Model Checking the Information Flow Security of Real-Time Systems

Christopher Gerking[1], David Schubert[2], and Eric Bodden[1,2]

[1] Paderborn University, Heinz Nixdorf Institute, Germany,
christopher.gerking@upb.de
[2] Fraunhofer IEM, Paderborn, Germany

Abstract. Cyber-physical systems are processing large amounts of sensitive information, but are increasingly often becoming the target of cyber attacks. Compared to smartphones and PCs, updating and maintaining cyber-physical systems is a major cost factor. Thus, it is essential to verify upfront the absence of unauthorized information flow. Our paper addresses this problem by proposing a novel approach to model-check the information flow security of cyber-physical systems represented by timed automata. We describe the transformation into so-called test automata, reducing the verification to a reachability test carried out using the off-the-shelf model checker Uppaal. Opposed to related work, we analyze the real-time behavior of systems, allowing software engineers to precisely identify timing channels that would enable attackers to draw conclusions from the system's response times. We illustrate the approach by example of a simplified model of a cyber-manufacturing system.

Keywords: model checking, information flow, security, real time

1 Introduction

Cyber-physical systems [35] are entrusted a fast-growing amount of sensitive data, but are inherently vulnerable to security breaches such as manipulation or leakage of information [25,15]. One subtle attack vector are *timing channels* [11], allowing attackers to infer sensitive information by observing the system's response times. In the worst case, such hidden flows of information could even be exploited to manipulate the physical behavior and compromise the safety of systems. Opposed to smartphones and PCs, updating cyber-physical systems is a major cost factor. Thus, it is essential to verify their *information flow security* early on, before the systems get deployed, to make them secure by design [40].

Model-driven engineering is widely used to secure cyber-physical systems [38], making the software under development accessible to formal verification. A well-established formal definition of information flow security is *noninterference* [26], restricting systems such that their publicly observable behavior must never depend on sensitive behavior. A well-known verification approach for noninterference is *bisimulation*, checking the observational equivalence between a system and a restricted variant of itself that is known to be secure by definition [19].

Nevertheless, checking the information flow security of cyber-physical systems is a challenging problem for software engineers because they are faced with *real-time systems*, restricted by hard timing constraints imposed by their physical environment [12]. Therefore, models of cyber-physical systems are based on formalisms like *timed automata* [5], and the real-time behavior needs to be taken into account during verification to detect leaks such as timing channels precisely. Verification techniques such as model checking of timed automata are available [4], but involve sophisticated tools that are hard to implement from scratch. In this paper, we therefore address the problem of applying existing verification tools for checking the information flow security of real-time systems.

Bisimulation of real-time systems is known to be decidable by existing verification techniques [14]. Nevertheless, previous approaches towards applied verification of information flow security have not taken into account real-time behavior [3,32,20,17]. Therefore, they fail to detect leaks such as timing channels precisely. Related work on the information flow security of timed automata exists [6,13,34,8,39], but has spent little effort on how to apply tool-supported model checking techniques in practice. Thus, in summary, none of the previous approaches fully combines real-time behavior with applied model checking.

In this paper, we fill in these gaps by reducing the check for noninterference of timed automata to a *refinement* check, adapting the work by Heinzemann et al. [30] to the application field of information flow security. This check is based on model transformations to construct a *test automaton* [1], introducing a dedicated location that is only reachable when violating a bisimulation between the original automaton and a variant that is secure by definition. By model checking the reachability of these dedicated locations using the off-the-shelf tool Uppaal [9], we obtain a novel verification technique for the information flow security of real-time systems. In contrast to related approaches, our work is based on timed automata, taking into account the real-time behavior of cyber-physical systems. Unlike other related work, we focus on applied model checking to meet the needs of software engineering practitioners.

We illustrate the approach using a simplified model of a cyber-manufacturing system that interacts with a cloud-based service market. Since market participants are potential attackers, the system must not allow them to draw any conclusions about a production plan that is classified as a business secret.

In summary, this paper makes the following contributions:

- We propose a model transformation process, reducing the check for information flow security of real-time systems to a model checking problem.
- At the core of this process, we illustrate the construction of test automata to check noninterference of timed automata.
- We give a proof of concept by detecting a timing channel in a simplified model of a cyber-manufacturing system.

Paper Organization: We introduce fundamentals in Section 2, and discuss related work in Section 3. In Section 4, we describe our approach of checking information flow security of real-time systems. We give a proof of concept in Section 5, before concluding in Section 6.

2 Fundamentals

In this section, we recall timed automata (cf. Section 2.1), timed bisimulation (cf. Section 2.2), and noninterference (cf. Section 2.3). Based on these fundamental concepts, we introduce our motivating example in Section 2.4.

2.1 Timed Automata

The formalism of *timed automata* [5] is used to model real-time behavior of stateful systems. A timed automaton is essentially a directed graph containing a finite set of locations, connected by a finite set of labeled edges. We utilize the definition of timed automata by Bengtsson and Yi [10]. Timed automata extend finite automata by real-valued variables that represent *clocks*. Clocks are initialized with zero, increase at the same rate, and may be set back to zero by using *resets* that can be assigned to edges.

Clock constraints restrict the behavior of an automaton with respect to the valuation of its clocks. A clock constraint is a conjunction of atomic constraints of the form $x \sim n$ or $x - y \sim n$, where x and y are clocks, $n \in \mathbb{N}$, and $\sim \in \{\leq, <, =, >, \geq\}$ [10]. Clock constraints are used as *invariants* and *time guards*. Invariants are assigned to locations. An active location is forced to be left by firing an edge before the location's invariant expires. Therefore, invariants have to be downward closed, i.e., only $\leq$ and $<$ operators are permitted. Time guards are assigned to edges. An edge may fire (i.e., it is *enabled*) only if its time guard evaluates to true. In addition, edges are labeled with actions, whereas firing an edge represents the execution of the action its is labelled with. To represent edges without an action, we refer to τ as the empty action.

Assuming a set $\mathcal{C}$ of clocks, a set $\mathcal{B}(\mathcal{C})$ of clock constraints, and an alphabet Σ of actions, the syntax of a timed automaton is defined as follows [10]:

Definition 1. *A timed automaton A is a tuple $\langle L, l_0, E, I \rangle$ where*

- *L is a finite set of locations,*
- *$l_0 \in L$ is the initial location,*
- *$E \subseteq L \times \mathcal{B}(\mathcal{C}) \times \Sigma \times 2^{\mathcal{C}} \times L$ is the set of edges where $\rho \in \mathcal{B}(\mathcal{C})$ is the time guard, $\mu \in \Sigma$ the action, and $\lambda \in 2^{\mathcal{C}}$ is the set of clock resets,*
- *$I : L \to \mathcal{B}(\mathcal{C})$ assigns invariants to locations.*

Uppaal[3] [9] is a well-established model checker for timed automata that is commonly applied by software engineering practitioners, and frequently integrated into domain-specific verification tools [44]. In the scope of this paper, we apply Uppaal to verify the information flow security of timed automata.

[3] `http://uppaal.org`

2.2 Timed Bisimulation

Bisimulation is a notion of equivalence that requires the observable behavior of two systems to be indistinguishable. *Timed bisimulation* is an extension of bisimulation for real-time systems, which is known to be decidable [14]. Intuitively, two systems are equivalent in terms of timed bisimulation if they execute the same sequences of observable actions in the same time intervals. We refer the reader to [10] for a formal definition in the context of timed automata.

There are two variants of timed bisimulation. *Strong* timed bisimulation is more restrictive as it considers all actions of a system as being observable, including τ actions. In the context of our paper, this assumption is too strong because we consider τ actions as internal and, therefore, not observable. In contrast to this, *weak* timed bisimulation ignores the execution of internal τ actions [14]. In the following, we consider only this weak variant of timed bisimulation between two timed automata A and B (denoted by $A \approx B$), and refer to it as timed bisimulation for brevity.

2.3 Noninterference

Noninterference was introduced by Goguen and Meseguer [26] to define information flow security for deterministic finite automata, such that the publicly observable behavior does never depend on any sensitive behavior. If so, the public observations never enable an unprivileged actor to distinguish whether or not sensitive actions are executed. In particular, no conclusions are possible about which sensitive actions actually occur. To characterize the sensitivity of behavior, the concept of noninterference is based on a separation between sensitive (or *high*) actions Σ_H and public (or *low*) actions Σ_L with $\Sigma_H, \Sigma_L \subseteq \Sigma$.

For nondeterministic systems such as timed automata, noninterference is frequently defined on the basis of bisimulation [19,34]. Noninterference holds if the publicly observable behavior of a system cannot be distinguished from a restricted behavior that is secure by definition. To define this property more precisely, we distinguish between *input actions* received from the environment, and *output actions* sent to the environment. Based on this distinction, noninterference reduces to a bisimulation of the publicly observable behavior between

1. the original system, and
2. a secure-by-definition system with all sensitive input actions *disabled*.

Disabling sensitive input actions ensures that the secure-by-definition system behaves as if no sensitive information is present. By disabling only inputs, we assume that all sensitive information is received from the environment, and never generated internally without depending on sensitive inputs [29].

To identify deviations only in the publicly observable behavior, non-public actions need to be *hidden* from the bisimulation, i.e., treated as non-observable. In the following definition, $\setminus_I$ denotes the disabling of inputs, and / the hiding of actions, wheras $\Sigma_{\bar{L}} = \Sigma \setminus \Sigma_L$ is the set of non-public actions.

Definition 2. *Timed noninterference holds for a timed automaton A, if and only if $A \,/\, \Sigma_{\bar{L}} \approx (A \setminus_I \Sigma_H) \,/\, \Sigma_{\bar{L}}$.*

2.4 Motivating Example

In our approach, we assume timed automata to be embedded in component-based software architectures, which are commonly used for the software design of cyber-physical systems [16]. In Figure 1, we show a software component named ManufacturingSystem as a model of our example announced in Section 1. The component embeds a timed automaton that drives the application-layer communication between the system and its environment. The communication is carried out by means of asynchronous message passing, whereas the set of messages corresponds to the alphabet Σ. Accordingly, when messages are received, they are buffered until they are processed by an input action of the automaton. Asynchronous communication is a characteristic property of cyber-physical systems because they are often spatially divided and dynamically interconnected over wireless networks. In Figure 1, we use / to separate input from output actions.

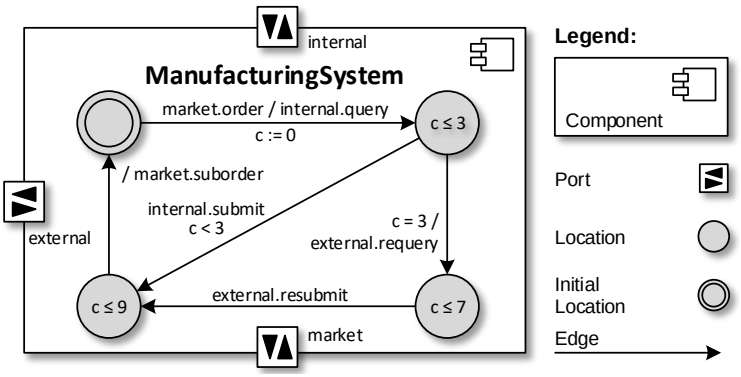


Fig. 1. Example timed automaton of a cyber-manufacturing system

Our example component uses three ports to exchange messages. The market port is used to interact with the public service market. Whenever a production order is received from the market, the internal port is the preferred way to access the production plan, provided that it is available as an in-house resource. Alternatively, access to the plan may also be purchased from a business partner over the external port. The messages exchanged over the market port are public (represented by the set Σ_L), whereas the messages exchanged over the internal port are sensitive (represented by Σ_H). In our example, messages exchanged over the external port are not further characterized as public or sensitive because information flow needs to be detected only from the internal to the market port.

In the following, we describe the behavior of the automaton from Figure 1. Whenever the automaton is in its initial location and receives an order message from the market, it sends a query for the production plan over the internal port. At the same time, it sets the clock c to zero which acts as a timeout. If the internal port does not provide the plan in terms of a submit message within three

time units (i.e., the company does not own the plan), a requery message is sent over the external port to purchase the plan from a business partner. In this case, the example system assumes that the plan is provided in terms of a resubmit message within four further time units ($c \leq 7$), i.e., a deadlock caused by an overdue message can never occur. Finally, if the plan is delivered from either the internal or external port, the system orders the corresponding subproducts from the market by sending a suborder message in the time interval $c \leq 9$.

However, the system violates timed noninterference because the effective timing of the public suborder message depends on whether or not the plan is owned by the company. If this is the case, the plan is provided by the sensitive internal port, and suborder may be sent when $c < 3$. Otherwise, when provided by the external port, the suborder message can only be sent when $c \geq 3$. This deviation represents a timing channel that allows market participants to infer whether the company owns the production plan or not. This knowledge is sensitive information that could be exploited in a future attack to gain access to the plan. Due to the subtleties of real time, such leaks can easily remain undetected during software design, and thus require a tool-supported verification technique.

3 Related Work

In Section 3.1, we recall general approaches towards checking information flow security, which are complementary to our work. In Section 3.2, we discuss related work on the information flow security of time-dependent systems.

3.1 Complementary Approaches

Unwinding [27] is a traditional verification technique to infer global information flow security from local properties of individual system actions (e.g., state transitions). In the context of real-time systems, this approach is hindered by the infinite, real-valued state space which makes such local properties hard to identify. *Language-based security* [43] is concerned with secure information flow at the level of programming languages, thus using a different model of computation compared to our automata-based approach. In this area, *type systems* are often used to enforce information flow security of programs statically. Furthermore, a technique called *self-composition* has been proposed [7], reducing language-based security to a logical formulation that is amenable to automated verification, similar to our approach in the context of automata-based systems. We refer the reader to [21,37] for a comparison of information flow security under different models of computation. Another complementary approach is the one by Finkbeiner et al. [18] on model checking *hyperproperties* which go beyond standard *safety* or *liveness* properties, as it is the case for security. Whereas this approach involves developing a completely new theory of specification and verification, our focus is on applied verification using existing techniques.

3.2 Timed Information Flow Security

In Table 3.2, we compare related work on the information flow security of time-dependent systems against the core characteristics of our approach, which combines dense real-time behavior with applied verification. Furthermore, we build on automata as the underlying model of computation, which are commonly used as a natural, well-established modeling approach [36]. Finally, according to our example given in Section 2.4, we focus on application-layer modeling, i.e., we abstract from responsibilities like scheduling.

Table 1. Comparative overview of approaches on timed information flow security

	Dense Real-Time	Automata-Based	Application Layer	Applied Verification
Evans & Schneider [17]	✗	✗	✓	✓
Focardi et al. [20]	✗	✗	✓	✓
Akella et al. [3]	✗	✗	✓	✓
Agat [2]	✗	✗	✓	✗
Giacobazzi & Mastroeni [24]	✗	✗	✓	✗
Rafnsson et al. [41]	✗	✗	✓	✗
Köpf & Basin [32]	✗	✓	✗	✓
Roscoe & Huang [42]	✓	✗	✓	✗
Son & Alves-Foss [45]	✓	✗	✗	✗
Kashyap et al. [31]	✓	✗	✗	✗
Cassez [13]	✓	✓	✓	✗
Lanotte et al. [34]	✓	✓	✓	✗
Benattar et al. [8]	✓	✓	✓	✗
Nielson et al. [39]	✓	✓	✓	✗
Barbuti & Tesei [6]	✓	✓	✓	(✓)
This paper	✓	✓	✓	✓

The work by Evans and Schneider [17], Focardi et al. [20], and Akella et al. [3] is based on process algebra. Existing verification techniques like theorem proving [17] or partial model checking [20] are applied, even in the context of cyber-physical systems [3]. By using process algebra, the authors differ from our automata-based approach in terms of their model of computation. In the context of language-based security, the work by Agat [2], Giaccobazzi and Mastroeni [24], as well as Rafnsson et al. [41] uses imperative programs as yet another model of computation. In contrast, the work by Köpf and Basin [32] on synchronous systems is automata-based and also amenable to applied verification. However, all of the above approaches are limited to discrete time, which is insufficient to capture the real-time behavior of cyber-physical systems. In contrast, other approaches consider dense real-time behavior. Among them, Roscoe and Huang [42] use process algebra as model of computation, differing from our automata-based approach. Son and Alves-Foss [45] as well as Kashyap et al. [31] both focus on scheduling of real-time tasks, i.e., are located beyond the application layer.

In the context of timed automata, Cassez [13] presents a real-time security property called *timed opacity* as a generalization of noninterference. The author proves the undecidability of the verification problem, i.e., is not concerned with applied verification. Lanotte et al. [33] consider real-time privacy properties of timed automata, and reduce the verification of such properties to a reachability analysis [28], similar to this paper. In [34], the same authors consider noninterference of timed automata extended by probabilistic behavior. However, the application of existing model checking techniques is beyond the scope of their approach. Benattar et al. [8] address the synthesis of controllers that ensure noninterference of timed automata. According to this constructive approach, they do not consider applied verification as well. Nielson et al. [39] combine timed automata with language-based security. The authors propose a type system to enforce noninterference in an imperative language, and translate this language into timed automata. Again, applied verification is beyond their scope.

Barbuti and Tesei [6] focus on noninterference of timed automata. Similar to our approach, they reduce the verification to a reachability analysis using model checking. However, their verification approach only checks that sensitive behavior does not influence the reachability of locations. Since this approximation gives rise to both *false positive* and *false negative* errors, it is not able to provide any security guarantee. Nevertheless, the approach by Barbuti and Tesei [6] is the only one that resembles the core characteristics of our paper (cf. Table 3.2).

4 Checking Noninterference of Timed Automata

In the following, we describe our approach of checking the information flow security of real-time systems. In Section 4.1, we give an overview on our approach, and describe the construction of the underlying test automata in Section 4.2.

4.1 Refinement Checking

To check timed noninterference, we reduce the verification to a refinement check for real-time systems as proposed by Heinzemann et al. [30]. The aforementioned work allows to verify refinement relations between real-time systems to check that an abstract behavior is correctly refined by a concrete behavior. The authors reduce the verification to a reachability test [1] that is carried out using model checking techniques. One possible refinement definition is timed bisimulation, as also used to define timed noninterference. Thus, in this paper, we adopt the notion of refinement to check timed noninterference. In Figure 2, we give an overview on our approach as an extension of the work by Heinzemann et al. [30].

In step ①, we transform a timed automaton A, as described in Section 2.4, into an auxiliary automaton A_{sec} that is secure by definition because sensitive inputs are disabled. This restriction corresponds to the automaton $A \setminus_I \Sigma_H$ from Definition 2. We disable sensitive inputs by removing the corresponding edges from A. In the context of the motivating example, Figure 3a depicts in red the removal of the edge that processes the sensitive input internal.submit.

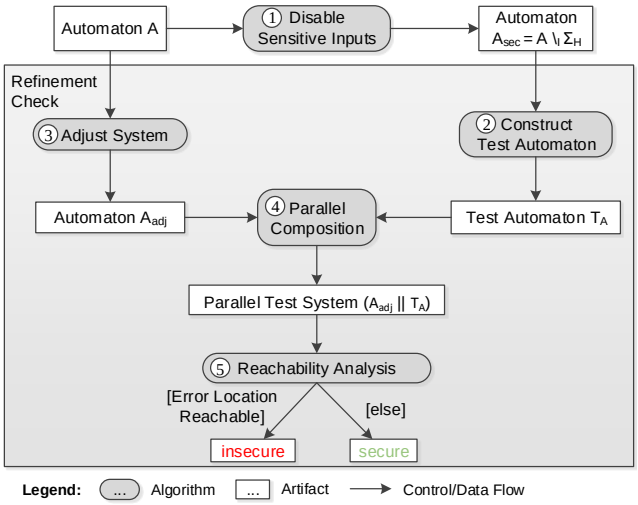
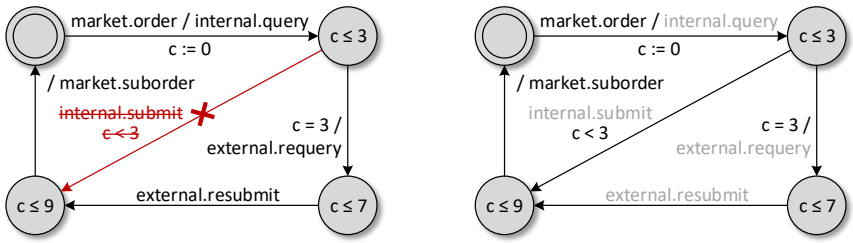


Fig. 2. Reduction of the noninterference check to a refinement check

The resulting automaton enables us to execute a specialized version of the refinement check for timed bisimulation, as proposed by Heinzemann et al. [30]. This approach allows to detect cases where a timed bisimulation between the original automaton A and the secure-by-definition automaton A_{sec} is violated because A deviates from the publicly observable behavior of A_{sec}. At the core of the approach is a *test automaton* [1] that acts as an oracle for the information flow security of the original automaton. In particular, the test automaton detects cases in which the behavior of the original automaton violates timed noninterference. To this end, step ② transforms the automaton A_{sec} into a test automaton T_A, introducing a designated *error* location that is reachable when violating timed noninterference by deviating from the secure-by-definition behavior.

According to Definition 2, we need to hide all non-public actions during the check for timed bisimulation because only deviations in the public behavior are violations of timed noninterference. In Figure 3b, we depict in grey those actions that need to be hidden. A natural approach to hide an action is to remove it from the corresponding edge [6,34], i.e., replacing it by a τ action. However, removing input actions may lead to an increased nondeterminism. The reason is that this approach potentially produces multiple τ transitions that are all executable on the same condition because they are no longer distinguishable by their input actions. However, the refinement check proposed by Heinzemann et al. is restricted to systems with a deterministic transition function [30], i.e., at most one edge can fire on a certain condition. Therefore, in contrast to the default test automata for timed bisimulation [30], our test automata must take responsibility for hiding non-public actions. We describe the construction of these specialized test automata in the upcoming Section 4.2.



(a) Disabling sensitive inputs (b) Hiding non-public actions

Fig. 3. Disabling and hiding of actions in the motivating example

To ensure that the test automaton acts as the oracle for the original automaton, we need to couple both automata with each other. Therefore, step ③ creates an adjusted automaton A_{adj} that has the same behavior as A. However, it synchronizes with the test automaton whenever both execute the same action. Furthermore, A_{adj} supports the hiding of non-public actions in the same fashion as the test automaton. In step ④, we compose both T_A and A_{adj} in parallel to enable synchronized execution of the automata. In the final step ⑤, the check for timed noninterference reduces to analyzing the resulting parallel test system for reachability of the *error* location. This reachability test [1] is carried out using the real-time model checker Uppaal [9], using its parallel composition operator || to enable the synchronizations between both automata [10]. In the end, the automaton A is secure in terms of timed noninterference, if and only if the *error* location is unreachable on all execution paths.

4.2 Test Automata Construction

To generate test automata, we adjust the construction schema for timed bisimulation proposed by Heinzemann et al. [30] such that it hides non-public communication, as demanded by Definition 2. We adopt the notion of a designated *error* location (named *Err* in our case) that is reachable if and only if timed noninterference is violated. Figure 4 illustrates our construction schema including the *Err* location. We apply this schema for each edge $S \rightarrow S'$ of A_{sec}. Our construction must ensure that T_A will ❶ accept secure communication (allowed by A_{sec} and correctly reflected by the original automaton A), ❷ reject insecure communication (i.e., public communication that is not allowed by A_{sec} but incorrectly present in A), and ❸ detect the absence of communication (i.e., public communication that is allowed by A_{sec} but incorrectly absent in A). Before going into details about the three cases, we introduce additional notation used in Figure 4. Synchronous input and output actions are denoted by $\mu?$ and $\mu!$ respectively. The set $\Theta(S)$ includes all actions that are *allowed* in a location S, i.e., all actions of outgoing edges of S. The invariant of S is denoted by $I(S)$. In accordance with Definition 1, ρ is a time guard, and λ is a set of clock resets.

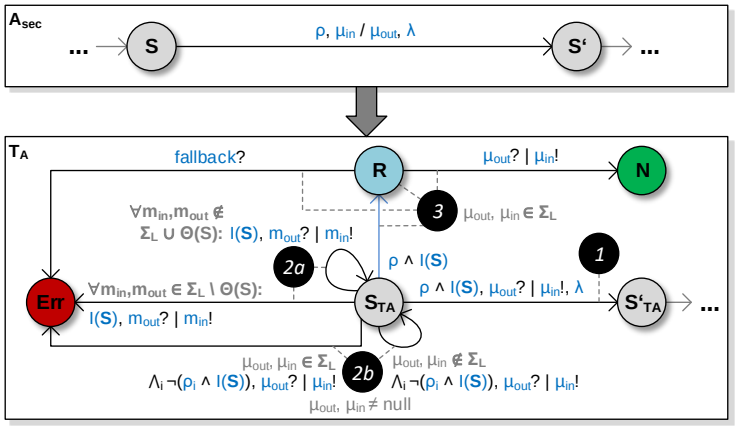


Fig. 4. Construction schema for test automata

Accepting Secure Communication. The edge labeled with ❶ in Figure 4 ensures that secure communication is accepted by T_A. For each edge $S \to S'$ in A_{sec}, we add an edge $S_{TA} \to S'_{TA}$ to T_A. Asynchronous actions of $S \to S'$ are transformed into synchronous actions to enable synchronization with A_{adj} when processing the same messages during parallel composition. Here, a synchronous output action $\mu_{in}!$ is created from an asynchronous input action μ_{in}, or a synchronous input action $\mu_{out}?$ is created from an asynchronous output action μ_{out}. Furthermore, we add the set of clock resets λ of $S \to S'$. Finally, we need to preserve the time intervals in which input/output actions are executed. Thus, the time guard of the edge is the conjunction of the original time guard ρ and the invariant $I(S)$ of S. Thus, $S_{TA} \to S'_{TA}$ is enabled whenever $S \to S'$ is enabled.

Rejecting Insecure Communication. Insecure communication includes two cases: ❷a executing a public action that is not allowed in a location S because no outgoing edge is labelled with it, and ❷b executing a public action that is allowed in S, but violates the timing defined in A_{sec}. For case ❷a, Figure 4 introduces one edge $S_{TA} \to Err$ for each action that is public (i.e., in Σ_L) but not in the set $\Theta(S)$ of messages allowed in S. The time guard $I(S)$ of these edges ensures that Err is reachable only by actions executed during the activity of S. In contrast to the construction by Heinzemann et al. [30], only communication using public actions makes Err reachable. For the opposite case of non-public communication, we add one self-edge $S_{TA} \to S_{TA}$ for each message that is not allowed, but also not in Σ_L. Thereby, instead of switching to the Err location, T_A hides all non-public actions that are not allowed.

To handle the timing violations of case ❷b, we create one more edge labelled with the allowed action μ_{out} or μ_{in} of $S \to S'$. However, this edge is enabled exactly at those times when A_{sec} does not allow its execution. To this end, the time guard of the edge is the conjunction of the negated enabling conditions for

all edges $S \to S^i$ executing the same action as $S \to S'$, resulting in $\bigwedge_i \neg(\rho_i \wedge I(S))$ where ρ_i is the time guard of $S \to S^i$. Thereby, we ensure that the edge can fire when exceeding the upper bounds of the time intervals in which the action is allowed to execute, or when falling below the lower bounds. In contrast to Heinzemann et al. [30], the edge leads to Err only in case of public messages from Σ_L. For non-public messages, analogously to case 2a, we construct a self-edge $S_{TA} \to S_{TA}$. This edge has the same enabling conditions, however, hides the corresponding non-public action instead of switching to the Err location.

Detecting Absent Communication. Timed noninterference demands that all public actions executable by A_{sec} are executable by the original automaton A in the same time intervals. To check such restrictions, Heinzemann et al. [30] add the constructs labeled with 3 in Figure 4. We adopt this construction only if μ_{out} or μ_{in} is public (i.e., in Σ_L). In this case, the location R represents a check for required communication. It is reachable during the full time interval in which the edge $S \to S'$ is enabled.

If A preserves the time interval in which A_{sec} can execute the public action, then the edge $R \to N$ is enabled whenever R is entered. The location N represents a neutral state of the analysis that is reachable whenever the required public action is properly executed. N has no outgoing edges because the execution does not have to be further explored from here. Instead, the location S'_{TA} is always reachable when N is reachable and ensures regular execution.

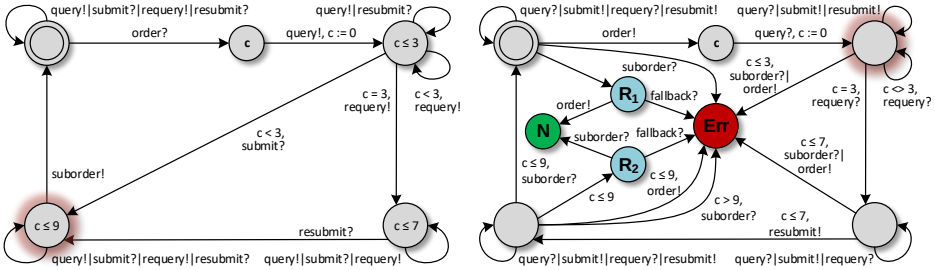
If at some time during its designated interval, the required public action can not be executed (because A lowers the upper bound or raises its lower bound), then $R \to N$ is not enabled. In this case, the edge $R \to Err$ fires by synchronizing over an auxiliary channel named $fallback$. Synchronization over this channel is always enabled, however, it has the lowest priority compared to all other channels used for communication. Thereby, Err is only reachable when the required public communication is absent.

5 Proof of Concept

In this section, we showcase the utility of our approach in the scope of the example given in Section 2.4. To this end, we demonstrate that our technique outperforms the related work by detecting a timing channel that remains undetected using the approach by Barbuti and Tesei [6]. We show that our technique rejects the insecure system, but accepts a mitigated system that is noninterferent.

To verify the information flow security of timed automata, Barbuti and Tesei check that disabling sensitive actions does not affect the reachability of locations. In particular, the authors verify that disabling edges with sensitive actions leads to the same amount of reachable locations compared to hiding sensitive actions. In the scope of our example, the corresponding transformation was shown in Figure 3a. Clearly, disabling the input internal.submit does not affect the amount of reachable locations compared to Figure 3b, and thus represents a *false negative* because the timing channel described in Section 2.4 can not be detected.

In contrast, Figure 5 illustrates the parallel composition of the adjusted automaton A_{adj} (Figure 5a) and the test automaton T_a (Figure 5b) as proposed in this paper. As an artifact for reproduction, we provide a corresponding model that is verifiable by the UPPAAL model checker [23]. Due to lack of space, Figure 5 merges multiple edges between the same source and target locations into a single edge with alternative synchronization labels. Furthermore, we omit the names of ports over which messages are sent or received. Finally, since UPPAAL does not allow edges with a conjunction of actions, we use a *committed* location [10] (labeled with c in Figure 5) to divide the order and query actions into two consecutive edges. Figure 5 also depicts the additional self-edges added to both automata for hiding non-public communication, as described for the test automata in Section 4.2.

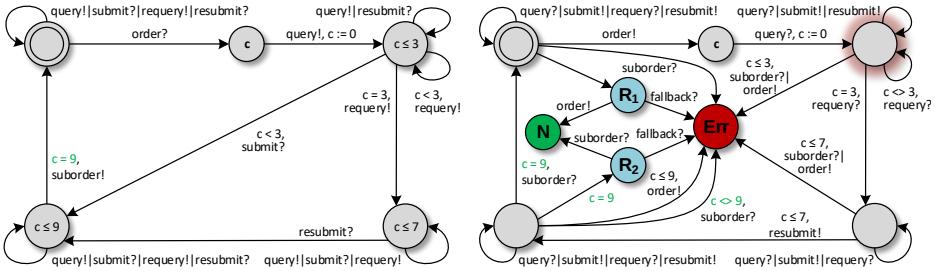


(a) Adjusted automaton A_{adj} (b) Test automaton T_a

Fig. 5. Parallel test system for the motivating example

In the situation depicted in Figure 5, the system has already processed the message sequence order, query, submit (in the time interval $c < 3$). Since sensitive inputs are disabled in T_a, it can not change its location when processing the sensitive submit message. Next, the public suborder message to be sent by A_{adj} corresponds to case 2a of our construction in Section 4.2. Thus, T_a will regard the message as insecure, and reject it by switching to the Err location. The reason for this violation is that A_{adj} sends the public suborder message too early, i.e., in the time interval $c \in [0, 3]$. Thus, in its current location (cf. Figure 5b), T_a regards the message as not allowed at this time, and switches to Err.

Figure 6 shows the countermeasures taken to mitigate the detected timing channel. In Figure 6a, we depict in green the time guard $c = 9$ added to delay the suborder message. Consequently, the timing of the message does no longer depend on whether or not the production plan was provided over the sensitive internal port. Figure 6b depicts in green the consequent changes of the test automaton. Since the timing of the suborder message is now fixed, the Err location is no longer reachable. Accordingly, our construction correctly identifies the mitigated system as noninterferent.



(a) Adjusted automaton A_{adj} (b) Test automaton T_a

Fig. 6. Mitigation of the timing channel in the motivating example

6 Conclusions & Future Work

This paper proposes a novel check for the information flow security of real-time systems given in the form of timed automata. Our approach is based on noninterference as a well-established definition of secure information flow. To provide a verification technique that applies existing tools and takes into account real-time behavior, we adapt the work on refinement checking by Heinzemann et al. [30] to the application field of security. We describe the transformation of timed automata into test automata, introducing a designated location that indicates violations of noninterference whenever it is reachable during execution. Thereby, we reduce the problem to a reachability test that is supported by model checking techniques used in software engineering practice. In particular, we apply the well-established Uppaal model checker for timed automata as our underlying verification engine. Our proof of concept demonstrates the advantages of our approach by detecting a timing channel that would remain undetected using the most closely related work.

The proposed idea provides software engineering practitioners with a tool-supported verification technique for the information flow security of timed automata, taking into account specific characteristics of cyber-physical systems like real-time behavior and asynchronous communication. Thereby, we enable engineers to identify information leaks such as timing channels early, and ensure security by design. For cyber-physical systems, this is of vital importance to avoid product recalls or even safety-critical attacks.

The proposed approach is part of ongoing work on tracing information flow security in cyber-physical systems engineering [22]. In future work, we will provide an in-depth validation of our approach in the context of a model-driven software design method for cyber-physical systems. In particular, to check the information flow security of hierarchical component architectures, our work needs to be extended to a compositional verification approach. Thereby, we seek to preserve security when composing overall software systems from single secure components.

References

1. Aceto, L., Burgueño, A., Larsen, K.G.: Model checking via reachability testing for timed automata. In: TACAS 1998. LNCS, vol. 1384, pp. 263–280. Springer (1998)
2. Agat, J.: Transforming out timing leaks. In: POPL 2000. pp. 40–53. ACM (2000)
3. Akella, R., Tang, H., McMillin, B.M.: Analysis of information flow security in cyber-physical systems. International Journal of Critical Infrastructure Protection 3(3-4), 157–173 (2010)
4. Alur, R., Courcoubetis, C., Dill, D.L.: Model-checking in dense real-time. Information and Computation 104(1), 2–34 (1993)
5. Alur, R., Dill, D.L.: A theory of timed automata. Theoretical Computer Science 126(2), 183–235 (1994)
6. Barbuti, R., Tesei, L.: A decidable notion of timed non-interference. Fundamenta Informaticae 54(2-3), 137–150 (2003)
7. Barthe, G., D'Argenio, P.R., Rezk, T.: Secure information flow by self-composition. Mathematical Structures in Computer Science 21(6), 1207–1252 (2011)
8. Benattar, G., Cassez, F., Lime, D., Roux, O.H.: Control and synthesis of non-interferent timed systems. International Journal of Control 88(2), 217–236 (2015)
9. Bengtsson, J., Larsen, K., Larsson, F., Pettersson, P., Yi, W.: Uppaal – a tool suite for automatic verification of real-time systems. In: Hybrid Systems III, pp. 232–243. Springer (1996)
10. Bengtsson, J., Yi, W.: Timed automata: Semantics, algorithms and tools. In: Lectures on Concurrency and Petri Nets, pp. 87–124. Springer (2004)
11. Biswas, A.K., Ghosal, D., Nagaraja, S.: A survey of timing channels and countermeasures. ACM Computing Surveys 50(1), 6:1–6:39 (2017)
12. Broman, D., Derler, P., Eidson, J.: Temporal issues in cyber-physical systems. Journal of Indian Institute of Science 93(3) (2013)
13. Cassez, F.: The dark side of timed opacity. In: ISA 2009. LNCS, vol. 5576, pp. 21–30. Springer (2009)
14. Cerans, K.: Decidability of bisimulation equivalences for parallel timer processes. In: CAV 1992. LNCS, vol. 663, pp. 302–315. Springer (1992)
15. Chattopadhyay, A., Prakash, A., Shafique, M.: Secure cyber-physical systems: Current trends, tools and open research problems. In: DATE 2017. pp. 1104–1109. IEEE (2017)
16. Crnkovic, I., Malavolta, I., Muccini, H., Sharaf, M.: On the use of component-based principles and practices for architecting cyber-physical systems. In: CBSE 2016. pp. 23–32. IEEE (2016)
17. Evans, N., Schneider, S.: Analysing time dependent security properties in CSP using PVS. In: ESORICS 2000. LNCS, vol. 1895, pp. 222–237. Springer (2000)
18. Finkbeiner, B., Rabe, M.N., Sánchez, C.: Algorithms for model checking HyperLTL and HyperCTL*. In: CAV 2015. LNCS, vol. 9206, pp. 30–48. Springer (2015)
19. Focardi, R., Gorrieri, R.: A taxonomy of security properties for process algebras. Journal of Computer Security 3(1), 5–34 (1995)
20. Focardi, R., Gorrieri, R., Martinelli, F.: Real-time information flow analysis. IEEE Journal on Selected Areas in Communications 21(1), 20–35 (2003)
21. Focardi, R., Rossi, S., Sabelfeld, A.: Bridging language-based and process calculi security. In: FOSSACS 2005. LNCS, vol. 3441, pp. 299–315. Springer (2005)
22. Gerking, C.: Traceability of information flow requirements in cyber-physical systems engineering. In: DocSym@MoDELS 2016. CEUR Workshop Proceedings, vol. 1735 (2016)

23. Gerking, C.: Detection of a timing channel in an Uppaal model of a cyber-manufacturing system (2018), https://doi.org/10.5281/zenodo.1035200
24. Giacobazzi, R., Mastroeni, I.: Timed abstract non-interference. In: FORMATS 2005. LNCS, vol. 3829, pp. 289–303. Springer (2005)
25. Giraldo, J., Sarkar, E., Cárdenas, A.A., Maniatakos, M., Kantarcioglu, M.: Security and privacy in cyber-physical systems: A survey of surveys. IEEE Design & Test 34(4), 7–17 (2017)
26. Goguen, J.A., Meseguer, J.: Security policies and security models. In: IEEE S&P. pp. 11–20. IEEE (1982)
27. Goguen, J.A., Meseguer, J.: Unwinding and inference control. In: IEEE S&P. pp. 75–87. IEEE (1984)
28. Gorrieri, R., Lanotte, R., Maggiolo-Schettini, A., Martinelli, F., Tini, S., Tronci, E.: Automated analysis of timed security. International Journal of Information Security 2(3-4), 168–186 (2004)
29. Guttman, J.D., Nadel, M.E.: What needs securing. In: CSFW. pp. 34–57. MITRE Corporation Press (1988)
30. Heinzemann, C., Brenner, C., Dziwok, S., Schäfer, W.: Automata-based refinement checking for real-time systems. Computer Science - R&D 30(3-4), 255–283 (2015)
31. Kashyap, V., Wiedermann, B., Hardekopf, B.: Timing- and termination-sensitive secure information flow. In: IEEE S&P. pp. 413–428. IEEE (2011)
32. Köpf, B., Basin, D.A.: Timing-sensitive information flow analysis for synchronous systems. In: ESORICS 2006. LNCS, vol. 4189, pp. 243–262. Springer (2006)
33. Lanotte, R., Maggiolo-Schettini, A., Tini, S.: Privacy in real-time systems. Electronic Notes in Theoretical Computer Science 52(3), 295–305 (2001)
34. Lanotte, R., Maggiolo-Schettini, A., Troina, A.: Time and probability-based information flow analysis. IEEE Transactions on Software Engineering 36(5), 719–734 (2010)
35. Lee, E.A.: CPS foundations. In: DAC 2010. pp. 737–742. ACM (2010)
36. van der Meyden, R., Zhang, C.: Algorithmic verification of noninterference properties. Electronic Notes in Theoretical Computer Science 168, 61–75 (2007)
37. van der Meyden, R., Zhang, C.: A comparison of semantic models for noninterference. Theoretical Computer Science 411(47), 4123–4147 (2010)
38. Nguyen, P.H., Ali, S., Yue, T.: Model-based security engineering for cyber-physical systems. Information & Software Technology 83, 116–135 (2017)
39. Nielson, F., Nielson, H.R., Vasilikos, P.: Information flow for timed automata. In: Models, Algorithms, Logics and Tools. LNCS, vol. 10460, pp. 3–21. Springer (2017)
40. Peisert, S., Margulies, J., Nicol, D.M., Khurana, H., Sawall, C.: Designed-in security for cyber-physical systems. IEEE Security & Privacy 12(5), 9–12 (2014)
41. Rafnsson, W., Jia, L., Bauer, L.: Timing-sensitive noninterference through composition. In: POST 2017. LNCS, vol. 10204, pp. 3–25. Springer (2017)
42. Roscoe, A.W., Huang, J.: Checking noninterference in timed CSP. Formal Aspects of Computing 25(1), 3–35 (2013)
43. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. IEEE Journal on Selected Areas in Communications 21(1), 5–19 (2003)
44. Schivo, S., Yildiz, B.M., Ruijters, E., Gerking, C., Kumar, R., Dziwok, S., Rensink, A., Stoelinga, M.: How to efficiently build a front-end tool for Uppaal: A model-driven approach. In: SETTA 2017. LNCS, vol. 10606, pp. 319–336. Springer (2017)
45. Son, J., Alves-Foss, J.: A formal framework for real-time information flow analysis. Computers & Security 28(6), 421–432 (2009)