

Systematic Classification of Attackers via Bounded Model Checking

Eric Rothstein-Morris¹, Sun Jun², and Sudipta Chattopadhyay¹

¹Singapore University of Technology and Design
{eric.rothstein,sudipta.chattopadhyay}@sutd.edu.sg
² Singapore Management University
junsun@smu.edu.sg



Abstract. In this work, we study the problem of verification of systems in the presence of attackers using bounded model checking. Given a system and a set of security requirements, we present a methodology to generate and classify attackers, mapping them to the set of requirements that they can break. A naive approach suffers from the same shortcomings of any large model checking problem, i.e., memory shortage and exponential time. To cope with these shortcomings, we describe two sound heuristics based on cone-of-influence reduction and on learning, which we demonstrate empirically by applying our methodology to a set of hardware benchmark systems.

1 Introduction

Problem Context. Some systems are designed to provide security guarantees in the presence of attackers. For example, the Diffie-Hellman key agreement protocol guarantees perfect forward secrecy [21, 27] (PFS), i.e., that the session key remains secret even if the long-term keys are compromised. These security guarantees are only valid in the context of the attacker models for which they were proven; more precisely, those guarantees only hold for attackers that fit those or weaker attacker models. For instance, PFS describes an attacker model (i.e., an attacker that can compromise the long-term keys, *and only those*), and a property that is guaranteed in the presence of an attacker that fits the model (i.e., confidentiality of the session keys). However, if we consider an attacker model that is stronger (e.g., an attacker that can directly compromise the session key), then Diffie-Hellman can no longer guarantee the confidentiality of the session keys. Clearly, it is difficult to provide any guarantees against an attacker model that is too capable, so it is in the interest of the system designer to choose an adequate attacker model that puts the security guarantees of the system in the context of realistic and relevant attackers.

Consider the following research question: **RQ1**) given a system and a list of security requirements, how do we systematically generate attackers that can potentially break these requirements, and how do we verify if they are successful? We approach this question at a high level for a system S with a set C of n components, and a set of security requirements R as follows. Let A be a subset

of C ; the set A models an attacker that can interact with S by means of each component c in it. More precisely, for every component c in A , the attacker can change the value of c at any time and any number of times during execution, possibly following an attack strategy. Considering the exponential size of the set of attackers (i.e., 2^n), a brute-force approach to checking whether each of those attackers breaks each requirement in R is inefficient for two reasons: 1) an attacker A may only affect an isolated part of the system, so requirements that refer to other parts of the system should not be affected by the presence of A , and 2) if some attacker B affects the system in a similar way to A (e.g., if they control a similar set of components), then the knowledge we obtain while verifying the system in the presence of A may be useful when verifying the system in the presence of B . These two reasons motivate a second research question: **RQ2**) which techniques can help us to *classify* attackers, i.e., to map each attacker to the set of requirements that it breaks?

To answer these two research questions in a more concrete and practical context, we study systems modelled by *And-inverter Graphs (AIGs)* (see [23, 25]). AIGs describe hardware models at the bit-level [10], and have attracted the attention of industry partners including IBM and Intel [16]. Due to being systems described at bit-level, AIGs present a convenient system model to study the problem of attacker classification, because the range of actions that attackers have over components is greatly restricted: either the attacker leaves the value of the component as it is, or the attacker negates its current value. However, this approach can be generalised to other systems by considering non-binary ranges for components, and by allowing attackers to choose any value in those ranges.

Contributions. In this paper, we provide:

- a formalisation of attackers of AIGs and how they interact with systems,
- a methodology to perform bounded model checking while considering the presence of attackers,
- a set of heuristics that characterise attacker frontiers for invariant properties using bounded model checking,
- experimental evidence of the effectiveness of the proposed methodology and heuristics.

2 Preliminaries

In this section, we provide the foundation necessary to formally present the problem of model checking And-inverter Graph (AIGs) in the presence of attackers. Let $\mathbb{B} = \{0, 1\}$ be the set of booleans/bits. An *And-inverter Graph* models a system of equations that has m boolean inputs, n boolean state variables and o boolean gates. The elements in the set $W = \{w_1, \dots, w_m\}$ represent the *inputs*, the elements in $V = \{v_1, \dots, v_n\}$ represent the *latches*, and the elements in $G = \{g_1, \dots, g_o\}$ represent the *and-gates*. We assume that W , V and G are pairwise disjoint, and we define the set of *components* C by $C \triangleq W \cup V \cup G$.

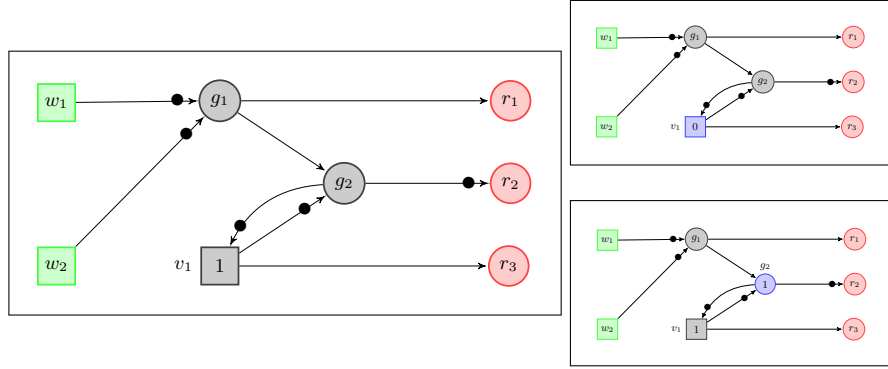


Fig. 1. Left: And-inverter graph describing a system with two inputs w_1 and w_2 (green boxes), one latch v_1 with initial value 1 (grey box), two gates g_1 and g_2 (gray circles), and three invariant requirements $r_1 = \Box g_1$, $r_2 = \Box \neg g_2$ and $r_3 = \Box v_1$ (red circles). Arrows represent logical dependencies, and bullets in the arrows imply negation.

Right above: an attacker that controls latch v_1 can set its initial value to 0 to break r_2 and r_3 in 0 steps.

Right below: an attacker that controls gate g_2 can set its value to 1 at time 0 to break r_2 in 0 steps and r_3 in 1 step, because the value of v_1 at time 1 is 0.

An *expression* e is described by the grammar $e ::= 0 \mid 1 \mid c \mid \neg c$, where $c \in C$. The set of all expressions is E . We use discrete time steps $t = 0, 1, \dots$ to describe the system of equations. To each latch $v \in V$ we associate a transition equation of the form $v(t+1) = e(t)$ and an initial equation of the form $v(0) = b$, where $e \in E$ and $b \in \mathbb{B}$. To each gate $g \in G$ we associate an equation of the form $g(t) = e_1(t) \wedge e_2(t)$, where $e_1, e_2 \in E$.

Example 1. Figure 1 shows an example AIG with $W = \{w_1, w_2\}$, $V = \{v_1\}$ and $G = \{g_1, g_2\}$. The corresponding system of equations is

$$\begin{aligned} v_1(0) &= 1, & v_1(t+1) &= \neg g_2(t), \\ g_1(t) &= \neg w_1(t) \wedge \neg w_2(t), & g_2(t) &= g_1(t) \wedge \neg v_1(t). \end{aligned}$$

The states of a system are all the different valuations of the variables in V ; formally, a state $\vec{v} : V \rightarrow \mathbb{B}$ is a map from V to bits. Similarly, the valuations of the variables in W are all the inputs to the systems; again, an input $\vec{w} : W \rightarrow \mathbb{B}$ is a map from W to bits. We refer to the set of all states by \vec{V} , and to the set of all inputs by \vec{W} . For $t = 0, 1, \dots$, we denote the state of the system at time t by $\vec{v}(t)$, with $\vec{v}(t) \triangleq \langle v_1(t), \dots, v_n(t) \rangle$. The initial state is $\vec{v}(0)$, defined by the initial equations for the latches. Similarly, we denote the input of the system at time t by $\vec{w}(t)$, with $\vec{w}(t) \triangleq \langle w_1(t), \dots, w_n(t) \rangle$. There are no restrictions or assumptions over $\vec{w}(t)$, so it can take any value in \vec{W} .

In Example 1, the states are $\vec{V} = \{ \langle (v_1, 0) \rangle, \langle (v_1, 1) \rangle \}$, the inputs are $\vec{W} = \{ \langle (w_1, 0), (w_2, 0) \rangle, \langle (w_1, 0), (w_2, 1) \rangle, \langle (w_1, 1), (w_2, 0) \rangle, \langle (w_1, 1), (w_2, 1) \rangle \}$, and the initial state is $\vec{v}(0) = \langle (v_1, 1) \rangle$.

Given an expression $e \in E$, the invariant $\Box e$ is the property that requires $e(t)$ to be true for all $t \geq 0$. The system S *fails* the invariant $\Box e$ iff there exists a finite sequence of inputs $\langle \vec{w}_0, \dots, \vec{w}_t \rangle$ such that, if we assume $\vec{w}(t) = \vec{w}_t$, then $e(t)$ is false. The system *satisfies* the invariant $\Box e$ if no such sequence of inputs exists. Every expression e represents a boolean predicate over the state of the latches of the system, and can be used to characterise states that are (un)safe. These expressions are particularly useful in safety-critical hardware, as they can signal the approach of a critical state.

In Example 1, we define three requirements: $r_1 \triangleq \Box g_1$, $r_2 \triangleq \Box \neg g_2$, and $r_3 \triangleq \Box v_1$. This system satisfies r_2 and r_3 , but it fails r_1 because $w_1 = 1$ and $w_2 = 0$ results in $g_1(0)$ being 0.

The *Cone-of-Influence* (COI) is a mapping from an expression to the components that can potentially influence its value. We obtain the COI of an expression $e \in E$, denoted $\nabla(e)$, by transitively tracing its dependencies to inputs, latches and gates. More precisely,

- $\nabla(0) = \emptyset$ and $\nabla(1) = \emptyset$;
- if $e = \neg c$ for $c \in C$, then $\nabla(e) = \nabla(c)$;
- if $e = w$ and w is an input, then $\nabla(e) = \{w\}$;
- if $e = v$ and v is a latch whose transition equation is $l(t+1) = e'(t)$, then $\nabla(e) = \{v\} \cup \nabla(e')$;
- if $e = g$ and g is a gate whose equation is $g(t) = e_1(t) \wedge e_2(t)$ then $\nabla(e) = \{g\} \cup \nabla(e_1) \cup \nabla(e_2)$.

The COI of a requirement $r = \Box e$ is $\nabla(r) \triangleq \nabla(e)$.

In Example 1, the COI for the requirements are $\nabla(\Box g_1) = \{g_1, w_1, w_2\}$, and $\nabla(\Box g_2) = \nabla(\Box g_3) = \{g_1, g_2, v_1, w_1, w_2\}$.

The set of *sources* of an expression $e \in E$, denoted $\text{src}(e)$ is the set of latches and inputs in the COI of e ; formally, $\text{src}(e) \triangleq \nabla(e) \cap (V \cup W)$. The *Jaccard index* of two expressions e_1 and e_2 is equal to $\frac{|\text{src}(e_1) \cap \text{src}(e_2)|}{|\text{src}(e_1) \cup \text{src}(e_2)|}$. This index provides a measure of how similar the sources of e_1 and e_2 are.

The *dual cone-of-influence* (IOC) of a component $c \in C$, denoted $\blacktriangle(c)$, is the set of components influenced by c ; more precisely $\blacktriangle(c) \triangleq \{c' \in C \mid c \in \nabla(c')\}$.

3 Motivational Example

In this section, we provide a motivational example of the problem of model checking compromised systems, and we illustrate how to classify attackers given a list of security requirements. Consider a scenario where an attacker A controls the gate g_2 of Example 1. By controlling g_2 , we mean that A can set the value of $g_2(t)$ at will for all $t \geq 0$. Since $r_2 = \Box \neg g_2$, it is possible for A to break r_2 by setting $g_2(0)$ to 1. We note that the same strategy works to break both




Fig. 2. Left: classification of attackers for requirements r_2 and r_3 . Right: classification of attackers for requirement r_1 . A green attacker cannot break the requirement, while a red attacker can.

requirements, but it need not be in the general case; i.e., an attacker may have one strategy to break one requirement, and a different strategy to break another. A can also break $r_3 = \Box v_1$, because, if A sets $g_2(0)$ to 1, then $v_1(1)$ is equal to 0. Since the original system fails to enforce r_1 , we say that A has the *power to break the requirements* r_1 , r_2 and r_3 . Now, consider a different attacker B which only controls the gate g_1 . No matter what value B chooses for $g_1(t)$ for all t , it is impossible for B to break r_2 or r_3 , so we say that B only has the power to break r_1 .

If we allow attackers to control any number of components, then there are 8 different attackers, described by the subsets of $\{v_1, g_1, g_2\}$. We do not consider attackers that control inputs, because the model checking of invariant properties requires the property to hold for all inputs, so giving control of inputs to an attacker does not make it more powerful (i.e. the attacker cannot break more requirements than it already could without the inputs). Figure 2 illustrates the classification of attackers depending on whether they can break a given requirement or not. Based on it, we can provide the following security guarantees: 1) the system cannot enforce r_1 , and 2) that the system can only enforce r_2 and r_3 in the presence of attackers that are as capable to interact with the system as $\{g_1\}$ (i.e. they only control g_1 or nothing).

According to the classification, attacker $\{g_2\}$ is as powerful as the attacker $\{v_1, g_1, g_2\}$, since both attackers can break the same requirements r_1 , r_2 and r_3 . This information may be useful to the designer of the system, because it may prioritise attackers that control less components but are as powerful as attackers that control more when deploying defensive mechanisms.

4 Bounded Model Checking of Compromised Systems

We recall the research questions that motivate this work: **RQ1**) given a system and a list of security requirements, how do we systematically generate attackers that can potentially break these requirements, and how do we verify if they are

successful? and **RQ2**) which techniques can help us *classify* attackers, i.e., to map each attacker to the set of requirements that it breaks? In this section, we aim to answer these research questions on a theoretical level by formalising the problem of attacker classification via bounded model checking AIGs in the presence of attackers. More precisely, to answer **RQ1**, we formalise attackers and their interactions with systems, and we show how to systematically generate bounded model checking problems that solve whether some given attackers can break some given requirements. We then propose two methods for the classification of attackers: 1) a brute-force method that creates a model checking problem for each attacker-requirement pair, and 2) a method that incrementally empowers attackers to find “minimal attackers,” since minimal attackers represent large portions of the universe of attackers thanks to a monotonicity relation between the set of components controlled by the attacker and the set of requirements that the attacker can break. The latter method is a theoretical approach to answer **RQ2**, while its practical usefulness is evaluated in Section 5.

4.1 Attackers and Compromised Systems

Since an AIG describes a system of equations, to incorporate the actions of an attacker A into the system, we modify the equations that are associated to the components controlled by A . Let $S = (W, V, G)$ be a system described by an AIG, let $R = \{r_1, \dots, r_n\}$ be a set of invariant requirements for S , and let $C = W \cup V \cup G$ be the set of components of S . By definition, an *attacker* A is any subset of C . If a component c belongs to an attacker A , then A has the *capability to interact with S through c* . We modify the equations of every latch $v \in V$ to be parametrised by an attacker A as follows: the original transition equation $v(t+1) = e(t)$ and the initial equation $v(0) = b$ changes to

$$v(t+1) = \begin{cases} e(t), & \text{if } v \notin A; \\ A_v(t+1), & \text{otherwise,} \end{cases} \quad v(0) = \begin{cases} b, & \text{if } v \notin A; \\ A_v(0), & \text{otherwise.} \end{cases} \quad (1)$$

where $A_v(t)$ is a value chosen by the attacker A at time t . Similarly, we modify the equation of gate $g \in G$ as follows: the original equation $g(t) = e_1(t) \wedge e_2(t)$ changes to

$$g(t) = \begin{cases} e_1(t) \wedge e_2(t), & \text{if } g \notin A; \\ A_g(t), & \text{otherwise,} \end{cases} \quad (2)$$

where $A_g(t)$ is, again, a value chosen by the attacker A at time t . We use $A[S]$ to denote the system S under the influence of attacker A ; i.e., $A[S]$ is the modified system of equations.

An *attack* $\vec{a} : A \rightarrow \mathbb{B}$ is a map of components in A to booleans. An *attack strategy* is a finite sequence of attacks $(\vec{a}_0, \vec{a}_1, \dots, \vec{a}_t)$ that fixes the values of all $A_c(k)$ (used in the equations above) by $A_c(k) = \vec{a}_k(c)$, with $c \in A$ and $0 \leq k \leq t$.

Definition 1 (Broken Requirement). *Given a requirement $r \in R$ with $r = \Box e$, we say that A breaks the requirement r if and only if there exists a sequence of inputs of length k and an attack strategy of length k such that $e(k)$ is false. We denote the set of requirements that A breaks by $A[R]$.*

Finally, we define two partial orders for attackers: **1)** an attacker A_i is strictly less *capable* (to interact with the system) than an attacker A_j in the context of S iff $A_i \subseteq A_j$ and $A_i \neq A_j$. The attacker A_i is equally capable to attacker A_j iff $A_i = A_j$; and **2)** an attacker A_i is strictly less *powerful* than an attacker A_j in the context of S and R iff $A_i[R] \subseteq A_j[R]$ and $A_i[R] \neq A_j[R]$. Similarly, attacker A_i is equally powerful to attacker A_j iff $A_i[R] = A_j[R]$. We simply state that A_i is less capable than A_j if S is clear from the context. Similarly, we simply say that A_i is less powerful than A_j if S and R are clear from the context.

We can now properly present the problem of *attacker classification*.

Definition 2 (Attacker Classification via Model Checking). *Given a system S , a set of requirements R , and a set of h attackers $\{A_1, \dots, A_h\}$, for every attacker A , we compute the set $A[R]$ of requirements that A can break by performing model checking of each requirement in R on the compromised system $A[S]$.*

Definition 2 assumes that exhaustive model checking is possible for S and the compromised versions $A[S]$ for all attackers A . However, if exhaustive model checking is not possible (e.g., due to time limitations or memory restrictions), we consider an alternative formulation for *Bounded Model Checking* (BMC):

Definition 3 (Attacker Classification via Bounded Model Checking). *Let S be a system, R be a set of requirements, and t be a natural number. Given a set of attackers $\{A_1, \dots, A_h\}$, for each attacker A , we compute the set $A[R]$ of requirements that A can break using a strategy of length up to t on the compromised system $A[S]$.*

In the following, we show how to construct a SAT formula that describes the attacker classification problem via bounded model checking.

4.2 A SAT Formula for BMC up to t Steps

For a requirement $r = \Box e$ and a time step $t \geq 0$, we are interested in finding an assignment of sources and attacker actions (i.e., an attack strategy) such that, for $0 \leq k \leq t$, the value of $e(k)$ is false. We define the proposition $\text{goal}(r, t)$ by

$$\text{goal}(\Box e, t) \triangleq \bigvee_{k=0}^t \neg e(k), \quad (3)$$

We must inform the SAT solver of the equalities and dependencies between expressions given by the definition of the AIG (e.g., that $e(k) \Leftrightarrow \neg v_1(k)$). Inspired by the work of Biere *et al.* [8], we transform the equations into a Conjunctive Normal Form formula (CNF) that the SAT solver can work with using

Tseitin encoding [31]. Each equation of the form

$$v(0) = \begin{cases} b, & \text{if } v \notin A; \\ A_v(0), & \text{otherwise,} \end{cases}$$

$$\text{becomes } (v^\downarrow \vee (v(0) \Leftrightarrow b)) \wedge (\neg v^\downarrow \vee (v(0) \Leftrightarrow A_v(0)))$$

where v^\downarrow is a literal that marks whether the latch v is an element of the attacker A currently being checked; i.e., we assume that v^\downarrow is true if $v \in A$, and we assume that v^\downarrow is false if $v \notin A$. Consequently, if $v \notin A$, then $v(0) \Leftrightarrow b$ must be true, and if $v \in A$, then $v(0) \Leftrightarrow A_v(0)$ must be true. We denote this new proposition by $\mathbf{encode}(v, 0)$, and it characterises the initial state of v .

Similarly, for $0 \leq k < t$, each equation of the form

$$v(k+1) = \begin{cases} e(k), & \text{if } v \notin A; \\ A_v(k+1), & \text{otherwise,} \end{cases}$$

$$\text{becomes } (v^\downarrow \vee (v(k+1) \Leftrightarrow e(k))) \wedge (\neg v^\downarrow \vee (v(k+1) \Leftrightarrow A_v(k+1))).$$

We denote this new proposition by $\mathbf{encode}(v, k)$. Finally, for $0 \leq k \leq t$, each equation of the form

$$g(k) = \begin{cases} e_1(k) \wedge e_2(k), & \text{if } g \notin A; \\ A_g(k), & \text{otherwise,} \end{cases}$$

$$\text{becomes } (g^\downarrow \vee (g(k) \Leftrightarrow e_1(k) \wedge e_2(k))) \wedge (\neg g^\downarrow \vee (g(k) \Leftrightarrow A_g(k))),$$

where g^\downarrow is a literal that marks whether the gate g is an element of the attacker A currently being checked in a similar way that the literal v^\downarrow works for the latch v . We denote this new proposition by $\mathbf{encode}(g, k)$.

To perform SAT solving, we need to find an assignment of inputs in W and attacker actions for each component c in A over t steps; thus, we need to assign at least $|W \times A| \times t$ literals. The SAT problem for checking whether requirement r is safe up to t steps, denoted $\mathbf{check}(r, t)$, is defined by

$$\mathbf{check}(r, t) \triangleq \mathbf{goal}(r, t) \wedge \bigwedge_{c \in (V \cup G)} \left(\bigwedge_{k=0}^t \mathbf{encode}(c, k) \right). \quad (4)$$

Proposition 1. *For a given attacker A and a requirement $r = \square e$, if we assume the literal c^\downarrow for all $c \in A$ and we assume $\neg x^\downarrow$ for all $x \notin A$ (i.e., $x \in (V \cup G) - A$), then A can break the requirement r in t steps (or less) if and only if $\mathbf{check}(r, t)$ is satisfiable.*

Proof. We first show that if A can break the requirement in t steps or less, then $\mathbf{check}(r, t)$ is satisfiable. Since A breaks $r = \square e$ in t steps or less, then, by Definition 1, there exists an assignment of inputs $(\vec{w}_0, \dots, \vec{w}_k)$ and an attacker strategy $(\vec{a}_0, \dots, \vec{a}_k)$ which causes $e(k)$ to be false for some $k \leq t$; this means that $\mathbf{goal}(r, t)$ is satisfiable, which, in turn, makes $\mathbf{check}(r, t)$ satisfiable.

Data: system $S = (W, V, G)$, a time step $t \geq 0$, a set of requirements R .

Result: A map that maps the attacker A to $A[R]$.

```

1 Map  $\mathcal{H}$ ;
2 foreach  $r \in R$  do
3   | foreach  $A$  such that  $A \subseteq (V \cup G)$  do
4   |   | if  $\text{check}(r, t)$  is satisfiable while assuming  $c^\downarrow$  for all  $c \in A$  then
5   |   |   | insert  $r$  in  $\mathcal{H}(A)$ ;
6   |   |   end
7   |   end
8 end
9 return  $\mathcal{H}$ ;

```

Algorithm 1: Naive attacker classification algorithm.

We now show that if $\text{check}(r, t)$ is satisfiable, then A can break the requirement. If $\text{check}(r, t)$ is satisfiable then $\text{goal}(r, t)$ is satisfiable, and $e(k)$ is false for some $k \leq t$. Consequently, there is an assignment of inputs $\vec{w}(k)$ and attacker actions $A_c(k)$, such that the $\text{encode}(c, k)$ propositions are satisfied for all $c \in A$. With $\vec{a}_k(c) = A_c(k)$ and $\vec{w}_k = \vec{w}(k)$, we obtain a witness input sequence and a witness attack strategy which proves that A can break r in k steps (i.e., in t steps or less since $k \leq t$). \square

Algorithm 1 describes a naive strategy to compute the sets $A[R]$ for each attacker A ; i.e. the set of requirements that A breaks in t steps (or less). Algorithm 1 works by solving, for each of the $2^{|V \cup G|}$ different attackers, a set of $|R|$ SAT problems, each of which has a size of at least $\mathcal{O}(|C| \times t)$ on the worst case.

In the rest of the section, we propose two sound heuristics in an attempt to improve Algorithm 1: the first technique aims to reduce the size of the SAT formula, while the other aims to record and propagate the results of verifications among the set of attackers so that some calls to the SAT solver can be avoided.

4.3 Isolation and Monotonicity

The first strategy involves relying on *isolation* to prove that it is impossible for a given attacker to break some requirements. To formally capture this notion, we first extend the notion of IOC to attackers. The IOC of an attacker A , denoted $\blacktriangle(A)$, is defined by the union of IOCs of the components in A ; more precisely, $\blacktriangle(A) \triangleq \bigcup \{ \blacktriangle(c) \mid c \in A \}$.

Informally, isolation happens whenever the IOC of A is disjoint from the COI of r , implying that A cannot interact with r .

Proposition 2 (Isolation). *Let A be an attacker and r be a requirement that is satisfied in the absence of A . If $\blacktriangle(A) \cap \blacktriangledown(r) = \emptyset$, then A cannot break r .*

Proof. For the attacker A to break the requirement r , there must be a component $c \in \blacktriangledown(r)$ whose behaviour was affected by the presence of A , and whose change

of behaviour caused r to fail. However, for A to affect the behaviour of c , there must be a dependency between the variables directly controlled by A and c , since A only chooses actions over the components it controls; implying that $c \in \blacktriangle(A)$. This contradicts the premise that the IOC of A and the COI of r are disjoint, so the component c cannot exist. \square

Isolation reduces the SAT formula by dismissing attackers that are outside the COI of the requirement to be verified. Isolation works similarly to *COI reduction* (see [13, 15, 7, 18, 14]), and it transforms Equation 4 into

$$\text{check}(r, t) \triangleq \text{goal}(r, t) \wedge \bigwedge_{c \in (\blacktriangledown(r) - W)} \left(\bigwedge_{k=0}^t \text{encode}(c, k) \right) \quad (5)$$

The second strategy uses *monotonicity* relation between capabilities and power of attackers.

Proposition 3 (Monotonicity). *For attackers A and B and a set of requirements R , if $A \subseteq B$, then $A[R] \subseteq B[R]$.*

Proof. If A is a subset of B , then attacker B can always choose the same attack strategies that A used to break the requirements in $A[R]$; thus, $A[R]$ must be a subset of $B[R]$. \square

Monotonicity allows us to define the notion of minimal (successful) attackers for a requirement r : attacker A is a *minimal attacker* for requirement r if and only if A breaks r , and there is no attacker $B \subset A$ such that B also breaks r . In the remainder of this section, we expand on this notion, and we describe a methodology for attacker classification that focuses on the identification of these minimal attackers.

4.4 Minimal (Successful) Attackers

The set of minimal attackers for a requirement r partitions the set of attackers into those that break r and those who do not. Any attacker that is more capable than a minimal attacker is guaranteed to break r by monotonicity (cf. Proposition 3), and any attacker that is less capable than a minimal attacker cannot break r ; otherwise, this less capable attacker would be a minimal attacker. Consequently, we can reduce the problem of attacker classification to the problem of finding the minimal attackers for all requirements.

Existence of a Minimal Attacker. Thanks to isolation (cf. Proposition 2) we can guarantee that a requirement r that is safe in the absence of an attacker A remains safe in the presence of A if $\blacktriangledown(r) \cap \blacktriangle(A)$ is empty. Thus, for each requirement $r \in R$, the set of attackers that could break r is $\mathcal{P}(\blacktriangledown(r) - W)$. Out of all the attackers of r , the most capable attacker is $\blacktriangledown(r) - W$, so we can test whether there *exists* any attacker that can break r in t steps by solving $\text{check}(r, t)$ against attacker $\blacktriangledown(r) - W$. For succinctness, we henceforth denote the attacker $\blacktriangledown(r) - W$ by r^{max} .

Corollary 1. *From monotonicity and isolation (cf. Propositions 3 and 2), if attacker r^{max} cannot break the requirement r , then there are no minimal attackers for r . Equivalently, if r^{max} cannot break r , then r does not belong to any set of broken requirements $A[R]$.*

```

Data: system  $S$ , a requirement  $r$ , and a time step  $t \geq 0$ .
Result: set  $M$  of minimal attackers for  $r$ , bounded by  $t$ .
1 if  $check(r, t)$  is not satisfiable while assuming  $c^\downarrow$  for all  $c \in r^{max}$  then
2   | return  $\emptyset$ ;
3 end
4 Set:  $P = \{\emptyset\}$ ,  $M = \emptyset$ ; //( $P$  contains the empty attacker  $\emptyset$ )
5 while  $P$  is not empty do
6   | extract  $A$  from  $P$  such that the size of  $A$  is minimal;
7   | if not (exists  $B \in M$  such that  $B \subseteq A$ ) then
8     |   if  $check(r, t)$  is satisfiable when assuming  $c^\downarrow$  for all  $c \in A$  then
9       |   | insert  $A$  in  $M$ ;
10      |   else
11        |   | foreach  $c \in (r^{max} - A)$  do
12          |   | | insert  $A \cup \{c\}$  in  $P$ ;
13          |   | end
14        |   end
15      | end
16 end
17 return  $M$ ;

```

Algorithm 2: The MinimalAttackers algorithm.

Finding Minimal Attackers. After having confirmed that at least one minimal attacker for r exists, we can focus on finding them. Our strategy consist of systematically increasing the capabilities of attackers that fail to break the requirement r until they do. Algorithm 2 describes this empowering procedure to computes the set of minimal attackers for a requirement r , which we call **MinimalAttackers**. As mentioned, we first check to see if a minimal attacker exists (Lines 1-3); then we start evaluating attackers in an orderly fashion by always choosing the smallest attackers in the set of pending attackers P (Lines 5-16). Line 7 uses monotonicity to discard the attacker A if there is a successful attacker B with $B \subseteq A$. Line 8 checks if the attacker A can break r in t steps (or less); if so, then A is a minimal attacker for r and is included in M (Line 9); otherwise, we empower A with a new component c , and we add these new attackers to P (Lines 11-13). We note that Line 11 relies on isolation, since we only add components that belong to the COI of r .

We recall the motivational example from Section 3. Consider the computation of **MinimalAttackers** for requirement r_2 . In this case, r_2^{max} is

Data: system S , a time step $t \geq 0$, and a set of requirements R .

Result: Set of all minimal attackers \mathcal{M} and an initial classification map \mathcal{H} .

```

1 Set:  $\mathcal{M} = \emptyset$ ;
2 Map:  $\mathcal{H}$ ;
3 foreach  $r \in R$  do
4   | foreach  $A \in \text{MinimalAttackers}(S, t, r)$  do
5   |   | insert  $r$  in  $\mathcal{H}(A)$ ;
6   |   | insert  $A$  in  $\mathcal{M}$ ;
7   |   end
8 end
9 return  $(\mathcal{M}, \mathcal{H})$ ;

```

Algorithm 3: The AllMinimalAttackers algorithm.

$\{g_1, g_2, v_1\}$, which is able to break r_2 , confirming the existence of (at least) a minimal attacker (Lines 1-3). We start to look for minimal attackers by checking the attacker \emptyset (Lines 5-8); after we see that it fails to break r_2 , we conclude that \emptyset is not a minimal attacker and that we need to increase its capabilities. We then derive the attackers $\{g_1\}$, $\{g_2\}$ and $\{v_1\}$ by adding one non-isolated component to \emptyset , and we put them into the set of pending attackers (Lines 11-13). For attackers $\{v_1\}$ and $\{g_2\}$, we know that they can break the requirement r_2 , so they get added to the set of minimal attackers, and are not empowered (Line 9); however, for attacker $\{g_1\}$, since it fails to break r_2 , we increase its capabilities and we generate attackers $\{v_1, g_1\}$ and $\{g_1, g_2\}$. Finally, for these two latter attackers, since the minimal attackers $\{v_1\}$ and $\{g_2\}$ have already been identified, the check in Line 7 fails, and they are dismissed from the set of pending attackers, since they cannot be minimal. The algorithm finishes with $M = \{\{v_1\}, \{g_2\}\}$.

Algorithm 3 applies Algorithm 2 to each requirement; it collects all minimal attackers in the set \mathcal{M} and initialises the attacker classification map \mathcal{H} . Finally, Algorithm 4 exploits monotonicity to compute the classification of each attacker A by aggregating the requirements broken by the minimal attackers that are subsets of A .

For the motivational example in Section 3, Algorithm 3 returns $\mathcal{M} = \{\emptyset, \{v_1\}, \{g_2\}\}$ and $\mathcal{H} = \{(\emptyset, \{r_1\}), (\{v_1\}, \{r_2, r_3\}), (\{g_2\}, \{r_2, r_3\})\}$. From there, Algorithm 4 completes the map \mathcal{H} , and returns

$$\begin{aligned} \mathcal{H} = & \{(\emptyset, \{r_1\}), (\{v_1\}, \{r_1, r_2, r_3\}), (\{g_1\}, \{r_1\}), (\{g_2\}, \{r_1, r_2, r_3\}), \\ & (\{v_1, g_2\}, \{r_1, r_2, r_3\}), (\{v_1, g_2\}, \{r_1, r_2, r_3\}), \\ & (\{g_1, g_2\}, \{r_1, r_2, r_3\}), (\{v_1, g_1, g_2\}, \{r_1, r_2, r_3\})\} \end{aligned}$$

4.5 On Soundness and Completeness

Just like any bounded model checking problem, if the time parameter t is below the *completeness threshold* (see [24]), the resulting attacker classification up to t

Data: system $S = (W, V, G)$, a time step $t \geq 0$, a set of requirements R .

Result: A map \mathcal{H} that maps the attacker A to $A[R]$.

```
1  $(\mathcal{M}, \mathcal{H}) = \text{AllMinimalAttackers}(S, t, R)$ ;  
2 foreach  $A \subseteq (V \cup G)$  do  
3   | foreach  $A' \in \mathcal{M}$  do  
4   |   | if  $A' \subseteq A$  then  
5   |   |   | insert all elements of  $\mathcal{H}(A')$  in  $\mathcal{H}(A)$ ;  
6   |   | end  
7   | end  
8 end  
9 return  $\mathcal{H}$ ;
```

Algorithm 4: Improved classification algorithm. We assume that \mathcal{H} initially maps every A to the empty set.

steps could be *incomplete*. More precisely, an attacker classification up to t steps may prove that an attacker A cannot break some requirement r with a strategy up to t steps, while in reality A can break r by using a strategy whose length is strictly greater than t . There are practical reasons that justify the use of a time parameter that is lower than the completeness threshold: 1) computing the exact completeness threshold is often as hard as solving the model-checking problem [18], so an approximation is taken instead; and 2), the complexity of the classification problem grows exponentially with t in the worst case, since the size of the SAT formulae grow with t , and there is an exponential number of attackers that need to be classified by making calls to the SAT solver. A classification that uses a t below the completeness threshold, while possibly incomplete, is *sound*, i.e., it does not falsely report that an attacker can break a requirement when in reality it cannot. In Section 6 we discuss possible alternatives to overcome incompleteness, but we leave a definite solution as future work.

We also consider the possibility of limiting the maximum size of minimal attackers to approximate the problem of attacker classification. The result of a classification whose minimal sets are limited to have up to z elements is also sound but incomplete, since does not identify minimal attackers that have more than z elements. We show in Section 5 that, even with restricted minimal attackers, it is possible to obtain a high coverage of the universe of attackers.

4.6 Requirement Clustering

Property clustering [13, 7, 14] is a state-of-the-art technique for the model checking of multiple properties. Clustering allows the SAT solver to reuse information when solving a similar instance of the same problem, but under different assumptions. To create clusters for attacker classification, we combine the SAT problems whose COI is similar (i.e., requirements that have a Jaccard index close to 1), and incrementally enable and disable properties during verification. More precisely, to use clustering, instead of computing $\text{goal}(r, t)$ for a single

requirement, we compute $\text{goal}(Y, t)$ for a cluster Y of requirements, defined by

$$\text{goal}(Y, t) \triangleq \bigwedge_{r \in Y} (\neg r^\downarrow \vee \text{goal}(r, t)). \quad (6)$$

where r^\downarrow is a new literal that plays a similar role to the ones used for gates and latches; i.e., we assume r^\downarrow when we want to find the minimal attackers for r , and we assume $\neg y^\downarrow$ for all other requirements $y \in Y$.

The SAT problem for checking whether the cluster of requirements Y is safe up to t steps is

$$\text{check}(Y, t) \triangleq \text{goal}(Y, t) \wedge \bigwedge_{c \in (\nabla(Y) - W)} \left(\bigwedge_{k=0}^t \text{encode}(c, k) \right), \quad (7)$$

where $\nabla(Y) = \bigcup \{ \nabla(r) \mid r \in Y \}$.

5 Evaluation

In this section, we perform experiments to evaluate how effective is the use of isolation and monotonicity for the classification of attackers, and we evaluate the completeness of partial classifications for different time steps.

For evaluation, we use a sample of AIG benchmarks from past Hardware Model-Checking Competitions (see [1, 2]), from their multiple-property verification track. Each benchmark has an associated list of invariants to be verified which, for the purposes of this evaluation, we interpret as the set of security requirements. As of 2014, the benchmark set was composed of 230 different instances, coming from both academia and industrial settings [16]. We quote from [16]:

“Among industrial entries, 145 instances belong to the SixthSense family (6s*, provided by IBM), 24 are Intel benchmarks (intel*), and 24 are Oski benchmarks. Among the academic related benchmarks, the set includes 13 instances provided by Robert (Bob) Brayton (bob*), 4 benchmarks coming from Politecnico di Torino (pdt*) and 15 Beem (beem*). Additionally, 5 more circuits, already present in previous competitions, complete the set.”

All experiments are performed on a quad core MacBook with 2.9 GHz Intel Core i7 and 16GB RAM, and we use the SAT solver CaDiCaL version 1.0.3 [3]. The source code of the artefact is available at [4].

We separate our evaluation in two parts: 1) a comparative study where we evaluate the effectiveness of using of monotonicity and isolation for attacker classification in several benchmarks, and 2) a case study, where we apply our classification methodology to a single benchmark `-pdtvsarmultip-` and we study the results of varying the time parameters for partial classification.

5.1 Evaluating Methodologies

Given a set of competing classification methodologies $\mathcal{M}_1, \dots, \mathcal{M}_n$ (e.g., Algorithm 1 and Algorithm 4), each methodology is given the same set of benchmarks S_1, \dots, S_m , each with its respective set of requirements R_1, \dots, R_m . To evaluate a methodology \mathcal{M} on a benchmark $S = (W, V, G)$ with a set of requirements R , we allow \mathcal{M} to “learn” for about 10 minutes per requirement by making calls to the SAT solver, and produce a (partial) attacker classification \mathcal{H} . Afterwards, we compute the coverage metric obtained by \mathcal{M} , defined as follows.

Definition 4 (Coverage). *Let $\mathcal{P}(V)$ be the set of all attackers, and let \mathcal{H} be the attacker classification produced by the methodology \mathcal{M} . We recall that \mathcal{H} is a map that maps each attacker A to a set of requirements, and in the ideal case, $\mathcal{H}(A) = A[R]$, for each attacker A . The attacker coverage obtained by methodology \mathcal{M} for a requirement r is the percentage of attackers $A \in \mathcal{P}(V)$ for which we can correctly determine whether A breaks r by computing $r \in \mathcal{H}(A)$ (i.e., we do not allow guessing and we do not allow making new calls to the SAT solver).*

We also measure the execution time of the classification per requirement. More precisely, the time it takes for the methodology to find minimal attackers, capped at about 10 minutes per requirement. We force stop the classification for each requirement if a timeout occurs, but not while the SAT solver is running (i.e., we do not interrupt the SAT solver), which is why sometimes the reported time exceeds 10 minutes.

5.2 Effectiveness of Isolation and Monotonicity

To test the effectiveness of isolation and monotonicity, we selected a small sample of seven benchmarks. For each benchmark, we test four variations of our methodology:

1. (+IS, +MO): Algorithm 4, which uses both isolation and monotonicity
2. (+IS, -MO): Algorithm 4 but removing the check for monotonicity on Algorithm 2, Line 7;
3. (-IS, +MO): Algorithm 4 but using Equation 5 instead of Equation 7 to remove isolation while preserving monotonicity; and
4. (+IS, +MO) Algorithm 1, which does not use isolation nor monotonicity.

The benchmarks we selected have an average of 173 inputs, 8306 gates, 517 latches, and 80 requirements. Under our formulation of attackers, these benchmarks have on average 2^{8823} attackers. However, since an attacker that controls a gate g can be emulated by an attacker that controls all latches in the sources of g , we restrict attackers to be comprised of only latches; reducing the size of the set of attackers from 2^{8823} to 2^{517} on average per benchmark. Furthermore, we arbitrarily restrict the number of components that minimal attackers may control to a maximum of 3, which implies that, on a worst case scenario, we need

to make a maximum of $80 \times \sum_{k=0}^3 \binom{517}{k}$ calls to the SAT solver per benchmark. We also arbitrarily define the time step parameter t to be 10.

Figure 3 illustrates the average coverage for the four different methodologies, for each of the seven benchmarks. The exact coverage values are reported in the Appendix of [29] (an extended version of this article). We see that our methodology consistently obtains the best coverage of all the other methodologies, with the exception of benchmark **6s155**, where the methodology that removes isolation triumphs over ours. We attribute this exception to the way the SAT solver reuses knowledge when working incrementally; it seems that, for $(-IS, +MO)$, the SAT solver can reuse more knowledge than for $(+IS, +MO)$, which is why $(-IS, +MO)$ can discover more minimal attackers in average than $(+IS, +MO)$.

We observe that the most significant element in play to obtain a high coverage is the use of monotonicity. Methodologies that use monotonicity always obtain better results than their counterparts without monotonicity. Isolation does not show a trend for increasing coverage, but has an impact in terms of classification time. Figure 4 presents the average classification time per requirement for the benchmarks under the different methodologies. We note that removing isolation often increases the average classification time of classification methodologies; the only exception –benchmark **6s325**– reports a smaller time because the SAT solver ran out of memory during SAT solving about 50% of the time, which caused an early termination of the classification procedure. This early termination also reflects on the comparatively low coverage for the method $(-IS, +MO)$ in this benchmark, reported on Figure 3.

5.3 Partial Classification of the `pdtsarmultip` Benchmark

The benchmark `pdtsarmultip` has 17 inputs, 130 latches, 2743 gates, and has an associated list of 33 invariant properties, out of which 31 are unique and we interpret as the list of security requirements. Since we are only considering attackers that control latches, there are a total of 2^{130} attackers that need to be classified for the 31 security requirements.

We consider 6 scenarios for partial classification up to t , with t taking values in $\{0, 1, 5, 10, 20, 30\}$. For each requirement, we obtain the execution time of classification (ms), the size of the set of source latches for the requirement ($\#C$), the number of minimal attackers found ($\#Min.$), the total number of calls to the SAT solver ($\#SAT$), the average number of components per minimal attacker ($\#C./Min$) and the coverage for the requirement (Cov.). We present the average of these measures in Table 1.

Normally, the attacker classification behaves in a similar way to what is reported for requirement $\square\neg g_{2177}$, shown in Table 2. More precisely, coverage steadily increases and stabilises as we increase t . However, we like to highlight two interesting phenomena that may occur: 1) coverage may *decrease* as we increase the time step (e.g., as shown in Table 3), and 2) the number of minimal attackers decreases while the coverage increases, as shown in Tables 2 and 3.

Case 1) occurs because the set of attackers that can effectively interact with the system at time 0 is rather small, i.e., 2^6 , while the set of attackers that can




Fig. 3. Average requirement coverage per benchmark. A missing bar indicates a value that is approximately 0.




Fig. 4. Average classification time per requirement per benchmark. A missing bar indicates a value that is approximately 0.

affect the system at times 0 and 1 has size 2^{26} . The size of this set increases with time until it stabilises at 2^{66} , which is the size of the set of attackers that cannot be dismissed by isolation.

Case 2) occurs because the minimal attackers that are found for smaller time steps represent a small percentage of the set of attackers that can affect the system, so there is very little we can learn by using monotonicity. More precisely, those minimal attackers control a relatively large set of components, which they need to be successful in breaking requirements, as shown in Step 5, column #C./Min in Tables 2 and 3. By considering more time steps, we are allowing attackers that control less components to further propagate their actions through the system, which enables attack strategies that were unsuccessful for smaller choices of time steps.

By taking an average over all requirements, we observe that coverage seems to steadily increase as we increase the number of steps for the classification, as reported in Table 1, column Cov. The low coverage for small t is due to the restriction on the size of minimal attackers. More precisely, for small t , attackers can only use short strategies, which limits their interaction with the system; we expect attackers to control a large number of components if they want to successfully influence a requirement in this single time step, and since we restricted our search to attackers of size 3 maximum, these larger minimal attackers are not found (e.g., as reported in Table 2 for Step 0).

Table 1. Average measures for all requirements per time steps.

Steps	ms	#C.	#Min.	#SAT	#C./Min.	Cov.
0	683.1290323	34.96774194	2.451612903	16328.77419	1.4	0.527277594
1	2387.548387	46.22580645	6.387096774	24420	1.650232484	0.572533254
5	5229.935484	58.93548387	44.90322581	28355.29032	1.639645689	0.84956949
10	24967.12903	58.93548387	151.1935484	25566.54839	1.460869285	0.918973269
20	13632.51613	58.93548387	17.67741935	20849.70968	1.176272506	0.979354259
30	12208.25806	58.93548387	15.93548387	20798.16129	1.104563895	0.979354274

Table 2. Coverage for $\Box\neg g_{2177}$.

$\Box\neg g_{2177}$						
Steps	ms	#C	#Min.	#SAT	#C./Min.	Cov.
0	895	59	0	34281	-	5.94E-14
1	2187	66	10	47378	2	0.499511
5	1735	66	205	12476	1.912195	0.999997
10	968	66	27	9948	1	0.999999
20	1275	66	27	9948	1	0.999999
30	1819	66	27	9948	1	0.999999

Table 3. Coverage for $\Box\neg g_{2220}$.

$\Box\neg g_{2220}$						
Steps	ms	#C.	#Min.	#SAT	#C./Min.	Cov.
0	1	6	1	28	1	0.90625
1	86	26	1	2628	1	0.500039
5	4511	67	17	47664	2.588235	0.852539
10	3226	67	6	37889	1	0.984375
20	3355	67	6	37889	1	0.984375
30	3562	67	6	37889	1	0.984375

We conclude that experimental evidence favours the use of both monotonicity and isolation for the classification of attackers, although some exceptions may occur for the use of isolation. Nevertheless, these two techniques help our classification methodology (+*IS*, +*MO*) consistently obtain significantly better coverage when compared to the naive methodology (-*IS*, -*MO*).

6 Related Work

On Defining Attackers. Describing an adequate attacker model to contextualise the security guarantees of a system is not a trivial task. Some attacker models may be adequate to provide guarantees over one property (e.g. confidentiality), but not for a different one (e.g., integrity). Additionally, depending on the nature of the system and the security properties being studied, it is sensible to describe attackers at different levels of abstraction. For instance, in the case of security protocols, Basin and Cremers define attackers in [6] as combinations of compromise rules that span over three dimensions: *whose* data is compromised, *which* kind of data it is, and *when* the compromise occurs. In the case of Cyber-physical Systems (CPS), works like [19, 30] model attackers as sets of components (e.g., some sensors or actuators), while other works like [33, 17, 32] model attackers that can arbitrarily manipulate any control inputs and any sensor measurements at will, as long as they avoid detection. In the same context of CPS, Rocchetto and Tippenhauer [28] model attackers more abstractly as combinations of quantifiable traits (e.g., insider knowledge, access to tools, and financial support), which, when provided a compatible system model, ideally fully define how the attacker can interact with the system.

Our methodology for the definition of attackers combines aspects from [6, 19] and [30]. The authors of [6] define symbolic attackers and a set of rules

that describe how the attackers affect the system, which is sensible since many cryptographic protocols are described symbolically. Our methodology describes attackers as sets of components (staying closer to the definitions of attackers in [19] and [30]), and has a lower level of abstraction since we describe the semantics of attacker actions in terms of how they change the functional behaviour of the AIG, and not in terms of what they ultimately represent. This lower level of abstraction lets us systematically and exhaustively generate attackers by simply having a benchmark description, but it limits the results of the analysis to the benchmark; Basin and Cremers can compare among different protocol implementations, because attackers have the same semantics even amongst different protocols. If we had an abstraction function from sets of gates and latches to symbolic notions (e.g., “components in charge of encryption”, or “components in charge of redundancy”), then it could be possible to compare results amongst different AIGs.

On Efficient Classification. The works by Cabodi, Camurati and Quer [15], Cabodi et. al [13], and Cabodi and Nocco [14] present several useful techniques that can be used to improve the performance of model checking when verifying multiple properties, including COI reduction and property clustering. We also mention the work by Goldberg et al. [20] where they consider the problem of efficiently checking a set of safety properties P_1 to P_k by individually checking each property while assuming that all other properties are valid. Ultimately, all these works inspired us to incrementally check requirements in the same cluster, helping us transform Equation 4 into Equation 7. Nevertheless, we note that all these techniques are described for model checking systems in the absence of attackers, which is why we needed to introduce the notions of isolation and monotonicity to account for them. Additionally, it may be possible to use or incorporate other techniques that improve the efficiency of BMC in general (e.g., interpolation [26]).

On Completeness. As mentioned in Section 4.5, if the time parameter for the classification is below the completeness threshold, the resulting attacker classification is most likely incomplete. To guarantee completeness, it may be possible to adapt existing termination methods (see [5]) to consider attackers. Alternatively, methods that compute a good approximation of the completeness threshold (see [24]) which guarantee the precision of resulting the coverage should help improve the completeness of attacker classifications. If we consider alternative verification techniques, then IC3 [11, 12] and PDR [11], which have seen some success in hardware model checking, may address the limitation of boundedness. Finally, interpolation [26] could also help finding a guarantee of completeness.

On Verifying Non-Safety Properties. In this work, we focused our analysis exclusively on safety properties of the form $\Box e$. However, we believe that it is possible to extend this methodology to other types of properties, it is possible to

efficiently encode Linear Temporal Logic formulae for bounded model checking [8, 9]. The formulations of the SAT problem change for the different nature of the formulae, but both isolation and monotonicity should remain valid heuristics, since they ultimately refer to how strategies of attackers can be inferred, not how they are constructed.

7 Conclusion and Future Work

In this work, we present a methodology to model check systems in the presence of attackers with the objective of mapping each attacker to the list of security requirements that it breaks. This mapping of attackers creates a classification for them, defining equivalence classes of attackers by the set of requirements that they can break. The system can then be considered safe in the presence of attackers that cannot break any requirement. While it is possible to perform a classification of attackers by exhaustively performing model checking, the exponential size of the set of attackers renders this naive approach impractical. Thus, we rely on ordering relations between attackers to efficiently classify a large percentage of them, and we demonstrate empirically by applying our methodology to a set of benchmarks that describe hardware systems at a bit level.

In our view, ensuring the completeness of the attacker classification is the most relevant direction for future work. Unlike complete classifications, incomplete classifications cannot provide guarantees that work in the general case if minimal attackers are not found. We also note that the effectiveness of monotonicity for classification is directly related to finding minimal attackers. Consequently, our methodology may benefit from any other method that helps in the identification of those minimal attackers. In particular, we are interested in checking the effectiveness of an approach where, instead of empowering attackers, we try to reduce successful attackers into minimal attackers by removing unnecessary capabilities. This, formally, is an *actual causality analysis* [22] of successful attackers.

References

1. (November 2011), <http://fmv.jku.at/hwmcc11/index.html>
2. (October 2013), <http://fmv.jku.at/hwmcc13/index.html>
3. (September 2019), <http://fmv.jku.at/cadical/>
4. (November 2020), <https://gitlab.com/asset-sutd/public/aig-ac>
5. Awedh, M., Somenzi, F.: Proving more properties with bounded model checking. In: Alur, R., Peled, D.A. (eds.) *Computer Aided Verification*. pp. 96–108. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
6. Basin, D., Cremers, C.: Know your enemy: Compromising adversaries in protocol analysis. *ACM Trans. Inf. Syst. Secur.* 17(2), 7:1–7:31 (Nov 2014), <http://doi.acm.org/10.1145/2658996>
7. Biere, A., Biere, A., Heule, M., van Maaren, H., Walsh, T.: *Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications*. IOS Press, Amsterdam, The Netherlands, The Netherlands (2009)

8. Biere, A., Cimatti, A., Clarke, E., Zhu, Y.: Symbolic model checking without bdds. In: Cleaveland, W.R. (ed.) *Tools and Algorithms for the Construction and Analysis of Systems*. pp. 193–207. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
9. Biere, A., Heljanko, K., Junttila, T., Latvala, T., Schuppan, V.: Linear Encodings of Bounded LTL Model Checking. *Logical Methods in Computer Science* Volume 2, Issue 5 (November 2006)
10. Biere, A., Heljanko, K., Wieringa, S.: AIGER 1.9 and beyond. Tech. rep., FMV Reports Series, Institute for Formal Models and Verification, Johannes Kepler University, Altenbergerstr. 69, 4040 Linz, Austria (2011)
11. Bradley, A.R.: Sat-based model checking without unrolling. In: Jhala, R., Schmidt, D. (eds.) *Verification, Model Checking, and Abstract Interpretation*. pp. 70–87. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
12. Bradley, A.R.: Understanding ic3. In: *Proceedings of the 15th International Conference on Theory and Applications of Satisfiability Testing*. pp. 1–14. SAT’12, Springer-Verlag, Berlin, Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-31612-8_1
13. Cabodi, G., Camurati, P.E., Loiacono, C., Palena, M., Pasini, P., Patti, D., Quer, S.: To split or to group: from divide-and-conquer to sub-task sharing for verifying multiple properties in model checking. *International Journal on Software Tools for Technology Transfer* 20(3), 313–325 (Jun 2018), <https://doi.org/10.1007/s10009-017-0451-8>
14. Cabodi, G., Nocco, S.: Optimized model checking of multiple properties. In: *2011 Design, Automation Test in Europe*. pp. 1–4 (March 2011)
15. Cabodi, G., Camurati, P., Quer, S.: A graph-labeling approach for efficient cone-of-influence computation in model-checking problems with multiple properties. *Software: Practice and Experience* 46(4), 493–511 (2016), <https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.2321>
16. Cabodi, G., Loiacono, C., Palena, M., Pasini, P., Patti, D., Quer, S., Vendraminetto, D., Biere, A., Heljanko, K.: Hardware model checking competition 2014: An analysis and comparison of model checkers and benchmarks. *Journal on Satisfiability, Boolean Modeling and Computation* 9, 135–172 (January 2015)
17. Cárdenas, A.A., Amin, S., Lin, Z.S., Huang, Y.L., Huang, C.Y., Sastry, S.: Attacks against process control systems: Risk assessment, detection, and response. In: *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*. pp. 355–366. ASIACCS ’11, ACM, New York, NY, USA (2011), <http://doi.acm.org/10.1145/1966913.1966959>
18. Clarke, E.M., Henzinger, T.A., Veith, H., Bloem, R. (eds.): *Handbook of Model Checking*. Springer (2018)
19. Giraldo, J., Urbina, D., Cardenas, A., Valente, J., Faisal, M., Ruths, J., Tippenhauer, N.O., Sandberg, H., Candell, R.: A survey of physics-based attack detection in cyber-physical systems. *ACM Comput. Surv.* 51(4), 76:1–76:36 (Jul 2018), <http://doi.acm.org/10.1145/3203245>
20. Goldberg, E., Gdemann, M., Kroening, D., Mukherjee, R.: Efficient verification of multi-property designs (the benefit of wrong assumptions). In: *2018 Design, Automation Test in Europe Conference Exhibition (DATE)*. pp. 43–48 (March 2018)
21. Günther, C.G.: An identity-based key-exchange protocol. In: Quisquater, J.J., Vandewalle, J. (eds.) *Advances in Cryptology — EUROCRYPT ’89*. pp. 29–37. Springer Berlin Heidelberg, Berlin, Heidelberg (1990)
22. Halpern, J.Y.: *Actual Causality*. MIT Press (2016), <http://www.jstor.org/stable/j.ctt1f5g5p9>

23. HELLERMAN, L.: A catalog of three-variable or-invert and and-invert logical circuits. *IEEE Transactions on Electronic Computers* EC-12(3), 198–223 (June 1963)
24. Kroening, D., Strichman, O.: Efficient computation of recurrence diameters. In: Zuck, L.D., Attie, P.C., Cortesi, A., Mukhopadhyay, S. (eds.) *Verification, Model Checking, and Abstract Interpretation*. pp. 298–309. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
25. Kuehlmann, A., Paruthi, V., Krohm, F., Ganai, M.K.: Robust boolean reasoning for equivalence checking and functional property verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 21(12), 1377–1394 (Dec 2002)
26. McMillan, K.L.: Interpolation and sat-based model checking. In: Hunt, W.A., Somenzi, F. (eds.) *Computer Aided Verification*. pp. 1–13. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
27. Menezes, A.J., Vanstone, S.A., Oorschot, P.C.V.: *Handbook of Applied Cryptography*. CRC Press, Inc., Boca Raton, FL, USA, 1st edn. (1996)
28. Rocchetto, M., Tippenhauer, N.O.: On attacker models and profiles for cyber-physical systems. In: Askoxylakis, I., Ioannidis, S., Katsikas, S., Meadows, C. (eds.) *Computer Security – ESORICS 2016*. pp. 427–449. Springer International Publishing, Cham (2016)
29. Rothstein-Morris, E., Sun, J., Chattopadhyay, S.: Systematic Classification of Attackers via Bounded Model Checking (Extended Version). arXiv:1911.05808 (2019), <http://arxiv.org/abs/1911.05808>
30. Rothstein, E., Murguia, C.G., Ochoa, M.: Design-time quantification of integrity in cyber-physical systems. In: *Proceedings of the 2017 Workshop on Programming Languages and Analysis for Security*. pp. 63–74. PLAS '17, ACM, New York, NY, USA (2017), <http://doi.acm.org/10.1145/3139337.3139347>
31. Tseitin, G. S.: On the Complexity of Derivation in Propositional Calculus. In: *Automation of Reasoning: 2: Classical Papers on Computational Logic (1967–1970)*
32. Urbina, D.I., Giraldo, J.A., Cardenas, A.A., Tippenhauer, N.O., Valente, J., Faisal, M., Ruths, J., Candell, R., Sandberg, H.: Limiting the impact of stealthy attacks on industrial control systems. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. pp. 1092–1105. CCS '16, ACM, New York, NY, USA (2016), <http://doi.acm.org/10.1145/2976749.2978388>
33. Weerakkody, S., Sinopoli, B., Kar, S., Datta, A.: Information flow for security in control systems. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. pp. 5065–5072 (Dec 2016)