

# SecIC3: Customizing IC3 for Hardware Security Verification

Qinhan Tan, Akash Gaonkar, Yu-Wei Fan, Aarti Gupta, and Sharad Malik  
*Princeton University, Princeton, New Jersey, USA*  
 {qinhant, agaonkar, yf9172, aartig, sharad}@princeton.edu

**Abstract**—Recent years have seen significant advances in using formal verification to check hardware security properties. Of particular practical interest are checking confidentiality and integrity of secrets, by checking that there is no information flow between the secrets and observable outputs. A standard method for checking information flow is to translate the corresponding non-interference hyperproperty into a safety property on a *self-composition* of the design, which has two copies of the design composed together. Although prior efforts have aimed to reduce the size of the self-composed design, there are no state-of-the-art model checkers that exploit their special structure for hardware security verification. In this paper, we propose SecIC3, a hardware model checking algorithm based on IC3 that is customized to exploit this self-composition structure. SecIC3 utilizes this structure in two complementary techniques: *symmetric state exploration* and *adding equivalence predicates*. We implement SecIC3 on top of two open-source IC3 implementations and evaluate it on a non-interference checking benchmark consisting of 10 designs. The experiment results show that SecIC3 significantly reduces the time for finding security proofs, with up to 49.3x proof speedup compared to baseline implementations.

**Index Terms**—Hardware Security, Non-Interference Property, Model Checking

## I. INTRODUCTION

Hardware security has received increasing attention in recent years as malicious attackers have used hardware vulnerabilities to compromise secret information. Breaches to *confidentiality*, where attackers steal secrets, and *integrity*, where attackers modify them, can result in economic loss [1] and threaten public safety [2]. One defensive approach that provides strong hardware security guarantees is *formal hardware security verification*, which uses formal verification techniques to systematically check a hardware implementation against a given attack model. Formal verification uniquely provides *exhaustive* security guarantees by constructing a rigorous security proof reasoning over all possible executions of a hardware design, and potentially identifying vulnerabilities that may have been missed by other (simulation-based) approaches, such as fuzzing [3], [4], [5]. However, current work in formal security verification [6], [7], [8], [9] has been limited to small hardware systems and does not scale to real-world designs. In this paper, we propose to address this scaling gap by customizing the model checking algorithm commonly used for proving security properties, particularly confidentiality and integrity.

In the literature [10], both confidentiality and integrity specify the absence of information flow: for confidentiality, changing a secret value should not change any public output, and for integrity, changing a public input should not




Fig. 1: Non-Interference Checking via Self-Composition.

change any protected data. More formally, both are forms of a hyperproperty called *non-interference*, in which every pair of executions that differ only in the values of given source variables produce the same values of given sink variables. We use the terminology of confidentiality throughout the paper. Seminal work by Barthe et al. [11] showed that the non-interference hyperproperty can be verified by constructing a *self-composition*, as depicted in Fig. 1. Two copies of a system are composed together, such that only the secret inputs are allowed to differ between the two copies. This simulates a pair of executions in parallel, so that one can prove non-interference by showing that corresponding outputs between the two copies are always equal — a safety property over the self-composition.

To our knowledge, prior work in formal hardware security verification [7], [12] has discharged the non-interference verification condition using generic model checking algorithms [13], treating the self-composition as a black-box transition system without taking advantage of its symmetric structure. Beyond security verification, previous work in general model checking (e.g., [14], [15]) has benefited from symmetry reduction. These were based on traditional model checking algorithms, and while their techniques do not translate directly to our context, are inspirational to our work.

This paper proposes SecIC3, a customization of the state-of-the-art IC3 model checking algorithm [13], [16] to make it more effective at verifying non-interference. As part of SecIC3 we introduce two complementary techniques that improve solver performance by exploiting the special structure of the non-interference self-composition:

a) *Exploring symmetric state*: Since a self-composition consists of two identical copies of the same design, its space of reachable states is highly symmetric. In particular, a copy-relabeling mapping allows us to take any reachable state, swap corresponding variables of the two copies, and produce another reachable state. We augment the model checker with

this relabeling mapping and use it to strengthen the blocking step of IC3 by learning additional lemmas via relabeling.

b) *Adding equivalence predicates*: For large designs, existing model checkers may struggle to decompose the proof objective relating corresponding outputs of the two copies into subtasks relating corresponding internal subcomponents/variables as this correspondence is not made explicit. We aid this decomposition by extending the self-composition with auxiliary *equivalence predicates* that indicate whether corresponding subcomponents between the two copies will output the same value. We modify the IC3 algorithm to better utilize these equivalence predicates, proposing three heuristics for determining which predicates to use during lemma learning, each with their own tradeoffs.

We have implemented two prototypes of SecIC3, extending the following two open-source IC3 implementations.

- ABC-PDR [16]: a classic IC3 implementation widely used as a baseline in model checking [17], [18], [19];
- rIC3 [20]: the winner of Hardware Model Checking Competition 2025 on both bit-level and word-level tracks.

We also create a new benchmark for non-interference checking, consisting of 10 different open-source hardware designs. Our evaluations on the benchmark show that SecIC3 can improve the performance of finding security proofs by up to 49.3x times compared to the baseline implementations.

### Summary of Our Contributions

- We propose SecIC3, an IC3-based model checking algorithm customized for verifying non-interference. In the algorithm, we introduce two new techniques, *exploring symmetric state* and *adding equivalence predicates*, that use symmetry within a non-interference verification problem to guide lemma learning within the model checker.
- We construct a benchmark suite comprising 10 diverse hardware designs specifically for evaluating non-interference checkers.
- We implement and evaluate SecIC3 on two open-source IC3 implementations, demonstrating that SecIC3 achieves up to 49.3x times speedup over baseline implementations in proving non-interference for this benchmark suite.

## II. BACKGROUND: IC3/PDR

IC3 [13], also known as *Property Directed Reachability (PDR)* [16] is a state-of-the-art hardware model checking algorithm for unbounded proofs of safety properties. We will assume some familiarity with the IC3 algorithm, but briefly review the parts relevant to our work. A more complete exposition can be found in Een et al. [16].

Given a transition system  $(Init, Tr)$  with state variables  $x$ , initial states  $Init(x)$ , and transition relation  $Tr(x, x')$ , suppose we wish to check a safety property  $P(x)$ . (We follow the standard notation of using a variable  $x$  to refer to current state, and  $x'$  as the next state.) As depicted in Fig. 2, IC3 tries to construct a safe inductive invariant proving  $P$  by maintaining an *inductive trace*, a sequence of formulas,  $F_0, F_1, \dots, F_n$ , called *frames*. Each frame  $F_k$  represents an over-approximation of the reachable states within  $k$  transitions, with all but the latest frame proven safe with respect to  $P$ . Formally,




Fig. 2: The frames of an IC3 inductive trace when checking a property  $P$ . IC3 refines frames by *blocking* bad states: the state is shown to be part of an unreachable cube  $c$ , and is thereby excluded from the latest frame  $F'_n = F_n \wedge \neg c$ .

- 1)  $F_0(x) = Init(x)$
- 2)  $\forall k < n, F_k(x) \implies F_{k+1}(x)$
- 3)  $\forall k < n, F_k(x) \wedge Tr(x, x') \implies F_{k+1}(x')$
- 4)  $\forall k < n, F_k(x) \implies P(x)$

IC3 incrementally refines and adds frames until either the latest frame becomes inductive ( $F_n \iff F_{n-1}$ ) or a bad state that violates  $P$  is found to be reachable from  $Init$ . The former case completes the safety proof, where  $F_n$  is an inductive invariant that holds on all reachable states of  $(Init, Tr)$  and that implies  $P(x)$ . In the latter case, an execution trace from an initial state to the bad state is returned as a counterexample to the safety of  $P$ .

If the current frame is not yet safe for  $P$ , IC3 needs to refine frames during verification. It identifies bad states or their predecessors that can be shown to be unreachable from the previous frame. As shown in Fig. 2, an unreachable bad state is generalized into a conjunction of literals (a variable or its negation), called a *cube*. This cube is then *blocked* in frames where it is provably unreachable by updating the frame's formula to include its negation. In the figure,  $F'_n = F_n \wedge \neg c$  updates  $F_n$  to incorporate the fact that cube  $c$  is unreachable within the first  $n$  transitions. This refinement procedure, called **BLOCK**, is the key step studied in our paper.

Algorithm 1 details how **BLOCK** checks the reachability of and potentially blocks a cube  $s$  on frame  $k$ . Line 3 uses a SAT query to check if  $s$  is unreachable in frame  $k$ . This step is referred to as a *reachability test* in the rest of this paper. If  $s$  is unreachable (UNSAT), a **GENERALIZE** procedure expands  $s$  into a cube  $c$  describing an unreachable superset of states (essentially by removing literals from  $s$ ). Cube  $c$  is then blocked in Line 6 by adding  $\neg c$  to frames up to  $k$ . On the other hand, if  $s$  is reachable (SAT), a predecessor cube  $t$  is given by the SAT solver. IC3 recursively tries to block  $t$  in the previous frame  $k-1$  (Line 10), returning to the current frame if successful (Line 11) or exiting with a failing execution trace if frame 0 is reached (Line 2).

## III. SECIC3 DESIGN

In this section, we detail our customizations to IC3 that use the symmetry of a self-composed design to improve

---

**Algorithm 1** BLOCK( $s, k, Tr$ )

---

**Require:** cube  $s$ , frame number  $k$ , and transition relation  $Tr$

- 1: **if**  $k = 0$  :
- 2:   **return false**   ▷ reaches  $F_0 = Init$ : real counterexample
- 3: **if** UNSAT( $F_{k-1} \wedge \neg s(x) \wedge Tr \wedge s(x')$ ) : ▷ reachability test
- 4:    $c \leftarrow \text{GENERALIZE}(s, k)$    ▷ a generalized cube
- 5:   **for**  $j \leftarrow 1$  **to**  $k$  :
- 6:      $F_j \leftarrow F_j \wedge \neg c$
- 7:   **return true**
- 8: **else** :
- 9:    $t \leftarrow$  predecessor cube for  $s$  in frame  $k - 1$
- 10:   **if** BLOCK( $t, k-1, Tr$ ) :
- 11:     **return** BLOCK( $s, k, Tr$ )   ▷ retry  $s$  after predecessor
- 12:   **else** :
- 13:     **return false**

---




Fig. 3: IC3 Example on Checking Non-Interference performance when verifying non-interference properties.

#### A. Overview

As an illustrative example, consider the self-composition shown in Fig. 3, used to check non-interference in a circuit containing a multiplexer and two registers. Call the registers  $s$ , for selector (1-bit), and  $r$ , for result ( $n$ -bit, for some  $n$ ), and suppose also that some input assumptions or combinational logic ensure  $s = 0$  always. We will use this example to point out some inefficiencies in the traditional IC3 algorithm that motivate our symmetry-based techniques.

When we use IC3 to prove  $output_1 == output_2$ , the algorithm repeatedly calls BLOCK on various cubes, one at a time, in order to construct a safe inductive invariant.

**Inefficiency 1 (Repeated reasoning).** For our example, one cube IC3 must block is  $(s_1)$ , because if  $s_1$  is high then in the coming cycles  $secret_1$  will leak into  $output_1$  producing a bad state. By symmetry, the same logic applies to cube  $(s_2)$ , but IC3 will likely separately repeat a very similar sequence of bad states and sat queries in order to block this cube also.

**Inefficiency 2 (Equalities are expensive).** After blocking the cubes  $(s_1)$  and  $(s_2)$ , IC3 must prove  $r_1 == r_2$  always holds by showing that  $r_1 \neq r_2$  is unreachable. Crucially, a cube (a conjunction of boolean literals) cannot individually capture complex relational constraints, such as register (in-)equality. To block all states where the  $n$ -bit registers  $r_1, r_2$  are unequal, IC3 must block two cubes for each bit  $i \in [0, n)$ :  $(r_1[i] \wedge \neg r_2[i])$  and  $(\neg r_1[i] \wedge r_2[i])$ . This requires (at least)  $2n$  calls to BLOCK, itself potentially recursive. A similar cost is paid for every equality IC3 must prove. In this example,

---

**Algorithm 2** SecIC3 modification on BLOCK (marked as blue)

---

- 1: **if** UNSAT( $F_{k-1} \wedge \neg s(x) \wedge Tr \wedge s(x')$ ) : ▷ reachability test
- 2:    $c \leftarrow \text{GENERALIZE}(s, k)$    ▷ a generalized cube
- 3:    $c \leftarrow \text{PREDICATE\_REPLACE}(c)$
- 4:    $c_{sym} \leftarrow \text{SYMMETRIC\_CUBE}(c)$
- 5:   **for**  $j \leftarrow 1$  **to**  $k$  :
- 6:      $F_j \leftarrow F_j \wedge \neg c \wedge \neg c_{sym}$
- 7:   **return true**

---




Fig. 4: Equivalence Predicate Construction

IC3 proves  $r_1 == r_2$  and soon after concludes that the non-interference property holds. In general, IC3 may need many equalities relating corresponding components in (the two copies of) a self-composition in order to prove non-interference for a large design.

SecIC3 addresses these inefficiencies through two techniques: *Symmetric State Exploration* and *Adding Equivalence Predicates*, respectively. Algorithm 2 shows the fragment of the BLOCK procedure that we modify to incorporate these techniques. The code snippet corresponds to lines 3-7 in Alg. 1, omitting the rest of BLOCK as it is unchanged. After IC3 generalizes an unreachable cube  $c$ , SecIC3 will call an additional generalization procedure PREDICATE\_REPLACE to update  $c$  to use equivalence predicates where appropriate (Line 3). It then calls SYMMETRIC\_CUBE to construct a second cube  $c_{sym}$  based on the symmetry of self-composition that is also guaranteed to be unreachable (Line 4). Both cubes  $c$  and  $c_{sym}$  are blocked in the current frame (Line 6). We further detail the novel techniques following their program order in Algorithm 2.

#### B. Adding Equivalence Predicates

To help IC3 prove equalities (addressing Inefficiency 2), we augment the self-composition with auxiliary boolean variables, called (*in*)*equivalence predicates*, that IC3 can block as part of a cube in order to guarantee a pair of corresponding registers between the two copies are equivalent. (We regret the double-negative, IC3 blocks inequivalence predicates.) We detail our augmentation and the PREDICATE\_REPLACE procedure that uses these predicates to further generalize a cube.

1) *Introducing Word-Level Inequivalence Predicates:* For every pair of registers  $r_1, r_2$  between the two copies in a self-composition, a 1-bit register,  $neq_r$ , is added to represent whether these two registers are unequal. We update the self-composition as shown in Fig. 4, adding a word-level not-equals comparison between the next state values of  $r_1$  and  $r_2$ , and store the result in  $neq_r$ . If IC3 blocks  $neq_r$ , then it has proven (in a single cube!) that the two registers  $r_1, r_2$  are always equal, regardless of the value of each copy's secret.

2) *Predicate Replacement in the Generalized Cube:* Given an augmented self-composition, IC3 is fully capable of reasoning about inequivalence predicates out-of-the-box. However,




Fig. 5: A lattice (partial order) of possible predicate replacements given the cube at the bottom of the lattice.

we found that in practice, many opportunities to block an inequivalence predicate were missed by the standard IC3 generalization procedure. SecIC3 adds an additional PREDICATE\_REPLACE step (Algorithm 2, Line 3) on each unreachable cube  $c$  after generalization, seeking to smartly replace registers in  $c$  with their corresponding  $neq$  inequivalence predicates while still preserving unreachability. Continuing our running example (Fig. 3), if IC3 finds an unreachable, generalized cube  $c := (r_1[0] \wedge \neg r_2[0])$ , it may be prudent to replace it with  $(neq_r)$ , rather than waiting to block the other half of  $r_1[0] \neq r_2[0]$ , and then incrementally blocking each bit of  $r$  ( $2n$  calls to BLOCK). We are careful, as part of this replacement, to use additional reachability tests (Algorithm 1, Line 3) to ensure the new cube remains unreachable.

In PREDICATE\_REPLACE, we use the pattern  $x_a[i] \wedge \neg x_b[i]$  (for some register  $x$ ; copies  $a, b$ ; bit  $i$ ) to identify opportunities to use an inequivalence predicate  $neq_x$ , as the pattern is a common part of cubes that facilitate equality proofs in IC3. Given a cube, we call each subformula that matches this pattern an *inequivalence group*, and choose whether or not to replace it with its corresponding predicate heuristically. We discuss three representative heuristics below, and compare their performance in our evaluation Section IV.

**All-or-Nothing Replacement** replaces all inequivalence groups with their corresponding  $neq$  variables and does one reachability test. For example, the cube at the bottom of Fig. 5, would be generalized by *All-or-Nothing Replacement* into the cube at top of the lattice. PREDICATE\_REPLACE would check its reachability, returning the top if it is unreachable, and otherwise the original cube at the bottom.

**Maximal Replacement** cumulatively replaces the inequivalence groups one by one, skipping groups if the replacement causes the cube to be reachable. In Fig. 5, *Maximal Replacement* would traverse from the bottom of the lattice, moving up when the cube is unreachable and to the right (another parent of the last unreachable cube) when the cube is not. PREDICATE\_REPLACE returns the last unreachable cube, defaulting to the original, at the bottom. Note that different orders of inequivalence groups can result in different final cubes, but each is optimal in that no additional inequivalence group can be replaced without becoming reachable.

**Maximum Replacement** identifies the set of all cubes that could be returned by Maximal Replacement, across all orderings of inequivalence groups. To do this efficiently, *Maximum Replacement* would traverse Fig. 5 by starting with a singleton set  $M$ , containing the cube at the top of the lattice. For each

cube added to  $M$ , we perform a reachability test. Unreachable cubes are left as-is, but reachable cubes are removed from the  $M$  and replaced with their lower-cover (all direct children in the lattice), filtering out cubes already subsumed by a cube in  $M$ . PREDICATE\_REPLACE returns the final cubes in set  $M$ , so that they can all be blocked in Algorithm 2.

These heuristics represent a tradeoff between using fewer SAT queries or blocking more states. For a cube with  $m$  inequivalence groups: *All-or-Nothing Replacement* costs one query, but gives up immediately; *Maximal Replacement* costs  $m$  queries, but is affected by the order in which inequivalence groups are replaced; *Maximum Replacement* costs up to  $2^m$  queries (in the worst case), but exhaustively blocks as much of the unreachable state-space as possible.

### C. Symmetric State Exploration

Given any reachability property of a self-composition (e.g., that a given cube is unreachable), it is always possible to construct a symmetric reachability property that swaps variables of the two copies and continues to hold on the same self-composition. To see this, note the symmetry of the space of reachable states: Suppose a state  $(copy_1, copy_2)$  can be reached with input  $(pub, sec_1, sec_2)$ , where  $pub$  corresponds to public inputs and  $sec_1, sec_2$  correspond to the secret inputs of each copy (over multiple cycles). Then the symmetric state  $(copy_2, copy_1)$  must be reachable through input that swaps the secrets of each copy,  $(pub, sec_2, sec_1)$ , because the two copies of the self-composition are identical.

In SecIC3, when we construct the self-composition, we record pairs of corresponding state variables (registers) between the two copies. During verification, whenever a cube is proved unreachable, we construct the symmetric unreachable cube by substituting variables with their paired counterpart. The one subtlety is that after predicate replacement, unreachable cubes may contain some equivalence predicates, which are auxiliary variables not part of either copy of the self-composition. Because these predicates check equality they are their own symmetric counterpart and thus are retained as is in the symmetric unreachable cube.

## IV. EVALUATION

We evaluate the performance of our SecIC3 implementation against the two open-source IC3 implementations it modifies, ABC-PDR [21] and rIC3 [20]. As part of our evaluation, we collected and organized several open-source hardware designs into a benchmark for information flow checking in hardware. We contribute this benchmark to the community and briefly overview its contents here. The implementation and evaluation code are open-sourced (<https://github.com/qinhant/SecIC3>).

### A. Implementation

SecIC3 is implemented as a Python frontend connected to custom versions of ABC-PDR (C) and rIC3 (Rust). Given a self-composition circuit and a non-interference property, the frontend (i) (optionally) adds inequivalence predicates, and (ii) uses the Yosys open synthesis suite [22] to convert the circuit into an And-Inverter-Graph (AIG) for verification. This AIG is verified by ABC-PDR or rIC3, which we have customized with additional options to enable our two verification techniques:

Multiplier	64-bit shift-and-add integer multiplier 194 inputs, 530 registers, 7926 gates
Modexp [23]	64-bit modular exponentiation calculator 258 inputs, 132 registers, 1300 gates
GCD [24]	8-bit greatest common divisor calculator 27 inputs, 38 registers, 836 gates
FP_ADD [25]	Single-precision floating pointer adder 197 inputs, 488 registers, 8695 gates
FP_MUL [25]	Single-precision floating pointer multiplier 197 inputs, 512 registers, 18128 gates
FP_DIV [25]	Single-precision floating pointer divider 197 inputs, 836 registers, 9531 gates
SecEnclave [26]	Secure enclave with AES encryption/decryption and integer arithmetic 524 inputs, 2866 registers, 1199964 gates
Cache [27]	4-way 4-word cache 270 inputs, 718 registers, 5630 gates
Sodor [28]	5-stage in-order RISCv core (RV32I) 132 inputs, 4004 registers, 51063 gates
Rocket [29]	5-stage in-order RISCv core (RV64GC) 521 inputs, 9026 registers, 123774 gates

TABLE I: Non-Interference Checking Benchmark

- *Symmetry* enables symmetric state exploration (§ III-C)
- *All-or-nothing/Maximal/Maximum* enables predicate replacement following the three heuristics in § III-B.

As the two techniques are independent, we will evaluate all configurations of symmetry and predicate replacement.

#### B. Hardware Non-Interference Checking Benchmark

We address the lack of a standard benchmark for evaluating information flow in hardware designs through constructing a benchmark with 10 representative open-source hardware designs including arithmetic modules, memory modules and processor cores, as shown in Table I. The statistics of the self-composition circuits are also shown in the table. The only output of these circuits is the property output.

For all designs, the non-interference property we are proving models *constant-time execution*, a form of confidentiality that is key in preventing microarchitectural timing side channels. For arithmetic modules, the source is the input operands and the sink is the `out_valid` signal, which serves as a proxy for execution time. For the cache module, the source is the input address and the sink is the output `hit` signal. For two processor cores, the source is the register file and the sink is the instruction `retire` signal. The register file is initialized with a free secret variable and operates normally after initialization. All 10 designs are common in that they do not satisfy constant-time execution when all inputs are free, but they are constant-time with constraints on the inputs. This is consistent with real practice. For example, the Sodor processor has constant-time and non-constant-time instructions, and we want a security proof for the constant-time instructions. Thus, we constrain the instruction memory to have only R-type and I-type instructions, with no branch instructions allowed.

#### C. Experiment Setup

We use the default IC3 mode in ABC-PDR and rIC3 as two baselines and evaluate SecIC3 on each of them with




Fig. 6: Proof Speedup Statistics

7 configurations: *Symmetry*, *All-or-Nothing*, *Maximal*, *Maximum*, *Symmetry+All-or-Nothing*, *Symmetry+Maximal*, *Symmetry+Maximum*. The time limit for each model checking task is 1 hour. Experiments are run on a machine with 10 M1-MAX cores (8 3.2Hz cores and 2 2.1Hz cores) and 32GB memory.

#### D. Experiment Results

Comparisons with ABC-PDR are shown in Table II. Relative to the baseline, SecIC3 can prove non-interference for two more designs, FP\_ADD and SecEnclave. For others, SecIC3 significantly improves the proof time or proof bounds (A proof bound  $n$  shows the property holds in the first  $n$  cycles.).

Comparisons with rIC3 are shown in Table III. While SecIC3 does not verify more designs than baseline, it does improve the proof bound of FP\_DIV from 62 to 103, and the proof bound of FP\_MUL from 14 to 15. For other cases, e.g., Cache, SecIC3 significantly improves proof time.

Among different configurations of SecIC3, no option is dominant over others in either table. However, we note that for some designs (and solvers), such as Cache in Table II (ABC-PDR) or SecEnclave in Table III (rIC3), using predicate replacement alone can significantly deteriorate performance relative to baseline. Therefore, *Symmetry* along with predicate replacement should be the default configuration of SecIC3.

Fig. 6 shows speedup statistics for each successfully verified design, dividing the best result among the two baselines by the best result among all SecIC3 configurations (across both backends). SecIC3 significantly improves the proof efficiency for all designs in the benchmark, with speedups ranging from 2.0x to 49.3x. The average speedup is 16.5x among these 8 cases. For the timeout cases FP\_MUL and FP\_DIV (not shown in Fig. 6), SecIC3 can improve the proof bound to 15 and 103 cycles, respectively, while the best among the two baselines can only prove 14 and 62 cycles.

## V. RELATED WORK

### A. Self-Composition

Self-composition was first proposed by Barthe et al. [11] to check program security properties. Many subsequent efforts focused on effective construction of self-composed programs (e.g., [30], [31], [32]) to improve verification performance, but they largely used model checkers as black boxes during safety verification of the self-composition.

Self-composition has been widely used in verifying hardware security properties such as constant-time execution [6], [33] and speculative leakage contracts [7], [34], [9]. These properties are all based on non-interference but they differ in their assumptions. For constant-time execution properties (as in our benchmarks), assumptions are typically logical expressions on input variables. In contrast, for speculative leakage

	ABC-PDR	Sym	AoN	Maximal	Maximum	Sym+AoN	Sym+Maximal	Sym+Maximum
Multiplier	11.5	3.0	0.1	0.1	0.1	0.1	0.1	0.1
Modexp	1.5	0.6	1.8	1.6	0.7	1.6	1.5	0.4
GCD	74.3	75.1	12.7	39.3	9.3	35.1	131.9	50.7
FP_ADD	TO(12)	TO(11)	466.2	TO(11)	150.3	143.7	186.0	TO(13)
FP_MUL	TO(12)	TO(9)	TO(11)	TO(13)	TO(12)	TO(12)	TO(12)	TO(13)
FP_DIV	TO(61)	TO(59)	TO(57)	TO(52)	TO(43)	TO(59)	TO(65)	TO(36)
SecEnclave	TO(7)	608.4	2.5	3.6	1.9	2.4	2.1	1.9
Cache	1419.1	370.5	783.0	TO(18)	TO(15)	1112.5	162.0	220.6
Sodor	5.0	2.7	0.9	0.8	0.9	0.4	0.4	0.4
Rocket	4.0	2.3	3.5	2.8	2.7	1.9	1.8	1.8

TABLE II: Evaluation results on ABC-PDR. Sym stands for Symmetry and AoN stands for All-or-Nothing. Time is shown in seconds. For timeout cases (TO), the proof bound is shown in ().

	rIC3	Sym	AoN	Maximal	Maximum	Sym+AoN	Sym+Maximal	Sym+Maximum
Multiplier	2.7	2.7	0.2	0.2	0.3	0.2	0.2	0.3
Modexp	0.6	0.3	0.7	0.8	1.9	0.6	1.0	1.7
GCD	7.1	13.5	0.8	3.3	26.2	3.7	1.4	2.2
FP_ADD	249.0	576.3	1925.4	62.4	80.5	44.8	69.0	114.5
FP_MUL	TO(14)	TO(14)	TO(14)	TO(14)	TO(14)	TO(14)	TO(15)	TO(15)
FP_DIV	TO(62)	TO(77)	TO(103)	TO(60)	TO(60)	TO(103)	TO(67)	TO(56)
SecEnclave	92.1	97.7	457.5	496.0	764.5	98.7	93.4	130.5
Cache	2970.2	44.0	207.3	71.4	94.5	112.1	87.2	88.4
Sodor	1.5	1.9	2.0	2.1	3.3	2.0	1.9	4.0
Rocket	6.3	6.5	6.9	7.5	9.3	7.3	8.4	12.0

TABLE III: Evaluation results on rIC3. Headings, units, and notation match Table II.

contracts, the assumptions can themselves be non-interference properties — we hope to address solver customization for such properties in future work.

### B. Hardware Information Flow Tracking

Aside from self-composition, *information flow tracking* has also been used to prove non-interference in hardware [35], [4], [36]. In this technique, the design is extended with additional “taint” signals that track if a wire/register either (i) is a secret or (ii) may depend on secrets based on the taint signals of its fan-in. This *taint propagation* is conservative [37]: proving that the taint signals of public outputs are always low guarantees non-interference, but counterexamples where output taint signals are high may be spurious (i.e., they do not leak secrets). Thus, while information flow tracking provides an alternative trade-off point in checking information flow, self-composition remains the most precise way to do this.

### C. IC3 Variants

Since the proposal of IC3 in 2011 [13], various variants have emerged. Word-level IC3 such as IC3ia [19] and IC3sa [18] intends to create word-level abstraction of hardware design and reason about it in first-order logic using SMT solvers. This avoids the expensive bit-blasting in bit-level reasoning, but additionally requires abstraction refinement, which may be more expensive. The equivalence predicates in SecIC3 also add some word-level reasoning ability to IC3, but it does not introduce abstraction and thus avoids abstraction refinement.

PDRER [38] dynamically adds auxiliary variables to the circuit during model checking based on the lemmas IC3 has learned so far. Those variables can then be used in IC3 to speed up verification. This idea is similar to the predicate variables in SecIC3, but the equivalence predicates in SecIC3 are added based on the knowledge of the self-composition design rather than the patterns in the learned lemmas.

## VI. CONCLUSION

In this paper, we propose SecIC3, a customized hardware model checker for hardware security verification based on the IC3 model checking algorithm. SecIC3 leverages the special structure of the self-composition circuit used in proving the non-interference property to extend IC3 with two key techniques, *symmetric state exploration* and *adding equivalence predicates* to guide the model checker. We implemented SecIC3 on top of (i) ABC-PDR, a widely used open-source IC3 implementation and (ii) rIC3, the winner of the 2025 Hardware Model Checking Competition in both the bit-level and word-level tracks, and evaluated its performance on a new security verification benchmark developed as part of this project consisting of 10 hardware designs. The evaluation results have shown that SecIC3 significantly improves the proof performance in security verification and can bring up to 49.3x proof speedup or increase the proof bound within 1 hour by at most 41 cycles compared to state-of-the-art baseline implementations.

## REFERENCES

- [1] IBM, “Cost of a data breach report 2024,” <https://www.ibm.com/reports/data-breach>, 2024, [Accessed 01-07-2025].
- [2] J. Shield, B. D. Hopkins, M. R. Beaumont, and C. North, “Hardware trojans—a systemic threat.” in *AISC*, 2015, pp. 45–51.
- [3] J. Hur, S. Song, S. Kim, and B. Lee, “Specdoctor: Differential fuzz testing to find transient execution vulnerabilities,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1473–1487.
- [4] F. Solt, B. Gras, and K. Razavi, “CellIFT: Leveraging cells for scalable and precise dynamic information flow tracking in RTL,” in *31st USENIX Security Symposium*, 2022.
- [5] J. Xu, Y. Zhou, X. Zhang, Y. Li, Q. Tan, Y. Zhang, Y. Zhou, R. Chang, and W. Shen, “Dejavuzz: Disclosing transient execution bugs with dynamic swappable memory and differential information flow tracking assisted processor fuzzing,” *arXiv preprint arXiv:2504.20934*, 2025.
- [6] K. v. Gleissenthall, R. G. Kıcı, D. Stefan, and R. Jhala, “IODINE: Verifying Constant-Time execution of hardware,” in *28th USENIX Security Symposium*, 2019.
- [7] Q. Tan, Y. Yang, T. Bourgeat, S. Malik, and M. Yan, “Rtl verification for secure speculation using contract shadow logic,” in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1*, 2025, pp. 970–986.
- [8] Y. Yang, T. Bourgeat, S. Lau, and M. Yan, “Pensieve: Microarchitectural modeling for security evaluation,” in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 2023, pp. 1–15.
- [9] Z. Wang, G. Mohr, K. von Gleissenthall, J. Reineke, and M. Guarnieri, “Specification and verification of side-channel security for open-source processors via leakage contracts,” in *Proceedings of the 30th ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS 2023. ACM, 2023, pp. 2128–2142.
- [10] J. A. Goguen and J. Meseguer, “Security policies and security models,” in *1982 IEEE Symposium on Security and Privacy*, 1982, pp. 11–11.
- [11] G. Barthe, P. R. D’Argenio, and T. Rezk, “Secure information flow by self-composition,” in *Proceedings. 17th IEEE Computer Security Foundations Workshop*. IEEE Computer Society, 2004, pp. 100–100.
- [12] Q. Tan, Y. Fisseha, S. Chen, L. Biernacki, J.-B. Jeannin, S. Malik, and T. Austin, “Security verification of low-trust architectures,” in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023, pp. 945–959.
- [13] A. R. Bradley, “Sat-based model checking without unrolling,” in *International Workshop on Verification, Model Checking, and Abstract Interpretation*. Springer, 2011, pp. 70–87.
- [14] E. A. Emerson and A. P. Sistla, “Symmetry and model checking,” in *Computer Aided Verification CAV, Proceedings*, ser. Lecture Notes in Computer Science, vol. 697. Springer, 1993, pp. 463–478.
- [15] D. Tang, S. Malik, A. Gupta, and C. N. Ip, “Symmetry reduction in sat-based model checking,” in *International Conference on Computer Aided Verification*. Springer, 2005, pp. 125–138.
- [16] N. Eén, A. Mishchenko, and R. Brayton, “Efficient implementation of property directed reachability,” in *2011 Formal Methods in Computer-Aided Design (FMCAD)*. IEEE, 2011, pp. 125–134.
- [17] G. Hu, J. Tang, C. Yu, W. Zhang, and H. Zhang, “Deepic3: Guiding ic3 algorithms by graph neural network clause prediction,” in *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2024, pp. 262–268.
- [18] A. Goel and K. Sakallah, “Model checking of verilog rtl using ic3 with syntax-guided abstraction,” in *NASA Formal Methods: 11th International Symposium, NFM 2019, Houston, TX, USA, May 7–9, 2019, Proceedings II*. Springer, 2019, pp. 166–185.
- [19] A. Cimatti, A. Griggio, S. Mover, and S. Tonetta, “Infinite-state invariant checking with ic3 and predicate abstraction,” *Formal Methods in System Design*, vol. 49, no. 3, pp. 190–218, 2016.
- [20] Y. Su, Q. Yang, Y. Ci, T. Bu, and Z. Huang, “The ric3 hardware model checker,” in *International Conference on Computer Aided Verification*. Springer, 2025, pp. 185–199.
- [21] R. Brayton and A. Mishchenko, “Abc: An academic industrial-strength verification tool,” in *International Conference on Computer Aided Verification*. Springer, 2010, pp. 24–40.
- [22] C. Wolf *et al.*, “Yosys open synthesis suite,” 2016.
- [23] A. Singh, “Rsa verilog implementation,” <https://github.com/Rajandeep/RSA-CRYPTOSYSTEM-using-verilog.git>, [Accessed 28-08-2024].
- [24] O. Muhammad, “Gcd verilog implementation,” <https://github.com/Oarmuhammadmu/GCD-FSMD>, [Accessed 28-08-2024].
- [25] J. Dawson, “Fpu verilog implementation,” <https://github.com/dawsonjon/fpu.git>, [Accessed 28-08-2024].
- [26] S. Chen, “Secure enclave verilog implementation,” <https://github.com/SE-and-verification/SE-prototype/tree/leakyMult>, [Accessed 28-08-2024].
- [27] E. Ogawa, “Cache verilog implementation,” <https://github.com/airin711/Verilog-caches>, [Accessed 28-08-2024].
- [28] C. Celio and Et al, “riscv-sodor github repository,” <https://github.com/ucb-bar/riscv-sodor>, 2024, [Accessed 17-09-2024].
- [29] K. Asanović, R. Avizienis, J. Bachrach, S. Beamer, D. Biancolin, C. Celio, H. Cook, D. Dabbelt, J. Hauser, A. Izraelevitz, S. Karandikar, B. Keller, D. Kim, J. Koenig, Y. Lee, E. Love, M. Maas, A. Magyar, H. Mao, M. Moreto, A. Ou, D. A. Patterson, B. Richards, C. Schmidt, S. Twigg, H. Vo, and A. Waterman, “The rocket chip generator,” EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2016-17, 2016. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-17.html>
- [30] T. Terauchi and A. Aiken, “Secure information flow as a safety problem,” in *Static Analysis Symposium SAS, Proceedings*, ser. Lecture Notes in Computer Science, vol. 3672. Springer, 2005, pp. 352–367.
- [31] G. Barthe, J. M. Crespo, and C. Kunz, “Relational verification using product programs,” in *Formal Methods FM, Proceedings*, ser. Lecture Notes in Computer Science, vol. 6664. Springer, 2011, pp. 200–214.
- [32] R. Shemer, A. Gurfinkel, S. Shoham, and Y. Vizel, “Property directed self composition,” in *Computer Aided Verification: 31st International Conference, CAV 2019, New York City, NY, USA, July 15–18, 2019, Proceedings, Part I 31*. Springer, 2019, pp. 161–179.
- [33] K. v. Gleissenthall, R. G. Kıcı, D. Stefan, and R. Jhala, “Solver-aided constant-time hardware verification,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 429–444.
- [34] T. Jauch, A. Wezel, M. R. Fadiheh, P. Schmitz, S. Ray, J. M. Fung, C. W. Fletcher, D. Stoffel, and W. Kunz, “Secure-by-construction design methodology for cpus: Implementing secure speculation on the rtl,” in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE, 2023, pp. 1–9.
- [35] M. Tiwari, H. M. Wassel, B. Mazloom, S. Mysore, F. T. Chong, and T. Sherwood, “Complete information flow tracking from the gates up,” in *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2009.
- [36] A. Ardeshricham, W. Hu, J. Marxen, and R. Kastner, “Register transfer level information flow tracking for provably secure hardware design,” in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2017.
- [37] W. Hu, J. Oberg, A. Irturk, M. Tiwari, T. Sherwood, D. Mu, and R. Kastner, “Theoretical fundamentals of gate level information flow tracking,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 30, no. 8, pp. 1128–1140, 2011.
- [38] A. Luka and Y. Vizel, “Property directed reachability with extended resolution,” in *International Conference on Computer Aided Verification*. Springer, 2025, pp. 258–280.