



HAL
open science

strategFTO: Untimed control for timed opacity

Étienne André, Shapagat Bolat, Engel Lefauchaux, Dylan Marinho

► **To cite this version:**

Étienne André, Shapagat Bolat, Engel Lefauchaux, Dylan Marinho. strategFTO: Untimed control for timed opacity. 8th International Workshop on Formal Techniques for Safety-Critical Systems (FTSCS 2022), Cyrille Artho; Peter Ölveczky, Dec 2022, Auckland, New Zealand. pp.27-33, 10.1145/3563822.3568013 . hal-03874435

HAL Id: hal-03874435



<https://hal.science/hal-03874435v1>

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.

strategFTO: Untimed control for timed opacity*

Étienne André^{1,2}, Shapagat Bolat², Engel Lefaucheur², and Dylan Marinho²

¹LIPN, CNRS UMR 7030, Université Sorbonne Paris Nord

²Université de Lorraine, CNRS, Inria, LORIA, Nancy, France

Abstract

We introduce a prototype tool `strategFTO` addressing the verification of a security property in critical software. We consider a recent definition of timed opacity where an attacker aims to deduce some secret while having access only to the total execution time. The system, here modelled by timed automata, is deemed opaque if for any execution time, there are either no corresponding runs, or both public and private corresponding runs. We focus on the untimed control problem: exhibiting a controller, i. e., a set of allowed actions, such that the system restricted to those actions is fully timed-opaque. We first show that this problem is not more complex than the full timed opacity problem, and then we propose an algorithm, implemented and evaluated in practice.

Keywords— opacity, timing leak, timed automata, security, control, IMITATOR

1 Introduction

We address here the control of timed systems to avoid timing leaks, i. e., the leakage of private information

that can be deduced from time. We use as underlying model timed automata (TAs) [AD94], an extension of finite-state automata with real-valued clocks.

Context Opacity is a key security property requiring that an external user should not be able to deduce whether the execution of a system contains a secret behavior through its observation. This property was first formalized for labeled transition systems [Bry+08], by specifying a subset of secret paths and requiring that, for any secret path, there is a non-secret one with the same observation. Opacity raises challenging research issues such as

1. specifying formally opacity in various frameworks [HS04; Bry+08],

*This manuscript is the author (and slightly extended) version of the manuscript of the same name published in the proceedings of the 8th International Workshop on Formal Techniques for Safety-Critical Systems (FTSCS 2022). The final authenticated version is available at dl.acm.org. This work is partially supported by the ANR-NRF French-Singaporean research program ProMiS (ANR-19-CE25-0015 / 2019 ANR NRF 0092) and the ANR research program BisoUS.

Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several universities as well as other organizations (see <https://www.grid5000.fr>).

2. verifying opacity properties [Maz04; Bry+08], and
3. developing mechanisms to design a system satisfying opacity while preserving functionality and performance [Bar+12; BHL17].

Franck Cassez proposed in [Cas09] a first definition of *timed opacity* asking whether an attacker can deduce a secret by observing a set of observable actions together with their timestamp. He proved that opacity is undecidable for TAs, mainly from the undecidability of the language inclusion problem for TAs [AD94]. The opacity problem is also undecidable for the restricted subclass of event-recording automata [AFH99]. Based on this definition of opacity, some decidable subclasses were proposed, for real-time automata [WZ18; WZA18] (a severely restricted subclass of TAs with a single clock), or over bounded-time [Amm+21].

In [And+22], we proposed a definition of opacity where the attacker only has access (in addition to the model knowledge) to the system *execution time*, i. e., the time from the initial location to a given location. The timed opacity problem therefore asks “for which execution times is the attacker unable to deduce whether a private location was visited?” The *full* timed opacity problem asks whether the system is timed-opaque for all execution times, i. e., the attacker is never able to deduce whether the private location was visited by an execution. We proved in [And+22] that this latter problem is decidable (in 3EXPTIME), and we proposed a practical algorithm using a parametric version of TAs [AHV93], implemented in IMITATOR [And21].

Contribution If a system is not fully timed-opaque, there may be ways to tune it to enforce opacity. For instance, one could change internal delays, or add some `sleep()` or `wait()` statements in the program (see e. g., [And+22]). In this paper, we consider a static (untimed) form of control of the system. This indicates whether there is a way of restricting the behavior of users to ensure full timed opacity. With that mindset, we assume the set of actions of the TA is partitioned into a set of *controllable* ac-

tions (that can be disabled) and a set of *uncontrollable* actions (that cannot be disabled). We address the following goal: exhibit a controller (i. e., a subset of the system controllable actions to be kept in addition to the uncontrollable actions, while other controllable actions are disabled) guaranteeing the system to be fully timed-opaque. We propose an algorithm exhibiting a set of controllers ensuring opacity, implemented into a tool `strategFTO`, calling IMITATOR [And21] for computing suitable opaque execution times, and POLYOP [BHZ08] for additional polyhedra operations.

Related works It is well known that observing the time taken by a system to finish some operation is a potential way to get information out of it (see e. g., [Koc96]). As such, identifying which information is released by the timing of a system has been studied both from a security and a safety perspective.

From the security point of view, beyond the works related to timed opacity and TAs [Cas09; WZ18; WZA18; Amm+21; And+22], the notion of non-interference has been widely studied. A first definition of timed non-interference was proposed for TAs in [Bar+02; BT03]. This notion is extended to PTAs in [AK20], with a semi-algorithm implemented using IMITATOR [And21]. In [GMR07], another notion of timed interference called timed strong non-deterministic non-interference (SNNI) which was based on timed language equivalence between the automaton with hidden low-level actions and the automaton with removed low-level actions was developed. This notion is in some aspects stronger than the opacity notion we consider, and is undecidable. SNNI was adapted in [VNN18] to allow some intentional information leakage and a form of control aimed at ensuring it was presented in [Ben+15]. Their framework gives to the attacker more information than the total execution time, and their control differs from ours to include that knowledge.

The *diagnosis* of TAs is one of the dominant research directions aimed at analysing information leakage from a safety perspective. Its goal is to detect, by observing the system, whether some faulty behavior occurred. As such, it is some form of dual

to opacity. Diagnosis was first introduced for TAs in [Tri02]. Diagnosability of a system is shown there to be decidable, though the actual diagnoser may be quite complex (see [BCD05] for subclasses of TAs allowing simpler diagnoser, see also [CT13] for a summary of the main results on the diagnosis of TAs and [Cas10] for a diagnosability focused control of TAs).

2 Preliminaries

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.

A clock guard g is a constraint over \mathbb{X} defined by a conjunction of inequalities of the form $x \bowtie d$, with $d \in \mathbb{Z}$ and $\bowtie \in \{<, \leq, =, \geq, >\}$. Given g , we write $\mu \models g$ if the expression obtained by replacing each x with $\mu(x)$ in g evaluates to true.

Definition 1 (TA [AD94]). A TA \mathcal{A} is a tuple $\mathcal{A} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E)$, where:

1. Σ is a finite set of actions,
2. L is a finite set of locations,
3. $\ell_0 \in L$ is the initial location,
4. $\ell_{priv} \in L$ is the private location,
5. $\ell_f \in L$ is the final location,
6. \mathbb{X} is a finite set of clocks,
7. I is the invariant, assigning to every $\ell \in L$ a clock 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 clock guard.

Example 1. Consider the TA in Fig. 1a, using one clock x . ℓ_1 is the initial location, while we assume that ℓ_f is the *final* location, i. e., a location in which an attacker can measure the execution time from the initial location. ℓ_2 is the private location, i. e., a secret to be preserved: the attacker should not be able to deduce whether it was visited or not. ℓ_2 has an invariant $x \leq 3$ (boxed); other locations invariants are true.

Definition 2 (Semantics of a TA [AD94]). Given a TA $\mathcal{A} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E)$, the semantics of \mathcal{A} is given by the timed transition system (TTS) $T_{\mathcal{A}} = (S, s_0, \rightarrow)$, with

- $S = \{(\ell, \mu) \in L \times \mathbb{R}_{\geq 0}^H \mid \mu \models I(\ell)\}$, $s_0 = (\ell_0, \vec{0})$,
- \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 $\mu' = [\mu]_R \models I(\ell')$, and $\mu \models g$.
 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 \mathcal{A} with semantics (S, s_0, \rightarrow) , a *run* of \mathcal{A} is an alternating sequence of states of $T_{\mathcal{A}}$ and pairs of delays and edges starting from the initial state s_0 of the form $s_0, (d_0, e_0), s_1, \dots$ where for all i , $e_i \in E$, $d_i \in \mathbb{R}_{\geq 0}$ and $s_i \xrightarrow{(d_i, e_i)} s_{i+1}$. 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 $dur(\rho) = \sum_{0 \leq j \leq i-1} d_j$.

2.1 Timed opacity definitions

We recall here the notion of timed opacity defined in [And+22].¹

¹We slightly modify the definitions from [And+22] by incorporating ℓ_{priv} within the definition of \mathcal{A} , and removing ℓ_{priv} and ℓ_f from the definition of $DReach^{priv}(\mathcal{A})$ and $DReach^{-priv}(\mathcal{A})$ to simplify the reading.




Figure 1: Running example

Given $\mathcal{A} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E)$, and a run ρ , we say that ℓ_{priv} is reached on the way to ℓ_f in ρ if ρ is of the form $(\ell_0, \mu_0), (d_0, e_0), (\ell_1, \mu_1), \dots, (\ell_m, \mu_m), (d_m, e_m), \dots, (\ell_n, \mu_n)$ for some $m, n \in \mathbb{N}$ such that $\ell_m = \ell_{priv}, \ell_n = \ell_f$ and $\forall 0 \leq i \leq m-1, \ell_i \neq \ell_f$. We denote by $\text{Reach}_{\ell_{priv}}^{\mathcal{A}}(\ell_f)$ the set of those runs, and refer to them as *private* runs. Conversely, we say that ℓ_{priv} is avoided on the way to ℓ_f in ρ if ρ is of the form $(\ell_0, \mu_0), (d_0, e_0), (\ell_1, \mu_1), \dots, (\ell_n, \mu_n)$ with $\ell_n = \ell_f$ and $\forall 0 \leq i < n, \ell_i \notin \{\ell_{priv}, \ell_f\}$. We denote the set of those runs by $\text{Reach}_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f)$, and refer to them as *public* runs.

While we model the secret behavior of the system using a private location ℓ_{priv} here, note that one could easily adapt these definitions if the secret is, for example, a set of locations, an action (this will be the case in our case study) or the value of a variable.

$D\text{Reach}^{priv}(\mathcal{A})$ (resp. $D\text{Reach}^{\neg priv}(\mathcal{A})$) is the set of all the durations of the runs for which ℓ_{priv} is reached (resp. avoided) on the way to ℓ_f . Formally: $D\text{Reach}^{priv}(\mathcal{A}) = \{d \in \mathbb{R}_{\geq 0} \mid \exists \rho \in \text{Reach}_{\ell_{priv}}^{\mathcal{A}}(\ell_f) \text{ such that } d = \text{dur}(\rho)\}$ and $D\text{Reach}^{\neg priv}(\mathcal{A}) = \{d \in \mathbb{R}_{\geq 0} \mid \exists \rho \in \text{Reach}_{\neg \ell_{priv}}^{\mathcal{A}}(\ell_f) \text{ such that } d = \text{dur}(\rho)\}$.

Definition 3 (full timed opacity). Given a TA \mathcal{A} , we say that \mathcal{A} is *fully timed-opaque* if $D\text{Reach}^{priv}(\mathcal{A}) = D\text{Reach}^{\neg priv}(\mathcal{A})$.

That is, a system is fully timed-opaque if, for any execution time d , there exists a run of duration d that reaches ℓ_f after going through ℓ_{priv} iff there exists another run of duration d that reaches ℓ_f without going through ℓ_{priv} . Hence, the attacker cannot deduce from the execution time whether ℓ_{priv} was visited or not.

Example 2. Consider again the TA in Fig. 1a. Recall that ℓ_2 is the private location. We have $D\text{Reach}^{priv}(\mathcal{A}) = [1, 5]$ and $D\text{Reach}^{\neg priv}(\mathcal{A}) = [1, 3] \cup [4, 4] \cup (5, +\infty)$. Since $D\text{Reach}^{priv}(\mathcal{A}) \neq D\text{Reach}^{\neg priv}(\mathcal{A})$, the system is *not* fully timed-opaque.

3 Untimed control for full timed opacity

In this section, we introduce an untimed control for controlling timed opacity. We assume $\Sigma = \Sigma_c \uplus \Sigma_u$ where Σ_c (resp. Σ_u) denote controllable (resp. uncontrollable) actions.

A (static, untimed) *strategy* of a TA \mathcal{A} is a set of actions $\sigma \subseteq \Sigma$ that contains at least all uncontrollable actions (i.e., $\Sigma_u \subseteq \sigma \subseteq \Sigma$). A strategy induces a restriction of \mathcal{A} where only the edges labeled by actions of σ are allowed:

Definition 4 (Controlled TA). Given $\mathcal{A} = (\Sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E)$ with $\Sigma = \Sigma_u \uplus \Sigma_c$ and a strategy $\sigma \subseteq \Sigma$, the *control* of \mathcal{A} using σ is the TA $\mathcal{A}' = Control(\mathcal{A}, \sigma) = (\sigma, L, \ell_0, \ell_{priv}, \ell_f, \mathbb{X}, I, E')$ where $E' = \{(\ell, g, a, R, \ell') \in E \mid a \in \sigma\}$.

Example 3. Consider again the TA \mathcal{A} in Fig. 1a. Fix $\sigma = \{u, a\}$. Then $Control(\mathcal{A}, \sigma)$ is in Fig. 1b.

Strategies represent some modifications of the system that can be implemented to ensure full timed opacity.

Definition 5 (fully timed-opaque strategy). A strategy σ is *fully timed-opaque* if $Control(\mathcal{A}, \sigma)$ is fully timed-opaque.

A strategy (even a maximal one) might achieve full timed opacity by blocking all runs (both private or public) from reaching the target. If reaching the target means completing a task, this might not be something one would desire. We call a strategy allowing to reach the target for at least some durations an *effective* strategy.

We define two slightly different problems: taking a TA \mathcal{A} as input, the **full timed (resp. effective full time) opacity control emptiness problem** asks whether the set of fully (resp. effective fully) timed-opaque strategies for \mathcal{A} is empty.

Full timed opacity control emptiness problem:

INPUT: A TA \mathcal{A}

PROBLEM: is the set of fully timed-opaque strategies for \mathcal{A} empty?

Effective full timed opacity control emptiness problem:

INPUT: A TA \mathcal{A}

PROBLEM: is the set of effective fully timed-opaque strategies for \mathcal{A} empty?

Note that, due to the presence of uncontrollable actions, the first problem (full timed opacity control emptiness) is not trivial. (If uncontrollable actions were not part of our definitions, choosing $\sigma = \emptyset$ would always yield an acceptable fully timed-opaque strategy.)

We will also refine those problems by considering a notion of *maximal* (i.e., most permissive) strategy w.r.t. full timed opacity based on the number of actions belonging to the strategy: given \mathcal{A} , a fully timed-opaque strategy σ is maximal if $\forall \sigma'$, if σ' is fully timed-opaque then $|\sigma'| \leq |\sigma|$. We define similarly *minimal* strategies (least permissive, i.e., disabling as many actions as possible) as well as maximal (resp. minimal) effective fully timed-opaque strategies, i.e., the set of largest (resp. smallest) effective fully timed-opaque strategies.

Example 4. Consider again the TA \mathcal{A} in Fig. 1a. Assume $\Sigma_u = \{u\}$ and $\Sigma_c = \{a, b, c, d, e, f\}$. Fix $\sigma_1 = \{u, b, c\}$. We have $DReach^{priv}(Control(\mathcal{A}, \sigma_1)) = [2, 5]$ while $DReach^{\neg priv}(Control(\mathcal{A}, \sigma_1)) = [4, 4]$; therefore, σ_1 is not fully timed-opaque. Now fix $\sigma_2 = \{u, a, f\}$. We have $DReach^{priv}(Control(\mathcal{A}, \sigma_2)) = DReach^{\neg priv}(Control(\mathcal{A}, \sigma_2)) = [1, 3]$; therefore, σ_2 is fully timed-opaque.

In fact, it can be shown that the set of effective fully timed-opaque strategies for \mathcal{A} is $\{\{u, a\}, \{u, a, e\}, \{u, a, f\}\}$; therefore, $\{u, a\}$ is the only minimal strategy, while $\{u, a, e\}, \{u, a, f\}$ are the two maximal strategies. In addition, $\{u, f\}$ is an example of a strategy that is not effective, as ℓ_f is always unreachable, whether ℓ_{priv} is visited or not.

3.1 Complexity

Proposition 1 (complexity). *One can compute the set of fully timed-opaque strategies over a TA \mathcal{A} in 3EXPTIME.*

Proof. The full timed opacity decision problem (i.e., checking if a given TA is fully timed-opaque) is decidable for TAs in (at most) 3EXPTIME [And+22]. Moreover, reachability of the final state can be decided in PSPACE [AD94]. Thus, for any given strategy, one can check in triple exponential time whether it is (effective) fully timed-opaque.

Computing the list of (effective) fully timed-opaque strategies can be done naively by testing each possible strategy one by one and keeping the ones that satisfy the property we want. As there is an exponential number of possible strategies and repeating exponen-

tially many times a 3EXPTIME algorithm remains in 3EXPTIME, this algorithm is in 3EXPTIME. \square

As a corollary of the above, the (effective) full timed opacity control emptiness problem is in 3EXPTIME as well. More precisely, the above proof establishes that the complexity class of the (effective) full timed opacity control emptiness problem is the maximum between PSPACE and the complexity of the full timed opacity problem. As the latter is PSPACE-hard (being trivially harder than reachability), the two problems lie in the same complexity class. From a theoretical point of view, one thus cannot do better than the naive enumeration approach described here to solve the control problem.

Finding the maximal (resp. minimal) strategies can be done slightly more efficiently by starting from the set with every (resp. no) controllable action and enumerating the potential strategies by decreasing (resp. increasing) order as one could then potentially stop before full enumeration. In the worst case, this will however have the same complexity as the full enumeration.

4 Implementation and experiments

4.1 Implementation in stratFTO

We implemented our strategy generation in stratFTO, an entirely automated open-source tool written in Java.² Our tool iteratively constructs strategies, then checks full timed opacity following Algorithm 1.

We give our strategy synthesis algorithm in Algorithm 1. The exhibition of these execution times ($DReach^{\neg priv}(\mathcal{A})$ and $DReach^{priv}(\mathcal{A})$, line 4) is done in our implementation by an automated model modification (following the procedure described in [And+22], but which was not entirely automated in [And+22]) followed by a synthesis problem

²Source code is available at <https://github.com/DylanMarinho/Controlling-TA>. Models and experiment results are available at [10.5281/zenodo.7181848](https://zenodo.org/record/7181848).

Algorithm 1: $\text{synthCtrl}(\mathcal{A})$ Exhibit all timed-opaque strategies

```

1  $\mathcal{S} \leftarrow \emptyset$ 
2 foreach  $s \subseteq \Sigma_c$  do
3    $\sigma \leftarrow s \cup \Sigma_u$ 
4   /* Compute execution times */
5    $\lambda_1 \leftarrow DReach^{\neg priv}(\text{Control}(\mathcal{A}, \sigma))$ 
6    $\lambda_2 \leftarrow DReach^{priv}(\text{Control}(\mathcal{A}, \sigma))$ 
7   /* Check for full timed opacity */
8   if  $\lambda_1 = \lambda_2$  then  $\mathcal{S} \leftarrow \mathcal{S} \cup \{s\}$ ;
9 return  $\mathcal{S}$ 

```

using a parametric extension of TAs [AHV93]. The synthesis of the execution times itself is done by a call to an external tool—IMITATOR 3.3 “*Cheese Caramel au beurre salé*” [And21].

stratFTO then checks whether both sets of execution times are equal; this is done by a call to another external tool—POLYOP 1.2³, that performs polyhedral operations as a simple interface for the Parma polyhedra library [BHZ08].

Algorithms We implement not only the exhibition of all timed-opaque strategies (denoted by $\text{synthCtrl}(\mathcal{A})$, in Algorithm 1), but also the following variants:

1. $\text{synthMaxCtrl}(\mathcal{A})$: synthesize all maximal strategies for \mathcal{A} ;
2. $\text{synthMinCtrl}(\mathcal{A})$: synthesize all minimal strategies;
3. $\text{witnessMaxCtrl}(\mathcal{A})$: witness *one* maximal strategy;
4. $\text{witnessMinCtrl}(\mathcal{A})$: witness *one* minimal strategy.

We implemented these other algorithms by changing the exploration order of the strategies, and/or by triggering immediate termination upon the first exhibition of a strategy.

³<https://github.com/etienneandre/PolyOp>

Input model The input TA model is given in the IMITATOR input syntax; while we presented a restricted setting in this paper for sake of clarity, our implementation in `strategFTO` is much more permissive, by allowing significant extensions of TAs with global (integer or Boolean) variables, multiple automata with synchronization, multi-rate clocks (including stopwatches), etc.

4.2 Proof of concept benchmark

As a proof of concept, we consider the TA model of an ATM (given in Fig. 2). The idea is that (as per our definition of timed opacity) the attacker only has access to the execution time, i. e., the time from the beginning of the program to reaching the end state. The secret is whether the ATM user has actually obtained cash (action *takeCash*).⁴ The TA uses two clocks: x for “local” actions, and y for a global time measurement. First, the user starts the process (action *start*), then the ATM displays a welcome screen for 3 time units, followed by another screen requesting the password (action *askPwd*). Then, the user can submit a correct (action *correctPwd*) or incorrect (*incorrectPwd*) password; if no password is input within 10 time units, the system moves to a cancelling phase. The same happens if 3 incorrect password have been input. After inputting the correct password, the user has the choice between a fixed-amount quick withdrawal (*quickWithdraw*), a normal withdrawal (*normalWithdraw*) or a balance request (*reqBalance*).

The quick withdrawal triggers a 15-time unit preparation followed by the availability of the money, which the user can take immediately (action *takeCash*), thus terminating the procedure. If the user does not take the money, the system moves to the cancelling phase.

The normal withdrawal asks the user to input the desired amount; similar to the password, after 3 wrong amounts (action *incorrectAmount*), or upon timeout, the system moves to cancelling phase. After the user retrieves cash (action *takeCash*), they are

⁴`strategFTO` allows not only private locations, but also actions.

asked whether they would like to perform another operation; if so (action *restart*), the system goes back to the choice location. Otherwise (action *pressFinish*), or unless a 10-time unit timeout is reached, the system moves to the terminating location. The balance request triggers the balance display, from which the user can immediately terminate the process (action *pressOK*), or go back to the choice menu.

The rationale is that, in the regular terminating and cancelling phases, the ATM terminates after constant time (invariant $y \leq 100$), avoiding leaking information. However, some actions may lead to quicker termination (quick withdrawal) or slower termination (multiple choices).

The uncontrollable actions are most of the user actions: *correctAmount*, *incorrectAmount*, *correctPwd*, *incorrectPwd*, *pressFinish*, *takeCash*. The controllable actions are the system actions (*askPwd*, *start*, *finish*) and some of the users actions that can be controlled by disabling the associated choice (*reqBalance*, *pressOK*, *quickWithdraw*, *restart*).

4.3 Experiments

We first exhibit in Table 1 controllers for our benchmark from Fig. 2 as computed by `strategFTO`, for all our algorithms. For space concern, we tabulate the actions to *disable*; the strategy is therefore Σ minus these actions. Also note that, for `witnessMaxCtrl` and `witnessMinCtrl`, the *order* in which we compute the subsets of Σ in Algorithm 1 has an impact on the result, as the algorithm stops as soon as *one* strategy is found. According to Table 1, the maximal strategies (i. e., the most permissive, disabling the least number of actions) are to disable either *restart* and *pressOK*, or *restart* and *reqBalance*. This is natural, as *restart* allows the user to restart a second operation, thus violating the constant-time nature of Fig. 2, while *pressOK* and *reqBalance*, if enabled together, allow a quick exit, shorter than a cash withdrawal operation—thus giving hint to the attacker that the *takeCash* secret did *not* occur.

Scalability Then, we test the scalability of `strategFTO` w.r.t. the number of actions. We modify Fig. 2 by adding an increasingly large numbers




Figure 2: ATM benchmark

Table 1: Strategy synthesis for Fig. 2

Actions to disable	synthMinCtrl	witnessMinCtrl	synthMaxCtrl	witnessMaxCtrl	synthCtrl
<i>restart, pressOK</i>			✓	✓	✓
<i>restart, reqBalance</i>			✓		✓
<i>restart, pressOK, quickWithdraw</i>					✓
<i>restart, pressOK, reqBalance</i>					✓
<i>restart, quickWithdraw, reqBalance</i>					✓
<i>restart, pressOK, quickWithdraw, reqBalance</i>	✓	✓			✓

of controllable actions; these actions do not play a role in the control (we basically add unguarded self-loops) but they will impact the computation time, as we will need to consider an increasingly (and exponentially) larger number of subsets of actions, from Algorithm 1. We add from 1 to 40 such actions, resulting (by adding the actions in Fig. 2) in a model with a number of controllable actions from 11 to 50. We plot these results in Fig. 3. (Raw results are in Table 2 in Appendix A.) From our results in Fig. 3, we see that, without surprise, the execution time for synthCtrl is exponential in the number of actions. However, synthMaxCtrl and synthMinCtrl

behave much better, by remaining respectively below 15 minutes and three minutes, even for up to 50 controllable actions. In addition, it is important to notice that witnessMaxCtrl and witnessMinCtrl do not decrease the time very much compared to the full versions synthMaxCtrl and synthMinCtrl. This is because, at a given size, the number of strategies to be tested remains relatively small.




Figure 3: Execution times for scalability (in seconds; TO set at 1,800s)

5 Conclusion

We introduced a prototype tool `strategFTO` implementing an algorithm to exhibit strategies to guarantee the full timed opacity of a system modeled by a timed automaton where the attacker only has access to the computation time. Even though relying on a simple enumeration of the subsets, our tool `strategFTO` shows good performance for synthesizing maximal or minimal strategies, with very reasonable times, even for several dozens of controllable actions.

Future works We plan to further optimize our implementation by maintaining a set of non-effective strategies, i. e., for which ℓ_f is unreachable: any strategy strictly included into a known non-effective strategy will necessarily be non-effective too, and therefore no full timed-opacity analysis is needed for this strategy. An option to efficiently represent this strategies set could be to store it using BDDs.

We also plan to strengthen strategies so that their choice may depend on how long has passed since the start of the execution. As these strategies still need a finite representation to be handled, this requires establishing exactly what strategies need to remember to chose optimally.

Our ultimate goal will be to extend timed automata to *parametric* timed automata [AHV93], and

use automated parameter synthesis techniques (e. g., [JLR15; And+21; AMP21]), with a parametric timed controller [JLR19; Gol21].

References

- [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](https://doi.org/10.1016/0304-3975(94)90010-8) (cit. on pp. 1–3, 5).
- [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](https://doi.org/10.1016/S0304-3975(97)00173-4) (cit. on p. 2).
- [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](https://doi.org/10.1145/167088.167242) (cit. on pp. 2, 6, 9).
- [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](https://doi.org/10.1109/ICECCS51672.2020.00012) (cit. on p. 2).

- [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. 2).
- [AMP21] Étienne André, Dylan Marinho, and Jaco van de Pol. “A Benchmarks Library for Extended Timed Automata”. In: *TAP* (June 21–25, 2021). Ed. by Frédéric Loulergue and Franz Wotawa. Vol. 12740. Lecture Notes in Computer Science. virtual: Springer, 2021, pp. 39–50. DOI: [10.1007/978-3-030-79379-1_3](https://doi.org/10.1007/978-3-030-79379-1_3) (cit. on p. 9).
- [And+21] Étienne André, Jaime Arias, Laure Petrucci, and Jaco van de Pol. “Iterative Bounded Synthesis for Efficient Cycle Detection in Parametric Timed Automata”. In: *TACAS* (Mar. 27–Apr. 1, 2021). Ed. by Jan Friso Groote and Kim G. Larsen. Vol. 12651. Lecture Notes in Computer Science. Virtual: Springer, 2021, pp. 311–329. DOI: [10.1007/978-3-030-72016-2_17](https://doi.org/10.1007/978-3-030-72016-2_17) (cit. on p. 9).
- [And+22] Étienne André, Didier Lime, Dylan Marinho, and Jun Sun. “Guaranteeing timed opacity using parametric timed model checking”. In: *ACM Transactions on Software Engineering and Methodology* 31.4 (Oct. 2022). To appear, pp. 1–36. DOI: [10.1145/3502851](https://doi.org/10.1145/3502851) (cit. on pp. 2, 3, 5, 6).
- [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. 2, 6).
- [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. 2).
- [Bar+12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. “On the (im)possibility of obfuscating programs”. In: *Journal of the ACM* 59.2 (2012), 6:1–6:48. DOI: [10.1145/2160158.2160159](https://doi.org/10.1145/2160158.2160159) (cit. on p. 2).
- [BCD05] Patricia Bouyer, Fabrice Chevalier, and Deepak D’Souza. “Fault Diagnosis Using Timed Automata”. In: *FoSSaCS* (Apr. 4–8, 2005). Ed. by Vladimiro Sassone. Vol. 3441. Lecture Notes in Computer Science. Edinburgh, UK: Springer, 2005, pp. 219–233. DOI: [10.1007/978-3-540-31982-5_14](https://doi.org/10.1007/978-3-540-31982-5_14) (cit. on p. 3).
- [Ben+15] 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 p. 2).
- [BHL17] Béatrice Bérard, Serge Haddad, and Engel Lefauchaux. “Probabilistic Disclosure: Maximisation vs. Minimisation”. In: *FSTTCS* (Dec. 11–15, 2017). Ed. by Satya V. Lokam and R. Ramanujam. Vol. 93. LIPIcs. Kanpur, India: Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 13:1–13:14. DOI: [10.4230/LIPIcs.FSTTCS.2017.13](https://doi.org/10.4230/LIPIcs.FSTTCS.2017.13) (cit. on p. 2).
- [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. 2, 6).
- [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. 1, 2).
- [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. 2).
- [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 p. 2).
- [Cas10] Franck Cassez. “Dynamic observers for fault diagnosis of timed systems”. In: *CDC* (Dec. 15–17, 2010). Atlanta, Georgia, USA: IEEE, 2010, pp. 4359–4364. DOI: [10.1109/CDC.2010.5717696](https://doi.org/10.1109/CDC.2010.5717696) (cit. on p. 3).
- [CT13] Franck Cassez and Stavros Tripakis. “Fault Diagnosis of Timed Systems”. In: *Communicating Embedded Systems*. Ed. by Claude Jard and Olivier H. Roux. Wiley, 2013, pp. 107–138. DOI: [10.1002/9781118558188.ch4](https://doi.org/10.1002/9781118558188.ch4) (cit. on p. 3).

- [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 p. 2).
- [Gol21] Ebru Aydin Gol. “Control Synthesis for Parametric Timed Automata under Unavoidability Specifications”. In: *ECC* (June 29–July 2, 2021). Virtual Event / Delft, The Netherlands: IEEE, 2021, pp. 740–745. DOI: [10.23919/ECC54610.2021.9655222](https://doi.org/10.23919/ECC54610.2021.9655222) (cit. on p. 9).
- [HS04] Dominic J. D. Hughes and Vitaly Shmatikov. “Information Hiding, Anonymity and Privacy: A Modular Approach”. In: *Journal of Computer Security* 12.1 (2004), pp. 3–36. DOI: [10.3233/jcs-2004-12102](https://doi.org/10.3233/jcs-2004-12102) (cit. on p. 1).
- [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 p. 9).
- [JLR19] Aleksandra Jovanović, Didier Lime, and Olivier H. Roux. “A game approach to the parametric control of real-time systems”. In: *International Journal of Control* 92.9 (2019), pp. 2025–2036. DOI: [10.1080/00207179.2018.1426883](https://doi.org/10.1080/00207179.2018.1426883) (cit. on p. 9).
- [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. 2).
- [Maz04] Laurent Mazaré. “Decidability of Opacity with Non-Atomic Keys”. In: *FAST* (Aug. 22–27, 2004). Ed. by Theodosios Dimitrakos and Fabio Martinelli. Vol. 173. IFIP. Toulouse, France: Springer, 2004, pp. 71–84. DOI: [10.1007/0-387-24098-5_6](https://doi.org/10.1007/0-387-24098-5_6) (cit. on p. 2).
- [Tri02] Stavros Tripakis. “Fault Diagnosis for Timed Automata”. In: *FTRTFT* (Sept. 9–12, 2002). Ed. by Werner Damm and Ernst-Rüdiger Olderog. Vol. 2469. Lecture Notes in Computer Science. Oldenburg, Germany: Springer, 2002, pp. 205–224. DOI: [10.1007/3-540-45739-9_14](https://doi.org/10.1007/3-540-45739-9_14) (cit. on p. 3).
- [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 p. 2).
- [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. 2).
- [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. 2).

Appendix

A Experiments: scalability test

We performed a sample scalability test on our benchmark. The plot is given in [Fig. 3](#).

We tabulate our full results in [Table 2](#).

Table 2: Execution times for scalability (in seconds; TO set at 1,800s)

Number of added actions	synthMinCtrl		witnessMinCtrl		synthMaxCtrl		witnessMaxCtrl		synthCtrl
	-find min	-find min	-find min	-witness	-find max	-find max	-find max	-witness	
1	2.89	2.04	2.04		12.61	5.92	5.92		22.98
2	3.19	2.44	2.44		17.81	11.30	11.30		44.68
3	3.84	2.74	2.74		23.99	17.58	17.58		87.07
4	4.43	2.85	2.85		31.15	24.92	24.92		172.26
5	4.90	3.77	3.77		39.69	10.34	10.34		342.95
6	6.07	4.09	4.09		48.78	11.12	11.12		683.77
7	7.02	4.54	4.54		59.14	12.35	12.35		1,362.48
8	8.34	4.69	4.69		70.09	13.46	13.46		TO
9	9.32	5.63	5.63		82.45	14.52	14.52		TO
10	10.51	5.91	5.91		95.86	15.65	15.65		TO
11	12.04	7.66	7.66		111.16	30.49	30.49		TO
12	13.99	9.43	9.43		126.11	46.00	46.00		TO
13	15.54	10.94	10.94		143.15	62.50	62.50		TO
14	17.58	12.83	12.83		160.41	80.32	80.32		TO
15	19.88	15.03	15.03		180.24	85.64	85.64		TO
16	21.94	17.38	17.38		199.34	105.15	105.15		TO
17	24.39	17.63	17.63		221.04	146.98	146.98		TO
18	27.64	20.53	20.53		241.70	168.79	168.79		TO
19	30.49	23.65	23.65		264.72	191.16	191.16		TO
20	33.43	26.59	26.59		287.85	215.33	215.33		TO
21	36.58	28.60	28.60		313.34	239.52	239.52		TO
22	40.46	30.92	30.92		339.24	265.90	265.90		TO
23	44.31	33.92	33.92		366.07	292.51	292.51		TO
24	48.52	36.43	36.43		395.95	322.16	322.16		TO
25	53.21	38.30	38.30		423.86	350.92	350.92		TO
26	58.02	41.21	41.21		453.59	368.86	368.86		TO
27	62.36	43.57	43.57		484.28	400.39	400.39		TO
28	68.56	46.60	46.60		517.03	419.43	419.43		TO
29	74.39	49.74	49.74		551.58	436.35	436.35		TO
30	80.64	53.15	53.15		586.03	453.37	453.37		TO
31	86.89	55.72	55.72		621.83	472.63	472.63		TO
32	92.91	59.14	59.14		656.75	492.10	492.10		TO
33	100.67	67.00	67.00		693.82	528.71	528.71		TO
34	111.66	76.17	76.17		730.93	564.68	564.68		TO
35	120.37	85.48	85.48		768.87	604.42	604.42		TO
36	128.35	94.59	94.59		809.36	645.01	645.01		TO
37	137.28	102.77	102.77		856.98	685.03	685.03		TO
38	147.68	112.83	112.83		897.39	728.48	728.48		TO
39	157.42	121.77	121.77		940.98	771.68	771.68		TO
40	168.74	132.45	132.45		984.65	818.25	818.25		TO