



HAL
open science

Guaranteeing Timed Opacity using Parametric Timed Model Checking

Étienne André, Didier Lime, Dylan Marinho, Jun Sun

► **To cite this version:**

Étienne André, Didier Lime, Dylan Marinho, Jun Sun. Guaranteeing Timed Opacity using Parametric Timed Model Checking. ACM Transactions on Software Engineering and Methodology, 2022, 31 (4), pp.1-36. 10.1145/3502851 . hal-03798157

HAL Id: hal-03798157





<https://hal.science/hal-03798157v1>

Submitted on 9 Oct 2024

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.

Guaranteeing Timed Opacity using Parametric Timed Model Checking*

Étienne André ¹, Didier Lime ², Dylan Marinho ¹ and Jun Sun ³

¹Université de Lorraine, CNRS, Inria, LORIA, Nancy, France

²École Centrale de Nantes, LS2N, UMR CNRS 6004, Nantes

³School of Information Systems, Singapore Management University

Abstract

Information leakage can have dramatic consequences on systems security. Among harmful information leaks, the timing information leakage occurs whenever an attacker successfully deduces confidential internal information. In this work, we consider that the attacker has access (only) to the system execution time. We address the following timed opacity problem: given a timed system, a private location and a final location, synthesize the execution times from the initial location to the final location for which one cannot deduce whether the system went through the private location. We also consider the full timed opacity problem, asking whether the system is opaque for all execution times. We show that these problems are decidable for timed automata (TAs) but become undecidable when one adds parameters, yielding parametric timed automata (PTAs). We identify a subclass with some decidability results. We then devise an algorithm for synthesizing PTAs parameter valuations guaranteeing that the resulting TA is opaque. We finally show that our method can also apply to program analysis.

Keywords— opacity, timed automata, IMITATOR, parameter synthesis

1 Introduction

Timed systems often combine hard real-time constraints with other complications such as concurrency. Information leakage can have dramatic consequences on the security of

*This is the author version of the manuscript of the same name published in ACM Transactions on Software Engineering and Methodology (ToSEM). The final version is available at [acm.org](https://arxiv.org/abs/2206.05438). We would like to thank Sudipta Chattopadhyay for helpful suggestions, Nicolas Markey for an interesting discussion on the proof of [Proposition 1](#), Jiaying Li for his help with preliminary model conversion, and an anonymous reviewer of ATVA 2019 for suggesting [Remark 1](#). This work is partially supported by the ANR national research program PACS (ANR-14-CE28-0002), by the ANR-NRF research program ProMiS (ANR-19-CE25-0015), and by ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603), JST.

such systems. Among harmful information leaks, the *timing information leakage* is the ability for an attacker to deduce internal information depending on timing information. In this work, 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.

Opacity In its most general form on partially observed labeled transitions systems, given a set of runs that reveal a secret (e. g., they perform a secret action or visit a secret state), *opacity* states that if there exists a run of the system that reveals the secret (i. e., belongs to the given secret set), there exists another run, with the same observation, that does not reveal that secret [Bry+08]. This secret is completely generic and, depending on its actual definition, properties and their decidability can differ.

In our setting, we define a form of opacity in which the observation is only *the time to reach a designated location*.

Contributions for TAs In this work, we consider the setting of timed automata (TAs), which is a popular extension of finite-state automata with clocks [AD94]. We consider the following version of timed opacity: given a TA, a private location denoting the execution of some secret behavior and a final location denoting the completion of the execution, the TA is opaque for a given execution time d (i. e., the time of a run from the initial location to the final location) if there exist two runs of duration d from the initial location to the final location, one going through the private location, and another run *not* going through the private location. That is, for this particular execution time, the system is opaque if one cannot deduce whether the system went through the private location. Such a notion of timed opacity can be used to capture many interesting security problems: for instance, it is possible to deduce whether a secret satisfies a certain condition based on whether a certain branch is visited or not.

To be explicit, the attacker knows a TA model of the system, and can observe the execution time from the system start until it reaches some particular final location. No other actions can be observed. Then, the system is timed opaque if the attacker cannot deduce whether the system has visited some particular private location. From a higher-level point of view, this means that the attacker cannot deduce some private information, such as whether some location has been visited, or whether some branch of a given program was visited, by only observing the execution time. In practice, this corresponds to a setting where the attacker may interact with some computational process on a remote machine (e. g., a server) and receives the responses only at the end of the process (e. g., a server message is received).

We consider two problems based on this notion of timed opacity:

1. a computation problem: the computation of the set of possible execution times for which the system is timed opaque; and
2. a decision problem: whether the TA is timed opaque for all execution times (referred to as full timed opacity).

We first prove that these problems can be effectively solved for TAs. We implement our procedure and apply it to a set of benchmarks containing notably a set of Java programs known for their (absence of) timing information leakage.

Contributions for parametric TAs As a second setting, we consider a higher-level version of these problems by allowing (internal) timing parameters in the system, which can model uncertainty or unknown constants at early design stage. The setting becomes parametric timed automata (PTAs) [AHV93].

On the theoretical side, we answer an existential parametric version of the two aforementioned problems, that is, the existence of (at least) one parameter valuation for which the TA is (fully) timed opaque. Although we show that these problems are in general undecidable, we exhibit a subclass with some decidability results.

Then, we address a practical problem: given a timed system with timing parameters, a private location and a final location, synthesize the timing parameters and the execution times for which one cannot deduce whether the system went through the private location. We devise a general procedure not guaranteed to terminate, but that behaves well on examples from the literature.

Summary of the contributions To sum up, this manuscript proposes the following contributions:

1. a notion of timed opacity, and a notion of full timed opacity for TAs;
2. a procedure to solve the timed opacity computation problem for TAs, and a procedure to answer the full timed opacity decision problem for TAs;
3. a study of two theoretical decision problems extending the two aforementioned problems to the parametric setting, and exhibition of a decidable subclass;
4. a practical algorithm to synthesize parameter valuations and execution times for which the TA is guaranteed to be opaque;
5. a set of experiments on a set of benchmarks, including PTAs translations from Java programs.

This manuscript is an extension of [AS19] with the following improvements.

- We provide all proofs of the results published in [AS19].
- We extend the theoretical part, by considering not one problem (as in [AS19]) but two versions (timed opacity w.r.t. a set of execution times, and full timed opacity), both for TAs and PTAs (including the subclass of L/U-PTAs).
- We propose a more elegant proof of [Proposition 1](#) (formerly [AS19, Proposition 1]), based on RA arithmetic [Wei99].
- On the practical side, we give hints to extend our construction to a richer framework ([Section 9.3](#)).

Outline After reviewing related works in [Section 2](#), [Section 3](#) recalls necessary concepts and [Section 4](#) introduces the problem. [Section 5](#) addresses timed opacity for timed automata. We then address the parametric version of timed opacity, with theory studied in [Sections 6](#) and [7](#), algorithmic in [Section 8](#) and experiments in [Section 9](#). [Section 10](#) concludes the paper.

2 Related works

Opacity and timed automata This work is closely related to the line of work on defining and analyzing information flow in timed automata. It is well-known (see

e. g., [Koc96; FS00; BB07; Kot+13; Ben+15a]) that time is a potential attack vector against secure systems. That is, it is possible that a non-interferent (secure) system can become interferent (insecure) when timing constraints are added [GMR07].

In *non-interference*, actions are partitioned into two levels of privilege, *high* and *low*, and we require that the system in which high-level actions are removed is equivalent to the system in which they are hidden (i. e., replaced by an unobservable action). Different equivalences lead to different flavors of non-interference. In [Bar+02; BT03], a first notion of *timed* non-interference is proposed for TAs. This notion is extended to PTAs in [AK20], with a semi-algorithm.

In [GMR07], Gardey *et al.* define timed strong non-deterministic non-interference (SNNI) based on timed language equivalence between the automaton with hidden low-level actions and the automaton with removed low-level actions. Furthermore, they show that the problem of determining whether a timed automaton satisfies SNNI is undecidable. In contrast, timed cosimulation-based SNNI, timed bisimulation-based SNNI and timed state SNNI are decidable. Classical SNNI is the one corresponding to the equality of the languages of the two systems. As such it is clearly a special case of opacity in which the secret runs are those containing a high-level action [Bry+08]. Other equivalence relations (namely (timed) cosimulation, (timed) bisimulation, sets of states) are not as easily relatable to opacity. No implementation is provided in [GMR07].

In [Cas09], it is proved that it is undecidable whether a TA is opaque, for the following definition of opacity: the system is opaque if an attacker cannot deduce whether some set of actions was performed, by only observing a given set of observable actions together with their timestamp. This problem is proved undecidable even for the restricted class of event-recording automata [AFH99], which is a subclass of TAs. No implementation nor procedure is provided. In contrast, our definition of opacity is decidable for TAs, notably because in our setting the attacker power is more restricted (they can only observe the “execution time”); in addition, our definition of opacity has some practical relevance nonetheless, when an attacker is able to interact remotely with the system under attack, and is therefore able to measure the response time.

In [Amm+21], the authors consider a *time-bound* opacity, where the attacker has to disclose the secret before an upper bound, using a partial observability. The authors prove that this problem is decidable for TAs. A construction and an algorithm are also provided to solve it; a case study is verified using SPACEEX [Fre+11]. In contrast, our definition of opacity only assumes observation of the execution time, does not assume any time-bounded setting, and our most general problem is parametric.

In [NNV17], the authors propose a type system dealing with non-determinism and (continuous) real-time, the adequacy of which is ensured using non-interference. We share the common formalism of TAs; however, we mainly focus on leakage as execution time, and we *synthesize* internal parts of the system (clock guards), in contrast to [NNV17] where the system is fixed.

In [VNN18], Vasilikos *et al.* define the security of timed automata in term of information flow using a bisimulation relation over a set of observable nodes and develop an algorithm for deriving a sound constraint for satisfying the information flow property locally based on relevant transitions.

In [GSB18], Gerking *et al.* study non-interference properties with input, high and low actions and provide a resolution method reducing a secure behavior to an unreachability construction. The proof-of-concept consists in the exhibition of a test automaton with a dedicated location that indicates violations of noninterference whenever it is reachable during execution. Then, UPPAAL [LPY97] is used to obtain the answer.

In [Ben+15a], Benattar *et al.* study the control synthesis problem of timed automata for SNNI. That is, given a timed automaton, they propose a method to automatically generate a (largest) sub-system such that it is non-interferent, if possible. Different from the above-mentioned work, our work considers *parametric* timed automata, i. e., timed systems with unknown design parameters, and focuses on synthesizing parameter valuations which guarantee information flow property. Compared to [Ben+15a], our approach is more realistic as it does not require change of program structure. Rather, our result provides guidelines on how to choose the timing parameters (e. g., how long to wait after certain program statements) for avoiding information leakage.

In [WZ18; WZA18], Wang *et al.* investigate interesting opacity problems for real-time automata. These works come with a dedicated Python implementation. Although their definition shares similarities with ours, real-time automata are a severely restricted formalism compared to TAs. Indeed, timed aspects are only considered by interval restrictions over the total elapsed time along transitions. Real-time automata can be seen as a subclass of TAs with a single clock, reset at each transition. Also, parameters are not considered in their work.

To the best of our knowledge, our approach is the first work on parametric model checking for timed automata for information flow property. In addition, and in contrast to most of the aforementioned works, our approach comes with an implementation.

Execution times and timed automata In this paper, we need to compute execution times in timed automata, i. e., the durations of all runs reaching the final state. Despite its natural aspect, this problem seems sparsely investigated in the literature. Both [BDR08; Ros19] deal with the computation of duration sets; we do reuse some of the reasoning from [BDR08] in our proof of [Proposition 1](#). Conversely, results from [BR07] cannot be used in our work: while the equality must be forbidden in PTCTL formulae to make the problems decidable, equality constraints in a PTCTL formula would be required for such an approach to answer our problems. Furthermore, the results presented in [Bou+21] cannot be applied to our study, since they concern one-clock timed automata.

Mitigating information leakage Complex systems may exhibit security problems through information leakage due to the presence of unintended communication media, called side channels. An example is time side channels in which measuring, e. g., execution times, gives information on some sensitive information. In 2018, the Spectre vulnerability [Koc+20] exploited speculative execution to bring secret information into the cache; subsequently, cache-timing attacks were launched to exfiltrate these secrets. Therefore, mitigation of timing attacks is of utmost importance.

Our work is related to work on mitigating information leakage through those time side channels [Aga00; Mol+05; Cop+09; WS17; Wu+18]. In [Aga00], Agat *et al.* proposed to eliminate time side channels through type-driven cross-copying. In [Mol+05], Molnar *et al.* proposed, along the program counter model, a method for mitigating side channel through merging branches. A similar idea was proposed in [BRW06]. Coppens *et al.* [Cop+09] developed a compiler backend for removing such leaks on x86 processors. In [WS17], Wang *et al.* proposed to automatically generate masking code for eliminating side channels through program synthesis. In [Wu+18], Wu *et al.* proposed to eliminate time side channels through program repair. Different from the above-mentioned works, we reduce the problem of mitigating time side channels as a

parametric model checking problem and solve it using parametric reachability analysis techniques.

This work is related to work on identifying information leakage through timing analysis [SPW18; CR11; Abb+16; Zha+18; DSF16; Doy+15; GWW18]. In [CR11], Chattopadhyay and Roychoudhury applied model checking to perform cache timing analysis. In [CJM16], Chu *et al.* performed similar analysis through symbolic execution. In [Abb+16], Abbasi *et al.* apply the NuSMV model checker to verify integrated circuits against information leakage through side channels. In [Doy+15], a tool is developed to identify time side channel through static analysis. In [Zha+18], Sung *et al.* developed a framework based on LLVM for cache timing analysis.

3 Preliminaries

In this work, we assume a system is modeled in the form of a parametric timed automaton (PTA). In Section 9.2, we discuss how we can model programs with unknown design parameters (e.g., a Java program with a statement `Thread.sleep(n)` where `n` is unknown) as PTA.

3.1 Clocks, parameters and guards

We assume a set $\mathbb{X} = \{x_1, \dots, x_H\}$ of *clocks*, i.e., real-valued variables that all evolve over time at the same rate. A clock valuation is a function $\mu : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$. We write $\vec{0}$ for the clock valuation assigning 0 to all clocks. Given $d \in \mathbb{R}_{\geq 0}$, $\mu + d$ denotes the valuation s.t. $(\mu + d)(x) = \mu(x) + d$, for all $x \in \mathbb{X}$. Given $R \subseteq \mathbb{X}$, we define the *reset* of a valuation μ , denoted by $[\mu]_R$, as follows: $[\mu]_R(x) = 0$ if $x \in R$, and $[\mu]_R(x) = \mu(x)$ otherwise.

We assume a set $\mathbb{P} = \{p_1, \dots, p_M\}$ of *parameters*, i.e., unknown constants. A parameter *valuation* v is a function $v : \mathbb{P} \rightarrow \mathbb{Q}_+$. We assume $\bowtie \in \{<, \leq, =, \geq, >\}$. A guard g is a constraint over $\mathbb{X} \cup \mathbb{P}$ defined by a conjunction of inequalities 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 g , we write $\mu \models v(g)$ if the expression obtained by replacing each x with $\mu(x)$ and each p with $v(p)$ in g evaluates to true.

3.2 Parametric timed automata

Parametric timed automata (PTAs) extend timed automata with parameters within guards and invariants in place of integer constants [AHV93].

3.2.1 Syntax

Definition 1 (PTA). A PTA \mathcal{A} is a tuple $\mathcal{A} = (\Sigma, L, \ell_0, \ell_f, \mathbb{X}, \mathbb{P}, I, E)$, 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. $\ell_f \in L$ is the (unique) final location,
5. \mathbb{X} is a finite set of clocks,
6. \mathbb{P} is a finite set of parameters,




Figure 1: A PTA example

7. I is the invariant, assigning to every $\ell \in L$ a guard $I(\ell)$,
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 guard.

Example 1. Consider the PTA in Fig. 1 (inspired by [GMR07, Fig. 1b]), using one clock x and two parameters p_1 and p_2 . ℓ_0 is the initial location, while we assume that ℓ_1 is the (only) *final* location, i. e., a location in which an attacker can measure the execution time from the initial location.

L/U-PTAs For some theoretical problems solved in Sections 6 and 7, we will consider the subclass of PTAs called “lower-bound/upper-bound parametric timed automata” (L/U-PTAs), introduced in [Hun+02].

Definition 2 (L/U-PTA [Hun+02]). 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 2. The PTA in Fig. 1 is an L/U-PTA with $\{p_1, p_2\}$ as lower-bound parameters, and \emptyset as upper-bound parameters.

The PTA in Fig. 5 is not an L/U-PTA, because p is compared to cl both as a lower-bound (in “ $p \times 32^2 \leq cl$ ”) and as an upper-bound (“ $cl \leq p \times 32^2 + \epsilon$ ”).

Given a parameter valuation v , we denote by $v(\mathcal{A})$ the non-parametric structure where all occurrences of any parameter p_i have been replaced by $v(p_i)$. We denote as a *timed automaton* any structure $v(\mathcal{A})$, by assuming a rescaling of the constants: by multiplying all constants in $v(\mathcal{A})$ by the least common multiple of their denominators, we obtain an equivalent (integer-valued) TA, as defined in [AD94].

Synchronized product of PTAs The *synchronous product* (using strong broadcast, i. e., synchronization on a given set of actions), or *parallel composition*, of several PTAs gives a PTA.

Definition 3 (synchronized product of PTAs). Let $N \in \mathbb{N}$. Given a set of PTAs $\mathcal{A}_i = (\Sigma_i, L_i, (\ell_0)_i, (\ell_f)_i, \mathbb{X}_i, \mathbb{P}_i, I_i, E_i)$, $1 \leq i \leq N$, and a set of actions Σ_s , the *synchronized product* of \mathcal{A}_i , $1 \leq i \leq N$, denoted by $\mathcal{A}_1 \parallel_{\Sigma_s} \mathcal{A}_2 \parallel_{\Sigma_s} \cdots \parallel_{\Sigma_s} \mathcal{A}_N$, is the tuple $(\Sigma, L, \ell_0, \ell_f, \mathbb{X}, \mathbb{P}, I, E)$, where:

1. $\Sigma = \bigcup_{i=1}^N \Sigma_i$,
2. $L = \prod_{i=1}^N L_i$,
3. $\ell_0 = ((\ell_0)_1, \dots, (\ell_0)_N)$,

4. $\ell_f = ((\ell_f)_1, \dots, (\ell_f)_N)$,
5. $\mathbb{X} = \bigcup_{1 \leq i \leq N} \mathbb{X}_i$,
6. $\mathbb{P} = \bigcup_{1 \leq i \leq N} \mathbb{P}_i$,
7. $I((\ell_1, \dots, \ell_N)) = \bigwedge_{i=1}^N I_i(\ell_i)$ for all $(\ell_1, \dots, \ell_N) \in L$,

and E is defined as follows. For all $a \in \Sigma$, let ζ_a be the subset of indices $i \in 1, \dots, N$ such that $a \in \Sigma_i$. For all $a \in \Sigma$, for all $(\ell_1, \dots, \ell_N) \in L$, for all $(\ell'_1, \dots, \ell'_N) \in L$, $((\ell_1, \dots, \ell_N), g, a, R, (\ell'_1, \dots, \ell'_N)) \in E$ if:

- if $a \in \Sigma_s$, then
 1. for all $i \in \zeta_a$, there exist g_i, R_i such that $(\ell_i, g_i, a, R_i, \ell'_i) \in E_i$, $g = \bigwedge_{i \in \zeta_a} g_i$, $R = \bigcup_{i \in \zeta_a} R_i$, and,
 2. for all $i \notin \zeta_a$, $\ell'_i = \ell_i$.
- otherwise (if $a \notin \Sigma_s$), then there exists $i \in \zeta_a$ such that
 1. there exist g_i, R_i such that $(\ell_i, g_i, a, R_i, \ell'_i) \in E_i$, $g = g_i$, $R = R_i$, and,
 2. for all $j \neq i$, $\ell'_j = \ell_j$.

That is, synchronization is only performed on Σ_s , and other actions are interleaved.

3.2.2 Concrete semantics of TAs

Let us now recall the concrete semantics of TA.

Definition 4 (Semantics of a TA). Given a PTA $\mathcal{A} = (\Sigma, L, \ell_0, \ell_f, \mathbb{X}, \mathbb{P}, I, E)$, and a parameter valuation v , the semantics of $v(\mathcal{A})$ is given by the timed transition system (TTS) [HMP91] $T_{v(\mathcal{A})} = (S, s_0, \rightarrow)$, with

- $S = \{(\ell, \mu) \in L \times \mathbb{R}_{\geq 0}^H \mid \mu \models v(I(\ell))\}$,
- $s_0 = (\ell_0, \vec{0})$,
- \rightarrow consists of the discrete and (continuous) delay transition relations:
 1. discrete transitions: $(\ell, \mu) \xrightarrow{e} (\ell', \mu')$, if $(\ell, \mu), (\ell', \mu') \in S$, and there exists $e = (\ell, g, a, R, \ell') \in E$, such that $\mu' = [\mu]_R$, and $\mu \models v(g)$.
 2. delay transitions: $(\ell, \mu) \xrightarrow{d} (\ell, \mu + d)$, with $d \in \mathbb{R}_{\geq 0}$, if $\forall d' \in [0, d], (\ell, \mu + d') \in S$.

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 $v(\mathcal{A})$ with concrete semantics (S, s_0, \rightarrow) , we refer to the states of $T_{v(\mathcal{A})}$ as the *concrete states* of $v(\mathcal{A})$. A *run* of $v(\mathcal{A})$ is a (finite or infinite) alternating sequence of concrete states of $v(\mathcal{A})$ and pairs of delays and edges starting from the initial state s_0 of the form $s_0, (d_0, e_0), s_1, \dots$ with $i = 0, 1, \dots$, $e_i \in E$, $d_i \in \mathbb{R}_{\geq 0}$ and $s_i \xrightarrow{(d_i, e_i)} s_{i+1}$.

Given a state $s = (\ell, \mu)$, we say that s is reachable in $v(\mathcal{A})$ if s appears in a run of $v(\mathcal{A})$. By extension, we say that ℓ is reachable; and by extension again, given a set L_T of locations, we say that L_T is reachable if there exists $\ell \in L_T$ such that ℓ is reachable

in $v(\mathcal{A})$.¹ Given $\ell, \ell' \in L$ and a run ρ , we say that ℓ is reached on the way to ℓ' 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 $\ell_m = \ell$, $\ell_n = \ell'$ and $\forall 0 \leq i \leq m-1, \ell_i \neq \ell'$. Conversely, ℓ is avoided on the way to ℓ' 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 $\ell_n = \ell'$ and $\forall 0 \leq i \leq n, \ell_i \neq \ell$. Given $\ell, \ell' \in L$, we say that ℓ is *reachable on the way to ℓ'* in $v(\mathcal{A})$ if there exists a run ρ of $v(\mathcal{A})$ for which ℓ is reached on the way to ℓ' in ρ . Otherwise, ℓ is *unreachable on the way to ℓ'* .

The *duration* of a finite run $\rho : s_0, (d_0, e_0), s_1, \dots, (d_{i-1}, e_{i-1}), (\ell_i, \mu_i)$ is $\text{dur}(\rho) = \sum_{0 \leq j \leq i-1} d_j$. We also say that ℓ_i is reachable in time $\text{dur}(\rho)$.

Example 3. Consider again the PTA \mathcal{A} in Fig. 1, and let v be such that $v(p_1) = 1$ and $v(p_2) = 2$. Consider the following run ρ of $v(\mathcal{A})$: $(\ell_0, x = 0), (1.4, e_2), (\ell_2, x = 1.4), (1.3, e_3), (\ell_1, x = 2.7)$, where e_2 is the edge from ℓ_0 to ℓ_2 in Fig. 1, and e_3 is the edge from ℓ_2 to ℓ_1 . We write “ $x = 1.4$ ” instead of “ μ such that $\mu(x) = 1.4$ ”. We have $\text{dur}(\rho) = 1.4 + 1.3 = 2.7$. In addition, ℓ_2 is reached on the way to ℓ_1 in ρ .

3.2.3 Timed automata regions

Let us next recall the concept of regions and the region graph [AD94].

Given a TA $v(\mathcal{A})$, for a clock x_i , we denote by c_i the largest constant to which x_i is compared within the guards and invariants of $v(\mathcal{A})$ (that is, $c_i = \max_i(\{d_i \mid x \bowtie d_i \text{ appears in a guard or invariant of } v(\mathcal{A})\})$). Given a clock valuation μ and a clock x_i , let $\lfloor \mu(x_i) \rfloor$ and $\text{fract}(\mu(x_i))$ denote respectively the integral part and the fractional part of $\mu(x_i)$.

Example 4. Consider again the PTA in Fig. 1, and let v be such that $v(p_1) = 2$ and $v(p_2) = 4$. In the TA $v(\mathcal{A})$, the clock x is compared to the constants in $\{2, 3, 4\}$. In that case, $c = 4$ is the largest constant to which the clock x is compared.

Definition 5 (Region equivalence). We say that two clock valuations μ and μ' are equivalent, denoted $\mu \approx \mu'$, if the following three conditions hold for any clocks x_i, x_j :

1. either
 - (a) $\lfloor \mu(x_i) \rfloor = \lfloor \mu'(x_i) \rfloor$ or
 - (b) $\mu(x_i) > c_i$ and $\mu'(x_i) > c_i$
2. $\text{fract}(\mu(x_i)) \leq \text{fract}(\mu(x_j))$ iff $\text{fract}(\mu'(x_i)) \leq \text{fract}(\mu'(x_j))$
3. $\text{fract}(\mu(x_i)) = 0$ iff $\text{fract}(\mu'(x_i)) = 0$

The equivalence relation \approx is extended to the states of $T_{v(\mathcal{A})}$: if $s = (\ell, \mu), s' = (\ell', \mu')$ are two states of $T_{v(\mathcal{A})}$, we write $s \approx s'$ iff $\ell = \ell'$ and $\mu \approx \mu'$.

We denote by $[s]$ the equivalence class of s for \approx . A *region* is an equivalence class $[s]$ of \approx . The set of all regions is denoted $\mathcal{R}_{v(\mathcal{A})}$. Given a state $s = (\ell, \mu)$ and $d \geq 0$, we write $s + d$ to denote $(\ell, \mu + d)$.

Definition 6 (Region graph [BDR08]). The *region graph* $\mathcal{RG}_{v(\mathcal{A})} = (\mathcal{R}_{v(\mathcal{A})}, \mathcal{F}_{v(\mathcal{A})})$ is a finite graph with:

¹We use an existential quantification over the set L_T of locations, i. e., the set of locations is reachable if at least one target location is reachable. This is a standard definition for reachability synthesis (see, e. g., [JLR15]), and our Algorithm 1 uses a singleton set for L_T anyway.




Figure 2: A TA example

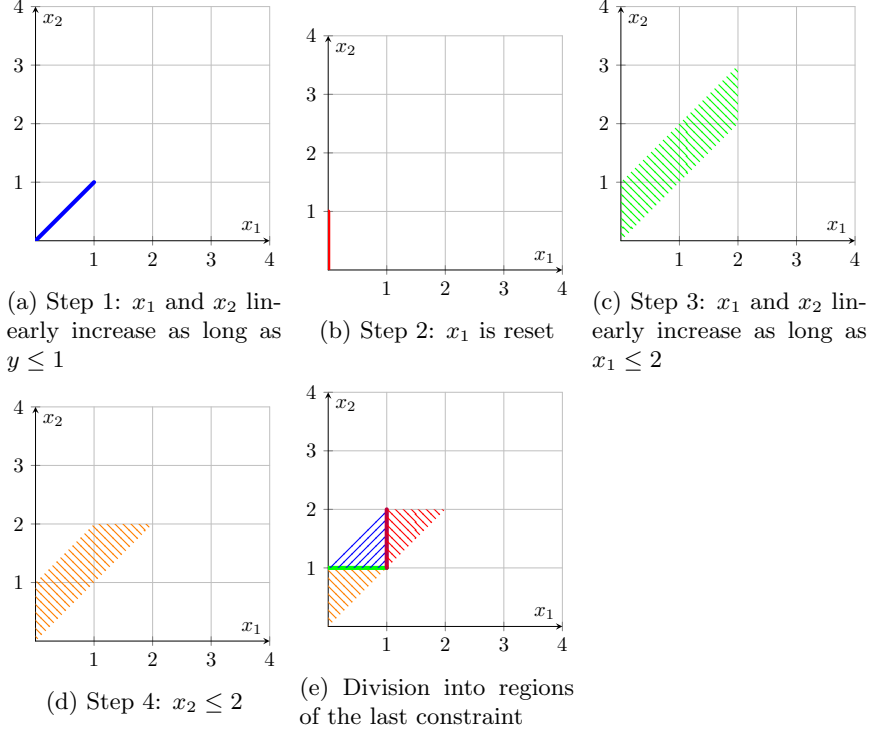


Figure 3: Region abstraction of the TA in Fig. 2

- $\mathcal{R}_{v(\mathcal{A})}$ as the set of vertices
- given two regions $r = [s], r' = [s'] \in \mathcal{R}_{v(\mathcal{A})}$, we have $(r, r') \in \mathcal{F}_{v(\mathcal{A})}$ if one of the following holds:
 - if $s \xrightarrow{e} s' \in T_{v(\mathcal{A})}$ for some $e \in E$ (*discrete instantaneous transition*);
 - if r' is a time successor of r : $r \neq r'$ and there exists d such that $s + d \in r'$ and $\forall d' < d, s + d' \in r \cup r'$ (*delay transition*);
 - $r = r'$ is unbounded: $s = (\ell, \mu)$ with $\mu(x_i) > c_i$ for all x_i (*equivalent unbounded regions*).

Example 5 (Region abstraction). Consider the TA in Fig. 2. We present in Fig. 3 some of its regions. For example, the region depicted in Fig. 3b is the successor (by taking the discrete transition from ℓ_0 to ℓ_1 , resetting x_1) of the region depicted in Fig. 3a.

Consider the valuations μ_1 and μ_2 such that $\mu_1(x_1) = 0$, $\mu_1(x_2) = 0.35$, $\mu_2(x_1) = 0$, and $\mu_2(x_2) = 0.75$. These two valuations satisfy all three conditions of [Definition 5](#), and therefore $\mu_1 \approx \mu_2$. Note that they both belong to the region depicted in [Fig. 3b](#), and any pair of valuations in this region is (by definition) equivalent.

3.3 Symbolic semantics

Let us now recall the symbolic semantics of PTAs (see e. g., [[Hun+02](#); [And+09](#)]).

Constraints We first need to define operations on constraints. A linear term over $\mathbb{X} \cup \mathbb{P}$ is of the form $\sum_{1 \leq i \leq H} \alpha_i x_i + \sum_{1 \leq j \leq M} \beta_j p_j + d$, with $x_i \in \mathbb{X}$, $p_j \in \mathbb{P}$, and $\alpha_i, \beta_j, d \in \mathbb{Z}$. A *constraint* C (i. e., a convex polyhedron²) over $\mathbb{X} \cup \mathbb{P}$ is a conjunction of inequalities of the form $lt \bowtie 0$, where lt is a linear term.

Given a parameter valuation v , $v(C)$ denotes the constraint over \mathbb{X} obtained by replacing each parameter p in C with $v(p)$. Likewise, given a clock valuation μ , $\mu(v(C))$ denotes the expression obtained by replacing each clock x in $v(C)$ with $\mu(x)$. We write $\mu \models v(C)$ whenever $\mu(v(C))$ evaluates to true. We say that v *satisfies* C , denoted by $v \models C$, if the set of clock valuations satisfying $v(C)$ is nonempty. We say that C is *satisfiable* if $\exists \mu, v$ s.t. $\mu \models v(C)$.

We define the *time elapsing* of C , denoted by C^\nearrow , as the constraint over \mathbb{X} and \mathbb{P} obtained from C by delaying all clocks by an arbitrary amount of time. That is,

$$\mu' \models v(C^\nearrow) \text{ if } \exists \mu : \mathbb{X} \rightarrow \mathbb{R}_+, \exists d \in \mathbb{R}_+ \text{ s.t. } \mu \models v(C) \wedge \mu' = \mu + d.$$

Given $R \subseteq \mathbb{X}$, we define the *reset* of C , denoted by $[C]_R$, as the constraint obtained from C by resetting the clocks in R to 0, and keeping the other clocks unchanged. That is,

$$\mu' \models v([C]_R) \text{ if } \exists \mu : \mathbb{X} \rightarrow \mathbb{R}_+ \text{ s.t. } \mu \models v(C) \wedge \forall x \in \mathbb{X} \begin{cases} \mu'(x) = 0 & \text{if } x \in R \\ \mu'(x) = \mu(x) & \text{otherwise.} \end{cases}$$

We denote by $C \downarrow_{\mathbb{P}}$ the projection of C onto \mathbb{P} , i. e., obtained by eliminating the variables not in \mathbb{P} (e. g., using Fourier-Motzkin [[Sch99](#)]).

Definition 7 (Symbolic state). A symbolic state is a pair (ℓ, C) where $\ell \in L$ is a location, and C its associated parametric zone.

Definition 8 (Symbolic semantics). Given a PTA $\mathcal{A} = (\Sigma, L, \ell_0, \ell_f, \mathbb{X}, \mathbb{P}, I, E)$, the symbolic semantics of \mathcal{A} is the labeled transition system called *parametric zone graph* $\mathcal{PZG} = (E, \mathbf{S}, \mathbf{s}_0, \Rightarrow)$, with

- $\mathbf{S} = \{(\ell, C) \mid C \subseteq I(\ell)\}$,
- $\mathbf{s}_0 = (\ell_0, (\bigwedge_{1 \leq i \leq H} x_i = 0)^\nearrow \wedge I(\ell_0))$, and
- $((\ell, C), e, (\ell', C')) \in \Rightarrow$ if $e = (\ell, g, a, R, \ell') \in E$ and

$$C' = (([C \wedge g]_R \wedge I(\ell'))^\nearrow \wedge I(\ell'))$$

with C' satisfiable.

²Strictly speaking, we manipulate *polytopes*, while polyhedra refer to 3-dimensional polytopes. However, for sake of consistency with the parametric timed model checking literature, and with the Parma polyhedra library (among others), we refer to these geometric objects as *polyhedra*.




Figure 4: Parametric zone graph of Fig. 1

That is, in the parametric zone graph, nodes are symbolic states, and arcs are labeled by *edges* of the original PTA.

If $(s, e, s') \in \Rightarrow$, we write $\text{Succ}(s, e) = s'$. By extension, we write $\text{Succ}(s)$ for $\cup_{e \in E} \text{Succ}(s, e)$.

In the non-parametric timed automata setting, a zone can be seen as a *convex union* of regions. In the parametric setting, a parametric zone constrains both clocks and parameters in such a way that, for each admissible parameter valuation, the resulting projection on clocks is a zone.

Example 6. Consider again the PTA \mathcal{A} in Fig. 1. The parametric zone graph of \mathcal{A} is given in Fig. 4, where e_1 is the edge from ℓ_0 to ℓ_1 in Fig. 1, e_2 is the edge from ℓ_0 to ℓ_2 , and e_3 is the edge from ℓ_2 to ℓ_1 . In addition, the symbolic states are:

$$\begin{aligned} \mathbf{s}_0 &= (\ell_0, 0 \leq x \leq 3 \wedge p_1 \geq 0 \wedge p_2 \geq 0) \\ \mathbf{s}_1 &= (\ell_1, x \geq p_2 \wedge 0 \leq p_2 \leq 3 \wedge p_1 \geq 0) \\ \mathbf{s}_2 &= (\ell_2, 3 \geq x \geq p_1 \wedge 0 \leq p_1 \leq 3 \wedge p_2 \geq 0) \\ \mathbf{s}_3 &= (\ell_1, x \geq p_1 \wedge 0 \leq p_1 \leq 3 \wedge p_2 \geq 0) . \end{aligned}$$

3.4 Reachability synthesis

We use reachability synthesis to solve the problems defined in Section 4. This procedure, called EFsynth , takes as input a PTA \mathcal{A} and a set of target locations L_T , and attempts to synthesize all parameter valuations v for which L_T is reachable in $v(\mathcal{A})$. $\text{EFsynth}(\mathcal{A}, L_T)$ was formalized in e. g., [JLR15] and is a procedure that may not terminate, but that computes an exact result (sound and complete) if it terminates. EFsynth traverses the *parametric zone graph* of \mathcal{A} .

Example 7. Consider again the PTA \mathcal{A} in Fig. 1. $\text{EFsynth}(\mathcal{A}, \{\ell_1\}) = p_1 \leq 3 \vee p_2 \leq 3$. Intuitively, it corresponds to all parameter constraints in the parametric zone graph in Fig. 4 associated to symbolic states with location ℓ_1 .

We finally recall the correctness of EFsynth .

Lemma 1 ([JLR15]). *Let \mathcal{A} be a PTA, and let L_T be a subset of the locations of \mathcal{A} . Assume $\text{EFsynth}(\mathcal{A}, L_T)$ terminates with result K . Then $v \models K$ iff L_T is reachable in $v(\mathcal{A})$.*

4 Timed opacity problems

4.1 Definitions

Let us first introduce two key concepts to define our notion of opacity. $D\text{Reach}_\ell^{v(\mathcal{A})}(\ell')$ (resp. $D\text{Reach}_{-\ell}^{v(\mathcal{A})}(\ell')$) is the set of all the durations of the runs for which ℓ is reachable

(resp. unreachable) on the way to ℓ' . Formally:

$$DReach_{\ell}^{v(\mathcal{A})}(\ell') = \{d \mid \exists \rho \text{ in } v(\mathcal{A}) \text{ such that } d = dur(\rho) \wedge \ell \text{ is reached on the way to } \ell' \text{ in } \rho\}$$

and

$$DReach_{-\ell}^{v(\mathcal{A})}(\ell') = \{d \mid \exists \rho \text{ in } v(\mathcal{A}) \text{ such that } d = dur(\rho) \wedge \ell \text{ is avoided on the way to } \ell' \text{ in } \rho\}.$$

These concepts can be seen as the set of execution times from the initial location ℓ_0 to a target location ℓ' while passing (resp. not passing) by a private location ℓ . Observe that, from the definition of $dur(\rho)$, this “execution time” does not include the time spent in ℓ' .

Example 8. Consider again the PTA in Fig. 1, and let v be such that $v(p_1) = 1$ and $v(p_2) = 2$. We have $DReach_{\ell_2}^{v(\mathcal{A})}(\ell_1) = [1, 3]$ and $DReach_{-\ell_2}^{v(\mathcal{A})}(\ell_1) = [2, 3]$.

We now introduce the concept of “timed opacity w.r.t. a set of durations (or execution times) D ”: a system is opaque w.r.t. a given location ℓ_{priv} on the way to ℓ_f for execution times D whenever, for any duration in D , it is not possible to deduce whether the system went through 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 9 (timed opacity w.r.t. D). Given a TA $v(\mathcal{A})$, a private location ℓ_{priv} , a target location ℓ_f and a set of execution times D , we say that $v(\mathcal{A})$ is *opaque w.r.t. ℓ_{priv} on the way to ℓ_f for execution times D* if $D \subseteq DReach_{\ell_{priv}}^{v(\mathcal{A})}(\ell_f) \cap DReach_{-\ell_{priv}}^{v(\mathcal{A})}(\ell_f)$.

If one does not have the ability to tune the system (i. e., change internal delays, or add some `sleep()` or `wait()` statements in the program), one may be first interested in knowing whether the system is opaque for all execution times (i. e., the durations of the runs from the initial location to the target location ℓ_f). In other words, if a system is *fully opaque*, for any possible measured execution time, an attacker is not able to deduce anything on the system, in terms of visit of ℓ_{priv} .

Definition 10 (full timed opacity). Given a TA $v(\mathcal{A})$, a private location ℓ_{priv} and a target location ℓ_f , we say that $v(\mathcal{A})$ is *fully opaque w.r.t. ℓ_{priv} on the way to ℓ_f* if $DReach_{\ell_{priv}}^{v(\mathcal{A})}(\ell_f) = DReach_{-\ell_{priv}}^{v(\mathcal{A})}(\ell_f)$.

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

Remark 1. This definition is symmetric: a system is not opaque iff an attacker can deduce ℓ_{priv} or $\neg\ell_{priv}$. For instance, if there is no path through ℓ_{priv} to ℓ_f , but a path to ℓ_f , a system is not opaque w.r.t. Definition 10.

Example 9. Consider the PTA \mathcal{A} in Fig. 5 where cl is a clock, while ϵ, p are parameters. We use a slightly extended PTA syntax: `read?x` reads the value input on a given channel `read`, and assigns it to a (discrete, global) variable `x`. `secret` is a constant variable of arbitrary value. If both `x` and `secret` are finite-domain variables (e. g., bounded integers) then they can be seen as syntactic sugar for locations. Such variables are supported by most model checkers, including UPPAAL [LPY97] and IMITATOR [And21].




Figure 5: A Java program encoded in a PTA

This PTA encodes a server process and is a (manual) translation of a Java program from the DARPA Space/Time Analysis for Cybersecurity (STAC) library³, that compares a user-input variable with a given secret and performs different actions taking different times depending on this secret. The original Java program is *vulnerable*, and tagged as such in the DARPA library, because some sensitive information can be deduced from the timing information. The original Java code is given in [Appendix A](#).

In our encoding, a single instruction takes a time in $[0, \epsilon]$, while p is a (parametric) factor to one of the `sleep` instructions of the program. Note that in the original Java code in [Appendix A](#), at line 25, there is no parameter p but an integer 2; that is, the code is fixed to have $v(p) = 2$. For sake of simplicity, we abstract away instructions not related to time, and merge subfunctions calls. For this work, we simplify the problem and abstract in this way. Precisely modeling the timing behavior of a program is itself a complicated problem (due to caching, speculative execution, etc.) and we leave to future work.

Let v_1 such that $v_1(\epsilon) = 1$ and $v_1(p) = 2$. For this example, $DReach_{\ell_{priv}}^{v_1(\mathcal{A})}(\ell_f) = [1024, 1029]$ while $DReach_{-\ell_{priv}}^{v_1(\mathcal{A})}(\ell_f) = [2048, 2053]$. Therefore, $v_1(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for execution times $D = [1024, 1029] \cap [2048, 2053] = \emptyset$.

Let v_2 such that $v_2(\epsilon) = 2$ and $v_2(p) = 1.002$. $DReach_{\ell_{priv}}^{v_2(\mathcal{A})}(\ell_f) = [1024, 1034]$ while $DReach_{-\ell_{priv}}^{v_2(\mathcal{A})}(\ell_f) = [1026.048, 1036.048]$. Therefore, $v_2(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for execution times $D = [1026.048, 1034]$.

Obviously, neither $v_1(\mathcal{A})$ nor $v_2(\mathcal{A})$ are fully opaque w.r.t. ℓ_{priv} on the way to ℓ_f .

4.2 Decision and computation problems

We can now define the timed opacity computation problem, which consists in computing the possible execution times ensuring opacity w.r.t. a private location. In other words, the attacker model is as follows: the attacker knows the system model in the form of a TA, and can only observe the computation time between the start of the program and the time it reaches a given (final) location.

Timed opacity computation Problem:

INPUT: A TA $v(\mathcal{A})$, a private location ℓ_{priv} , a target location ℓ_f

PROBLEM: Compute the execution times D such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for these execution times D .

Let us illustrate that this computation problem is certainly not easy. For the TA \mathcal{A} in [Fig. 6](#), the execution times D for which \mathcal{A} is opaque w.r.t. ℓ_{priv} on the way to ℓ_f

³https://github.com/Apogee-Research/STAC/blob/master/Canonical_Examples/Source/Category1_vulnerable.java

is exactly \mathbb{N} ; that is, only integer times are opaque (as the system can only leave ℓ_{priv} and hence enter ℓ_f at an integer time).




Figure 6: TA for which the set of opaque execution times is \mathbb{N}

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

Timed opacity synthesis Problem:

INPUT: A PTA \mathcal{A} , a private location ℓ_{priv} , a target location ℓ_f

PROBLEM: Synthesize the parameter valuations v and the execution times D_v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for these execution times D_v .

Note that the execution times can depend on the parameter valuations.

We can also define the full timed opacity decision problem, which consists in answering whether a timed automaton is fully opaque w.r.t. a private location.

Full timed opacity decision Problem:

INPUT: A TA $v(\mathcal{A})$, a private location ℓ_{priv} , a target location ℓ_f

PROBLEM: Is $v(\mathcal{A})$ fully opaque w.r.t. ℓ_{priv} on the way to ℓ_f ?

Note that a last problem of interest would be *full timed opacity synthesis*, aiming at synthesizing (ideally the entire set of) parameter valuations v for which $v(\mathcal{A})$ is *fully opaque* w.r.t. ℓ_{priv} on the way to ℓ_f . This is left as future work.

5 Timed opacity problems for timed automata

In this section, we address the non-parametric problems defined in Section 4.2, i.e., the timed opacity computation problem (Section 5.2) and full timed opacity decision problem (Section 5.3). We show that both problems can be solved using a construction of the $DReach$ sets based on the RA arithmetic [Wei99] (Section 5.1).

In this section, let \mathcal{A} denote a (non-parametric) timed automaton.

5.1 Computing $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$ and $DReach_{-\ell_{priv}}^{\mathcal{A}}(\ell_f)$

We must be able to express the set of execution times, i.e., the durations of all runs from the initial location to a given target location. While the problem of expressing the set of execution times seems very natural for timed automata, it was barely addressed in the literature, with the exception of [BDR08; Ros19].

5.1.1 The RA arithmetic

We use the RA arithmetic, which is the set of first-order formulae, interpreted over the real numbers, of $\langle \mathbb{R}, +, <, \mathbb{N}, 0, 1 \rangle$ where \mathbb{N} is a unary predicate such that $\mathbb{N}(z)$ is true iff z is a natural number. This arithmetic has a decidable theory with a complexity of 3EXPTIME [Wei99].

5.1.2 Computing execution times of timed automata

With $r, r' \in \mathcal{R}_{\mathcal{A}}$, we denote by $\lambda_{r,r'}$ the set of durations d such that there exists a finite path $\rho = (s_i)_i$ in $T_{\mathcal{A}}$ such that $\text{dur}(\rho) = d$ and the associated path $\pi(\rho) = (r_k)_{0 \leq k \leq K}$ in the region graph satisfies $r_0 = r, r_K = r'$. It is shown in [BDR08, Proposition 5.3] that these sets $\lambda_{r,r'}$ can be defined in RA arithmetic. Moreover, they are definable by a disjunction of terms of the form $d = m, \exists z, \mathbb{N}(z) \wedge d = m + cz$ and $\exists z, \mathbb{N}(z) \wedge m + cz < d < m + cz + 1$, where $c, m \in \mathbb{N}$.

Let us give the main idea of the proof presented in [BDR08] (even though this explanation is not necessary to follow our reasoning). The idea of the proof of [BDR08] is to consider a TA \mathcal{A}^0 obtained from \mathcal{A} by adding a new clock x_0 which is reset to 0 each time it reaches the value 1 and to count all of the resets of x_0 . The construction of \mathcal{A}^0 ensures that each (finite) run ρ of $T_{\mathcal{A}}$ corresponds to a run ρ^0 of $T_{\mathcal{A}^0}$ (at each state, the value of x_0 is the fractional part of the total time elapsed), and conversely (erasing x_0). The authors propose next a classical automaton C as a particular subgraph of the region graph $\mathcal{R}_{\mathcal{A}^0}$, where the only action a denotes the reset of x_0 (all other transitions are labeled with the silent action). The conclusion follows because $t \in \lambda_{r,r'}$ if $\lfloor t \rfloor$ is the length of a path in C , the deterministic automaton obtained from C by subset construction.

Lemma 2 (Reachability-duration computation). *The sets $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$ and $DReach_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f)$ are computable and definable in RA arithmetic.*

Proof. Let \mathcal{A} be a TA. We aim at reducing the computation of the sets $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$ and $DReach_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f)$ to the computation of sets $\lambda_{r,r'}$.

First, let us compute $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$. From \mathcal{A} , we define a TA \mathcal{A}' by adding a Boolean discrete variable b , initially false. Recall that discrete variables over a finite domain are syntactic sugar for *locations*: therefore, ℓ_f with $b = \text{false}$ and ℓ_f with $b = \text{true}$ can be seen as two different locations. Then, we set $b \leftarrow \text{true}$ on any transition whose target location is ℓ_{priv} ; therefore, $b = \text{true}$ denotes that ℓ_{priv} has been visited. We denote by $\ell'_{f, \text{true}}$ the final state of \mathcal{A}' where $b = \text{true}$. $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$ is exactly the set of executions times in \mathcal{A}' between ℓ_0 and $\ell'_{f, \text{true}}$. For all the regions r'_i associated to $\ell'_{f, \text{true}}$, we can compute (using [BDR08]) λ_{r,r'_i} , where r is the region associated to ℓ_0 in \mathcal{A}' . Therefore, $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$ can be computed as the union of all the λ_{r,r'_i} (of which there is a finite number), which is definable in RA arithmetic.

Second, let us compute $DReach_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f)$. We define another TA \mathcal{A}'' obtained from \mathcal{A} by deleting all the transitions leading to ℓ_{priv} . Therefore, the set of durations reaching ℓ_f in \mathcal{A}'' is exactly the set of durations reaching ℓ_f in \mathcal{A} associated to runs not visiting ℓ_{priv} . $DReach_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f)$ is exactly the set of executions times in \mathcal{A}'' between ℓ_0 and ℓ_f . For all the regions r'_i associated to ℓ_f , we can compute λ_{r,r'_i} , where r is the region associated to ℓ_0 in \mathcal{A}'' . Therefore, $DReach_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f)$ can be computed as the union of all the λ_{r,r'_i} . \square

5.2 Answering the timed opacity computation problem

Proposition 1 (timed opacity computation). *The timed opacity computation problem is solvable for TAs.*

Proof. Let \mathcal{A} be a TA. From [Lemma 2](#), we can compute and define in RA arithmetic the sets $DReach_{-\ell_{priv}}^{\mathcal{A}}(\ell_f)$ and $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$.

By the decidability of RA arithmetic, the intersection of these sets is computable. Then, the set $D = DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f) \cap DReach_{-\ell_{priv}}^{\mathcal{A}}(\ell_f)$ is effectively computable (with a 3EXPTIME upper bound). \square

This positive result can be put in perspective with the negative result of [\[Cas09\]](#) that proves that it is undecidable whether a TA (and even the more restricted subclass of event-recording automata [\[AFH99\]](#)) 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 [\[Cas09\]](#).

5.3 Checking for full timed opacity

Proposition 2 (full timed opacity decision). *The full timed opacity decision problem is decidable for TAs.*

Proof. Let \mathcal{A} be a TA. From [Lemma 2](#), we can compute and define in RA arithmetic the sets $DReach_{-\ell_{priv}}^{\mathcal{A}}(\ell_f)$ and $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$.

From the decidability of RA arithmetic, the equality between these sets is decidable. Therefore, $DReach_{\ell_{priv}}^{\mathcal{A}}(\ell_f) \stackrel{?}{=} DReach_{-\ell_{priv}}^{\mathcal{A}}(\ell_f)$ is decidable. \square

From [\[Wei99\]](#) and [\[BDR08, Theorem 7.5\]](#), the computation of a set $\lambda_{r,r'}$ is 2EXPTIME and the RA arithmetic has a decidable theory with complexity 3EXPTIME. Therefore, our construction is 3EXPTIME, which is an upper-bound for the problem complexity. Note that, as in [\[BDR08\]](#), we did not compute a lower bound for the complexity of [Propositions 1](#) and [2](#). This remains to be done as future work.

Example 10. Consider again the PTA \mathcal{A} in [Fig. 1](#), and let v be such that $v(p_1) = 1$ and $v(p_2) = 2$. Recall from [Example 8](#) that $DReach_{\ell_2}^{v(\mathcal{A})}(\ell_1) = [1, 3]$ and $DReach_{-\ell_2}^{v(\mathcal{A})}(\ell_1) = [2, 3]$. Thus, $DReach_{\ell_2}^{v(\mathcal{A})}(\ell_1) \neq DReach_{-\ell_2}^{v(\mathcal{A})}(\ell_1)$ and therefore $v(\mathcal{A})$ is *not (fully) opaque* w.r.t. ℓ_2 on the way to ℓ_1 .

Now, consider v' such that $v'(p_1) = v'(p_2) = 1.5$. This time, $DReach_{\ell_2}^{v'(\mathcal{A})}(\ell_1) = DReach_{-\ell_2}^{v'(\mathcal{A})}(\ell_1) = [1.5, 3]$ and therefore $v'(\mathcal{A})$ is *(fully) opaque* w.r.t. ℓ_2 on the way to ℓ_1 .

6 The theory of parametric timed opacity w.r.t. execution times

We address in this section the parametric problems of timed opacity w.r.t. execution times.

Let us consider the following decision problem, i.e., the problem of checking the *emptiness* of the parameter valuations and execution times set guaranteeing timed




Figure 7: Reduction from reachability-emptiness

opacity w.r.t. execution times. The decision problem associated to *full* timed opacity will be considered in Section 7.

Timed opacity emptiness Problem:

INPUT: A PTA \mathcal{A} , a private location ℓ_{priv} , a target location ℓ_f

PROBLEM: Is the set of valuations v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for a non-empty set of execution times empty?

The negation of the timed opacity emptiness consists in deciding whether there exists at least one parameter valuation for which $v(\mathcal{A})$ is opaque for at least some execution time.

6.1 Undecidability in general

We prove undecidability results for a “sufficient” number of clocks and parameters. Put it differently, our proofs of undecidability require a minimum number of clocks and parameters to work; the problems are obviously undecidable for larger numbers, but the decidability is open for smaller numbers. This will be briefly discussed in Section 10.

With the rule of thumb that all non-trivial decision problems are undecidable for general PTAs [And19], the following result is not surprising, and follows from the undecidability of reachability-emptiness for PTAs.

Theorem 1 (undecidability). *The timed opacity emptiness problem is undecidable for general PTAs.*

Proof. We reduce from the reachability-emptiness problem, i. e., the existence of a parameter valuation for which there exists a run reaching a given location in a PTA, which is undecidable (e. g., [AHV93; Mil00; Doy07; JLR15; Ben+15b]). Consider an arbitrary PTA \mathcal{A} with initial location ℓ_0 and a given location ℓ_f . It is undecidable whether there exists a parameter valuation for which there exists a run reaching ℓ_f (proofs of undecidability in the literature generally reduce from the halting problem of a 2-counter machine which is undecidable [Min67], so one can see \mathcal{A} as an encoding of a 2-counter machine).

Now, add the following locations and transitions (all unguarded) as in Fig. 7: a new urgent⁴ initial location ℓ'_0 with outgoing transitions to ℓ_0 and to a new location ℓ_{pub} ; a new urgent location ℓ_{priv} with an incoming transition from ℓ_f ; a new final location ℓ'_f with incoming transitions from ℓ_{priv} and ℓ_{pub} . Also, ℓ_f is made urgent. Let \mathcal{A}' denote this new PTA.

⁴Where time cannot elapse (depicted in dotted yellow in our figures).

First note that, due to the unguarded transitions, ℓ'_f is reachable for any parameter valuation and for any execution time by runs going through ℓ_{pub} and not going through ℓ_{priv} . That is, for all v , $DReach_{-\ell_{priv}}^{v(\mathcal{A}')}(\ell'_f) = [0, \infty)$.

Assume there exists some parameter valuation v such that ℓ_f is reachable from ℓ_0 in $v(\mathcal{A})$ for some execution times D : then, due to our construction with additional urgent locations, ℓ_{priv} is reachable on the way to ℓ'_f in $v(\mathcal{A}')$ for the exact same execution times D . Therefore, $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ'_f for execution times D .

Conversely, if ℓ_f is not reachable from ℓ_0 in \mathcal{A} for any valuation, then ℓ_{priv} is not reachable on the way to ℓ'_f for any valuation in \mathcal{A}' . Therefore, there is no valuation v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ'_f for any execution time. Therefore, there exists a valuation v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ'_f iff ℓ_f is reachable in \mathcal{A} —which is undecidable. \square

Remark 2. Our proof reduces from the reachability-emptiness problem, for which several undecidability proofs were proposed (notably [AHV93; Mil00; Doy07; JLR15; Ben+15b]), with various flavors (numbers of parameters, integer- or dense-time, integer- or rational-valued parameters, etc.). See [And19] for a survey. Notably, this means (from [Ben+15b]) that [Theorem 1](#) holds for PTAs with at least 3 clocks and a single parameter.

6.2 A decidable subclass

We now show that the timed opacity emptiness problem is decidable for the subclass of PTAs called L/U-PTAs [Hun+02]. Despite early positive results for L/U-PTAs [Hun+02; BL09], more recent results (notably [JLR15; AL17; ALR18; ALM20]) mostly proved undecidable properties of L/U-PTAs, and therefore this positive result is welcome.

Theorem 2 (decidability). *The timed opacity emptiness problem is decidable for L/U-PTAs.*

Proof. We reduce to the timed opacity computation problem of a given TA, which is decidable ([Proposition 1](#)).

Let \mathcal{A} be an L/U-PTA. Let $\mathcal{A}_{0,\infty}$ denote the structure obtained as follows: any occurrence of a lower-bound parameter is replaced with 0, and any occurrence of a conjunct $x \triangleleft p$ (where p is necessarily an upper-bound parameter) is deleted, i.e., replaced with **true**.

Let us show that the set of valuations v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for a non-empty set of execution times is non empty iff the solution to the timed opacity computation problem for $\mathcal{A}_{0,\infty}$ is non-empty.

\Rightarrow Assume there exists a valuation v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for a non-empty set of execution. Therefore, the solution to the timed opacity computation problem for $\mathcal{A}_{0,\infty}$ is non-empty. That is, there exists a duration d such that there exists a run of duration d such that ℓ_{priv} is reachable on the way to ℓ_f , and there exists a run of duration d such that ℓ_{priv} is unreachable on the way to ℓ_f .

We now need the following monotonicity property of L/U-PTAs:

Lemma 3 ([Hun+02]). *Let \mathcal{A} be an L/U-PTA and v be a parameter valuation. Let v' be a valuation such that for each upper-bound parameter p^u , $v'(p^u) \geq v(p^u)$ and for each lower-bound parameter p^l , $v'(p^l) \leq v(p^l)$. Then any run of $v(\mathcal{A})$ is a run of $v'(\mathcal{A})$.*

Therefore, from Lemma 3, the runs of $v(\mathcal{A})$ of duration d such that ℓ_{priv} is reachable (resp. unreachable) on the way to ℓ_f are also runs of $\mathcal{A}_{0,\infty}$. Therefore, there exists a non-empty set of durations such that $\mathcal{A}_{0,\infty}$ is opaque, i. e., solution to the timed opacity computation problem for $\mathcal{A}_{0,\infty}$ is non-empty.

⇐ Assume the solution to the timed opacity computation problem for $\mathcal{A}_{0,\infty}$ is non-empty. That is, there exists a duration d such that there exists a run of duration d such that ℓ_{priv} is reachable on the way to ℓ_f in $\mathcal{A}_{0,\infty}$, and there exists a run of duration d such that ℓ_{priv} is unreachable on the way to ℓ_f in $\mathcal{A}_{0,\infty}$.

The result could follow immediately—if only assigning 0 and ∞ to parameters was a proper parameter valuation. From [Hun+02; BL09], if a location is reachable in the TA obtained by valuating lower-bound parameters with 0 and upper-bound parameters with ∞ , then there exists a sufficiently large constant C such that this run exists in $v(\mathcal{A})$ such that v assigns 0 to lower-bound and C to upper-bound parameters. Here, we can trivially pick $d+1$, as any clock constraint $x \leq d+1$ or $x < d+1$ will be satisfied for a run of duration d . Let v assign 0 to lower-bound and d to upper-bound parameters. Then, there exists a run of duration d such that ℓ_{priv} is reachable on the way to ℓ_f in $v(\mathcal{A})$, and there exists a run of duration d such that ℓ_{priv} is unreachable on the way to ℓ_f in $v(\mathcal{A})$. Therefore, the set of valuations v such that $v(\mathcal{A})$ is opaque w.r.t. ℓ_{priv} on the way to ℓ_f for a non-empty set of execution times is non empty—which concludes the proof. \square

Remark 3. The class of L/U-PTAs is known to be relatively meaningful, and many case studies from the literature fit into this class, including case studies proposed even before this class was defined in [Hun+02], e. g., a toy railroad crossing model and a model of Fischer mutual exclusion protocol given in [AHV93] (see [And19] for a survey). Even though the PTA in Fig. 5 does not fit in this class, it can easily be transformed into an L/U-PTA, by duplicating p into p^l (used in lower-bound comparisons with clocks) and p^u (used in upper-bound comparisons with clocks).

6.3 Intractability of synthesis for L/U-PTAs

Even though the *emptiness* problem for the timed opacity w.r.t. a set of execution times D is decidable for L/U-PTAs (Theorem 2), the *synthesis* of the parameter valuations remains intractable in general, as shown in the following Proposition 3. By intractable, we mean more precisely that the solution, if it can be computed, cannot (in general) 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 answer the question “is the computed solution intersected with an equality test between variables empty?” cannot be used to represent the solution. By empty, we mean emptiness of the valuations set. For example, let us question whether we could represent the solution of the timed opacity 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 [BHZ08] can typically compute the answer to

this question. Therefore, from the following [Proposition 3](#), finite unions of polyhedra cannot be used to represent the solution of the timed opacity 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 [\[JLR15, Theorem 2\]](#)).

Proposition 3 (intractability). *If it can be computed, the solution to the timed opacity synthesis problem for L/U-PTAs cannot in general be represented using any formalism for which the emptiness of the intersection with equality constraints is decidable.*

Proof. We reuse a reasoning inspired by [\[JLR15, Theorem 2\]](#), and we reduce from the undecidability of the timed opacity emptiness problem for general PTAs ([Theorem 1](#)). Assume the solution of the timed opacity synthesis problem for L/U-PTAs can be represented in a formalism for which the emptiness of the intersection with equality constraints is decidable.

Assume an arbitrary PTA \mathcal{A} with notably two locations ℓ_{priv} and ℓ_f . From \mathcal{A} , we define an L/U-PTA \mathcal{A}' as follow: for each parameter p_i that is used both as an upper bound and a lower bound, replace its occurrences as upper bounds by a fresh parameter p_i^u and its occurrences as lower bounds by a fresh parameter p_i^l .

By assumption, the solution of the synthesis problem $\Gamma = \{(v, D_v) \mid v(\mathcal{A}') \text{ is opaque w.r.t. } \ell_{priv} \text{ on the way to } \ell_f \text{ for times } D_v\}$ for \mathcal{A}' can be computed and represented in a formalism for which the emptiness of the intersection with equality constraints is decidable.

However, solving the emptiness of $\{(v, D_v) \in \Gamma \mid \bigwedge_i v(p_i^u) = v(p_i^l)\}$ (which is decidable by assumption), we can decide the timed opacity emptiness for the PTA \mathcal{A} (which is undecidable from [Theorem 1](#)). This leads to a contradiction, and therefore the solution of the timed opacity synthesis problem for L/U-PTAs cannot in general be represented in a formalism for which the emptiness of the intersection with equality constraints is decidable. \square

7 The theory of parametric full timed opacity

We address here the following decision problem, which asks about the emptiness of the parameter valuations set guaranteeing full timed opacity.

Full timed opacity Emptiness Problem:

INPUT: A PTA \mathcal{A} , a private location ℓ_{priv} , a target location ℓ_f

PROBLEM: Is the set of valuations v such that $v(\mathcal{A})$ is fully opaque w.r.t. ℓ_{priv} on the way to ℓ_f empty?

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

7.1 Undecidability in general

Considering that [Theorem 1](#) shows that the undecidability of the timed opacity emptiness problem, the undecidability of the emptiness problem for the full timed opacity is not surprising, but does not follow immediately.

To prove this result (that will be stated formally in [Theorem 3](#)), we first need the following lemma stating that the reachability-emptiness (hereafter sometimes referred




Figure 8: Undecidability of EF-emptiness over constant time

to as “EF-emptiness”) problem is undecidable in constant time, i. e., for a fixed time bound T . That is, the following lemma shows that, given a PTA \mathcal{A} , a target location ℓ_T and a time bound T , it is undecidable whether the set of parameter valuations for which there exists a run reaching ℓ_T in exactly T time units is empty or not.

Lemma 4 (Reachability in constant time). *The reachability-emptiness problem in constant time is undecidable for PTAs with 4 clocks and 2 parameters.*

Proof. In [ALM20, Theorem 3.12], we showed that the EF-emptiness problem is undecidable over bounded time for PTAs with (at least) 3 clocks and 2 parameters. That is, given a fixed bound T and a location ℓ_f , it is undecidable whether the set of parameter valuations for which at least one run reaches ℓ_f within T time units is empty or not.

We reduce the reachability in bounded time (i. e., in at most T time units) to the reachability in constant time (i. e., in exactly T time units). In this proof, we fix $T = 1$.

Assume a PTA \mathcal{A} with a location ℓ_f . We define a PTA \mathcal{A}' as in Fig. 8 by adding a new location ℓ'_f , and a transition from ℓ_f to ℓ'_f guarded by $x = 1$, where x is a new clock (initially 0), not used in \mathcal{A} and therefore never reset in the automaton.

Let us show that there is no valuation such that ℓ_f is reachable in at most 1 time unit iff there is no valuation such that ℓ'_f is reachable exactly in 1 time unit.

- \Leftarrow Assume EF-emptiness holds in constant time for \mathcal{A}' , i. e., there exists no parameter valuation for which ℓ'_f is reachable in $T = 1$ time units. Then, from the construction of \mathcal{A}' , no parameter valuation exists for which ℓ_f is reachable in $T \leq 1$ time units.
- \Rightarrow Conversely, assume EF-emptiness holds over bounded time for \mathcal{A} , i. e., there exists no parameter valuation for which ℓ_f is reachable in $T \leq 1$ time units. Then, from the construction of \mathcal{A}' , no parameter valuation exists for which ℓ'_f is reachable in $T = 1$ time units.

This concludes the proof of the lemma. □

We can now state and prove [Theorem 3](#).

Theorem 3 (undecidability). *The full timed opacity emptiness problem is undecidable for general PTAs with (at least) 4 clocks and 2 parameters.*

Proof. We reduce from the reachability-emptiness problem in constant time, which is undecidable ([Lemma 4](#)).

Consider an arbitrary PTA \mathcal{A} with (at least) 4 clocks and 2 parameters, with initial location ℓ_0 and a given location ℓ_f . We add the following locations and transitions in \mathcal{A} to obtain a PTA \mathcal{A}' , as in [Fig. 9](#): a new urgent initial location ℓ'_0 , with outgoing transition to ℓ_0 and to a new location ℓ_{pub} , a new urgent location ℓ_{priv} with an incoming




Figure 9: Reduction from reachability-emptiness

transition from ℓ_f , a new urgent and final location ℓ'_f with incoming transitions from ℓ_{priv} and ℓ_{pub} , and a guard $x = 1$ (with a new clock x , never reset) on the transition from ℓ_{pub} to ℓ'_f .

First, note that, due to the guarded transition, ℓ'_f is reachable for any parameter valuation and (only) for an execution time equal to 1 by runs going through ℓ_{pub} and not going through ℓ_{priv} . That is, for all v , $DReach_{-\ell_{priv}}^{v(\mathcal{A})}(\ell_f) = \{1\}$.

We show that there exists a valuation v such that $v(\mathcal{A}')$ is fully opaque w.r.t. ℓ_{priv} on the way to ℓ'_f iff there exists a valuation v such that ℓ_f is reachable in $v(\mathcal{A})$ for execution time equal to 1.

- \Leftarrow Assume there exists some valuation v such that ℓ_f is reachable from ℓ_0 in \mathcal{A} (only) for execution time equal to 1. Then, due to our construction, ℓ_{priv} is reachable on the way to ℓ'_f in $v(\mathcal{A}')$ for the exact same execution time 1. Therefore, $v(\mathcal{A}')$ is fully opaque w.r.t. ℓ_{priv} on the way to ℓ'_f .
- \Rightarrow Conversely, if ℓ_f is not reachable from ℓ_0 in \mathcal{A} for any valuation for execution time 1, then ℓ_{priv} is not reachable on the way to ℓ'_f for any valuation of \mathcal{A}' . Therefore, there is no valuation v such that $v(\mathcal{A}')$ is fully opaque w.r.t. ℓ_{priv} on the way to ℓ'_f for execution time 1.

Therefore, there exists a valuation v such that $v(\mathcal{A}')$ is fully opaque w.r.t. ℓ_{priv} on the way to ℓ'_f iff there exists a valuation v such that ℓ_f is reachable in $v(\mathcal{A})$ for execution time equal to 1—which is undecidable. This concludes the proof.

Let us briefly discuss the minimum number of clocks necessary to obtain undecidability using our proof (the case of smaller numbers of clocks remains open). Recall that [Lemma 4](#) needs 4 clocks; in the current proof of [Theorem 3](#), we add a new clock x which is never reset; however, since the proof of [Lemma 4](#) also uses a clock which is never reset, therefore we can reuse it, and our proof does not need an additional clock. So the result holds for 4 clocks and 2 parameters. \square

7.2 Undecidability for L/U-PTAs

Note that reasoning like in [Section 6.2](#), i.e., reducing the emptiness problem to a decision problem of the non-parametric $\mathcal{A}_{0,\infty}$, is not relevant. [Fig. 10](#) shows an L/U-PTA \mathcal{A} (and more precisely, a U-PTA, i.e., an L/U-PTA with an empty set of lower-bound parameters [[BL09](#)]) which is not fully opaque for any parameter valuation, but whose associated TA $\mathcal{A}_{0,\infty}$ is.

In addition, 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




Figure 10: $\mathcal{A}_{0,\infty}$ is not sufficient for the full timed opacity emptiness problem




Figure 11: No monotonicity for full timed opacity in L/U-PTAs

parameter preserves reachability”) as recalled in [Lemma 3](#), we can show in the following example that this is not the case for full timed opacity.

Example 11. Consider the PTA in [Fig. 11](#). First assume v such that $v(p) = 0.5$. Then, $v(\mathcal{A})$ is not full timed opaque: indeed, ℓ_f can be reached in 1 time unit via ℓ_{priv} , but not without going through ℓ_{priv} .

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

Finally, let us enlarge p further, and assume v'' such that $v''(p) = 2$. Then, $v''(\mathcal{A})$ becomes again not full timed opaque: indeed, ℓ_f can be reached in 2 time units not going through ℓ_{priv} , but cannot be reached in 2 time units by going through ℓ_{priv} .

As a side note, remark that this PTA is actually a U-PTA, that is, monotonicity for this problem does not even hold for U-PTAs.

In fact, we show that, while the timed opacity emptiness problem is decidable for L/U-PTAs ([Theorem 2](#)), the full timed opacity emptiness problem becomes undecidable for this same class (from 4 parameters). This confirms (after previous works in [[JLR15](#); [AL17](#); [ALR18](#)]) that L/U-PTAs stand at the frontier between decidability and undecidability.

Theorem 4. *The full timed opacity emptiness problem is undecidable for L/U-PTAs with (at least) 4 clocks and 4 parameters.*

Proof. Let us recall from [[ALM20](#), Theorem 3.12] that the EF-emptiness problem is undecidable over bounded time for PTAs with (at least) 3 clocks and 2 parameters. Assume a PTA \mathcal{A} with 3 clocks and 2 parameters, say p_1 and p_2 , and a target location ℓ_f . Take 1 as a time bound. From [[ALM20](#), Theorem 3.12], it is undecidable whether there exists a parameter valuation for which ℓ_f is reachable in time ≤ 1 .

The idea of our proof is that, as in [[JLR15](#); [And+19](#)], we “split” each of the two parameters used in \mathcal{A} 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 construction of the form $x < p_i$ (resp. $x \leq p_i$) is replaced

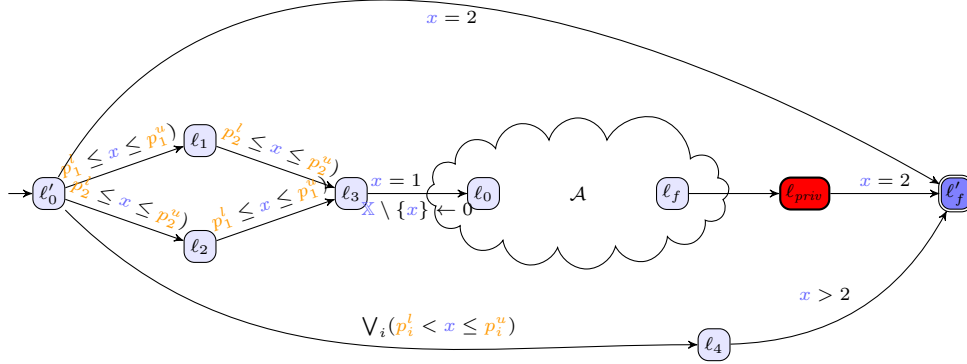


Figure 12: Undecidability of full timed opacity emptiness for L/U-PTAs

with $x < p_i^u$ (resp. $x \leq p_i^u$) while each construction 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{A} 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 [ALM20, Theorem 3.12].

Now, consider the extension of \mathcal{A} given in Fig. 12, and let \mathcal{A}' be this extension. We assume that x is an extra clock not used in \mathcal{A} . The syntax “ $\mathbb{X} \setminus \{x\} \leftarrow 0$ ” denotes that all clocks of the original PTA \mathcal{A} are reset—but not the new clock x . The guard on the lower 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 upper transition from l'_0 to l'_f at time 2, i. e., l'_f is always reachable in time 2 without going through location l_{priv} ;
2. the original automaton \mathcal{A} 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. if a run reaches l_{priv} on the way to l'_f , then its duration is necessarily 2;
4. from [ALM20, 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:

1. if $p_1^l > p_1^u$ 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{A} (and therefore to reach l_{priv} on the way to l'_f); since these valuations can still reach l'_f for some execution times (notably $x = 2$ through the upper transition from l'_0 to l'_f), then \mathcal{A}' is not fully opaque for any of these valuations.

2. if $p_1^l < p_1^u$ or $p_2^l < p_2^u$, then one of the lower transitions from ℓ'_0 to ℓ_4 can be taken, and therefore ℓ'_f is reachable in a time > 2 without going through ℓ_{priv} . Since no run can reach ℓ'_f while going through ℓ_{priv} for a duration $\neq 2$, then again \mathcal{A}' is not fully opaque for any of these valuations.
3. if $p_1^l = p_1^u$ and $p_2^l = p_2^u$, then the behavior of the modified \mathcal{A} (with duplicate parameters) is exactly the one of the original \mathcal{A} . Also, note that the lower transitions from ℓ'_0 to ℓ'_f (via ℓ_4) cannot be taken. In contrast, the upper transition from ℓ'_0 to ℓ'_f can still be taken, and therefore there exists a run of duration 2 reaching ℓ'_f without visiting ℓ_{priv} .

Now, assume there exists a parameter valuation for which there exists a run of \mathcal{A} of duration ≤ 1 reaching ℓ_f . And, as a consequence, ℓ_{priv} is reachable, and therefore there exists some run of duration 2 (including the 1 time unit to go from ℓ_0 to ℓ'_0) reaching ℓ'_f after visiting ℓ_{priv} . From the above reasoning, all runs reaching ℓ'_f have duration 2; in addition, we exhibited a run visiting ℓ_{priv} and a run not visiting ℓ_{priv} ; therefore the modified automaton \mathcal{A}' is fully opaque for such a parameter valuation.

Conversely, assume there exists no parameter valuation for which there exists a run of \mathcal{A} of duration ≤ 1 reaching ℓ_f . In that case, \mathcal{A}' is not fully opaque for any parameter valuation.

As a consequence, there exists a parameter valuation v' for which $v'(\mathcal{A}')$ is fully opaque iff there exists a parameter valuation v for which there exists a run in $v(\mathcal{A})$ of duration ≤ 1 reaching ℓ_f —which is undecidable from [ALM20, Theorem 3.12]. \square

8 Parameter synthesis for opacity

Despite the negative theoretical result of [Theorem 1](#), we now address the timed opacity 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 can be summarized as follows.

1. We enrich 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;
3. We perform reachability-synthesis using EFsynth on ℓ_f with contradictory values of b .

We detail each operation in the following.

In this section, we assume a PTA \mathcal{A} , a given private location ℓ_{priv} and a given final location ℓ_f .

8.1 Enriching the PTA

We first add a Boolean flag b initially set to false, and then set to true on any transition whose target location is ℓ_{priv} (in the line of the proof of [Proposition 1](#)). Therefore, $b = \text{true}$ denotes that ℓ_{priv} has been visited. Second, we add a synchronization action `finish` on any transition whose target location is ℓ_f . Third, we add a new clock x_{abs} (never reset) together with a new parameter p_{abs} , and we guard all transitions to ℓ_f




Figure 13: Transformed version of Fig. 5

with $x_{abs} = p_{abs}$. This will allow to measure the (parametric) execution time. Let $\text{Enrich}(\mathcal{A}, \ell_{priv}, \ell_f)$ denote this procedure.

Example 12. Fig. 13 shows the transformed version of the PTA in Fig. 5.

8.2 Self-composition

We use here the principle of *self-composition* [BDR11], i. e., composing the PTA with a copy of itself. More precisely, given a PTA $\mathcal{A}' = \text{Enrich}(\mathcal{A}, \ell_{priv}, \ell_f)$, we first perform an identical copy of \mathcal{A}' with *distinct variables*: that is, a clock x of \mathcal{A}' is distinct from a clock x in the copy of \mathcal{A}' —which can be trivially performed using variable renaming.⁵ Let $\text{Copy}(\mathcal{A}')$ denote this copy of \mathcal{A}' . We then compute $\mathcal{A}' \parallel_{\{\text{finish}\}} \text{Copy}(\mathcal{A}')$. That is, \mathcal{A}' and $\text{Copy}(\mathcal{A}')$ evolve completely independently due to the interleaving—except that they are forced to enter ℓ_f at the same time, thanks to the synchronization action *finish*.

8.3 Synthesis

Then, we apply reachability synthesis EFsynth (over all parameters, i. e., the “internal” timing parameters, but also the p_{abs} parameter) to the following goal location: the original \mathcal{A}' is in ℓ_f with $b = \text{true}$ while its copy $\text{Copy}(\mathcal{A}')$ is in ℓ'_f with $b' = \text{false}$ (primed variables denote variables from the copy). Intuitively, we synthesize timing parameters and execution times such that there exists a run reaching ℓ_f with $b = \text{true}$ (i. e., that has visited ℓ_{priv}) and there exists another run of same duration reaching ℓ_f with $b = \text{false}$ (i. e., that has not visited ℓ_{priv}).

Let $\text{SynthOp}(\mathcal{A}, \ell_{priv}, \ell_f)$ denote the entire procedure. We formalize SynthOp in Algorithm 1, where “ $\ell_f \wedge b = \text{true}$ ” denotes the location ℓ_f with $b = \text{true}$. Recall that p_{abs} is added by the enrichment step described in Section 8.1. The set of execution times D is therefore given by the possible valuations of p_{abs} ; these valuations may depend on the model timing parameters (in the form of a constraint). Finally note that EFsynth is called on a set made of a single location of $\mathcal{A}' \parallel_{\{\text{finish}\}} \text{Copy}(\mathcal{A}')$; by definition of the synchronous product, this location is a *pair* of locations, one from \mathcal{A}' (i. e., “ $\ell_f \wedge b = \text{true}$ ”) and one from $\text{Copy}(\mathcal{A}')$ (i. e., “ $\ell'_f \wedge b' = \text{false}$ ”).

⁵In fact, the fresh clock x_{abs} and parameter p_{abs} can be shared to save two variables, as x_{abs} is never reset, and both PTAs enter ℓ_f at the same time, therefore both “copies” of x_{abs} and p_{abs} always share the same values.

Algorithm 1: SynthOp($\mathcal{A}, \ell_{priv}, \ell_f$)

input : A PTA \mathcal{A} with parameters set \mathbb{P} , locations ℓ_{priv}, ℓ_f
output: Parameter constraint K over $\mathbb{P} \cup \{p_{abs}\}$

- 1 $\mathcal{A}' \leftarrow \text{Enrich}(\mathcal{A}, \ell_{priv}, \ell_f)$
 - 2 $\mathcal{A}'' \leftarrow \mathcal{A}' \parallel_{\{\text{finish}\}} \text{Copy}(\mathcal{A}')$
 - 3 **return** EFSynth($\mathcal{A}'', \{(\ell_f \wedge b = \text{true}, \ell'_f \wedge b' = \text{false})\}$)
-

Example 13. Consider again the PTA \mathcal{A} in Fig. 5: its enriched version \mathcal{A}' is given in Fig. 13. Fix $v(\epsilon) = 1$, $v(p) = 2$. We then perform the synthesis applied to the self-composition of \mathcal{A}' according to Algorithm 1. The result obtained with IMITATOR is: $p_{abs} = \emptyset$ (as expected from Example 9).

Now fix $v(\epsilon) = 2$, $v(p) = 1.002$. We obtain: $p_{abs} \in [1026.048, 1034]$ (again, as expected from Example 9).

Now let us keep all parameters unconstrained. The result of Algorithm 1 is the following 3-dimensional constraint:

$$\begin{aligned} & 5 \times \epsilon + 1024 \geq p_{abs} \geq 1024 \\ \wedge \quad & 1024 \times p + 5 \times \epsilon \geq p_{abs} \geq 1024 \times p \geq 0 \end{aligned}$$

8.4 Correctness

We will state below that, whenever SynthOp($\mathcal{A}, \ell_{priv}, \ell_f$) terminates, then its result is an exact (sound and complete) answer to the timed opacity synthesis problem.

Let us first prove a technical lemma used later to prove the soundness of SynthOp.

Lemma 5. *Assume SynthOp($\mathcal{A}, \ell_{priv}, \ell_f$) terminates with result K . For all $v \models K$, there exists a run ending in ℓ_f at time $v(p_{abs})$ in $v(\mathcal{A})$.*

Proof. From the construction of the procedure Enrich, we added a new clock x_{abs} (never reset) together with a new parameter p_{abs} , and we guarded all transitions to ℓ_f with $x_{abs} = p_{abs}$. Therefore, valuations of p_{abs} correspond exactly to the times at which ℓ_f can be reached in $v(\mathcal{A})$. \square

We can now prove soundness and completeness.

Proposition 4 (soundness). *Assume SynthOp($\mathcal{A}, \ell_{priv}, \ell_f$) terminates with result K . For all $v \models K$, there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is reachable on the way to ℓ_f in $v(\mathcal{A})$ and there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is avoided on the way to ℓ_f in $v(\mathcal{A})$.*

Proof. SynthOp($\mathcal{A}, \ell_{priv}, \ell_f$) is the result of EFSynth called on the self-composition of Enrich($\mathcal{A}, \ell_{priv}, \ell_f$). Recall that Enrich has enriched \mathcal{A} with the addition of a guard $x_{abs} = p_{abs}$ on the incoming transitions of ℓ_f , as well as a Boolean flag b that is true iff ℓ_{priv} was visited along a run. Assume $v \models K$. From Lemma 1, there exists a run of \mathcal{A}'' reaching $\ell_f \wedge b = \text{true}, \ell'_f \wedge b' = \text{false}$. From Lemma 5, this run takes $v(p_{abs})$ time units. From the self-composition that is made of interleaving only (except for the final synchronization), there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is reachable on the way to ℓ_f in $v(\mathcal{A})$ and there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is avoided on the way to ℓ_f in $v(\mathcal{A})$. \square

Proposition 5 (completeness). *Assume $\text{SynthOp}(\mathcal{A}, \ell_{priv}, \ell_f)$ terminates with result K . Assume v . Assume there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is reachable on the way to ℓ_f in $v(\mathcal{A})$ and there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is avoided on the way to ℓ_f in $v(\mathcal{A})$. Then $v \models K$.*

Proof. Assume $\text{SynthOp}(\mathcal{A}, \ell_{priv}, \ell_f)$ terminates with result K . Assume v . Assume there exists a run ρ of duration $v(p_{abs})$ such that ℓ_{priv} is reachable on the way to ℓ_f in $v(\mathcal{A})$ and there exists a run ρ' of duration $v(p_{abs})$ such that ℓ_{priv} is avoided on the way to ℓ_f in $v(\mathcal{A})$.

First, from **Enrich**, there exists a run ρ of duration $v(p_{abs})$ such that ℓ_{priv} is reachable (resp. avoided) on the way to ℓ_f in $v(\mathcal{A})$ implies that there exists a run ρ of duration $v(p_{abs})$ such that $\ell_f \wedge b = \text{true}$ (resp. $b = \text{false}$) is reachable in $v(\text{Enrich}(\mathcal{A}))$.

Since our self-composition allows any interleaving, runs ρ of $v(\mathcal{A}')$ and ρ' in $v(\text{Copy}(\mathcal{A}'))$ are independent—except for reaching ℓ_f . Since ρ and ρ' have the same duration $v(p_{abs})$, then they both reach ℓ_f at the same time and, from our definition of self-composition, they can simultaneously fire action `finish` and enter ℓ_f at time $v(p_{abs})$. Hence, there exists a run reaching $\ell_f \wedge b = \text{true}$, $\ell_f \wedge b' = \text{false}$ in $v(\mathcal{A}'')$.

Finally, from **Lemma 1**, $v \models K$. \square

Theorem 5 (correctness). *Assume $\text{SynthOp}(\mathcal{A}, \ell_{priv}, \ell_f)$ terminates with result K . Assume v . The following two statements are equivalent:*

1. *There exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is reachable on the way to ℓ_f in $v(\mathcal{A})$ and there exists a run of duration $v(p_{abs})$ such that ℓ_{priv} is avoided on the way to ℓ_f in $v(\mathcal{A})$.*
2. $v \models K$.

Proof. From **Propositions 4** and **5** \square

9 Experiments

9.1 Experimental environment

We use IMITATOR [And21], a parametric timed model checking tool taking as input networks of PTAs extended with several handful features such as shared global discrete variables, PTA synchronization through strong broadcast, non-timing rational-valued parameters, etc. IMITATOR supports various parameter synthesis algorithms, including reachability synthesis. IMITATOR represents symbolic states as polyhedra, relying on PPL [BHZ08]. IMITATOR is a leading tool for parameter synthesis for extensions of parametric timed automata. Related tools are Romeo [Lim+09] (which cannot be used here, as it does not support parametric timed automata, but extensions of Petri nets), SPACEEX [Fre+11] (which does not perform parameter synthesis), or UPPAAL [LPY97] (which cannot be used here, as our algorithm requires timing parameters, not supported by UPPAAL).

We ran experiments using IMITATOR 2.10.4 “Butter Jellyfish” (build 2477 HEAD/5b53333) on a Dell XPS 13 9360 equipped with an Intel[®] Core™ i7-7500U CPU @ 2.70GHz with 8 GiB memory running Linux Mint 18.3 64 bits.⁶

⁶Sources, models and results are available at doi.org/10.5281/zenodo.3251141 and imitator.fr/static/ATVA19/.

9.2 Translating programs into PTAs

We will consider case studies from the PTA community and from previous works focusing on privacy using (parametric) timed automata. In addition, we will be 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 timed automata is highly non-trivial due to complication of hardware pipelining, caching, OS scheduling, etc. The readers are referred to the rich literature in, for instance, [Lv+10]. In this work, we instead make the following simplistic assumption on execution time of a program statement and focus on solving the parameter synthesis problem. How to precisely model the execution time of programs is orthogonal and complementary to our work.

We assume that the execution time of a program statement other than `Thread.sleep(n)` is within a range $[0, \epsilon]$ where ϵ is a small integer constant (in milliseconds), whereas the execution time of statement `Thread.sleep(n)` is within a range $[n, n + \epsilon]$. 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.⁷ 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.

9.3 A richer framework

The symbolic representation of variables and parameters in IMITATOR allows us to reason *symbolically* concerning variables. That is, instead of enumerating all possible (bounded) values of `x` and `secret` in Fig. 5, we turn them to parameters (i. e., unknown constants), and IMITATOR performs a symbolic reasoning. Even better, the analysis terminates for this example even when no bound is provided on these variables. This is often not possible in (non-parametric) timed automata based model checkers, which usually have to enumerate these values. Therefore, in our PTA representation of Java programs, we turn all user-input variable and secret constant variables to non-timing rational-valued parameters, also supported by IMITATOR. Other local variables are implemented using IMITATOR discrete (shared, global) variables.

We also discuss how to enlarge the scope of our framework.

Multiple private locations This can be easily achieved by setting `b` to true along any incoming transition of one of these private locations.

Multiple final locations The technique used depends on whether these multiple final locations can be distinguished or not. If they are indistinguishable (i. e., the observer knows when the program has terminated, but not in which state), then it suffices to merge all these final locations in a single one, and our framework trivially applies. If they are distinguishable, then one analysis needs to be conducted on each of these locations (with a different parameter p_{abs} for each of these), and the obtained constraints must be intersected.

⁷<https://github.com/Apogee-Research/STAC/>




Figure 14: [VNN18, Fig. 5]

Access to high-level variables In the literature, a distinction is sometimes made between low-level (“public”) and high-level (“private”) variables. Opacity or non-interference can be defined in terms of the ability for an observer to deduce some information on the high-level variables.

Example 14. For example, in Fig. 14 (where cl is a clock and h a variable), if ℓ_2 is reachable in 20 time units, then it is clear that the value of the high-level variable h is negative.

Our framework can also be used to address this problem, e. g., by setting b to true, not on locations but on selected tests / valuations of such variables.

Example 15. For example, setting b to true on the upper transition from ℓ_1 to ℓ_2 in Fig. 14, the answer to the timed opacity computation problem is $D = (30, \infty)$, and the system is therefore not opaque since ℓ_2 can be reached for any execution time in $[0, \infty)$.

9.4 Experiments

9.4.1 Benchmarks

As a proof of concept, we applied our method to a set of examples from the literature. The first five models come from previous works from the literature [GMR07; Ben+15a; VNN18], also addressing non-interference or opacity in timed automata.⁸ In addition, we used two common models from the (P)TA literature, not necessarily linked to security: a toy coffee machine (*Coffee*) used as benchmark in a number of papers, and a model Fischer’s mutual exclusion protocol (*Fischer-HRSV02*) [Hun+02]. In both cases, we added manually a definition of private location (the number of sugars ordered, and the identity of the process entering the critical section, respectively), and we verified whether they are opaque w.r.t. these internal behaviors.

We also applied our approach to a set of Java programs from the aforementioned STAC library. We use identifiers of the form *STAC:1:n* where **1** denotes the identifier in the library, while **n** (resp. **v**) denotes non-vulnerable (resp. vulnerable). We manually translated these programs to PTAs, following the method described in Section 9.2. We used a representative set of programs from the library; however, some of them were too complex to fit in our framework, notably when the timing leaks come from calls to external libraries (*STAC:15:v*), when dealing with complex computations such as operations on matrices (*STAC:16:v*) or when handling probabilities (*STAC:18:v*). Proposing efficient and accurate ways to represent arbitrary programs into (parametric) timed automata is orthogonal to our work, and is the object of future works.

⁸As most previous works on opacity and timed automata do not come with an implementation nor with benchmarks, it is not easy to find larger models coming in the form of TAs.

Table 1: Experiments: timed opacity

Model			Transf. PTA			Result	
Name	$ \mathcal{A} $	$ \mathbb{X} $	$ \mathcal{A} $	$ \mathbb{X} $	$ \mathbb{P} $	Time (s)	Opaque?
[VNN18, Fig. 5]	1	1	2	3	3	0.02	(×)
[GMR07, Fig. 1b]	1	1	2	3	1	0.04	(×)
[GMR07, Fig. 2a]	1	1	2	3	1	0.05	(×)
[GMR07, Fig. 2b]	1	1	2	3	1	0.02	(×)
Web privacy problem [Ben+15a]	1	2	2	4	1	0.07	(×)
Coffee	1	2	2	5	1	0.05	✓
Fischer-HSRV02	3	2	6	5	1	5.83	(×)
STAC:1:n			2	3	6	0.12	(×)
STAC:1:v			2	3	6	0.11	×
STAC:3:n			2	3	8	0.72	✓
STAC:3:v			2	3	8	0.74	(×)
STAC:4:n			2	3	8	6.40	×
STAC:4:v			2	3	8	265.52	×
STAC:5:n			2	3	6	0.24	✓
STAC:11A:v			2	3	8	47.77	(×)
STAC:11B:v			2	3	8	59.35	(×)
STAC:12c:v			2	3	8	18.44	×
STAC:12e:n			2	3	8	0.58	×
STAC:12e:v			2	3	8	1.10	(×)
STAC:14:n			2	3	8	22.34	(×)

9.4.2 Timed opacity computation

First, we *verified* whether a given TA model is opaque, i. e., if for all execution times reaching a given final location, both an execution goes through a given private location and an execution does not go through this private location. To this end, we also answer the timed opacity computation problem, i. e., to synthesize all execution times for which the system is opaque. While this problem can be verified on the region graph (Proposition 1), we use the same framework as in Section 8, but without parameters in the original TA. That is, we use the Boolean flag b and the parameter p_{abs} to compute all possible execution times. In other words, we use a parametric analysis to solve a non-parametric problem.

We tabulate the experiments results in Table 1. We give from left to right the model name, the numbers of automata and of clocks in the original timed automaton (this information is not relevant for Java programs as the original model is not a TA), the numbers of automata⁹, of clocks and of parameters in the transformed PTA, the computation time in seconds (for the timed opacity computation problem), and the result. In the result column, “✓” (resp. “×”) denotes that the model is opaque (resp. is not opaque, i. e., vulnerable), while “(×)” denotes that the model is not opaque, but could be fixed. That is, although $DReach_{\ell_{priv}}^{v(\mathcal{A})}(\ell_f) \neq DReach_{\neg\ell_{priv}}^{v(\mathcal{A})}(\ell_f)$, their intersection is non-empty and therefore, by tuning the computation time, it may be possible to make the system opaque. This will be discussed in Section 9.5.

Even though we are interested here in timed opacity computation (and not in synthesis), note that all models derived from Java programs feature the parameter ϵ .

⁹As usual, it may be simpler to write PTA models as a network of PTAs. Recall from Definition 3 that a network of PTAs gives a PTA. In this case, $|\mathcal{A}|$ denotes the number of input PTA components.

The result is obtained by variable elimination, i. e., by existential quantification over the parameters different from p_{abs} . In addition, the number of parameters is increased by the parameters encoding the symbolic variables (such as x and $secret$ in Fig. 5).

Concerning the Java programs, we decided to keep the most abstract representation, by imposing that each instruction lasts for a time in $[0, \epsilon]$, with ϵ a parameter. However, fixing an identical (parametric) time ϵ for all instructions, or fixing an arbitrary time in a constant interval $[0, \epsilon]$ (for some constant ϵ , e. g., 1), or even fixing an identical (constant) time ϵ (e. g., 1) for all instructions, significantly speeds up the analysis. These choices can be made for larger models.

Discussion Overall, our method is able to answer the timed opacity computation problem for practical case studies, exhibiting which execution times are opaque (timed opacity computation problem), and whether *all* execution times indeed guarantee opacity (timed opacity problem).

In many cases, while the system is not opaque, we are able to *infer* the execution times guaranteeing opacity (cells marked “(×)”). This is an advantage of our method w.r.t. methods outputting only binary answers.

We observed some mismatches in the Java programs, i. e., some of the programs marked n (non-vulnerable) in the library are actually vulnerable according to our method. This mainly comes from the fact that the STAC library expect tools to use imprecise analyses on the execution times, while we use an exact method. Therefore, a very small mismatch between $DReach_{\ell_{priv}}^{v(\mathcal{A})}(\ell_f)$ and $DReach_{-\ell_{priv}}^{v(\mathcal{A})}(\ell_f)$ will lead our algorithm to answer “not opaque”, while some methods may not be able to differentiate this mismatch from imprecision (noise). This is notably the case of **STAC:14:n** where some action lasts either 5,010,000 or 5,000,000 time units depending on some secret, which our method detects to be different, while the library does not. For **STAC:1:n**, using our data, the difference in the execution time upper bound between an execution performing some secret action and an execution not performing it is larger than 1%, which we believe is a value which is not negligible, and therefore this case study might be considered as vulnerable.

STAC:4:n requires a more detailed discussion. This particular program is targeting vulnerabilities that can be detected easily *when they accumulate*, typically in loops. This program checks a number of times (10) a user-input password, and each password check is made in the most insecure way, i. e., by returning “incorrect” as soon as one character differs between the input password and the expected password. This way is very insecure because the execution time is proportional to the number of consecutive correct characters in the input password and, by observing the execution time, an attacker can guess how many characters are correct, and therefore using a limited number of tests, (s)he will eventually guess the correct password. The difference between the vulnerable (**STAC:4:v**) and the non-vulnerable (**STAC:4:n**) versions is that the non-vulnerable version immediately stops if the password is incorrect, and performs the 10 checks only if the password is correct. Therefore, while the computation time is very different between the correct input password and any incorrect input password, it is however very similar between an incorrect input password that would only be incorrect because, say, of the last character (e. g., “**kouignamaz**” while the expected password is “**kouignaman**”), and a completely incorrect input password differing as early as the first character (e. g., “**andouille**”). This makes the attacker’s task very difficult. The main reason for the STAC library to label **STAC:4:n** as a non-vulnerable program is because of the “very similar” nature of the computation times between an

Table 2: Experiments: timed opacity synthesis

Model				Transf. PTA			Result	
Name	$ \mathcal{A} $	$ \mathbb{X} $	$ \mathbb{P} $	$ \mathcal{A} $	$ \mathbb{X} $	$ \mathbb{P} $	Time (s)	Constraint
[VNN18, Fig. 5]	1	1	0	2	3	4	0.02	K
[GMR07, Fig. 1b]	1	1	0	2	3	3	0.03	K
[GMR07, Fig. 2]	1	1	0	2	3	3	0.05	K
Web privacy problem [Ben+15a]	1	2	2	2	4	3	0.07	K
Coffee	1	2	3	2	5	4	0.10	\top
Fischer-HSRV02	3	2	2	6	5	3	7.53	K
STAC:3:v			2	2	3	9	0.93	K

incorrect input password that would only be incorrect because of the last character, and a completely incorrect input password. (In contrast, the vulnerable version `STAC:4:v` is completely vulnerable because this time difference is amplified by the loop, here 10 times.) While “very similar” might be acceptable for most tools, in our setting based on formal verification, we *do* detect that testing “kouignamaz” or testing “kouignamzz” will yield a slightly faster computation time for the second input, because the first incorrect letter occurs earlier—and the program is therefore vulnerable.

9.4.3 Timed opacity synthesis

Then, we address the timed opacity synthesis problem. In this case, we *synthesize* both the execution time and the internal values of the parameters for which one cannot deduce private information from the execution time.

We consider the same case studies as for timed opacity computation; however, the Java programs feature no internal “parameter” and cannot be used here. Still, as a proof of concept, we artificially enriched one of them (`STAC:3:v`) as follows: in addition to the parametric value of ϵ and the execution time, we parameterized one of the `sleep` timers. The resulting constraint can help designers to refine this latter value to ensure opacity. Note that it may not be that easy to tune a Java program to make it non-opaque: while this is reasonably easy on the PTA level (restraining the execution times using an additional clock), this may not be clear on the original model: Making a program terminate slower than originally is easy with a `Sleep` statement; but making it terminate “earlier” is less obvious, as it may mean an abrupt termination, possibly leading to wrong results.

We tabulate the results in Table 2, where the columns are similar to Table 1. A difference is that the first $|\mathbb{P}|$ column denotes the number of parameters in the original model (without counting these added by our transformation). In addition, Table 2 does not contain a “vulnerable?” column as we *synthesize* the condition for which the model is non-vulnerable, and therefore the answer is non-binary. However, in the last column (“Constraint”), we make explicit whether no valuations ensure opacity (“ \top ”), all of them (“ K ”), or some of them (“ K ”).

Discussion An interesting outcome is that the computation time is comparable to the (non-parametric) timed opacity computation, with an increase of up to 20% only. In addition, for all case studies, we exhibit at least some valuations for which the system can be made opaque. Also note that our method always terminates for these models, and therefore the result exhibited is complete. Interestingly, `Coffee` is opaque for any valuation of the 3 internal parameters.

9.5 “Repairing” a non-opaque PTA

Our method gives a result in time of a union of polyhedra over the internal timing parameters and the execution time. On the one hand, we believe tuning the internal timing parameters should be easy: for a program, an internal timing parameter can be the duration of a `sleep`, for example. On the other hand, tuning the execution time of a program may be more subtle. A solution is to enforce a minimal execution time by adding a second thread in parallel with a `wait()` primitive to ensure a minimal execution time. Ensuring a *maximal* execution time can be achieved with an exception stopping the program after a given time; however there is a priori no guarantee that the result of the computation is correct.

10 Conclusion

In this work, we proposed an approach based on parametric timed model checking to not only decide whether the model of a timed system can be subject to timing information leakage, but also to *synthesize* internal timing parameters and execution times that render the system opaque. We implemented our approach in a framework based on IMITATOR, and performed experiments on case studies from the literature and from a library of Java programs.

We now discuss future works in the following.

Theory We proved decidability of the timed opacity computation problem ([Proposition 1](#)) and of the full timed opacity decision problem ([Proposition 2](#)) for TAs, but we only provided an upper bound (3EXPTIME) on the complexity. It can be easily shown that these problems are at least PSPACE, but the exact complexity remains to be exhibited.

In addition, the decidability of several “low-dimensional” problems (i.e., with “small” number of clocks or parameters) remains open. Among these, the one-clock case for parametric timed opacity emptiness ([Theorem 1](#)) remains open: that is, is the timed opacity emptiness problem decidable for PTAs using a single clock? Our method in [Section 8](#) consists in duplicating the automaton and adding a clock that is never reset, thus resulting in a PTA with 3 clocks, for which reachability-emptiness is undecidable [[AHV93](#)]. However, since one of the clocks is never reset, and since the automaton is structurally constrained (it is the result of the composition of two copies of the same automaton), decidability might be envisioned. Recall that the 2-clock reachability-emptiness problem is a famous open problem [[And19](#)], despite recent advances, notably over discrete time [[BO14](#); [GH21](#)]. The 1-clock question also remains open for full timed opacity emptiness ([Theorem 3](#)). The minimum number of parameters required for our proof of the undecidability of the full timed opacity emptiness problem for PTAs (resp. L/U-PTAs) to work is 2 (resp. 4), as seen in [Theorem 3](#) (resp. [Theorem 4](#)); it is open whether using less parameters can render these problems decidable.

Finally, concerning L/U-PTAs, we proved two negative results, despite the decidability of the timed opacity emptiness problem ([Theorem 2](#)): the undecidability of the full timed opacity emptiness ([Theorem 4](#)) and the intractability of timed opacity synthesis ([Proposition 3](#)). It remains open whether these results still apply to the more restrictive class of U-PTAs [[BL09](#)].

Full timed opacity synthesis We leave full timed opacity synthesis as future work; while we could certainly reuse partially our algorithm, this is not entirely trivial, as we need to select only the parameter valuations for which the whole set of execution times is exactly the set of opaque times.

Applications The translation of the STAC library required some non-trivial creativity: while the translation from programs to quantitative extensions of automata is orthogonal to our work, proposing automated translations of (possibly annotated) programs to timed automata dedicated to timing analysis is on our agenda.

Adding probabilities to our framework will be interesting, helping to quantify the execution times of “untimed” instructions in program with a finer grain than an interval; also note that some benchmarks make use of probabilities (notably [STAC:18:v](#)).

Finally, IMITATOR is a general model checker, not specifically aimed at solving the problem we address here. Notably, constraints managed by PPL contain all variables (clocks, timing parameters, and parameters encoding symbolic variables of programs), yielding an exponential complexity. Separating certain types of independent variables (typically parameters encoding symbolic variables of programs, and other variables) should increase efficiency.

References

- [Abb+16] Imran Hafeez Abbasi, Faiq Khalid Lodhi, Awais Mehmood Kamboh, and Osman Hasan. “Formal Verification of Gate-Level Multiple Side Channel Parameters to Detect Hardware Trojans”. In: *FTSCS* (Nov. 14, 2016). Ed. by Cyrille Artho and Peter Csaba Ölveczky. Vol. 694. Communications in Computer and Information Science. Tokyo, Japan, 2016, pp. 75–92. DOI: [10.1007/978-3-319-53946-1_5](#) (cit. on p. 6).
- [AD94] Rajeev Alur and David L. Dill. “A theory of timed automata”. In: *Theoretical Computer Science* 126.2 (Apr. 1994), pp. 183–235. DOI: [10.1016/0304-3975\(94\)90010-8](#) (cit. on pp. 2, 7, 9).
- [AFH99] Rajeev Alur, Limor Fix, and Thomas A. Henzinger. “Event-Clock Automata: A Determinizable Class of Timed Automata”. In: *Theoretical Computer Science* 211.1-2 (1999), pp. 253–273. DOI: [10.1016/S0304-3975\(97\)00173-4](#) (cit. on pp. 4, 17).
- [Aga00] Johan Agat. “Transforming Out Timing Leaks”. In: *POPL* (Jan. 19–21, 2000). Boston, Massachusetts, USA, 2000, pp. 40–53. DOI: [10.1145/325694.325702](#) (cit. on p. 5).
- [AHV93] Rajeev Alur, Thomas A. Henzinger, and Moshe Y. Vardi. “Parametric real-time reasoning”. In: *STOC* (May 16–18, 1993). Ed. by S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal. San Diego, California, United States: ACM, 1993, pp. 592–601. DOI: [10.1145/167088.167242](#) (cit. on pp. 3, 6, 18–20, 35).
- [AK20] Étienne André and Aleksander Kryukov. “Parametric non-interference in timed automata”. In: *ICECCS* (Mar. 4–6, 2021). Ed. by Yi Li and Alan Liew. Singapore, 2020, pp. 37–42. DOI: [10.1109/ICECCS51672.2020.00012](#) (cit. on p. 4).
- [AL17] Étienne André and Didier Lime. “Liveness in L/U-Parametric Timed Automata”. In: *ACSD* (June 25–30, 2017). Ed. by Alex Legay and Klaus Schneider. Zaragoza, Spain: IEEE, 2017, pp. 9–18. DOI: [10.1109/ACSD.2017.19](#) (cit. on pp. 19, 24).

- [ALM20] Étienne André, Didier Lime, and Nicolas Markey. “Language Preservation Problems in Parametric Timed Automata”. In: *Logical Methods in Computer Science* 16.1 (Jan. 2020). DOI: [10.23638/LMCS-16\(1:5\)2020](https://doi.org/10.23638/LMCS-16(1:5)2020) (cit. on pp. 19, 22, 24–26).
- [ALR18] Étienne André, Didier Lime, and Mathias Ramparison. “TCTL model checking lower/upper-bound parametric timed automata without invariants”. In: *FORMATS* (Sept. 4–6, 2018). Ed. by David N. Jansen and Pavithra Prabhakar. Vol. 11022. Lecture Notes in Computer Science. Beijing, China: Springer, 2018, pp. 1–17. DOI: [10.1007/978-3-030-00151-3_3](https://doi.org/10.1007/978-3-030-00151-3_3) (cit. on pp. 19, 24).
- [Amm+21] Ikhlass Ammar, Yamen El Touati, Moez Yeddes, and John Mullins. “Bounded opacity for timed systems”. In: *Journal of Information Security and Applications* 61 (Sept. 2021), pp. 1–13. DOI: [10.1016/j.jisa.2021.102926](https://doi.org/10.1016/j.jisa.2021.102926) (cit. on p. 4).
- [And+09] Étienne André, Thomas Chatain, Emmanuelle Encrenaz, and Laurent Fribourg. “An Inverse Method for Parametric Timed Automata”. In: *International Journal of Foundations of Computer Science* 20.5 (2009), pp. 819–836. DOI: [10.1142/S0129054109006905](https://doi.org/10.1142/S0129054109006905) (cit. on p. 11).
- [And+19] Étienne André, Vincent Bloemen, Laure Petrucci, and Jaco van de Pol. “Minimal-Time Synthesis for Parametric Timed Automata”. In: *TACAS, Part II* (Apr. 8–11, 2019). Ed. by Tomáš Vojnar and Lijun Zhang. Vol. 11428. Lecture Notes in Computer Science. Prague, Czech Republic: Springer, 2019, pp. 211–228. DOI: [10.1007/978-3-030-17465-1_12](https://doi.org/10.1007/978-3-030-17465-1_12) (cit. on p. 24).
- [And19] Étienne André. “What’s decidable about parametric timed automata?” In: *International Journal on Software Tools for Technology Transfer* 21.2 (Apr. 2019), pp. 203–219. DOI: [10.1007/s10009-017-0467-0](https://doi.org/10.1007/s10009-017-0467-0) (cit. on pp. 18–20, 35).
- [And21] Étienne André. “IMITATOR 3: Synthesis of timing parameters beyond decidability”. In: *CAV* (July 18–23, 2021). Ed. by Rustan Leino and Alexandra Silva. Vol. 12759. Lecture Notes in Computer Science. virtual: Springer, 2021, pp. 1–14. DOI: [10.1007/978-3-030-81685-8_26](https://doi.org/10.1007/978-3-030-81685-8_26) (cit. on pp. 13, 29).
- [AS19] Étienne André and Jun Sun. “Parametric Timed Model Checking for Guaranteeing Timed Opacity”. In: *ATVA* (Oct. 28–31, 2019). Ed. by Yu-Fang Chen, Chih-Hong Cheng, and Javier Esparza. Vol. 11781. Lecture Notes in Computer Science. Taipei, Taiwan: Springer, 2019, pp. 115–130. DOI: [10.1007/978-3-030-31784-3_7](https://doi.org/10.1007/978-3-030-31784-3_7) (cit. on p. 3).
- [Bar+02] Roberto Barbuti, Nicoletta De Francesco, Antonella Santone, and Luca Tesei. “A Notion of Non-Interference for Timed Automata”. In: *Fundamenta Informaticae* 51.1-2 (2002), pp. 1–11 (cit. on p. 4).
- [BB07] Andrew Bortz and Dan Boneh. “Exposing private information by timing Web applications”. In: *WWW* (May 8–12, 2007). Ed. by Carey L. Williamson, Mary Ellen Zurko, Peter F. Patel-Schneider, and Prashant J. Shenoy. Banff, Alberta, Canada: ACM, 2007, pp. 621–628. DOI: [10.1145/1242572.1242656](https://doi.org/10.1145/1242572.1242656) (cit. on p. 4).
- [BDR08] Véronique Bruyère, Emmanuel Dall’Olio, and Jean-Francois Raskin. “Durations and parametric model-checking in timed automata”. In: *ACM Transactions on Computational Logic* 9.2 (2008), 12:1–12:23. DOI: [10.1145/1342991.1342996](https://doi.org/10.1145/1342991.1342996) (cit. on pp. 5, 9, 15–17).
- [BDR11] Gilles Barthe, Pedro R. D’Argenio, and Tamara Rezk. “Secure information flow by self-composition”. In: *Mathematical Structures in Computer Science* 21.6 (2011), pp. 1207–1252. DOI: [10.1017/S0960129511000193](https://doi.org/10.1017/S0960129511000193) (cit. on p. 27).
- [Ben+15a] Gilles Benattar, Franck Cassez, Didier Lime, and Olivier H. Roux. “Control and synthesis of non-interferent timed systems”. In: *International Journal of Control* 88.2 (2015), pp. 217–236. DOI: [10.1080/00207179.2014.944356](https://doi.org/10.1080/00207179.2014.944356) (cit. on pp. 4, 5, 31, 32, 34).

- [Ben+15b] Nikola Beneš, Peter Bezděk, Kim Gulstrand Larsen, and Jiří Srba. “Language Emptiness of Continuous-Time Parametric Timed Automata”. In: *ICALP, Part II* (July 6–10, 2015). Ed. by Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann. Vol. 9135. Lecture Notes in Computer Science. Kyoto, Japan: Springer, July 2015, pp. 69–81. DOI: [10.1007/978-3-662-47666-6_6](https://doi.org/10.1007/978-3-662-47666-6_6) (cit. on pp. 18, 19).
- [BHZ08] Roberto Bagnara, Patricia M. Hill, and Enea Zaffanella. “The Parma Polyhedra Library: Toward a Complete Set of Numerical Abstractions for the Analysis and Verification of Hardware and Software Systems”. In: *Science of Computer Programming* 72.1–2 (2008), pp. 3–21. DOI: [10.1016/j.scico.2007.08.001](https://doi.org/10.1016/j.scico.2007.08.001) (cit. on pp. 20, 29).
- [BL09] Laura Bozzelli and Salvatore La Torre. “Decision problems for lower/upper bound parametric timed automata”. In: *Formal Methods in System Design* 35.2 (2009), pp. 121–151. DOI: [10.1007/s10703-009-0074-0](https://doi.org/10.1007/s10703-009-0074-0) (cit. on pp. 19, 20, 23, 35).
- [BO14] Daniel Bundala and Joël Ouaknine. “Advances in Parametric Real-Time Reasoning”. In: *MFCS, Part I* (Aug. 25–29, 2014). Ed. by Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik. Vol. 8634. Lecture Notes in Computer Science. Budapest, Hungary: Springer, 2014, pp. 123–134. DOI: [10.1007/978-3-662-44522-8](https://doi.org/10.1007/978-3-662-44522-8) (cit. on p. 35).
- [Bou+21] Patricia Bouyer, Léo Henry, Samy Jaziri, Thierry Jéron, and Nicolas Markey. “Diagnosing timed automata using timed markings”. In: *International Journal on Software Tools for Technology Transfer* 23.2 (2021), pp. 229–253. DOI: [10.1007/s10009-021-00606-2](https://doi.org/10.1007/s10009-021-00606-2) (cit. on p. 5).
- [BR07] Véronique Bruyère and Jean-François Raskin. “Real-Time Model-Checking: Parameters everywhere”. In: *Logical Methods in Computer Science* 3.1:7 (2007), pp. 1–30. DOI: [10.2168/LMCS-3\(1:7\)2007](https://doi.org/10.2168/LMCS-3(1:7)2007) (cit. on p. 5).
- [BRW06] Gilles Barthe, Tamara Rezk, and Martijn Warnier. “Preventing Timing Leaks Through Transactional Branching Instructions”. In: *Electronic Notes in Theoretical Computer Science* 153.2 (2006), pp. 33–55. DOI: [10.1016/j.entcs.2005.10.031](https://doi.org/10.1016/j.entcs.2005.10.031) (cit. on p. 5).
- [Bry+08] Jeremy W. Bryans, Maciej Koutny, Laurent Mazaré, and Peter Y. A. Ryan. “Opacity generalised to transition systems”. In: *International Journal of Information Security* 7.6 (2008), pp. 421–435. DOI: [10.1007/s10207-008-0058-x](https://doi.org/10.1007/s10207-008-0058-x) (cit. on pp. 2, 4).
- [BT03] Roberto Barbuti and Luca Tesei. “A Decidable Notion of Timed Non-Interference”. In: *Fundamenta Informaticae* 54.2-3 (2003), pp. 137–150 (cit. on p. 4).
- [Cas09] Franck Cassez. “The Dark Side of Timed Opacity”. In: *ISA* (June 25–27, 2009). Ed. by Jong Hyuk Park, Hsiao-Hwa Chen, Mohammed Atiquzzaman, Changhoon Lee, Tai-Hoon Kim, and Sang-Soo Yeo. Vol. 5576. Lecture Notes in Computer Science. Seoul, Korea: Springer, 2009, pp. 21–30. DOI: [10.1007/978-3-642-02617-1_3](https://doi.org/10.1007/978-3-642-02617-1_3) (cit. on pp. 4, 17).
- [CJM16] Duc-Hiep Chu, Joxan Jaffar, and Rasool Maghareh. “Precise Cache Timing Analysis via Symbolic Execution”. In: *RTAS* (Apr. 11–14, 2016). Vienna, Austria: IEEE Computer Society, 2016, pp. 293–304. DOI: [10.1109/RTAS.2016.7461358](https://doi.org/10.1109/RTAS.2016.7461358) (cit. on p. 6).
- [Cop+09] Bart Coppens, Ingrid Verbauwhede, Koen De Bosschere, and Bjorn De Sutter. “Practical Mitigations for Timing-Based Side-Channel Attacks on Modern x86 Processors”. In: *S&P* (May 17–20, 2009). Oakland, California, USA, 2009, pp. 45–60. DOI: [10.1109/SP.2009.19](https://doi.org/10.1109/SP.2009.19) (cit. on p. 5).

- [CR11] Sudipta Chattopadhyay and Abhik Roychoudhury. “Scalable and Precise Refinement of Cache Timing Analysis via Model Checking”. In: *RTSS* (Nov. 29–Dec. 2, 2011). Vienna, Austria, 2011, pp. 193–203. DOI: [10.1109/RTSS.2011.25](https://doi.org/10.1109/RTSS.2011.25) (cit. on p. 6).
- [Doy+15] Goran Doychev, Boris Köpf, Laurent Mauborgne, and Jan Reineke. “CacheAudit: A Tool for the Static Analysis of Cache Side Channels”. In: *ACM Transactions on Information and System Security* 18.1 (2015), 4:1–4:32. DOI: [10.1145/2756550](https://doi.org/10.1145/2756550) (cit. on p. 6).
- [Doy07] Laurent Doyen. “Robust Parametric Reachability for Timed Automata”. In: *Information Processing Letters* 102.5 (2007), pp. 208–213. DOI: [10.1016/j.ip1.2006.11.018](https://doi.org/10.1016/j.ip1.2006.11.018) (cit. on pp. 18, 19).
- [DSF16] Louise A. Dennis, Marija Slavkovic, and Michael Fisher. “‘How Did They Know?’ – Model-Checking for Analysis of Information Leakage in Social Networks”. In: *COIN@AAMAS* (May 9, 2016). Ed. by Stephen Cranefield, Samhar Mahmoud, Julian A. Padget, and Ana Paula Rocha. Vol. 10315. Lecture Notes in Computer Science. Singapore: Springer, 2016, pp. 42–59. DOI: [10.1007/978-3-319-66595-5_3](https://doi.org/10.1007/978-3-319-66595-5_3) (cit. on p. 6).
- [Fre+11] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. “SpaceX: Scalable Verification of Hybrid Systems”. In: *CAV* (July 14–20, 2011). Ed. by Ganesh Gopalakrishnan and Shaz Qadeer. Vol. 6806. Lecture Notes in Computer Science. Snowbird, UT, USA: Springer, 2011, pp. 379–395. DOI: [10.1007/978-3-642-22110-1_30](https://doi.org/10.1007/978-3-642-22110-1_30) (cit. on pp. 4, 29).
- [FS00] Edward W. Felten and Michael A. Schneider. “Timing attacks on Web privacy”. In: *CCS* (Nov. 1–4, 2000). Ed. by Dimitris Gritzalis, Sushil Jajodia, and Pierangela Samarati. Athens, Greece: ACM, 2000, pp. 25–32. DOI: [10.1145/352600.352606](https://doi.org/10.1145/352600.352606) (cit. on p. 4).
- [GH21] Stefan Göller and Mathieu Hilaire. “Reachability in Two-Parametric Timed Automata with One Parameter Is EXPSPACE-Complete”. In: *STACS* (Mar. 16–19, 2021). Ed. by Markus Bläser and Benjamin Monmege. Vol. 187. LIPIcs. Saarbrücken, Germany: Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 36:1–36:18. DOI: [10.4230/LIPIcs.STACS.2021.36](https://doi.org/10.4230/LIPIcs.STACS.2021.36) (cit. on p. 35).
- [GMR07] Guillaume Gardey, John Mullins, and Olivier H. Roux. “Non-Interference Control Synthesis for Security Timed Automata”. In: *Electronic Notes in Theoretical Computer Science* 180.1 (2007), pp. 35–53. DOI: [10.1016/j.entcs.2005.05.046](https://doi.org/10.1016/j.entcs.2005.05.046) (cit. on pp. 4, 7, 31, 32, 34).
- [GSB18] Christopher Gerking, David Schubert, and Eric Bodden. “Model Checking the Information Flow Security of Real-Time Systems”. In: *ESSoS* (June 26–27, 2018). Ed. by Mathias Payer, Awais Rashid, and Jose M. Such. Vol. 10953. Lecture Notes in Computer Science. Paris, France: Springer, 2018, pp. 27–43. DOI: [10.1007/978-3-319-94496-8_3](https://doi.org/10.1007/978-3-319-94496-8_3) (cit. on p. 4).
- [GWW18] Shengjian Guo, Meng Wu, and Chao Wang. “Adversarial symbolic execution for detecting concurrency-related cache timing leaks”. In: *ESEC/SIGSOFT FSE* (Nov. 4–9, 2018). Ed. by Gary T. Leavens, Alessandro Garcia, and Corina S. Pasareanu. Lake Buena Vista, FL, USA: ACM, 2018, pp. 377–388. DOI: [10.1145/3236024.3236028](https://doi.org/10.1145/3236024.3236028) (cit. on p. 6).
- [HMP91] Thomas A. Henzinger, Zohar Manna, and Amir Pnueli. “Temporal Proof Methodologies for Real-time Systems”. In: *POPL* (Jan. 21–23, 1991). Ed. by David S. Wise. Orlando, Florida, USA: ACM Press, 1991, pp. 353–366. DOI: [10.1145/99583.99629](https://doi.org/10.1145/99583.99629) (cit. on p. 8).
- [Hun+02] Thomas Hune, Judi Romijn, Mariëlle Stoelinga, and Frits W. Vaandrager. “Linear parametric model checking of timed automata”. In: *Journal of Logic and Algebraic Programming* 52-53 (2002), pp. 183–220. DOI: [10.1016/S1567-8326\(02\)00037-1](https://doi.org/10.1016/S1567-8326(02)00037-1) (cit. on pp. 7, 11, 19, 20, 31).

- [JLR15] Aleksandra Jovanović, Didier Lime, and Olivier H. Roux. “Integer Parameter Synthesis for Real-Time Systems”. In: *IEEE Transactions on Software Engineering* 41.5 (2015), pp. 445–461. DOI: [10.1109/TSE.2014.2357445](https://doi.org/10.1109/TSE.2014.2357445) (cit. on pp. 9, 12, 18, 19, 21, 24).
- [Koc+20] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. “Spectre attacks: Exploiting speculative execution”. In: *Communications of the ACM* 63.7 (2020), pp. 93–101. DOI: [10.1145/3399742](https://doi.org/10.1145/3399742) (cit. on p. 5).
- [Koc96] Paul C. Kocher. “Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems”. In: *CRYPTO* (Aug. 18–22, 1996). Ed. by Neal Koblitz. Vol. 1109. Lecture Notes in Computer Science. Santa Barbara, California, USA: Springer, 1996, pp. 104–113. DOI: [10.1007/3-540-68697-5_9](https://doi.org/10.1007/3-540-68697-5_9) (cit. on p. 4).
- [Kot+13] Robert Kotcher, Yutong Pei, Pranjal Jumde, and Collin Jackson. “Cross-origin pixel stealing: timing attacks using CSS filters”. In: *CCS* (Nov. 4–8, 2013). Ed. by Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung. Berlin, Germany: ACM, 2013, pp. 1055–1062. DOI: [10.1145/2508859.2516712](https://doi.org/10.1145/2508859.2516712) (cit. on p. 4).
- [Lim+09] Didier Lime, Olivier H. Roux, Charlotte Seidner, and Louis-Marie Traonouez. “Romeo: A Parametric Model-Checker for Petri Nets with Stopwatches”. In: *TACAS* (Mar. 22–29, 2009). Ed. by Stefan Kowalewski and Anna Philippou. Vol. 5505. Lecture Notes in Computer Science. York, United Kingdom: Springer, Mar. 2009, pp. 54–57. DOI: [10.1007/978-3-642-00768-2_6](https://doi.org/10.1007/978-3-642-00768-2_6) (cit. on p. 29).
- [LPY97] Kim Guldstrand Larsen, Paul Pettersson, and Wang Yi. “UPPAAL in a Nutshell”. In: *International Journal on Software Tools for Technology Transfer* 1.1-2 (1997), pp. 134–152. DOI: [10.1007/s1000900050010](https://doi.org/10.1007/s1000900050010) (cit. on pp. 4, 13, 29).
- [Lv+10] Mingsong Lv, Wang Yi, Nan Guan, and Ge Yu. “Combining Abstract Interpretation with Model Checking for Timing Analysis of Multicore Software”. In: *RTSS* (Nov. 30–Dec. 3, 2010). San Diego, California, USA: IEEE Computer Society, 2010, pp. 339–349. DOI: [10.1109/RTSS.2010.30](https://doi.org/10.1109/RTSS.2010.30) (cit. on p. 30).
- [Mil00] Joseph S. Miller. “Decidability and Complexity Results for Timed Automata and Semi-linear Hybrid Automata”. In: *HSCC* (Mar. 23–25, 2000). Ed. by Nancy A. Lynch and Bruce H. Krogh. Vol. 1790. Lecture Notes in Computer Science. Pittsburgh, PA, USA: Springer, 2000, pp. 296–309. DOI: [10.1007/3-540-46430-1_26](https://doi.org/10.1007/3-540-46430-1_26) (cit. on pp. 18, 19).
- [Min67] Marvin L. Minsky. *Computation: Finite and infinite machines*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1967. ISBN: 0-13-165563-9 (cit. on p. 18).
- [Mol+05] David Molnar, Matt Piotrowski, David Schultz, and David A. Wagner. “The Program Counter Security Model: Automatic Detection and Removal of Control-Flow Side Channel Attacks”. In: *ICISC*. Seoul, Korea, 2005, pp. 156–168. DOI: [10.1007/11734727_14](https://doi.org/10.1007/11734727_14) (cit. on p. 5).
- [NNV17] Flemming Nielson, Hanne Riis Nielson, and Panagiotis Vasilikos. “Information Flow for Timed Automata”. In: *Models, Algorithms, Logics and Tools*. Ed. by Luca Aceto, Giorgio Bacci, Giovanni Bacci, Anna Ingólfssdóttir, Axel Legay, and Radu Mardare. Vol. 10460. Lecture Notes in Computer Science. Springer, 2017, pp. 3–21. DOI: [10.1007/978-3-319-63121-9_1](https://doi.org/10.1007/978-3-319-63121-9_1) (cit. on p. 4).
- [Ros19] Amnon Rosenmann. “The Timestamp of Timed Automata”. In: *FORMATS* (Aug. 27–29, 2019). Ed. by Étienne André and Mariëlle Stoelinga. Vol. 11750. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, 2019, pp. 181–198. DOI: [10.1007/978-3-030-29662-9_11](https://doi.org/10.1007/978-3-030-29662-9_11) (cit. on pp. 5, 15).

- [Sch99] Alexander Schrijver. *Theory of linear and integer programming*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1999. ISBN: 978-0-471-98232-6 (cit. on p. 11).
- [SPW18] Chunga Sung, Brandon Paulsen, and Chao Wang. “CANAL: a cache timing analysis framework via LLVM transformation”. In: *ASE* (Sept. 3–Mar. 7, 2018). Montpellier, France, 2018, pp. 904–907. DOI: [10.1145/3238147.3240485](https://doi.org/10.1145/3238147.3240485) (cit. on p. 6).
- [VNN18] Panagiotis Vasilikos, Flemming Nielson, and Hanne Riis Nielson. “Secure Information Release in Timed Automata”. In: *POST* (Apr. 14–20, 2018). Ed. by Lujo Bauer and Ralf Küsters. Vol. 10804. Lecture Notes in Computer Science. Thessaloniki, Greece: Springer, 2018, pp. 28–52. DOI: [10.1007/978-3-319-89722-6_2](https://doi.org/10.1007/978-3-319-89722-6_2) (cit. on pp. 4, 31, 32, 34).
- [Wei99] Volker Weispfenning. “Mixed Real-Integer Linear Quantifier Elimination”. In: *ISSAC* (July 29–31, 1999). Ed. by Keith O. Geddes, Bruno Salvy, and Samuel S. Dooley. Vancouver, BC, Canada: Association for Computing Machinery, 1999, pp. 129–136. DOI: [10.1145/309831.309888](https://doi.org/10.1145/309831.309888) (cit. on pp. 3, 15–17).
- [WS17] Chao Wang and Patrick Schaumont. “Security by compilation: an automated approach to comprehensive side-channel resistance”. In: *SIGLOG News* 4.2 (2017), pp. 76–89. DOI: [10.1145/3090064.3090071](https://doi.org/10.1145/3090064.3090071) (cit. on p. 5).
- [Wu+18] Meng Wu, Shengjian Guo, Patrick Schaumont, and Chao Wang. “Eliminating timing side-channel leaks using program repair”. In: *ISSTA* (July 16–21, 2018). Amsterdam, The Netherlands, 2018, pp. 15–26. DOI: [10.1145/3213846.3213851](https://doi.org/10.1145/3213846.3213851) (cit. on p. 5).
- [WZ18] Lingtai Wang and Naijun Zhan. “Decidability of the Initial-State Opacity of Real-Time Automata”. In: *Symposium on Real-Time and Hybrid Systems - Essays Dedicated to Professor Chaochen Zhou on the Occasion of His 80th Birthday*. Ed. by Cliff B. Jones, Ji Wang, and Naijun Zhan. Vol. 11180. Lecture Notes in Computer Science. Springer, 2018, pp. 44–60. DOI: [10.1007/978-3-030-01461-2_3](https://doi.org/10.1007/978-3-030-01461-2_3) (cit. on p. 5).
- [WZA18] Lingtai Wang, Naijun Zhan, and Jie An. “The Opacity of Real-Time Automata”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37.11 (2018), pp. 2845–2856. DOI: [10.1109/TCAD.2018.2857363](https://doi.org/10.1109/TCAD.2018.2857363) (cit. on p. 5).
- [Zha+18] Jun Zhang, Pengfei Gao, Fu Song, and Chao Wang. “SCInfer: Refinement-Based Verification of Software Countermeasures Against Side-Channel Attacks”. In: *CAV, Part II* (July 14–17, 2018). Ed. by Hana Chockler and Georg Weissenbacher. Vol. 10982. Lecture Notes in Computer Science. Oxford, UK, 2018, pp. 157–177. DOI: [10.1007/978-3-319-96142-2_12](https://doi.org/10.1007/978-3-319-96142-2_12) (cit. on p. 6).

A The code of the Java example

```
1 import java.io.BufferedReader;
2 import java.io.IOException;
3 import java.io.InputStreamReader;
4 import java.io.PrintWriter;
5 import java.net.ServerSocket;
6 import java.net.Socket;
7
8 //Coefficients are Disregarded
9 public class Category1-vulnerable {
10     private static final int port = 8000;
11     private static final int secret = 1234;
12     private static final int n = 32;
13     private static ServerSocket server;
14
15     private static void checkSecret(int guess) throws InterruptedException {
16         if (guess <= secret) {
17             for (int i = 0; i < n; i++) {
18                 for (int t = 0; t < n; t++) {
19                     Thread.sleep(1);
20                 }
21             }
22         } else {
23             for (int i = 0; i < n; i++) {
24                 for (int t = 0; t < n; t++) {
25                     Thread.sleep(2);
26                 }
27             }
28         }
29     }
30
31     private static void startServer() {
32         try {
33             server = new ServerSocket(port);
34             System.out.println("Server Started Port: " + port);
35             Socket client;
36             PrintWriter out;
37             BufferedReader in;
38             String userInput;
39             int guess;
40             while (true) {
41                 client = server.accept();
42                 out = new PrintWriter(client.getOutputStream(), true);
43                 in = new BufferedReader(new InputStreamReader(client.getInputStream
44
45                 userInput = in.readLine();
46                 try {
47                     guess = Integer.parseInt(userInput);
48                     if (guess < 0) {
49                         throw new IllegalArgumentException();
50                     }
51                     checkSecret(guess);
52                     out.println("Process Complete");
53                 } catch (IllegalArgumentException | InterruptedException e) {
54                     out.println("Unable to Process Input");
55                 }
56                 client.shutdownOutput();
57                 client.shutdownInput();
58                 client.close();
59             } catch (IOException e) {
60                 System.exit(-1);
61             }
62         }
63     }
64
65     public static void main(String[] args) throws InterruptedException {
66         startServer();
67     }
68 }
```

Note that the two “for” loops featuring a `Thread.sleep(1)` (resp. 2) could be equivalently replaced with a simple `Thread.sleep(32*32)` (resp. `Thread.sleep(2*32*32)`) statement, but

1. this is the way the program is presented in the DARPA library, and
2. a (minor) difficulty may come from these loops instead of a simple `Thread.sleep(32*32)` statement.