is expected to jump to J (the exact offset is calculated in Section III-C1), but when the branch jumps to other instruction blocks or no jump occurs, the value of register $Result_Reg$ will be set to 0. If the range of the jump exceeds the length of the instruction queue in the scheduler, the instruction $Result_Reg \leftarrow 0$ is dispatched to indicate an error in the branch prediction unit.

D. Interfacing RTL Design and Model Checking

In this phase, the scheduler (see Section III-C2) and assertion modules are integrated into the DUV's RTL. Assertions are functions that evaluate a boolean value over the state-holding elements of the processor. The safety property (SP , refer to Section III-A) expressed as an assertion is inserted into the RTL. A feature of *TIUP* is its ability to reduce the burden of engineers in writing a plethora of tedious and error-prone assertions. The safety property to be verified is as follows:

$$Finish_Reg \rightarrow Result_Reg \quad (6)$$




Fig. 2. Dispatch instructions from scheduler

Upon completion of the execution and writeback of the instruction sequence under the mapping of tautologies, the register recording the results should be set to 1.

TIUP converts RTL code into a word-level model checking format - BTOR2 [21] by the open synthesis suite Yosys [22]. The model-checking tool is able to find a counterexample quickly if the DUV fails to comply with the required tautology.

IV. EVALUATION

We evaluate *TIUP* to assess its efficacy and practicality across two CPU designs. BIRISCV [23] is a dual-issue, seven-stage pipeline in-order processor, whilst RIDECORE [24] is a dual-issue, six-stage pipeline out-of-order processor. We injected 20 distinct logic anomaly scenarios from [1], [12], [14], [25] as well as some real bug scenarios into the two CPUs' RTL. Model checking was performed using the open-source tool Pono [26] on an Intel(R) Core(TM) i9-10900K CPU with 3.70GHz and 64GB of RAM.

We compared *TIUP* with *SQED* and *S²QED*, both of which are based on a universal property. In comparison with the original *SQED*, we only implemented the basic EDDI-V transformation, which targets self-consistency property checking. BMC was utilized as the proof mode for all three methods.

The results are summarized in Table I. *Category* refers to whether the anomaly is related to single-instruction or multiple-instruction. The state of being independent of previously executed instructions pertains to single-instruction anomalies, while multiple-instruction anomalies require specific instruction sequences to trigger that error scenario. *Processor* refers to the distribution of anomalies across the two processors. The anomalies *a01* and *a02* are not artificially injected into the processors. For the anomaly *a01*, *TIUP* identifies that BIRISCV does not support a division-by-zero exception. While division-by-zero exceptions are not mandatory for processor designs, they are generally used to ensure program stability and correctness. On the other hand, *TIUP* identifies that RIDECORE, as an incomplete demo, does not include division and modulo instructions because it only implements a subset of the RV32I.

The following observations can be made based on Table I:

Observation 1: As discussed in Section II, self-consistency can result in false positives. Because of the inherent limitations

of the self-consistency property, *SQED* and *S²QED* missed single-instruction anomalies, whereas *TIUP* can detect both.

Observation 2: *SQED* and *S²QED* are capable of efficiently detecting some bugs under the length of two instructions. The shortest computation spent of *TIUP* to detect a bug is longer than *SQED* and *S²QED* because the property (6) is only checked after all instructions corresponding to specified tautology have been committed. However, the longest computation time to detect a bug is shorter in *TIUP* than in *SQED* and *S²QED*. *TIUP* can verify all tautology specifications simultaneously, meaning that multiple tautology specifications may cover a bug, and we select the shortest bug-finding time. This is because the solver's search object is constrained by each tautology, thereby resulting in a shorter computation time than *SQED* and *S²QED* (Table I shows the longest computation for detecting an anomaly in *TIUP* was less than 1.5h of *SQED* and *S²QED*). Since the state space substantially influences the complexity of the proof, being able to divide the search space makes *TIUP* more scalable for large-scale designs.

Observation 3: *TIUP* was unable to detect the anomaly *a17*, where the source operand of an instruction is misidentified as 0. This anomaly, however, does not violate any of the tautologies verified in our experiments. In fact, our future research will focus on identifying the minimal set of universal properties that can comprehensively cover the behavior of the processor.

V. CONCLUSION

In this paper, we introduce *TIUP*, an end-to-end approach that employs tautologies as universal properties for formal verification of processors. *TIUP* enables engineers to verify processors at the architectural level without the need for in-depth analysis of the microarchitecture. By utilizing a set of universal properties, *TIUP* prevents the issues of false positives when using a single universal property and lets the verification problem be divided into a set of sub-problems, each of which is characterized through an abstract specification generated by a tautology. The experiment demonstrates the effective detection of anomalies in both in-order and out-of-order processors by *TIUP*.

ACKNOWLEDGEMENTS

This work was supported by Strategic Priority Research Program of Chinese Academy of Sciences under Grant No. XDC05020201. We also thank the anonymous reviewers for their feedback.

REFERENCES

- [1] F. Lonsing, K. Ganesan, M. Mann, S. S. Nuthakki, E. Singh, M. Srouji, Y. Yang, S. Mitra, and C. Barrett, "Unlocking the power of formal hardware verification with cosa and symbolic qed," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2019, pp. 1–8.
- [2] E. Clarke, "Jr., o. grumberg, and da peled," *Model Checking*, 1999.
- [3] E. M. Clarke Jr, O. Grumberg, D. Kroening, D. Peled, and H. Veith, *Model checking*. MIT press, 2018.
- [4] A. Reid, R. Chen, A. Deligiannis, D. Gilday, D. Hoyes, W. Keen, A. Pathirane, O. Shepherd, P. Vrabell, and A. Zaidi, "End-to-end verification of processors with isa-formal," in *International Conference on Computer Aided Verification*. Springer, 2016, pp. 42–58.

TABLE I
ANOMALY DETECTION RESULTS

No.	Synopsis	Category	Processor		Universal property-based method		
			RIDECORE	BIRISCV	SQED	S ² QED	TIUP
a01	No implementation of divide-by-zero exception (Return 0xffffffff)	single	x	✓	x	x	✓
a02	No implementation of the division-modulo execution unit	single	✓	x	x	x	✓
a03	Register target redirection	multiple	✓	✓	✓	✓	✓
a04	Register source redirection	multiple	✓	✓	✓	✓	✓
a05	Incorrect unsigned operand less-than compare	single	✓	✓	x	x	✓
a06	GPR0 can be assigned	multiple	✓	✓	✓	✓	✓
a07	Incorrect instruction fetched after dispatch stall	multiple	✓	✓	✓	✓	✓
a08	Incorrect instruction fetched after an LSU stall	multiple	✓	✓	✓	✓	✓
a09	One of the buggy RS-m entries corrupted the MULH/MULHU instruction	multiple	✓	x	✓	✓	✓
a10	Erroneous branch addresses	single	✓	✓	x	x	✓
a11	Erroneous branch directions	single	✓	✓	x	x	✓
a12	Error in decoding next instruction's operand	single	✓	✓	x	x	✓
a13	Processor incorrectly decodes the next instruction to a NOP instruction	multiple	✓	✓	✓	✓	✓
a14	The value of the next register read is corrupted to all 0's	multiple	✓	✓	x	x	✓
a15	Erroneous speculative instruction aren't flushed	single	✓	✓	✓	✓	✓
a16	Unsigned multiply operand converts to signed	single	✓	✓	x	x	✓
a17	Source operand is misidentified as 0	multiple	✓	✓	✓	✓	x
a18	ALU opcode does not match with the actual circuit	single	✓	✓	x	x	✓
a19	Error in determining whether instructions in decode queue have been popped	multiple	x	✓	✓	✓	✓
a20	Logical error in fetch instruction valid signal	multiple	✓	✓	✓	✓	✓
Detect single-instruction anomalies					x	x	✓
Detect multiple-instruction anomalies					✓	✓	✓
Runtime (with anomalies) [min, max]					[< 60s, > 1.5h]	[< 60s, > 1.5h]	[< 90s, < 988s]
Runtime (without anomalies)					Timeout	Timeout	< 988s
Counterexample length ([min, max] instructions)					[2, 14]	[2, 7]	[3, > 14]

- [5] J. Bormann, S. Beyer, A. Maggiore, M. Siegel, S. Skalberg, T. Blackmore, and F. Bruno, "Complete formal verification of tricore2 and other processors," in *Design and Verification Conference (DVCon)*, 2007.
- [6] S. Katz, O. Grumberg, and D. Geist, "have I written enough properties?" - A method of comparison between specification and implementation," in *Correct Hardware Design and Verification Methods, 10th IFIP WG 10.5 Advanced Research Working Conference, CHARME '99, Bad Herrenalb, Germany, September 27-29, 1999, Proceedings*, ser. Lecture Notes in Computer Science, L. Pierre and T. Kropf, Eds., vol. 1703. Springer, 1999, pp. 280–297. [Online]. Available: https://doi.org/10.1007/3-540-48153-2_21
- [7] C. Wolf, "Risc-v formal verification framework," <https://github.com/YosysHQ/riscv-formal>, 2018.
- [8] P. Subramanyan, Y. Vizel, S. Ray, and S. Malik, "Template-based synthesis of instruction-level abstractions for soc verification," in *2015 Formal Methods in Computer-Aided Design (FMCAD)*. IEEE, 2015, pp. 160–167.
- [9] B.-Y. Huang, H. Zhang, P. Subramanyan, Y. Vizel, A. Gupta, and S. Malik, "Instruction-level abstraction (ila) a uniform specification for system-on-chip (soc) verification," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 24, no. 1, pp. 1–24, 2018.
- [10] B.-Y. Huang, H. Zhang, A. Gupta, and S. Malik, "Ilang: a modeling and verification platform for socs using instruction-level abstractions," in *Tools and Algorithms for the Construction and Analysis of Systems: 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings, Part I* 25. Springer, 2019, pp. 351–357.
- [11] A. Reid, "Trustworthy specifications of arm® v8-a and v8-m system level architecture," in *2016 Formal Methods in Computer-Aided Design (FMCAD)*. IEEE, 2016, pp. 161–168.
- [12] D. Lin, T. Hong, Y. Li, S. Eswaran, S. Kumar, F. Fallah, N. Hakim, D. S. Gardner, and S. Mitra, "Effective post-silicon validation of system-on-chips using quick error detection," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2014.
- [13] D. Lin, E. Singh, C. Barrett, and S. Mitra, "A structured approach to post-silicon validation and debug using symbolic quick error detection," *International Test Conference*, 2015.
- [14] E. Singh, D. Lin, C. Barrett, and S. Mitra, "Logic bug detection and localization using symbolic quick error detection," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2018.
- [15] M. R. Fadiheh, J. Urdahl, S. S. Nuthakki, S. Mitra, C. Barrett, D. Stoffel, and W. Kunz, "Symbolic quick error detection using symbolic initial state for pre-silicon verification," *Design, Automation, and Test in Europe*, 2018.
- [16] K. Ganesan, F. Lonsing, S. S. Nuthakki, E. Singh, M. R. Fadiheh, W. Kunz, D. Stoffel, C. Barrett, and S. Mitra, "Effective pre-silicon verification of processor cores by breaking the bounds of symbolic quick error detection," *arXiv preprint arXiv:2106.10392*, 2021.
- [17] K. Devarajgowda, M. R. Fadiheh, E. Singh, C. Barrett, S. Mitra, W. Ecker, D. Stoffel, and W. Kunz, "Gap-free processor verification by s 2 qed and property generation," in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2020, pp. 526–531.
- [18] K. Devarajgowda and W. Ecker, "Meta-model based automation of properties for pre-silicon verification," in *2018 IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC)*. IEEE, 2018, pp. 231–236.
- [19] E. Singh, K. Devarajgowda, S. Simon, R. Schnieder, K. Ganesan, M. Fadiheh, D. Stoffel, W. Kunz, C. Barrett, W. Ecker *et al.*, "Symbolic qed pre-silicon verification for automotive microcontroller cores: Industrial case study," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 1000–1005.
- [20] E. W. Weisstein, "Tautology. From MathWorld—A Wolfram Web Resource," <https://mathworld.wolfram.com/Tautology.html>, last visited on 4/12/2023.
- [21] A. Niemetz, M. Preiner, C. Wolf, and A. Biere, "Btor2, btorc and boolector 3.0," in *Computer Aided Verification: 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I*. Springer, 2018, pp. 587–595.
- [22] C. Wolf, "Yosys open synthesis suite," 2016.
- [23] ultraembedded, "birisc-v - 32-bit dual issue risc-v cpu," <https://github.com/ultraembedded/biriscv>, 2021.
- [24] T. V. C. Masashi Fujinami, Susumu Mashimo and K. Kise, "Ridecore(risc-v dynamic execution core)," <https://github.com/ridecore/ridecore>, 2017.
- [25] R. Zhang, C. Deutschbein, P. Huang, and C. Sturton, "End-to-end automated exploit generation for validating the security of processor designs," in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2018, pp. 815–827.
- [26] M. Mann, A. Irfan, F. Lonsing, Y. Yang, H. Zhang, K. Brown, A. Gupta, and C. W. Barrett, "Pono: A flexible and extensible smt-based model checker," in *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part II*, ser. Lecture Notes in Computer Science, A. Silva and K. R. M. Leino, Eds., vol. 12760. Springer, 2021, pp. 461–474. [Online]. Available: https://doi.org/10.1007/978-3-030-81688-9_22