



HAL
open science

Configuring Timing Parameters to Ensure Execution-Time Opacity in Timed Automata

Étienne André, Engel Lefauchaux, Didier Lime, Dylan Marinho, Jun Sun

► **To cite this version:**

Étienne André, Engel Lefauchaux, Didier Lime, Dylan Marinho, Jun Sun. Configuring Timing Parameters to Ensure Execution-Time Opacity in Timed Automata. *Electronic Proceedings in Theoretical Computer Science*, 2023, 392, pp.1 - 26. 10.4204/eptcs.392.1 . hal-04312156

HAL Id: hal-04312156

<https://hal.science/hal-04312156>

Submitted on 28 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Configuring Timing Parameters to Ensure Execution-Time Opacity in Timed Automata*

Étienne André

Université Sorbonne Paris Nord, LIPN, CNRS UMR 7030
F-93430 Villetaneuse, France

Engel Lefauchaux

Université de Lorraine, CNRS, Inria, LORIA
F-54000 Nancy, France

Didier Lime

Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004
F-44000 Nantes, France

Dylan Marinho

Université de Lorraine, CNRS, Inria, LORIA
F-54000 Nancy, France

Jun Sun

School of Computing and Information Systems
Singapore Management University

Timing information leakage occurs whenever an attacker successfully deduces confidential internal information by observing some timed information such as events with timestamps. Timed automata are an extension of finite-state automata with a set of clocks evolving linearly and that can be tested or reset, making this formalism able to reason on systems involving concurrency and timing constraints. In this paper, we summarize a recent line of works using timed automata as the input formalism, in which we assume that the attacker has access (only) to the system execution time. First, we address the following execution-time opacity problem: given a timed system modeled by a timed automaton, given a secret location and a final location, synthesize the execution times from the initial location to the final location for which one cannot deduce whether the secret location was visited. This means that for any such execution time, the system is opaque: either the final location is not reachable, or it is reachable with that execution time for both a run visiting and a run not visiting the secret location. We also address the full execution-time opacity problem, asking whether the system is opaque for all execution times; we also study a weak counterpart. Second, we add timing parameters, which are a way to configure a system: we identify a subclass of parametric timed automata with some decidability results. In addition, we devise a semi-algorithm for synthesizing timing parameter valuations guaranteeing that the resulting system is opaque. Third, we report on problems when the secret has itself an expiration date, thus defining expiring execution-time opacity problems. We finally show that our method can also apply to program analysis with configurable internal timings.

1 Introduction

Complex timed systems often combine hard real-time constraints with concurrency. Information leakage, notably through side channels (see, e.g., [23, 31]), can have dramatic consequences on the security of such systems. Among harmful information leaks, the *timing information leakage* (see, e.g., [22, 25, 38, 33, 35]) is the ability for an attacker to deduce internal information depending on observable timing information. In this paper, we focus on timing leakage through the total execution time, i.e., when a system works as an almost black-box and the ability of the attacker is limited to know the model and observe the total execution time. We consider here the formalism of timed automata (TAs) [1], which is a popular extension of finite-state automata with clocks measuring time, i.e., variables evolving linearly

*This work is partially supported by the ANR-NRF French-Singaporean research program ProMiS (ANR-19-CE25-0015 / 2019 ANR NRF 0092) and by ANR BisoUS (ANR-22-CE48-0012).

at the same rate. Such clocks can be tested against integer constants in locations (“invariants”) or along transitions (“guards”), and can be reset to 0 when taking transitions.

Context and related works Franck Cassez proposed in [19] a first definition of *timed* opacity for TAs: the system is opaque if an attacker can never deduce whether some sequence of actions (possibly with timestamps) was performed, by only observing a given set of observable actions together with their timestamp. It is then proved in [19] that it is undecidable whether a TA is opaque, even for the restricted class of event-recording automata [2] (a subclass of TAs). This notably relates to the undecidability of timed language inclusion for TAs [1]. Security problems for TAs are surveyed in [13].

The aforementioned negative result leaves hope only if the definition or the setting is changed, which was done in three main lines of works. The different studied options were to reduce the expressiveness of the formalism [36, 37], to constrain the system to evolve in a time-bounded setting [4] or to consider a weaker attacker, who has access only to the *execution time* [9, 8], rather than to all observable actions with their timestamps. We present here a summary of our recent works in this latter setting [9, 8].

Contributions In the setting of TAs, we denote by *execution time* the time from the system start to the time a given (final) location is entered. Therefore, given a secret location, a TA is execution-time opaque (ET-opaque) for an execution time d if there exist at least two runs of duration d from the initial location to a final location: one visiting the secret location, and another one *not* visiting the secret location. In other words, if an attacker measures such an execution time from the initial location to the target location ℓ_f , then this attacker is not able to deduce whether the system visited ℓ_{priv} . Deciding whether at least one such d exists can be seen as an *existential* version of ET-opacity (called \exists -ET-opacity).

Then, a TA is *fully ET-opaque* if it is ET-opaque for *all execution times*: that is, for each possible execution time d , either the final location is unreachable, or the final location is reachable for at least two runs, one visiting the secret location, and another one not visiting it. We define a *weak* version of ET-opacity by only requiring that runs visiting the secret location on the way to the final location have a counterpart of the same duration not visiting the secret location on the way to the final location, but not necessarily the opposite: the TA is *weakly ET-opaque* if for each run visiting the secret location, there exists a run not visiting it with the same duration; the dual does not necessarily hold.

We also consider an *expiring version* of ET-opacity, where the secret is subject to an expiration date Δ . That is, we consider that an attack is successful only when the attacker can decide that the secret location was visited less than Δ time units before the system completion. Conversely, if the attacker exhibits an execution time d for which it is certain that the secret location was visited, but this location was visited strictly more than Δ time units prior to the system completion, then this attack is useless, and can be seen as a failed attack. The system is therefore *fully expiring ET-opaque* if the set of execution times for which the private location was visited within Δ time units prior to system completion (referred as “secret times”) is exactly equal to the set of execution times for which the private location was either not visited or visited more than Δ time units prior to system completion (referred as “non-secret times”). Moreover, it is *weakly expiring ET-opaque* when the inclusion of the secret times into the non-secret ones is verified—and not necessarily the dual.

Finally, we study the aforementioned problems for a *parametric* extension of TAs, i.e., parametric timed automata (PTAs) [3], where integer constants compared to clocks can be made (rational-valued) timing parameters, i.e., unknown constants. Interesting problems include emptiness problems, i.e., the emptiness of the parameter valuations set such that (expiring) ET-opacity holds, and synthesis, i.e., the synthesis of all parameter valuations such that (expiring) ET-opacity holds.

Table 1: Summary of the results for ET-opacity [9]

		\exists -ET-opaque	weakly ET-opaque	fully ET-opaque
Decision	TA	√(Proposition 2)	√(Proposition 4)	√(Proposition 3)
p-emptiness	L/U-PTA	√(Theorem 2)	×(Theorem 6)	×(Theorem 4)
	PTA	×(Theorem 1)	×(Theorem 5)	×(Theorem 3)
p-synthesis	L/U-PTA	×(Proposition 5)	×(Corollary 5)	×(Corollary 3)
	PTA	×(Corollary 1)	×(Corollary 4)	×(Corollary 2)

Table 2: Summary of the results for exp-ET-opacity [8]

		\exists -exp-ET-opaque	weakly exp-ET-opaque	fully exp-ET-opaque
Decision	TA	√(Theorem 9)	√(Theorem 8)	√(Theorem 8)
Δ -emptiness	TA	?	√(Corollary 6)	√(Theorem 11)
Δ -computation		?	√(Theorem 10)	?
Δ -p-emptiness	L/U-PTA	?	×(Theorem 12)	×(Theorem 12)
	PTA	?	×(Theorem 13)	×(Theorem 13)
Δ -p-synthesis	L/U-PTA	?	×(Corollary 7)	×(Theorem 12)
	PTA	?	×(Corollary 8)	×(Corollary 8)

About this manuscript This manuscript mainly summarizes results from two recent works, providing unified notations and concept names for the sake of consistency:

1. defining and studying ET-opacity problems [9] in TAs (Section 3) and PTAs (Section 4); these notions from [9] are presented differently (including the problem names) in this paper for sake of consistency; and
2. defining and studying expiring execution-time opacity (exp-ET-opacity) problems [8] in both TAs and PTAs (Section 5).

In addition, we prove a few original results on weak ET-opacity (that were not addressed in [9] because we had not yet defined the concept of weak ET-opacity when writing [9]) and on exp-ET-opacity. These original results are Propositions 2 and 4 and Theorems 5, 6 and 9.

In Tables 1 and 2, we summarize the decidability results recalled in this paper for ET-opacity and exp-ET-opacity. We denote a problem with a green check if it is decidable, with a red cross if it is undecidable, and with a yellow question mark if it is open (or not considered in the aforementioned papers [9, 8]). We emphasize using a bold font the original results of this paper. The p-emptiness (resp. p-synthesis) problem asks for the synthesis (resp. for the non-existence) of a parameter valuation for which ET-opacity is enforced. The Δ -p-synthesis (resp. emptiness) problem asks for the synthesis (resp. for the non-existence) of a parameter valuation and an expiring bound Δ for which the exp-ET-opacity is enforced. L/U-PTA denote the lower-bound/upper-bound parametric timed automata [27] subclass of PTAs. These notions will be formally defined in the paper.

Outline Section 2 recalls the necessary preliminaries, notably (parametric) timed automata. Section 3 defines and reviews execution-time opacity problems in timed automata. Section 4 defines and reviews execution-time opacity problems in timed automata. Section 5 defines and reviews *expiring* execution-time opacity problems in (parametric) timed automata. Section 6 briefly reports on our existing imple-

mentation of some of the problems using the parametric timed model checker IMITATOR [6]. Section 7 concludes the paper and reports on perspectives.

2 Preliminaries

We denote by $\mathbb{N}, \mathbb{Z}, \mathbb{Q}_{\geq 0}, \mathbb{R}_{\geq 0}$ the sets of non-negative integers, integers, non-negative rationals and non-negative reals, respectively.

2.1 Clocks, parameters and constraints

Clocks are real-valued variables that all evolve over time at the same rate. Throughout this paper, we assume a set $\mathbb{X} = \{x_1, \dots, x_H\}$ of *clocks*. A *clock valuation* is a function $\mu : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, assigning a non-negative value to each clock. We write $\vec{0}$ for the clock valuation assigning 0 to all clocks. Given a constant $d \in \mathbb{R}_{\geq 0}$, $\mu + d$ denotes the valuation s.t. $(\mu + d)(x) = \mu(x) + d$, for all $x \in \mathbb{X}$.

A (*timing*) *parameter* is an unknown rational-valued constant of a model. Throughout this paper, we assume a set $\mathbb{P} = \{p_1, \dots, p_M\}$ of *parameters*. A *parameter valuation* v is a function $v : \mathbb{P} \rightarrow \mathbb{Q}_{\geq 0}$.

As often, we choose *real-valued* clocks and *rational-valued* parameters, because irrational constants render reachability undecidable in TAs [30] (see [5] for a survey on the impact of these domains in (P)TAs).

We assume $\bowtie \in \{<, \leq, =, \geq, >\}$. A constraint C is a conjunction of inequalities over $\mathbb{X} \cup \mathbb{P}$ of the form $x \bowtie \sum_{1 \leq i \leq M} \alpha_i p_i + d$, with $p_i \in \mathbb{P}$, and $\alpha_i, d \in \mathbb{Z}$. Given C , we write $\mu \models v(C)$ if the expression obtained by replacing each x with $\mu(x)$ and each p with $v(p)$ in C evaluates to true.

2.2 Timed automata

A TA is a finite-state automaton extended with a finite set of real-valued clocks. We also add to the standard definition of TAs a special private location, which will be used to define our subsequent opacity concepts.

Definition 1 (Timed automaton [1]). A TA \mathcal{A} is a tuple $\mathcal{A} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E)$, where:

1. Σ is a finite set of actions,
2. L is a finite set of locations,
3. $\ell_0 \in L$ is the initial location,
4. $\ell_{priv} \in L$ is a special private location,
5. $\ell_f \in L$ is the final location,
6. \mathbb{X} is a finite set of clocks,
7. I is the invariant, assigning to every $\ell \in L$ a constraint $I(\ell)$ over \mathbb{X} (called *invariant*),
8. E is a finite set of edges $e = (\ell, g, a, R, \ell')$ where $\ell, \ell' \in L$ are the source and target locations, $a \in \Sigma$, $R \subseteq \mathbb{X}$ is a set of clocks to be reset, and g is a constraint over \mathbb{X} (called *guard*).

Example 1. In Fig. 1, we give an example of a TA with three locations ℓ_0, ℓ_1 and ℓ_2 , three edges, three actions $\{a, b, c\}$, and one clock x . ℓ_0 is the initial location, ℓ_2 is the private location, while ℓ_1 is the final location. ℓ_0 has an invariant $x \leq 3$ and the edge from ℓ_0 to ℓ_2 has a guard $x \geq 1$.




Figure 1: A TA example

Concrete semantics of timed automata We recall the concrete semantics of a TA using a timed transition system (TTS) [26].

Definition 2 (Semantics of a TA). Given a TA $\mathcal{A} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E)$, the semantics of \mathcal{A} is given by the TTS $\mathfrak{T}_{\mathcal{A}} = (\mathfrak{S}, \mathfrak{s}_0, \Sigma \cup \mathbb{R}_{\geq 0}, \rightarrow)$, with

1. $\mathfrak{S} = \{(\ell, \mu) \in L \times \mathbb{R}_{\geq 0}^H \mid \mu \models I(\ell)v\}$,
2. $\mathfrak{s}_0 = (\ell_0, \vec{0})$,
3. \rightarrow consists of the discrete and (continuous) delay transition relations:
 - (a) discrete transitions: $(\ell, \mu) \xrightarrow{e} (\ell', \mu')$, if $(\ell, \mu), (\ell', \mu') \in \mathfrak{S}$, and there exists $e = (\ell, g, a, R, \ell') \in E$, such that $\mu' = [\mu]_R$, and $\mu \models v(g)$.
 - (b) delay transitions: $(\ell, \mu) \xrightarrow{d} (\ell, \mu + d)$, with $d \in \mathbb{R}_{\geq 0}$, if $\forall d' \in [0, d], (\ell, \mu + d') \in \mathfrak{S}$.

Moreover we write $(\ell, \mu) \xrightarrow{(d,e)} (\ell', \mu')$ for a combination of a delay and discrete transition if $\exists \mu'' : (\ell, \mu) \xrightarrow{d} (\ell, \mu'') \xrightarrow{e} (\ell', \mu')$.

Given a TA \mathcal{A} with concrete semantics $(\mathfrak{S}, \mathfrak{s}_0, \Sigma \cup \mathbb{R}_{\geq 0}, \rightarrow)$, we refer to the states of \mathfrak{S} as the *concrete states* of \mathcal{A} . A *run* of \mathcal{A} is an alternating sequence of concrete states of \mathcal{A} and pairs of edges and delays starting from the initial state \mathfrak{s}_0 of the form $(\ell_0, \mu_0), (d_0, e_0), (\ell_1, \mu_1), \dots$ with $i = 0, 1, \dots, e_i \in E, d_i \in \mathbb{R}_{\geq 0}$ and $(\ell_i, \mu_i) \xrightarrow{(d_i, e_i)} (\ell_{i+1}, \mu_{i+1})$.

Definition 3 (Duration of a run). Given a finite run $\rho : (\ell_0, \mu_0), (d_0, e_0), (\ell_1, \mu_1), \dots, (d_{i-1}, e_{i-1}), (\ell_n, \mu_n)$, the *duration* of ρ is $dur(\rho) = \sum_{0 \leq i \leq n-1} d_i$. We also say that ℓ_n is reachable in time $dur(\rho)$.

Example 2. Consider again the TA \mathcal{A} in Fig. 1. Consider the following run ρ of \mathcal{A} : $(\ell_0, x = 0), (1.4, a), (\ell_2, x = 1.4), (0.4, b), (\ell_1, x = 1.8)$. Note that we write “ $x = 1.4$ ” instead of “ μ such that $\mu(x) = 1.4$ ”. We have $dur(\rho) = 1.4 + 0.4 = 1.8$.

2.3 Parametric timed automata

A PTA is a TA extended with a finite set of timing parameters allowing to model unknown constants.

Definition 4 (Parametric timed automaton [3]). A PTA \mathcal{P} is a tuple $\mathcal{P} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, \mathbb{P}, I, E)$, where:

1. Σ is a finite set of actions;
2. L is a finite set of locations;
3. $\ell_0 \in L$ is the initial location;
4. $\ell_{priv} \in L$ is a special private location,
5. $\ell_f \in L$ is the final location;




Figure 2: A PTA example

6. \mathbb{X} is a finite set of clocks;
7. \mathbb{P} is a finite set of parameters;
8. I is the invariant, assigning to every $\ell \in L$ a constraint $I(\ell)$ over $\mathbb{X} \cup \mathbb{P}$ (called *invariant*);
9. E is a finite set of edges $e = (\ell, g, a, R, \ell')$ where $\ell, \ell' \in L$ are the source and target locations, $a \in \Sigma$, $R \subseteq \mathbb{X}$ is a set of clocks to be reset, and g is a constraint over $\mathbb{X} \cup \mathbb{P}$ (called *guard*).

Example 3. In Fig. 2, we give an example of a PTA with three locations ℓ_0 , ℓ_1 and ℓ_2 , three edges, three actions $\{a, b, c\}$, one clock x and two parameters $\{p_1, p_2\}$. ℓ_0 is the initial location, ℓ_2 is the private location, while ℓ_1 is the final location. ℓ_0 has an invariant $x \leq 3$ and the edge from ℓ_0 to ℓ_2 has a guard $x \geq p_1$.

Definition 5 (Valuation of a PTA). Given a parameter valuation v , we denote by $v(\mathcal{P})$ the non-parametric structure where all occurrences of a parameter p_i have been replaced by $v(p_i)$.

Remark 1. We have a direct correspondence between the valuation of a PTA and the definition of a TA given in Definition 1. TAs were originally defined with integer constants in [1] (as done in Definition 1), while our definition of PTAs allows *rational*-valued constants. By assuming a rescaling of the constants (i.e., by multiplying all constants in a TA by the least common multiple of their denominators), we obtain an equivalent (integer-valued) TA, as defined in Definition 1. So we assume in the following that $v(\mathcal{P})$ is a TA.

Example 4. Consider again the PTA in Fig. 2 and let v be such that $v(p_1) = 1$ and $v(p_2) = 2$. Then $v(\mathcal{P})$ is the TA depicted in Fig. 1.

Lower/upper parametric timed automaton While most decision problems are undecidable for the general class of PTAs (see [5] for a survey), lower/upper parametric timed automata (L/U-PTAs) [27] is the most well-known subclass of PTAs with some decidability results: for example, reachability-emptiness (“the emptiness of the valuations set for which a given location is reachable”), which is undecidable for PTAs [3], becomes decidable for L/U-PTAs [27]. Various other results were studied for this subclass (e.g., [17, 28, 12]).

Definition 6 (Lower/upper parametric timed automaton [27]). An L/U-PTA is a PTA where the set of parameters is partitioned into lower-bound parameters and upper-bound parameters, where each upper-bound (resp. lower-bound) parameter p_i must be such that, for every guard or invariant constraint $x \bowtie \sum_{1 \leq i \leq M} \alpha_i p_i + d$, we have:

- $\bowtie \in \{\leq, <\}$ implies $\alpha_i \geq 0$ (resp. $\alpha_i \leq 0$), and
- $\bowtie \in \{\geq, >\}$ implies $\alpha_i \leq 0$ (resp. $\alpha_i \geq 0$).

Example 5. The PTA in Fig. 2 is an L/U-PTA with $\{p_1\}$ as lower-bound parameter set, and $\{p_2\}$ as upper-bound parameter set.

3 Execution-time opacity problems in timed automata

Throughout this paper, the attacker model is as follows: the attacker knows the TA modeling the system, and can only observe the execution time between the start of the system and the time it reaches the final location. The attacker cannot observe actions, nor the values of the clocks, nor whether some locations are visited. Its goal will be to deduce from its observations whether the private location was visited.

3.1 Defining the execution times

Let us first introduce two key concepts necessary to define our notion of execution-time opacity.

Given a TA \mathcal{A} and a run ρ , we say that ℓ_{priv} is *visited on the way to ℓ_f in ρ* if ρ is of the form

$$(\ell_0, \mu_0), (d_0, e_0), (\ell_1, \mu_1), \dots, (\ell_m, \mu_m), (d_m, e_m), \dots, (\ell_n, \mu_n)$$

for some $m, n \in \mathbb{N}$ such that $\ell_m = \ell_{priv}$, $\ell_n = \ell_f$ and $\forall 0 \leq i \leq n-1, \ell_i \neq \ell_f$. We denote by $Visit^{priv}(\mathcal{A})$ the set of those runs, and refer to them as *private runs*. We denote by $DVisit^{priv}(\mathcal{A})$ the set of all the durations of these runs.

Conversely, we say that ℓ_{priv} is *avoided on the way to ℓ_f in ρ* if ρ is of the form

$$(\ell_0, \mu_0), (d_0, e_0), (\ell_1, \mu_1), \dots, (\ell_n, \mu_n)$$

with $\ell_n = \ell_f$ and $\forall 0 \leq i < n, \ell_i \notin \{\ell_{priv}, \ell_f\}$. We denote the set of those runs by $\overline{Visit}^{priv}(\mathcal{A})$, referring to them as *public runs*, and by $D\overline{Visit}^{priv}(\mathcal{A})$ the set of all the durations of these public runs.

Therefore, $DVisit^{priv}(\mathcal{A})$ (resp. $D\overline{Visit}^{priv}(\mathcal{A})$) is the set of all the durations of the runs for which ℓ_{priv} is visited (resp. avoided) on the way to ℓ_f .

These concepts can be seen as the set of execution times from the initial location ℓ_0 to the final location ℓ_f while visiting (resp. not visiting) a private location ℓ_{priv} . Observe that, from the definition of the duration of a run (Definition 3), this “execution time” does not include the time spent in ℓ_f .

Example 6. Consider again the TA in Fig. 1. We have $DVisit^{priv}(\mathcal{A}) = [1, 2]$ and $D\overline{Visit}^{priv}(\mathcal{A}) = [0, 3]$.

3.2 Defining execution-time opacity

We now introduce formally the concept of “ET-opacity for a set of durations (or execution times) D ”: a system is *ET-opaque for execution times D* whenever, for any duration in D , it is not possible to deduce whether the system visited ℓ_{priv} or not. In other words, if an attacker measures an execution time within D from the initial location to the target location ℓ_f , then this attacker is not able to deduce whether the system visited ℓ_{priv} .

Definition 7 (Execution-time opacity (ET-opacity) for D). Given a TA \mathcal{A} and a set of execution times D , we say that \mathcal{A} is *execution-time opaque (ET-opaque) for execution times D* if $D \subseteq (DVisit^{priv}(\mathcal{A}) \cap D\overline{Visit}^{priv}(\mathcal{A}))$.

In the following, we will be interested in the existence of such an execution time. We say that a TA is \exists -ET-opaque if it is ET-opaque for a non-empty set of execution times.

Definition 8 (\exists -ET-opacity). A TA \mathcal{A} is \exists -ET-opaque if $(DVisit^{priv}(\mathcal{A}) \cap D\overline{Visit}^{priv}(\mathcal{A})) \neq \emptyset$.

If one does not have the ability to tune the system (i.e., change internal delays, or add some `Thread.sleep()` statements in a program), one may be first interested in knowing whether the system is ET-opaque for all execution times. In other words, if a system is *fully ET-opaque*, for any possible measured execution time, an attacker is not able to deduce whether ℓ_{priv} was visited or not.

Definition 9 (full ET-opacity). A TA \mathcal{A} is *fully ET-opaque* if $D\text{Visit}^{\text{priv}}(\mathcal{A}) = \overline{D\text{Visit}^{\text{priv}}}(\mathcal{A})$.

That is, a system is fully ET-opaque if, for any execution time d , a run of duration d reaches ℓ_f after visiting ℓ_{priv} iff another run of duration d reaches ℓ_f without visiting ℓ_{priv} .

Remark 2. This definition is symmetric: a system is not fully ET-opaque iff an attacker can deduce ℓ_{priv} or $\neg\ell_{\text{priv}}$. For instance, if there is no run to ℓ_f visiting ℓ_{priv} , but still a run to ℓ_f (not visiting ℓ_{priv}), a system is not fully ET-opaque w.r.t. Definition 9.

We finally define weak ET-opacity, not considered in [9], but defined in the specific context of expiring opacity [8]. We therefore reintroduce this definition in the “normal” opacity setting considered in this section, in the following:

Definition 10 (weak ET-opacity). A TA \mathcal{A} is *weakly ET-opaque* if $D\text{Visit}^{\text{priv}}(\mathcal{A}) \subseteq \overline{D\text{Visit}^{\text{priv}}}(\mathcal{A})$.

That is, a TA is weakly ET-opaque whenever, for any run reaching the final location after visiting the private location, there exists another run of the same duration reaching the final location but not visiting the private location; but the converse does not necessarily hold.

Remark 3. Our notion of weak ET-opacity may still leak some information: on the one hand, if a run indeed visits the private location, there exists an equivalent run not visiting it, and therefore the system is ET-opaque; *but* on the other hand, there may exist execution times for which the attacker can deduce that the private location was *not* visited. This remains acceptable in some cases, and this motivates us to define a weak version of ET-opacity. Also note that the “initial-state opacity” for real-time automata considered in [36] can also be seen as *weak* in the sense that their language inclusion is also unidirectional.

Example 7. Consider again the PTA \mathcal{P} in Fig. 2 and let v such that $v(p_1) = 1$ while $v(p_2) = 2$ (i.e., the TA in Fig. 1). Recall that $D\text{Visit}^{\text{priv}}(v(\mathcal{P})) = [1, 2]$ and $\overline{D\text{Visit}^{\text{priv}}}(v(\mathcal{P})) = [0, 3]$. Hence, it holds that $D\text{Visit}^{\text{priv}}(v(\mathcal{P})) \subseteq \overline{D\text{Visit}^{\text{priv}}}(v(\mathcal{P}))$ and therefore $v(\mathcal{P})$ is weakly ET-opaque. However, $D\text{Visit}^{\text{priv}}(v(\mathcal{P})) \neq \overline{D\text{Visit}^{\text{priv}}}(v(\mathcal{P}))$ and therefore $v(\mathcal{P})$ is not fully ET-opaque.

Now consider again the PTA \mathcal{P} in Fig. 2 and let v' such that $v'(p_1) = 0$ while $v'(p_2) = 3$. This time, $D\text{Visit}^{\text{priv}}(v'(\mathcal{P})) = \overline{D\text{Visit}^{\text{priv}}}(v'(\mathcal{P})) = [0, 3]$ and therefore $v'(\mathcal{P})$ is fully ET-opaque.

3.3 Decision and computation problems

3.3.1 Computation problem for ET-opacity

We can now define the ET-opacity t-computation problem, which consists in computing the possible execution times ensuring ET-opacity.

ET-opacity t-computation problem:


INPUT: A TA \mathcal{A}

PROBLEM: Compute the execution times D such that \mathcal{A} is ET-opaque for D .

Let us illustrate that this computation problem is certainly not easy. For the TA \mathcal{A} in Fig. 3, the execution times D for which \mathcal{A} is ET-opaque is exactly \mathbb{N} ; that is, only integer times ensure ET-opacity (as the system can only leave ℓ_{priv} and hence enter ℓ_f at an integer time), while non-integer times violate ET-opacity.

3.3.2 Decision problems

We define the three following decision problems:

Figure 3: TA for which the set of execution times ensuring ET-opacity is \mathbb{N} **\exists -ET-opacity decision problem:**INPUT: A TA \mathcal{A} PROBLEM: Is \mathcal{A} \exists -ET-opaque?**Full ET-opacity decision problem:**INPUT: A TA \mathcal{A} PROBLEM: Is \mathcal{A} fully ET-opaque?**Weak ET-opacity decision problem:**INPUT: A TA \mathcal{A} PROBLEM: Is \mathcal{A} weakly ET-opaque?

3.4 Answering the ET-opacity t-computation problem

Proposition 1 (Solvability of the ET-opacity t-computation problem [9, Proposition 5.2]). *The ET-opacity t-computation problem is solvable for TAs.*

This positive result can be put in perspective with the negative result of [19] that proves that it is undecidable whether a TA (and even the more restricted subclass of event-recording automata (ERAs) [2]) is opaque, in a sense that the attacker can deduce some actions, by looking at observable actions together with their timing. The difference in our setting is that only the global time is observable, which can be seen as a single action, occurring once only at the end of the computation. In other words, our attacker is less powerful than the attacker in [19].

3.5 Checking for \exists -ET-opacity

The following result was not strictly speaking proved in [9], and we provide here an original proof for it.

Proposition 2 (Decidability of the \exists -ET-opacity decision problem). *The \exists -ET-opacity decision problem is decidable in 5EXPTIME for TAs.*

Proof. Let \mathcal{A} be a TA. Suppose we add a Boolean variable $priv$ to \mathcal{A} which is initially false and set to true on every edge going into the location l_{priv} . This Boolean variable (not strictly part of the TA syntax) can also be simulated by adding a copy of \mathcal{A} instead, and jumping to that copy on edges going into location l_{priv} .

Then the \exists -ET-opacity decision problem amounts to checking the following parametric TCTL formula [18], with p a parameter:

$$\exists p (\exists \diamond_{=p} (l_f \wedge priv) \wedge \exists \diamond_{=p} (l_f \wedge \neg priv))$$

From [18], this can be checked in 5EXPTIME, since the size of the TA it is checked on is at most twice that of \mathcal{A} , and the size of the formula is constant w.r.t. the size of \mathcal{A} . \square

3.6 Checking for full ET-opacity

The following result matches [9, Proposition 5.3] but we provide an original proof, also fixing a complexity issue in [9, Proposition 5.3].

Proposition 3 (Decidability of the full ET-opacity decision problem). *The full ET-opacity decision problem is decidable in 5EXPTIME for TAs.*

Proof. As before, we can write a parametric TCTL formula for this problem, with p a parameter:

$$\forall p(\exists \diamond_{=p}(\ell_f \wedge \text{priv}) \Leftrightarrow \exists \diamond_{=p}(\ell_f \wedge \neg \text{priv}))$$

This formula can be checked in 5EXPTIME [18]. \square

3.7 Checking for weak ET-opacity

The *weak* notion of ET-opacity had not been defined in [9]. Nevertheless, the proof of Proposition 3 can be adapted in a very straightforward manner to prove its weak counterpart as follows:

Proposition 4 (Decidability of the weak ET-opacity decision problem). *The weak ET-opacity decision problem is decidable in 5EXPTIME for TAs.*

Proof. Let \mathcal{A} be a TA. As before, we can write a parametric TCTL formula for this problem, with p a parameter:

$$\forall p(\exists \diamond_{=p}(\ell_f \wedge \text{priv}) \Rightarrow \exists \diamond_{=p}(\ell_f \wedge \neg \text{priv}))$$

This formula can be checked in 5EXPTIME [18]. \square

4 Execution-time opacity problems in parametric timed automata

We now extend opacity problems to parametric timed automata. We first address the parametric problems related to \exists -ET-opacity in Section 4.1. The decision problems associated to full ET-opacity and weak ET-opacity will then be considered in Sections 4.2 and 4.3 respectively.

Following the usual concepts for parametric timed automata, we consider both *emptiness* and *synthesis* problems. An emptiness problem aims at *deciding* whether the set of parameter valuations for which a given property holds in the valuated TA is empty, while a synthesis problem aims at *synthesizing* the set of parameter valuations for which a given property holds in the valuated TA.

4.1 \exists -ET-opacity

4.1.1 Problems

Emptiness problem for \exists -ET-opacity Let us consider the following decision problem, i.e., the problem of checking the *emptiness* of the set of parameter valuations guaranteeing \exists -ET-opacity.

\exists -ET-opacity p-emptiness problem:INPUT: A PTA \mathcal{P} PROBLEM: Decide the emptiness of the set of parameter valuations v such that $v(\mathcal{P})$ is \exists -ET-opaque.

The negation of the \exists -ET-opacity p-emptiness problem consists in deciding whether there exists at least one parameter valuation for which $v(\mathcal{P})$ is \exists -ET-opaque.

Synthesis problem for \exists -ET-opacity The synthesis counterpart allows for a higher-level problem by also synthesizing the internal timings guaranteeing \exists -ET-opacity.

 \exists -ET-opacity p-synthesis problem:INPUT: A PTA \mathcal{P} PROBLEM: Synthesize the set V of parameter valuations such that $v(\mathcal{P})$ is \exists -ET-opaque, for all $v \in V$.**4.1.2 Undecidability in general**

With the rule of thumb that all non-trivial decision problems are undecidable for general PTAs [5], the following result is not surprising, and follows from the undecidability of reachability-emptiness for PTAs [3].

Theorem 1 (Undecidability of the \exists -ET-opacity p-emptiness problem [9, Theorem 6.1]). *The \exists -ET-opacity p-emptiness problem is undecidable for general PTAs.*

Since the emptiness problem is undecidable, the synthesis problem is immediately unsolvable as well.

Corollary 1. *The \exists -ET-opacity p-synthesis problem is unsolvable for general PTAs.*

Nevertheless, in [9] we proposed a procedure solving this problem. While this procedure is not guaranteed to terminate, its result is correct when termination can be achieved. See [9, Section 8] for details.

4.1.3 The subclass of L/U-PTAs

Decidability We now show that the \exists -ET-opacity p-emptiness problem is decidable for L/U-PTAs. Despite early positive results for L/U-PTAs [27, 17], more recent results (notably [28, 11, 12]) mostly proved undecidable properties of L/U-PTAs, and therefore this positive result is welcome.

Theorem 2 (Decidability of the \exists -ET-opacity p-emptiness problem [9, Theorem 6.2]). *The \exists -ET-opacity p-emptiness problem is decidable for L/U-PTAs.*

Intractability of synthesis for lower/upper parametric timed automata Even though the \exists -ET-opacity p-emptiness problem is decidable for L/U-PTAs (Theorem 2), the *synthesis* of the parameter valuations remains intractable in general, as shown in the following Proposition 5. By intractable we mean more precisely that the solution, if it can be computed, cannot (in general, i.e., for some sufficiently complex solutions) be represented using any formalism for which the emptiness of the intersection with equality constraints is decidable. That is, a formalism in which it is decidable to decide “the emptiness of the valuation set of the computed solution intersected with an equality test between variables” cannot be used to represent the solution. For example, let us question whether we could represent the solution of

the \exists -ET-opacity p-synthesis problem for L/U-PTAs using the formalism of a *finite union of polyhedra*: testing whether a finite union of polyhedra intersected with “equality constraints” (typically $p_1 = p_2$) is empty or not *is* decidable. The Parma polyhedra library [14] can typically compute the answer to this question. Therefore, from the following Proposition 5, finite unions of polyhedra cannot be used to represent the solution of the \exists -ET-opacity p-synthesis problem for L/U-PTAs. As finite unions of polyhedra are a very common formalism (not to say the *de facto* standard) to represent the solutions of various timing parameters synthesis problems, the synthesis is then considered to be infeasible in practice, or *intractable* (following the vocabulary used in [28, Theorem 2]).

Proposition 5 (Intractability of the \exists -ET-opacity p-synthesis problem [9, Proposition 6.4]). *In case a solution to the \exists -ET-opacity p-synthesis problem for L/U-PTAs can be computed, this solution may be not representable using any formalism for which the emptiness of the intersection with equality constraints is decidable.*

4.2 Parametric full ET-opacity

We address here the following decision problem, which asks about the emptiness of the parameter valuation set guaranteeing full ET-opacity. We also define the full ET-opacity p-synthesis problem, this time *synthesizing* the timing parameters guaranteeing full ET-opacity.

4.2.1 Problem definitions

Full ET-opacity p-emptiness problem:

INPUT: A PTA \mathcal{P}

PROBLEM: Decide the emptiness of the set of parameter valuations v such that $v(\mathcal{P})$ is fully ET-opaque.

Equivalently, we are interested in deciding whether there exists at least one parameter valuation for which $v(\mathcal{P})$ is fully ET-opaque.

We also define the *full ET-opacity p-synthesis problem*, aiming at synthesizing (ideally the entire set of) parameter valuations v for which $v(\mathcal{P})$ is fully ET-opaque.

Full ET-opacity p-synthesis problem:

INPUT: A PTA \mathcal{P}

PROBLEM: Synthesize the set V of parameter valuations such that $v(\mathcal{P})$ is fully ET-opaque, for all $v \in V$.

4.2.2 Undecidability for general PTAs

Considering that Theorem 1 shows the undecidability of the \exists -ET-opacity p-emptiness problem, the undecidability of the full ET-opacity p-emptiness problem is not surprising, but does not follow immediately.

Theorem 3 (Undecidability of the full ET-opacity p-emptiness problem [9, Theorem 7.2]). *The full ET-opacity p-emptiness problem is undecidable for general PTAs.*

The proof relies on a reduction from the problem of reachability-emptiness in constant time, a result proved itself undecidable in the same paper [9, Lemma 7.1].

Since the emptiness problem is undecidable, the synthesis problem is immediately unsolvable as well.

Corollary 2. *The full ET-opacity p-synthesis problem is unsolvable for PTAs.*




Figure 4: No monotonicity for full ET-opacity in L/U-PTAs

4.2.3 Undecidability for lower/upper parametric timed automata

Let us now study the full ET-opacity p-emptiness problem for L/U-PTAs. While it is well-known that L/U-PTAs enjoy a monotonicity for reachability properties (“enlarging an upper-bound parameter or decreasing a lower-bound parameter preserves reachability”) [27], we can show in the following example that this is not the case for full ET-opacity.

Example 8. Consider the PTA in Fig. 4. First assume v such that $v(p) = 0.5$. Then, $v(\mathcal{P})$ is not fully ET-opaque: indeed, l_f can be reached in 1 time unit by visiting l_{priv} , but not without visiting l_{priv} .

Second, assume v' such that $v'(p) = 1$. Then, $v'(\mathcal{P})$ is fully ET-opaque: indeed, l_f can be reached for any duration in $[0, 1]$ by runs both visiting and not visiting l_{priv} .

Finally, let us enlarge p further, and assume v'' such that $v''(p) = 2$. Then, $v''(\mathcal{P})$ becomes again not fully ET-opaque: indeed, l_f can be reached in 2 time units without visiting l_{priv} , but cannot be reached in 2 time units by visiting l_{priv} .

As a side note, remark that this PTA is actually an upper-bound parametric timed automaton (U-PTA) [17], that is, monotonicity for this problem does not even hold for U-PTAs.

In fact, we show that, while the \exists -ET-opacity p-emptiness problem is decidable for L/U-PTAs (Theorem 2), the full ET-opacity p-emptiness problem becomes undecidable for this same class. This confirms (after previous works in [17, 28, 11, 12]) that L/U-PTAs stand at the frontier between decidability and undecidability.

Theorem 4 (Undecidability of the full ET-opacity p-emptiness problem for L/U-PTAs [9, Theorem 7.4]). *The full ET-opacity p-emptiness problem is undecidable for L/U-PTAs.*

Since the emptiness problem is undecidable, the synthesis problem is immediately unsolvable as well.

Corollary 3. *The full ET-opacity p-synthesis problem is unsolvable for L/U-PTAs.*

Remark 4. Since L/U-PTAs are a subclass of PTAs (put it differently: “any L/U-PTA is a PTA”), the negative results proved for L/U-PTAs (Theorem 4 and Corollary 3) immediately imply those previously shown for general PTAs (Theorem 3 and Corollary 2). However, in [9], a smaller number of clocks and parameters is needed to prove the aforementioned negative results for general PTAs, which justifies the two versions of the proofs in [9].

4.3 Parametric weak ET-opacity

4.3.1 Problem definitions

Weak ET-opacity p-emptiness problem:

INPUT: A PTA \mathcal{P}

PROBLEM: Decide the emptiness of the set of parameter valuations v such that $v(\mathcal{P})$ is weakly ET-opaque.




Figure 5: Reduction from reachability-emptiness for the proof of Theorem 5

Weak ET-opacity p-synthesis problem:INPUT: A PTA \mathcal{P} PROBLEM: Synthesize the parameter valuations v such that $v(\mathcal{P})$ is weakly ET-opaque.**4.3.2 Undecidability for general PTAs**

We provide below an original result in the context of *weak* opacity, but partially inspired by the construction used in the proof of Theorem 3.

Theorem 5 (Undecidability of the weak ET-opacity p-emptiness problem). *The weak ET-opacity p-emptiness problem is undecidable for general PTAs.*

Proof. We reduce from the reachability-emptiness problem in bounded time, which is undecidable from [10, Theorem 3.12]. (This is different from the proof of [9, Theorem 7.2], which reduces from the reachability-emptiness problem in *constant* time, which is undecidable according to [9, Lemma 7.1].)

Consider an arbitrary PTA \mathcal{P} , with initial location l_0 and a final location l_f . We add the following locations and transitions in \mathcal{P} to obtain a PTA \mathcal{P}' , as in Fig. 5: (i) a new initial location l'_0 , with outgoing transitions in 0-time (due to their guard $x = 0$, where x is a new clock not belonging to \mathcal{P} , and never reset in \mathcal{P}') to l_0 and to a new location l_{priv} , (ii) a new location l_{pub} with an incoming transition from l_f guarded by $x = 1$, and (iii) a new final location l'_f with incoming transitions from l_{pub} and l_{priv} both guarded by $x = 1$.

First, note that, due to the guarded transitions, l'_f is reachable for any parameter valuation via runs visiting l_{priv} , (only) for an execution time equal to 1. That is, for all v , $DVisit^{priv}(v(\mathcal{P}')) = \{1\}$.

We now show that there exists a valuation v such that $v(\mathcal{P}')$ is weakly ET-opaque (with l_{priv} as private location, and l'_f as final location) iff there exists a valuation v such that l_f is reachable in $v(\mathcal{P})$ for an execution time ≤ 1 .

\Leftarrow Assume there exists some valuation v such that l_f is reachable from l_0 in \mathcal{P} for an execution time ≤ 1 . Then, due to our construction, l_{pub} is visited on the way to l'_f in $v(\mathcal{P}')$ (only) for the execution time 1. Therefore, $D\overline{Visit}^{priv}(v(\mathcal{P}')) = \{1\} = DVisit^{priv}(v(\mathcal{P}'))$ and then $v(\mathcal{P}')$ is weakly ET-opaque (and also fully ET-opaque, which plays no role here).

\Rightarrow Conversely, if l_f is not reachable from l_0 in \mathcal{P} for any valuation for an execution time ≤ 1 , then no run reaches l'_f in time 1 without visiting l_{priv} , for any valuation of \mathcal{P}' . Therefore, for any valuation v , $DVisit^{priv}(v(\mathcal{P}')) = \{1\} \not\subseteq D\overline{Visit}^{priv}(v(\mathcal{P}')) = \emptyset$. Therefore, there is no valuation v such that $v(\mathcal{P}')$ is weakly ET-opaque.




Figure 6: Undecidability of full ET-opacity p-emptiness problem for L/U-PTAs

Therefore, there exists a valuation v such that $v(\mathcal{P}')$ is weakly ET-opaque iff there exists a valuation v such that l_f is reachable in $v(\mathcal{P})$ for an execution time ≤ 1 —which is undecidable from [10, Theorem 3.12]. This concludes the proof. \square

Since the emptiness problem is undecidable, the synthesis problem is immediately unsolvable as well.

Corollary 4. *The weak ET-opacity p-synthesis problem is unsolvable for general PTAs.*

4.3.3 Undecidability for lower/upper parametric timed automata

We provide below another original result in the context of *weak* opacity, this time for L/U-PTAs, largely inspired by the proof of Theorem 4, even though our construction needed to be changed.

Theorem 6 (Undecidability of the weak ET-opacity p-emptiness problem for L/U-PTAs). *The weak ET-opacity p-emptiness problem is undecidable for L/U-PTAs.*

Proof. Let us recall from [10, Theorem 3.12] that the reachability-emptiness problem is undecidable over bounded time for PTAs with (at least) 3 clocks and 2 parameters. Assume a PTA \mathcal{P} with 3 clocks and 2 parameters, say p_1 and p_2 , and a final location l_f . Take 1 as a time bound. From [10, Theorem 3.12], it is undecidable whether there exists a parameter valuation for which l_f is reachable in \mathcal{P} in time ≤ 1 .

The idea of our proof is that, as in [28, 9], we “split” each of the two parameters used in \mathcal{P} into a lower-bound parameter (p_1^l and p_2^l) and an upper-bound parameter (p_1^u and p_2^u). Each constraint of the form $x < p_i$ (resp. $x \leq p_i$) is replaced with $x < p_i^u$ (resp. $x \leq p_i^u$) while each constraint of the form $x > p_i$ (resp. $x \geq p_i$) is replaced with $x > p_i^l$ (resp. $x \geq p_i^l$); $x = p_i$ is replaced with $p_i^l \leq x \leq p_i^u$.

The idea is that the PTA \mathcal{P} is exactly equivalent to our construction with duplicated parameters only when $p_1^l = p_1^u$ and $p_2^l = p_2^u$. The crux of the rest of this proof is that we will “rule out” any parameter valuation not satisfying these equalities, so as to use directly the undecidability result of [10, Theorem 3.12].

Now, consider the extension of \mathcal{P} given in Fig. 6, and let \mathcal{P}' be this extension. We assume that x is an extra clock not used in \mathcal{P} . The syntax “ $\mathbb{X} \setminus \{x\} \leftarrow 0$ ” denotes that all clocks of the original PTA \mathcal{P} are reset—but not the new clock x . The guard on the transition from l_0' to l_4 stands for 2 different transitions guarded with $p_1^l < x \leq p_1^u$, and $p_2^l < x \leq p_2^u$, respectively.

Let us first make the following observations:

1. for any parameter valuation, one can take the transition from ℓ_0' to ℓ_{priv} at time 2 and then to ℓ_f' in 0-time (i.e., at time 2), i.e., ℓ_f' is always reachable in time 2 while visiting location ℓ_{priv} ; put differently, $\{2\} \subseteq DVisit^{priv}(v(\mathcal{P}'))$ for any parameter valuation v ;
2. the original automaton \mathcal{P} can only be entered whenever $p_1^l \leq p_1^u$ and $p_2^l \leq p_2^u$; going from ℓ_0' to ℓ_0 takes exactly 1 time unit (due to the $x = 1$ guard);
3. to reach ℓ_f' without visiting ℓ_{priv} , a run must go through \mathcal{P} and visit ℓ_f , and its duration is necessarily 2; put differently, $D\overline{Visit}^{priv}(v(\mathcal{P}')) \subseteq \{2\}$ for any parameter valuation v ;
4. from [10, Theorem 3.12], it is undecidable whether there exists a parameter valuation for which there exists a run reaching ℓ_f from ℓ_0 in time ≤ 1 , i.e., reaching ℓ_f from ℓ_0' in time ≤ 2 .

Let us consider the following cases depending on the valuations:

1. for valuations v such that $p_1^l > p_1^u$ or $p_2^l > p_2^u$, then thanks to the transitions from ℓ_0' to ℓ_0 , there is no way to enter the original PTA \mathcal{P} (and therefore to reach ℓ_f' without visiting ℓ_{priv}); hence, $D\overline{Visit}^{priv}(v(\mathcal{P}')) = \emptyset$, and therefore $\{2\} \subseteq DVisit^{priv}(v(\mathcal{P}')) \not\subseteq D\overline{Visit}^{priv}(v(\mathcal{P}'))$, i.e., \mathcal{P}' is not weakly ET-opaque for any of these valuations.
2. for valuations v such that $p_1^l < p_1^u$ or $p_2^l < p_2^u$, then the transition from ℓ_0' to ℓ_4 can be taken, and therefore there exist runs reaching ℓ_f' after a duration > 2 (for example of duration 3) and visiting ℓ_{priv} . Since no run can reach ℓ_f' without visiting ℓ_{priv} for a duration $\neq 2$, then $\{3\} \subseteq DVisit^{priv}(v(\mathcal{P}')) \not\subseteq D\overline{Visit}^{priv}(v(\mathcal{P}')) \subseteq \{2\}$ and again \mathcal{P}' is not weakly ET-opaque for any of these valuations.
3. for valuations such that $p_1^l = p_1^u$ and $p_2^l = p_2^u$, then the behavior of the modified \mathcal{P} (with duplicate parameters) is exactly the one of the original \mathcal{P} . Also, note that the transition from ℓ_0' to ℓ_4 cannot be taken. In contrast, the transition from ℓ_0' to ℓ_{priv} can still be taken, and therefore there exists a run of duration 2 visiting ℓ_{priv} and reaching ℓ_f' . Hence, $DVisit^{priv}(v(\mathcal{P}')) = \{2\}$ for any such valuation v .

- Now, assume there exists such a parameter valuation v for which there exists a run of $v(\mathcal{P})$ of duration ≤ 1 reaching ℓ_f . And, as a consequence, there exists a run of $v(\mathcal{P}')$ of duration 2 (including the 1 time unit to go from ℓ_0' to ℓ_0) reaching ℓ_f' without visiting ℓ_{priv} . Hence, $D\overline{Visit}^{priv}(v(\mathcal{P}')) = \{2\}$. Therefore $D\overline{Visit}^{priv}(v(\mathcal{P}')) = DVisit^{priv}(v(\mathcal{P}')) = \{2\}$.

As a consequence, the modified automaton \mathcal{P}' is weakly ET-opaque (and actually fully ET-opaque—which plays no role in this proof) for such a parameter valuation.

- Conversely, assume there exists no parameter valuation for which there exists a run of \mathcal{P} of duration ≤ 1 reaching ℓ_f . In that case, ℓ_f' can never be reached without visiting ℓ_{priv} : $D\overline{Visit}^{priv}(v(\mathcal{P}')) = \emptyset$, and therefore $\{2\} \subseteq DVisit^{priv}(v(\mathcal{P}')) \not\subseteq D\overline{Visit}^{priv}(v(\mathcal{P}'))$, i.e., $v(\mathcal{P}')$ is not fully ET-opaque for any such parameter valuation v .

As a consequence, there exists a parameter valuation v' for which $v'(\mathcal{P}')$ is weakly ET-opaque iff there exists a parameter valuation v for which there exists a run in $v(\mathcal{P})$ of duration ≤ 1 reaching ℓ_f —which is undecidable from [10, Theorem 3.12]. \square

Corollary 5. *The weak ET-opacity p -synthesis problem is unsolvable for L/U-PTAs.*

Table 3: Summary of the definitions for ET-opacity and expiring ET-opacity [9, 8]

	Secret runs	Non-secret runs
ET-opacity	Runs visiting the private location (= private runs)	Runs not visiting the private location (= public runs)
exp-ET-opacity	Runs visiting the private location $\leq \Delta$ time units before the system completion	(i) Runs not visiting the private location and (ii) Runs visiting the private location $> \Delta$ time units before the system completion

The system is (resp. expiring)	if
ET-opaque	$\{\text{secret runs}\} \cap \{\text{non-secret runs}\} \neq \emptyset$
weakly ET-opaque	$\{\text{secret runs}\} \subseteq \{\text{non-secret runs}\}$
full ET-opacity	$\{\text{secret runs}\} = \{\text{non-secret runs}\}$

5 Expiring execution-time opacity problems

In [4], the authors consider a time-bounded notion of the opacity of [19], where the attacker has to disclose the secret before an upper bound, using a partial observability. This can be seen as a secrecy with an *expiration date*. The rationale is that retrieving a secret “too late” is useless; this is understandable, e.g., when the secret depends of the status of the memory; if the cache was overwritten since, then knowing the secret is probably useless in most situations. In addition, the analysis in [4] is carried over a time-bounded horizon; this means there are two time bounds in [4]: one for the secret expiration date, and one for the bounded-time execution of the system.

In this section, we review a recent work of ours [8] in which we incorporate this secret expiration date into our notion of ET-opacity: we only consider the former notion of time bound from [4] (the secret expiration date), and lift the assumption regarding the latter (the bounded-time execution of the system). More precisely, we consider an *expiring version of ET-opacity*, where the secret is subject to an expiration date; this can be seen as a combination of both concepts from [9] and [4]. That is, we consider that an attack is successful only when the attacker can decide that the secret location was entered less than Δ time units before the system completion. Conversely, if the attacker exhibits an execution time d for which it is certain that the secret location was visited, but this location was entered strictly more than Δ time units prior to the system completion, then this attack is useless, and can be seen as a failed attack. The system is therefore *fully exp-ET-opaque* if the set of execution times for which the private location was entered within Δ time units prior to system completion is exactly equal to the set of execution times for which the private location was either not visited or entered $> \Delta$ time units prior to system completion.

In addition, when the former (secret) set of execution times is *included* into the latter (non-secret) set of times, we say that the system is *weakly exp-ET-opaque*; this encodes situations when the attacker might be able to deduce that no secret location was visited, but is not able to confirm that the secret location *was* indeed visited.

On the one hand, our attacker model is *less powerful* than [4], because our attacker has only access to the execution time (and to the input model); in that sense, our attacker capability is identical to [9]. On the other hand, we lift the time-bounded horizon analysis from [4], allowing to analyze systems without any assumption on their execution time; therefore, we only import from [4] the notion of *expiring secret*.

We summarize in Table 3 our different notions of ET-opacity and expiring ET-opacity; we will define

formally expiring ET-opacity in the following.

5.1 Exp-ET-opacity

Let us first introduce some notions dedicated to expiring ET-opacity (hereafter referred to as exp-ET-opacity). Let $\mathbb{R}_{\geq 0}^{\infty} = \mathbb{R}_{\geq 0} \cup \{+\infty\}$. Given a TA \mathcal{A} and a finite run ρ in $\mathfrak{T}_{\mathcal{A}}$, the *duration* between two states of $\rho : \mathfrak{s}_0, (d_0, e_0), \mathfrak{s}_1, \dots, \mathfrak{s}_k$ is $dur_{\rho}(\mathfrak{s}_i, \mathfrak{s}_j) = \sum_{i \leq m \leq j-1} d_m$. We also define the *duration* between two locations ℓ_1 and ℓ_2 as the duration $dur_{\rho}(\ell_1, \ell_2) = dur_{\rho}(\mathfrak{s}_i, \mathfrak{s}_j)$ with $\rho : \mathfrak{s}_0, (d_0, e_0), \mathfrak{s}_1, \dots, \mathfrak{s}_i, \dots, \mathfrak{s}_j, \dots, \mathfrak{s}_k$ where \mathfrak{s}_j the first occurrence of a state with location ℓ_2 and \mathfrak{s}_i is the last state of ρ with location ℓ_1 before \mathfrak{s}_j . We choose this definition to coincide with the definitions of opacity that we will define in the following Definition 11. Indeed, we want to make sure that revealing a secret (ℓ_1 in this definition) is not a failure if it is done after a given time. Thus, as soon as the system reaches its final state (ℓ_2), we will be interested in knowing how long the secret has been present, and thus the last time it was entered (\mathfrak{s}_i).

Given $\Delta \in \mathbb{R}_{\geq 0}^{\infty}$, we define $Visit_{\leq \Delta}^{priv}(\mathcal{A})$ (resp. $Visit_{> \Delta}^{priv}(\mathcal{A})$) as the set of runs $\rho \in Visit^{priv}(\mathcal{A})$ s.t. $dur_{\rho}(\ell_{priv}, \ell_f) \leq \Delta$ (resp. $dur_{\rho}(\ell_{priv}, \ell_f) > \Delta$). We refer to the runs of $Visit_{\leq \Delta}^{priv}(\mathcal{A})$ as *secret runs*; their durations are denoted by $DVisit_{\leq \Delta}^{priv}(\mathcal{A})$. Similarly, the durations of the runs of $Visit_{> \Delta}^{priv}(\mathcal{A})$ are denoted by $DVisit_{> \Delta}^{priv}(\mathcal{A})$.

We define below two notions of ET-opacity w.r.t. a time bound Δ . We will compare two sets:

1. the set of execution times for which the private location was entered at most Δ time units prior to system completion; and
2. the set of execution times for which either the private location was not visited at all, or it was last entered more than Δ time units prior to system completion (which, in our setting, is somehow similar to *not* visiting the private location, in the sense that entering it “too early” is considered of little interest).

If both sets match, the system is fully ($\leq \Delta$)-ET-opaque. If the former is included into the latter, then the system is weakly ($\leq \Delta$)-ET-opaque.

Definition 11 (Expiring Execution-time opacity). Given a TA \mathcal{A} and a bound (i.e., an expiration date for the secret) $\Delta \in \mathbb{R}_{\geq 0}^{\infty}$ we say that \mathcal{A} is *fully exp-ET-opaque* w.r.t. the expiration date Δ , denoted by *fully ($\leq \Delta$)-ET-opaque*, if

$$DVisit_{\leq \Delta}^{priv}(\mathcal{A}) = DVisit_{> \Delta}^{priv}(\mathcal{A}) \cup \overline{DVisit}^{priv}(\mathcal{A}).$$

Moreover, \mathcal{A} is *weakly exp-ET-opaque* w.r.t. the expiration date Δ , denoted by *weakly ($\leq \Delta$)-ET-opaque*, if

$$DVisit_{\leq \Delta}^{priv}(\mathcal{A}) \subseteq DVisit_{> \Delta}^{priv}(\mathcal{A}) \cup \overline{DVisit}^{priv}(\mathcal{A}).$$

Finally, \mathcal{A} is *\exists -ET-opaque* w.r.t. the expiration date Δ , denoted by *\exists -($\leq \Delta$)-ET-opaque*, if

$$DVisit_{\leq \Delta}^{priv}(\mathcal{A}) \cap (DVisit_{> \Delta}^{priv}(\mathcal{A}) \cup \overline{DVisit}^{priv}(\mathcal{A})) \neq \emptyset.$$

Example 9. Consider again the PTA in Fig. 2; let v be such that $v(p_1) = 1$ and $v(p_2) = 2.5$. Fix $\Delta = 1$.

We have:

- $\overline{DVisit}^{priv}(v(\mathcal{P})) = [0, 3]$
- $DVisit_{> \Delta}^{priv}(v(\mathcal{P})) = (2, 2.5]$
- $DVisit_{\leq \Delta}^{priv}(v(\mathcal{P})) = [1, 2.5]$

Therefore, we say that $v(\mathcal{P})$ is:

- \exists - (≤ 1) -ET-opaque, as $[1, 2.5] \cap ((2, 2.5] \cup [0, 3]) \neq \emptyset$
- weakly (≤ 1) -ET-opaque, as $[1, 2.5] \subseteq ((2, 2.5] \cup [0, 3])$
- not fully (≤ 1) -ET-opaque, as $[1, 2.5] \not\subseteq ((2, 2.5] \cup [0, 3])$

As noted in Remark 3, despite the weak (≤ 1) -ET-opacity of \mathcal{A} , the attacker can deduce some information about the visit of the private location for some execution times. For example, if a run has a duration of 3 time units, it cannot be a private run, and therefore the attacker can deduce that the private location was not visited at all.

5.2 Exp-ET-opacity problems in timed automata

5.2.1 Problem definitions

We define seven different problems in the context of (non-parametric) TAs:

\exists -exp-ET-opacity decision problem:

INPUT: A TA \mathcal{A} and a bound $\Delta \in \mathbb{R}_{\geq 0}^{\infty}$

PROBLEM: Decide whether \mathcal{A} is \exists - $(\leq \Delta)$ -ET-opaque.

Full (resp. weak) exp-ET-opacity decision problem:

INPUT: A TA \mathcal{A} and a bound $\Delta \in \mathbb{R}_{\geq 0}^{\infty}$

PROBLEM: Decide whether \mathcal{A} is fully (resp. weakly) $(\leq \Delta)$ -ET-opaque.

Full (resp. weak) exp-ET-opacity Δ -emptiness problem:

INPUT: A TA \mathcal{A}

PROBLEM: Decide the emptiness of the set of bounds Δ such that \mathcal{A} is fully (resp. weakly) $(\leq \Delta)$ -ET-opaque.

Full (resp. weak) exp-ET-opacity Δ -computation problem:

INPUT: A TA \mathcal{A}

PROBLEM: Compute the maximal set \mathcal{D} of bounds such that \mathcal{A} is fully (resp. weakly) $(\leq \Delta)$ -ET-opaque for all $\Delta \in \mathcal{D}$.

Example 10. Consider again the PTA in Fig. 2; let v be such that $v(p_1) = 1$ and $v(p_2) = 2.5$ (as in Example 9). Let us exemplify some of the problems defined above.

- Given $\Delta = 1$, the weak exp-ET-opacity decision problem asks whether $v(\mathcal{P})$ is weakly (≤ 1) -ET-opaque—the answer is “yes” from Example 9.
- The answer to the weak exp-ET-opacity Δ -emptiness problem is therefore “no” because the set of bounds Δ such that $v(\mathcal{P})$ is weakly $(\leq \Delta)$ -ET-opaque is not empty.
- Finally, the weak exp-ET-opacity Δ -computation problem asks to compute all the corresponding bounds: in this example, the solution is $\Delta \in \mathbb{R}_{\geq 0}^{\infty}$, i.e., the solution is the set all possible (non-negative) values for Δ .

Relations with the ET-opacity problems Note that, when considering $\Delta = +\infty$, $DVisit_{>\Delta}^{priv}(\mathcal{A}) = \emptyset$ and all the execution times of runs visiting ℓ_{priv} are in $DVisit_{\leq \Delta}^{priv}(\mathcal{A})$. Therefore, full $(\leq +\infty)$ -ET-opacity matches the full ET-opacity. We can therefore notice that answering the full exp-ET-opacity decision

problem for $\Delta = +\infty$ is decidable (Proposition 3). However, the emptiness and computation problems cannot be reduced to full ET-opacity problems from Section 4.1.3.

Conversely, it is possible to answer the full ET-opacity decision problem by checking the full exp-ET-opacity decision problem with $\Delta = +\infty$. Moreover, the ET-opacity t-computation problem reduces to the full exp-ET-opacity Δ -computation problem: if $+\infty \in \mathcal{D}$, we get the answer.

Recall that we summarize our different definitions of (expiring) ET-opacity in Table 3.

5.2.2 Results

In general, the link between the full and weak notions of the three aforementioned problems is not obvious. However, for a fixed value of Δ , we establish the following theorem.

Theorem 7 ([8, Theorem 1]). *The full exp-ET-opacity decision problem reduces to the weak exp-ET-opacity decision problem.*

We can now study the aforementioned problems.

Theorem 8 (Decidability of full (resp. weak) exp-ET-opacity decision problem [8, Theorems 2 and 5]). *The full (resp. weak) exp-ET-opacity decision problem is decidable in NEXPTIME.*

Remark 5. In Proposition 3, we established that the full ($\leq +\infty$)-ET-opacity decision problem is in 5EXPTIME. Theorem 8 thus extends our former results in three ways:

1. by including the parameter Δ ,
2. by reducing the complexity and
3. by considering as well the *weak* notion of ET-opacity (considered separately in Proposition 4).

We complete these results from [8] with the following result analog to Proposition 2.

Theorem 9 (Decidability of \exists -exp-ET-opacity decision problem). *The \exists -exp-ET-opacity decision problem is decidable in PSPACE.*

Proof. The full (resp. weak) exp-ET-opacity decision problem was solved in [8] by building two non-deterministic finite automata whose languages represented the secret and the non-secret durations of the system, respectively. These automata being of exponential size and with a unary language, testing the equality or inclusion of languages led to the NEXPTIME algorithm quoted in Theorem 8. Similarly, the \exists -exp-ET-opacity decision problem can be decided by testing whether the intersection of the languages of these automata is empty. This can be done in NLOGSPACE in the size of the automata (classically, by first building the product between these two automata, and then by checking the reachability of a pair of final states), hence the PSPACE algorithm. \square

Theorem 10 (Solvability of weak exp-ET-opacity Δ -computation problem [8, Theorems 3 and 5]). *The weak exp-ET-opacity Δ -computation problem is solvable.*

Corollary 6 (Decidability of weak exp-ET-opacity Δ -emptiness problem [8, Corollary 1]). *The weak exp-ET-opacity Δ -emptiness problem is decidable.*

In contrast to the weak exp-ET-opacity Δ -computation problem, we only show below that the full exp-ET-opacity Δ -emptiness problem is decidable; the computation problem remains open.

Theorem 11 (Decidability of the full exp-ET-opacity Δ -emptiness problem [8, Theorems 4 and 5]). *The full exp-ET-opacity Δ -emptiness problem is decidable.*

5.3 Exp-ET-opacity in parametric timed automata

We now study exp-ET-opacity problems for PTAs: we will be interested in the synthesis and in the emptiness of the valuations set ensuring that a system is fully (resp. weakly) exp-ET-opaque.

5.3.1 Definitions

We define the following problems, where we ask for parameter valuations v and for valuations of Δ s.t. $v(\mathcal{P})$ is fully (resp. weakly) $(\leq \Delta)$ -ET-opaque.

Full (resp. weak) exp-ET-opacity Δ -p-emptiness problem:

INPUT: A PTA \mathcal{P}

PROBLEM: Decide whether the set of parameter valuations v and valuations of Δ such that $v(\mathcal{P})$ is fully (resp. weakly) $(\leq \Delta)$ -ET-opaque is empty

Full (resp. weak) exp-ET-opacity Δ -p-synthesis problem:

INPUT: A PTA \mathcal{P}

PROBLEM: Synthesize the set of parameter valuations v and valuations of Δ such that $v(\mathcal{P})$ is fully (resp. weakly) $(\leq \Delta)$ -ET-opaque

Example 11. Consider again the PTA \mathcal{P} in Fig. 2.

For this PTA, the answer to the weak exp-ET-opacity Δ -p-emptiness problem is false, as there exists such a valuation (e.g., the valuation given in Example 10).

Moreover, we can show that, for all Δ and v :

- $D\overline{Visit}^{priv}(v(\mathcal{P})) = [0, 3]$
- if $v(p_1) > 3$ or $v(p_1) > v(p_2)$, it is not possible to reach ℓ_f with a run visiting ℓ_{priv} and therefore $DVisit_{>\Delta}^{priv}(v(\mathcal{P})) = DVisit_{\leq\Delta}^{priv}(v(\mathcal{P})) = \emptyset$
- if $v(p_1) \leq 3$ and $v(p_1) \leq v(p_2)$
 - $DVisit_{>\Delta}^{priv}(v(\mathcal{P})) = (v(p_1) + \Delta, v(p_2)]$
 - $DVisit_{\leq\Delta}^{priv}(v(\mathcal{P})) = [v(p_1), \min(\Delta + 3, v(p_2))]$

Recall that the full exp-ET-opacity Δ -p-synthesis problem aims at synthesizing the valuations such that $DVisit_{\leq\Delta}^{priv}(v(\mathcal{P})) = DVisit_{>\Delta}^{priv}(v(\mathcal{P})) \cup D\overline{Visit}^{priv}(v(\mathcal{P}))$. The answer to this problem is therefore the set of valuations of timing parameters and of Δ s.t.:

$$v(p_1) = 0 \wedge \left((\Delta \leq 3 \wedge 3 \leq v(p_2) \leq \Delta + 3) \vee (v(p_2) < \Delta \wedge v(p_2) = 3) \right).$$

5.3.2 Results

The subclass of lower/upper parametric timed automata

Theorem 12 (Undecidability of full (resp. weak) exp-ET-opacity Δ -p-emptiness problem [8, Theorem 6]). *The full (resp. weak) exp-ET-opacity Δ -p-emptiness problem is undecidable for L/U-PTAs.*

The synthesis problems are therefore immediately unsolvable as well.

Corollary 7 ([8, Corollary 2]). *The full (resp. weak) exp-ET-opacity Δ -p-synthesis problem is unsolvable for L/U-PTAs.*

The full class of parametric timed automata The undecidability of the emptiness problems for L/U-PTAs proved above (Theorem 12) immediately implies undecidability for the larger class of PTAs. However, as in Remark 4, the full proof (given in [8]) of the result stated below uses less clocks and parameters than for L/U-PTAs (Theorem 12).

Theorem 13 (Undecidability of full (resp. weak) exp-ET-opacity Δ -p-emptiness problem [8, Theorem 7]). *The full (resp. weak) exp-ET-opacity Δ -p-emptiness problem is undecidable for general PTAs.*

Again, the synthesis problems are therefore immediately unsolvable as well.

Corollary 8 ([8, Corollary 3]). *The full (resp. weak) exp-ET-opacity Δ -p-synthesis problem is unsolvable for PTAs.*

6 Implementation and application to Java programs

A motivation for the works on ET-opacity (described in Sections 3 and 4) is the analysis of programs. More precisely, we are interested in deciding whether a program, e.g., written in Java, is ET-opaque, i.e., whether an attacker is incapable of deducing internal behavior by only looking at its execution time. A second motivation is the *configuration* of internal timing values from a program, e.g., changing some internal delays, or tuning some `Thread.sleep()` statements in the program, so that the program becomes ET-opaque—justifying notably the results in Section 4.

Semi-algorithm and implementation Despite the negative theoretical results (notably Theorem 1), we addressed in [9] the \exists -ET-opacity p-synthesis problem for the full class of PTAs. Our method may not terminate (due to the undecidability) but, if it does, its result is correct. Our workflow [9] can be summarized as follows.

1. We slightly modify the original PTA (by adding a Boolean flag b and a final synchronization action);
2. We perform *self-composition* (i.e., parallel composition with a copy of itself) of this modified PTA, a method commonly used in security analyses [32, 15];
3. We perform reachability-synthesis (i.e., the synthesis of parameter valuations for which a given location is reachable) on ℓ_f with contradictory values of b .

Reachability-synthesis is implemented in IMITATOR [6], a parametric timed model checker taking as inputs networks of (extensions of) parametric timed automata, and synthesizing parameter valuations for which a number of properties (including reachability) hold.

Analysis of Java programs In addition, we are interested in analyzing programs too. In order to apply our method to the analysis of programs, we need a systematic way of translating a program (e.g., a Java program) into a PTA. In general, precisely modeling the execution time of a program using models like TA is highly non-trivial due to complication of hardware pipelining, caching, OS scheduling, etc. The readers are referred to the rich literature in, e.g., [29, 20]. In [9], we instead make the following simplistic assumption on execution time of a program statement and focus on solving the parameter synthesis problem. We assume that the execution time of a program statement other than `Thread.sleep(n)` is within a range $[0, \varepsilon]$ where ε is a small integer constant (in milliseconds), whereas the execution time of statement `Thread.sleep(n)` is within a range $[n, n + \varepsilon]$. In fact, we choose to keep ε *parametric* to be as general as possible, and to not depend on particular architectures.

Our test subject is a set of benchmark programs from the DARPA Space/Time Analysis for Cybersecurity (STAC) program.¹ These programs are being released publicly to facilitate researchers to develop methods and tools for identifying STAC vulnerabilities in the programs. These programs are simple yet non-trivial, and were built on purpose to highlight vulnerabilities that can be easily missed by existing security analysis tools. We *manually* translated these programs to PTAs, following the method described above, and using a number of assumptions (such as collapsing loops with predefined duration).

In addition, we applied our method to a set of PTAs examples from the literature, notably from [27, 24, 16, 34].

Experiments reported in [9] show that we can decide whether these benchmarks (including the programs) are fully ET-opaque or \exists -ET-opaque. When adding timing parameters, we additionally answer the \exists -ET-opacity p-synthesis problem, i.e., we synthesize the parameter valuations v and the associated execution times D such that $v(\mathcal{P})$ is ET-opaque. Our method allows to exhibit cases when the system can never be made ET-opaque, including by tuning internal delays, or is always ET-opaque, or is ET-opaque only for some execution times and internal timing parameters.

To summarize, the following problems can be answered using our framework:

- \exists -ET-opacity decision problem
- full ET-opacity decision problem
- weak ET-opacity decision problem (not considered in our experiments in [9], but can be easily adapted)
- \exists -ET-opacity p-synthesis problem, but without guarantee of termination, due to the undecidability of Theorem 1.

However, our procedure cannot in its current form answer neither the full ET-opacity p-synthesis problem nor the weak ET-opacity p-synthesis problem. The expiring opacity problems in Section 5 were not addressed either.

7 Conclusion and perspectives

In this paper, we recalled (and proved a few original) results related to the ET-opacity in TAs. Our notion of ET-opacity consists in considering an attacker model that can only observe the execution time of the system, i.e., the time from the initial location to a final location. The secret consists in deciding whether a special private location was visited or not. In contrast to another notion of opacity with a more powerful attacker able to observe some actions together with their timestamps, which led to the undecidability of the decision problem for TAs [19], our notion of ET-opacity yields decidability results for TAs. Parameterizing the problems using timing parameters brings undecidability for PTAs, but the subclass of L/U-PTAs gives mildly positive results.

When in addition we consider that the secret has an expiration date, similarly to the concepts introduced in [4], we are able to not only *decide* problems for TAs, but also to *synthesize* valuations for the expiration date such that the TA is weakly exp-ET-opaque. However, problems extended with timing parameters all become undecidable.

Recall that we summarized in Tables 1 and 2 the decidability results recalled in this paper, with a bold emphasis on the original results of this paper.

¹<https://github.com/Apogee-Research/STAC/>

We also reported here on an implementation using IMITATOR, which is able to answer non-parametric problems (\exists -ET-opacity decision problem, full ET-opacity decision problem, weak ET-opacity decision problem), and also answering a parameter synthesis problem (\exists -ET-opacity p-synthesis problem) without guarantee of termination for the latter problem.

Perspectives The main theoretical future work is the open problems in Table 2 (mainly the full exp-ET-opacity Δ -computation problem): it is unclear whether we can *compute* the exact set of expiration dates Δ for which a TA is fully ($\leq \Delta$)-ET-opaque.

In terms of synthesis, we have so far no procedure able (whenever it terminates) to answer the full ET-opacity p-synthesis problem or the weak ET-opacity p-synthesis problem. Synthesis procedures to answer expiring opacity problems (defined in Section 5) for PTAs remain to be designed too. These procedures cannot be both exact and guaranteed to terminate due to the aforementioned undecidability results.

Exact analysis of opacity for programs, including a more precise modeling of the cache, is also on our agenda, following works such as [20, 21].

A different direction is that of *control*: can we turn a non-opaque system into an opaque system, by restraining its possible behaviors? A first step with our notion of ET-opacity was presented in [7], with only an *untimed* controller. In addition, in [24], Gardey *et al.* propose several definitions of non-interference, related to various notions of simulation: they consider not only the *verification* problem (“is the system non-interferent?”) but also the (timed) *control* problem (“synthesize a controller that will restrict the system in order to enforce non-interference”). Extending our current line works on ET-opacity to *timed* controllers remains to be done.

Acknowledgments We are grateful to Clemens Dubsclaff and Maurice ter Beek for the opportunity to give an invited talk at TiCSA 2023, and for useful suggestions on this manuscript.

References

- [1] Rajeev Alur & David L. Dill (1994): *A theory of timed automata*. *Theoretical Computer Science* 126(2), pp. 183–235, doi:10.1016/0304-3975(94)90010-8.
- [2] Rajeev Alur, Limor Fix & Thomas A. Henzinger (1999): *Event-Clock Automata: A Determinizable Class of Timed Automata*. *Theoretical Computer Science* 211(1-2), pp. 253–273, doi:10.1016/S0304-3975(97)00173-4.
- [3] Rajeev Alur, Thomas A. Henzinger & Moshe Y. Vardi (1993): *Parametric real-time reasoning*. In S. Rao Kosaraju, David S. Johnson & Alok Aggarwal, editors: *STOC*, ACM, New York, NY, USA, pp. 592–601, doi:10.1145/167088.167242.
- [4] Ikhlass Ammar, Yamen El Touati, Moez Yeddes & John Mullins (2021): *Bounded opacity for timed systems*. *Journal of Information Security and Applications* 61, pp. 1–13, doi:10.1016/j.jisa.2021.102926.
- [5] Étienne André (2019): *What’s decidable about parametric timed automata?* *International Journal on Software Tools for Technology Transfer* 21(2), pp. 203–219, doi:10.1007/s10009-017-0467-0.
- [6] Étienne André (2021): *IMITATOR 3: Synthesis of timing parameters beyond decidability*. In Rustan Leino & Alexandra Silva, editors: *CAV, Lecture Notes in Computer Science* 12759, Springer, pp. 1–14, doi:10.1007/978-3-030-81685-8_26.

- [7] Étienne André, Shapagat Bolat, Engel Lefauchaux & Dylan Marinho (2022): *strategFTO: Untimed control for timed opacity*. In Cyrille Artho & Peter Ölveczky, editors: *FTSCS*, ACM, pp. 27–33, doi:10.1145/3563822.3568013.
- [8] Étienne André, Engel Lefauchaux & Dylan Marinho (2023): *Expiring opacity problems in parametric timed automata*. In Yamine Ait-Ameur & Ferhat Khendek, editors: *ICECCS*. To appear.
- [9] Étienne André, Didier Lime, Dylan Marinho & Jun Sun (2022): *Guaranteeing timed opacity using parametric timed model checking*. *ACM Transactions on Software Engineering and Methodology* 31(4), pp. 1–36, doi:10.1145/3502851.
- [10] Étienne André, Didier Lime & Nicolas Markey (2020): *Language Preservation Problems in Parametric Timed Automata*. *Logical Methods in Computer Science* 16(1), doi:10.23638/LMCS-16(1:5)2020. Available at <https://lmcs.episciences.org/6042>.
- [11] Étienne André, Didier Lime & Mathias Ramparison (2018): *TCTL model checking lower/upper-bound parametric timed automata without invariants*. In David N. Jansen & Pavithra Prabhakar, editors: *FORMATS, Lecture Notes in Computer Science* 11022, Springer, pp. 1–17, doi:10.1007/978-3-030-00151-3_3.
- [12] Étienne André, Didier Lime & Olivier H. Roux (2022): *Reachability and liveness in parametric timed automata*. *Logical Methods in Computer Science* 18(1), pp. 31:1–31:41, doi:10.46298/lmcs-18(1:31)2022. Available at <https://lmcs.episciences.org/9070/pdf>.
- [13] Johan Arcile & Étienne André (2023): *Timed automata as a formalism for expressing security: A survey on theory and practice*. *ACM Computing Surveys* 55(6), pp. 1–36, doi:10.1145/3534967.
- [14] Roberto Bagnara, Patricia M. Hill & Enea Zaffanella (2008): *The Parma Polyhedra Library: Toward a Complete Set of Numerical Abstractions for the Analysis and Verification of Hardware and Software Systems*. *Science of Computer Programming* 72(1–2), pp. 3–21, doi:10.1016/j.scico.2007.08.001.
- [15] Gilles Barthe, Pedro R. D’Argenio & Tamara Rezk (2011): *Secure information flow by self-composition*. *Mathematical Structures in Computer Science* 21(6), pp. 1207–1252, doi:10.1017/S0960129511000193.
- [16] Gilles Benattar, Franck Cassez, Didier Lime & Olivier H. Roux (2015): *Control and synthesis of non-interferent timed systems*. *International Journal of Control* 88(2), pp. 217–236, doi:10.1080/00207179.2014.944356.
- [17] Laura Bozzelli & Salvatore La Torre (2009): *Decision problems for lower/upper bound parametric timed automata*. *Formal Methods in System Design* 35(2), pp. 121–151, doi:10.1007/s10703-009-0074-0.
- [18] Véronique Bruyère, Emmanuel Dall’Olio & Jean-Francois Raskin (2008): *Durations and parametric model-checking in timed automata*. *ACM Transactions on Computational Logic* 9(2), pp. 12:1–12:23, doi:10.1145/1342991.1342996.
- [19] Franck Cassez (2009): *The Dark Side of Timed Opacity*. In Jong Hyuk Park, Hsiao-Hwa Chen, Mohammed Atiqzaman, Changhoon Lee, Tai-Hoon Kim & Sang-Soo Yeo, editors: *ISA, Lecture Notes in Computer Science* 5576, Springer, pp. 21–30, doi:10.1007/978-3-642-02617-1_3.
- [20] Franck Cassez & Jean-Luc Béchenec (2013): *Timing Analysis of Binary Programs with UPPAAL*. In Josep Carmona, Mihai T. Lazarescu & Marta Pietkiewicz-Koutny, editors: *ACSD*, IEEE Computer Society, pp. 41–50, doi:10.1109/ACSD.2013.7.
- [21] Duc-Hiep Chu, Joxan Jaffar & Rasool Maghareh (2016): *Precise Cache Timing Analysis via Symbolic Execution*. In: *RTAS*, IEEE Computer Society, pp. 293–304, doi:10.1109/RTAS.2016.7461358.
- [22] Shuwen Deng, Wenjie Xiong & Jakub Szefer (2018): *Cache timing side-channel vulnerability checking with computation tree logic*. In Jakub Szefer, Weidong Shi & Ruby B. Lee, editors: *ISCA*, ACM, pp. 2:1–2:8, doi:10.1145/3214292.3214294.
- [23] Goran Doychev, Boris Köpf, Laurent Mauborgne & Jan Reineke (2015): *CacheAudit: A Tool for the Static Analysis of Cache Side Channels*. *ACM Transactions on Information and System Security* 18(1), pp. 4:1–4:32, doi:10.1145/2756550.

- [24] Guillaume Gardey, John Mullins & Olivier H. Roux (2007): *Non-Interference Control Synthesis for Security Timed Automata*. *Electronic Notes in Theoretical Computer Science* 180(1), pp. 35–53, doi:10.1016/j.entcs.2005.05.046.
- [25] Shengjian Guo, Meng Wu & Chao Wang (2018): *Adversarial symbolic execution for detecting concurrency-related cache timing leaks*. In Gary T. Leavens, Alessandro Garcia & Corina S. Pasareanu, editors: *ESEC/SIGSOFT FSE*, ACM, pp. 377–388, doi:10.1145/3236024.3236028.
- [26] Thomas A. Henzinger, Zohar Manna & Amir Pnueli (1992): *Timed Transition Systems*. In J. W. de Bakker, Cornelis Huizing, Willem P. de Roever & Grzegorz Rozenberg, editors: *REX, Lecture Notes in Computer Science* 600, Springer, pp. 226–251, doi:10.1007/BFb0031995.
- [27] Thomas Hune, Judi Romijn, Mariëlle Stoelinga & Frits W. Vaandrager (2002): *Linear parametric model checking of timed automata*. *Journal of Logic and Algebraic Programming* 52-53, pp. 183–220, doi:10.1016/S1567-8326(02)00037-1.
- [28] Aleksandra Jovanović, Didier Lime & Olivier H. Roux (2015): *Integer Parameter Synthesis for Real-Time Systems*. *IEEE Transactions on Software Engineering* 41(5), pp. 445–461, doi:10.1109/TSE.2014.2357445.
- [29] Mingsong Lv, Wang Yi, Nan Guan & Ge Yu (2010): *Combining Abstract Interpretation with Model Checking for Timing Analysis of Multicore Software*. In: *RTSS*, IEEE Computer Society, pp. 339–349, doi:10.1109/RTSS.2010.30.
- [30] Joseph S. Miller (2000): *Decidability and Complexity Results for Timed Automata and Semi-linear Hybrid Automata*. In Nancy A. Lynch & Bruce H. Krogh, editors: *HSCC, Lecture Notes in Computer Science* 1790, Springer, pp. 296–309, doi:10.1007/3-540-46430-1_26.
- [31] Quoc-Sang Phan, Lucas Bang, Corina S. Pasareanu, Pasquale Malacaria & Tevfik Bultan (2017): *Synthesis of Adaptive Side-Channel Attacks*. In: *CSF*, IEEE Computer Society, pp. 328–342, doi:10.1109/CSF.2017.8.
- [32] Tachio Terauchi & Alexander Aiken (2005): *Secure Information Flow as a Safety Problem*. In Chris Hankin & Igor Siveroni, editors: *Proceedings of the 12th International Symposium on Static Analysis (SAS 2005)*, *Lecture Notes in Computer Science* 3672, Springer, pp. 352–367, doi:10.1007/11547662_24.
- [33] Saeid Tizpaz-Niari, Pavol Cerný & Ashutosh Trivedi (2019): *Quantitative Mitigation of Timing Side Channels*. In Işıl Dillig & Serdar Tasiran, editors: *CAV, Part I, Lecture Notes in Computer Science* 11561, Springer, pp. 140–160, doi:10.1007/978-3-030-25540-4_8.
- [34] Panagiotis Vasilikos, Flemming Nielson & Hanne Riis Nielson (2018): *Secure Information Release in Timed Automata*. In Lujo Bauer & Ralf Küsters, editors: *POST, Lecture Notes in Computer Science* 10804, Springer, pp. 28–52, doi:10.1007/978-3-319-89722-6_2.
- [35] Panagiotis Vasilikos, Hanne Riis Nielson, Flemming Nielson & Boris Köpf (2019): *Timing Leaks and Coarse-Grained Clocks*. In: *CSF*, IEEE, pp. 32–47, doi:10.1109/CSF.2019.00010.
- [36] Lingtai Wang & Naijun Zhan (2018): *Decidability of the Initial-State Opacity of Real-Time Automata*. In Cliff B. Jones, Ji Wang & Naijun Zhan, editors: *Symposium on Real-Time and Hybrid Systems - Essays Dedicated to Professor Chaochen Zhou on the Occasion of His 80th Birthday*, *Lecture Notes in Computer Science* 11180, Springer, pp. 44–60, doi:10.1007/978-3-030-01461-2_3.
- [37] Lingtai Wang, Naijun Zhan & Jie An (2018): *The Opacity of Real-Time Automata*. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37(11), pp. 2845–2856, doi:10.1109/TCAD.2018.2857363.
- [38] Meng Wu, Shengjian Guo, Patrick Schaumont & Chao Wang (2018): *Eliminating timing side-channel leaks using program repair*. In Frank Tip & Eric Bodden, editors: *ISSTA*, ACM, pp. 15–26, doi:10.1145/3213846.3213851.