# Opacity Degree in Interval Labelled Markov Chains

Hugo Bazille[1], Eric Fabre[2], Kritin Garg[3], and Blaise Genest[4]

[1] Univ Rennes, Rennes 1, France hbazille@irisa.fr
[2] Univ Rennes, INRIA, France fabre@inria.fr
[3] IIT Bombay, Mumbai, India kritingarg0209@gmail.com,
[4] Univ Rennes, CNRS, France bgenest@irisa.fr

**Abstract.** Nowadays, almost every system is connected, which raises new security questions, such as "can we deduce some confidential information from observing the system". It is thus important to verify that the system is opaque, for different notions of opacity. In this paper, we consider a quantitative notion of opacity: to be useful, systems usually need to leak some information - hence the question is not "can some information be deduced", but "is it the case that only a limited amount of information is leaked". We quantify the degree of opacity in terms of probabilities of deducing some information, and we thus naturally model the system as a stochastic system, namely *Interval* Labelled Markov chains (also called *Interval* Hidden Markov Models), where probabilities do not need to be known exactly.

In this paper, we study weighted degrees, averaging the information leaked over all runs of the system, as well as the worst case degree, considering the observation revealing the most information about the system. We show that one cannot compute exactly nor approximate the worst-case degree in general, while computing exactly the average degree is easy. We show that one cannot compute exactly more refined weighted degrees, but one can approximate them, unlike the worst-case degree. Finally, we provide the exact complexity of deciding qualitative worst-case degree, i.e. whether it is positive or 1, with complexities ranging from NLOGSPACE to EXPTIME-complete.

## 1 Introduction

In our connected world, privacy concerns are becoming prominent. It is thus important to analyze different key concepts to measure information leakages. Very precise semantical notions have been developed, such as information flow (e.g. [10]). However, no general techniques exist yet to automatically compute information flow *over a whole system* beyond a particular finite run. On the other hand, syntactical notions, less precise but more easily computable, have also emerged, such as *opacity* in partially observable systems [17], based on a notion of secret and non-secret runs: a system is *opaque* when all secret runs are observationally equivalent with non-secret runs. It is decidable in PSPACE

whether a system is opaque [19], and one can also try to control the system to make it opaque [12].

This notion of opacity is purely qualitative. This is quite restrictive for two reasons. First, it is possible that a system is not opaque, but extremely few runs allows an attacker to obtain information, which could be perfectly fine for non-critical systems. Second, it is possible that a system is opaque using this qualitative definition, while an attacker could know that an observation corresponds to a secret run with very high probability, because equivalent non-secret runs could have a negligible probability to happen.

To handle the first issue, [20, 4] considered probabilistic systems and proposed to quantify the number of secret runs which are not observationally equivalent with non-secret runs. To compute it, they add a 0/1 signal to each observation (opaque or not), and count the probability to obtain non-opaque observations.

In this paper, we handle the second issue by quantifying the opacity degree at the level of each observation, comparing the probabilities of non-secret runs and of all runs corresponding to this observation. We then consider two problems: compute the worst possible opacity degree among all observations, or compute an average opacity degree by weighting observations with their probabilities and their degree. We show that the worst-case degree cannot be computed, and cannot even be approximated, by reducing an unapproximable problem on PFAs [16]. This contrasts with the average degree, which can be computed in PTIME.

Our main results concern cases between these two extremes: First, we consider weighted opacity degree giving more information than a pure average degree, by setting any threshold $\theta$ of probability under which a run is considered opaque: if the ratio between secret runs and all runs having a given observation is less than $\theta$, then we consider the observation opaque. Otherwise, we weight its non opacity w.r.t. the threshold $\theta$. We show that it is undecidable to compute this threshold degree exactly. However, we provide several results to $\varepsilon$-approximate the threshold degree, either with high confidence or full certainty (albeit with higher complexities). Then we show that qualitative worst-case degree is decidable (is the worst case degree $> 0$? is it $= 1$?), and provide the exact complexities. Further, because it is in practice hard to know the probabilities of all transitions, we introduce for the first time Interval Labelled Markov Chains, inspired by Interval MC and Labelled MC. We develop techniques to handle intervals in this partially-observable setting: for most problems, the blow-up in the precise complexity is limited wrt plain Labelled Markov Chains.

### 1.1 Related work

Opacity degrees have been considered in different frameworks [4, 3, 5]. The first difference is that they study these notions for Markov Decision Processes (MDP) while we consider Interval Labeled Markov Chains (ILMC). While ILMCs can be encoded as MDPs, the encoding needs a uncountable number of actions (one for each choice of probabilities), hence complexities cannot be directly related. In [4], Bérard and co-authors define *restrictive probability of opacity* similar to our opacity degree. They do not provide (un)decidability result on computing the

opacity degree of a system (except for the trivial case where the number of runs is bounded). The main focus of [3, 5] is *disclosure*, which is dual with opacity. There, one wants to know whether *all* words discloses fully the secret, while worst-case opacity asks whether *there exists* a word which discloses fully the secret. We do not consider disclosure in this paper. Concerning opacity degree, in [3], the authors mention an opacity degree similar to our weighted opacity degree, and prove an undecidability result on the existence of an observation with a given opacity degree, similar to the first half of our Proposition 1. As they do not consider approximability, they did not study the problem further, which we do in Section 5. Last [5] (Corollary 11) obtain easily the *decidability* of whether there exists a system for which no observation has an opacity degree 0 (qualitative opacity degree). This is related to our Theorem 4, where we provide much more precise result by providing tight complexity bounds (PSPACE-complete and EXPTIME-complete depending on the semantics of ILMCs, better than using the generic construction of [5]). Compared with this particular case, we provide tight results concerning the other qualitative cases as well (Theorem 3,5,6).

## 2 Preliminaries and Background

### 2.1 Labelled Markov Chain

Let $Q$ be a finite set of states. A probability distribution over $Q$ is a function $\mu : S \longrightarrow [0, 1]$ such that $\sum_{q \in Q} \mu(q) = 1$. Let $Dist(Q)$ be the set of distributions over Q. The *support* of $\mu$ is $\text{support}(\mu) = \{s \in Q | \mu(q) > 0\}$.

A *labelled (discrete-time) Markov Chain (LMC for short)* over finite alphabet $\Sigma$ is a tuple $A = (S, M, \mu_0)$ with $S$ its set of states, $\mu_0$ its initial distribution, and $M : S \times \Sigma \times S \to [0, 1]$, such that for all $s$, $\sum_{a \in \Sigma, s' \in S} M(s, a, s') = 1$.

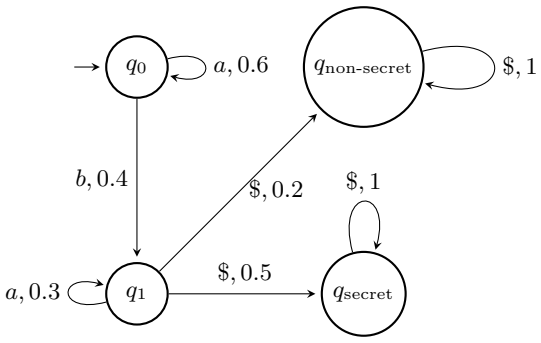


Fig. 1: Example of an LMC

We will model secret runs by using a secret state $q_{\text{secret}}$, and non-secret runs by using a non-secret state $q_{\text{non-secret}}$, as usual [19]. As we need to quantify over

runs, we do not want to count a run several times. [20] solves this by only considering runs of the same size. This is not possible in our context where we want to weight a degree over all runs. Instead, we consider that runs eventually stop, and only quantify over stopped runs. If a run stops in $q_{\text{secret}}$, it is secret, and otherwise it stops in $q_{\text{non-secret}}$ and it is non-secret. Further, the last action before stopping is labelled by the special symbol \$, such that stopping is observable. To obtain an LMC, we will formally allow \$ to loop over these states: $M(q_{\text{secret}}, \$, q_{\text{secret}}) = 1$ and $M(q_{\text{non-secret}}, \$, q_{\text{non-secret}}) = 1$. We will never consider runs past the first \$.

A path of $(S, M, \mu_0)$ is a sequence $\rho = s_0 a_1 s_1 \cdots a_n s_n$ in $S(\Sigma \times S)^*$, where $\mu_0(s_0) > 0$ and for all $1 \leq i < n$, $M(s_{i-1}, a_i, s_i) > 0$. The path is final if $a_n = \$$ (and thus $s_n \in \{q_{\text{secret}}, q_{\text{non-secret}}\}$). Given any path $\rho$, $s^-(\rho)$ represents the initial state of the path, $s^+(\rho)$ represents the final state of the path. The observation associated with $\rho$ is $obs(\rho) = a_1 a_2 a_3 ....a_n \in \Sigma^*$. We say that $w \in \Sigma^*$ is an observation if there exists a path $\rho$ associated with $w$. The language of an LMC $A$ is defined as the set of observations $L(A) = \{w \in \Sigma^* \mid \exists \text{ final } \rho, obs(\rho) = w\}$.

For a path $\rho$, we define its probability as $\mathcal{P}(\rho) = \mu(s_0) \cdot \prod_{i=1}^{n-1} M(s_i, a_{i+1}, s_{i+1})$. We define the probability of an observation $w \in \Sigma^*$ as $\mathcal{P}(w) = \sum_{\rho | obs(\rho) = w} \mathcal{P}(\rho)$ i.e. the sum of the probabilities over all paths such that $obs(\rho) = w$. We write $\mathcal{P}^A_{\mu_0}$ to denote probability in $A$ with initial distribution $\mu_0$, and use $\mathcal{P}^A_s$ for $\mu_0(s) = 1$. We define $\mathcal{P}_{\text{secret}}(w) = \sum_{\rho | obs(\rho) = w, s^+(\rho) = q_{\text{secret}}} \mathcal{P}(\rho)$ and $\mathcal{P}_{\text{non-secret}}(w) = \sum_{\rho | obs(\rho) = w, s^+(\rho) = q_{\text{non-secret}}} \mathcal{P}(\rho)$ i.e. the probability to observe $w$ and end up in $q_{\text{secret}}$ and $q_{\text{non-secret}}$ respectively. We also define $\mathcal{P}_{\text{stop}}(w) = \mathcal{P}_{\text{non-secret}}(w) + \mathcal{P}_{\text{secret}}(w)$. We assume that from each state $s \in S$, there exists a path from $s$ to $q_{\text{non-secret}}$ or to $q_{\text{secret}}$, because path should stop with probability 1. As the number of states is finite, this ensures that the probability to eventually reach $q_{\text{secret}}$ or $q_{\text{non-secret}}$ is one, that is all paths eventually terminate with probability 1.

*Example 1.* Figure 1 shows an LMC. The initial distribution $\mu_0$ is such that $\mu_0(q_0) = 1$. An example of a path is $\rho = q_0 b q_1 \$ q_{\text{secret}}$, with probability $\mathcal{P}(\rho) = 0.4 \cdot 0.5 = 0.2$. Its observation is $b\$$, with $\mathcal{P}(b\$) = 0.4 \cdot (0.5 + 0.2) = 0.28$. So $\mathcal{P}_{\text{secret}}(b\$) = 0.2$ and $\mathcal{P}_{\text{non-secret}}(b\$) = 0.08$.

## 3 Degree of Opacity

The notion of opacity focuses on characterizing the information flow from the system to an external observer and aims at determining whether a given system secret behavior (i.e., a subset of the behaviors of the system that is considered critical) is kept opaque to outsiders [21]: Given an LMC $A = (S, M, \mu_0)$, we say that $A$ is *opaque* if for all path $\rho$ s.t. $s^+(\rho) = q_{\text{secret}}$, there is a path $\rho'$ s.t. $obs(\rho') = obs(\rho)$ and $s^+(\rho') = q_{\text{non-secret}}$. Consider Figure 2: the LMC on the left (a) is opaque, while the LMC on the right (b) is not opaque, because observation $b\$$ can lead to both $q_{\text{secret}}$ and $q_{\text{non-secret}}$.

It is usually hard to make a system totally opaque, unless its functional part is severely impaired. Most of the time, the practical question is not: can

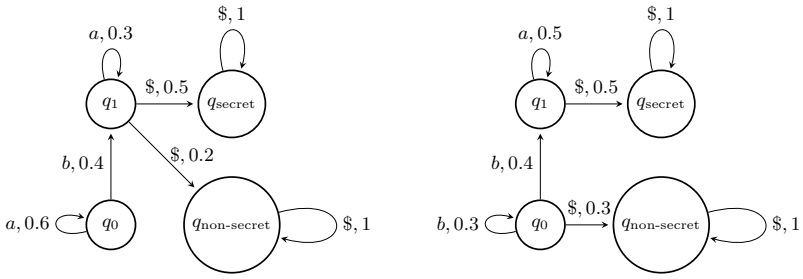


(a) Opaque LMC (b) Non-opaque LMC

Fig. 2: Two LMCs: LMC in (a) is opaque, while LMC in (b) is not.

the system leak some information, but rather how much information does the system leak. We will thus be considering the opacity degree of an observation as a measure about the opacity of a system, with a definition very similar to *restrictive probability of opacity* from [4], and of *$\varepsilon$-disclosure* in [3].

**Definition 1.** *Let $A = (S, M, \mu_0)$ be an LMC. We define the opacity degree of an observation $w$ as $d_A(w) = \mathcal{P}^A_{non\text{-}secret}(w)/\mathcal{P}^A_{stop}(w)$ for $\mathcal{P}^A_{stop}(w) > 0$, and 1 otherwise. We define the* average opacity degree *of $A$ as*

$$d_{avg}(A) = \sum_{w \in \Sigma^*} \mathcal{P}^A_{stop}(w) \cdot d_A(w).$$

We start with a negative result, regarding "worst-case" opacity, whose first part has been shown in [3] (Theorem 3):

**Proposition 1.** *It is undecidable to know, given $\alpha \in (0, 1)$ and an LMC $A$, whether for all observation $w$, $d_A(w) \geq \alpha$ [3]. Further, one cannot even approximate $d_{worst}(A) = \inf_{w \in \Sigma^*} d_A(w)$.*

*Proof.* (Sketch of.) The proof of undecidability is quite standard (see the appendix and [3]), reducing the emptiness problem in Rabin's probabilistic finite automaton (PFA) (defined in the appendix) to our problem.

To prove that $d_{\text{worst}}(A)$ is not approximable, we show in the appendix that the following problem on LMCs is undecidable: Given $\varepsilon > 0$ and an LMC $A$ such that either (1) $d_{\text{worst}}(A) < \varepsilon$, or (2) $d(w\$) > 1 - \varepsilon$ for all observations $w\$$. Deciding which case holds is undecidable, that is, one cannot approximate $d_{\text{worst}}(A)$. We prove this by reducing from the same problem on PFAs, which is undecidable [22]. □

This contrasts with the positive result regarding average opacity:

**Proposition 2.** *$d_{avg}(A)$ can be computed in polynomial time for an LMC $A$.*

*Proof.* We have $d_A(w) = \mathcal{P}_{\text{non-secret}}(w)/\mathcal{P}_{\text{stop}}(w)$. Thus $d_{avg} = \sum_w \mathcal{P}_{\text{non-secret}}(w)$ as $\mathcal{P}_{\text{non-secret}}(w) > 0$ implies that $w$ ends with \$, and thus $\mathcal{P}_{\text{stop}}(w) = \mathcal{P}(w)$: this is equal to the probability to reach state $q_{\text{non-secret}}$. Computing this probability can be done in PTIME [2]. □

Our paper will thus focus inbetween these two extreme cases: For worst-case opacity, we will consider the qualitative cases, that is the cases with $\alpha = 0, 1$.

On the other hand, we will also go beyond average opacity: while the average opacity degree is an interesting value, many very different situations can result in the same average: all paths can have the same opacity degree, or a small proportion of paths can have much smaller opacity degree than others. Consider the case where 10% of paths have much smaller opacity degree: while the overall average opacity degree is high, it means that the system is not opaque in a tangible number of cases (10%).

In order to quantify opacity in a finer way, we will resort to weighted threshold degree, considering paths that have opacity smaller than $\theta$, weighted by the difference of opacity with $\theta$. A sensible choice is to take $\theta$ being e.g. the average opacity, allowing to quantify the weighted average over paths with opacity degree lower than this average.

**Definition 2.** *Let $\theta \in (0, 1)$ and $A$ be an LMC. We define the weighted $\theta$-threshold degree as $d_\theta(A) = \sum_w \mathcal{P}^A(w) \cdot (\theta - min(d^A(w), \theta))$.*

For an observation $w$, we will denote $d_\theta(w) = \theta - min(d^A(w), \theta)$. That is, $d_\theta$ counts only observations $w$ for which the opacity degree is lower than $\theta$: the further away below $\theta$, the more $w$ contributes to the threshold degree, and this contribution is weighted by the probability to see $w$.

## 4 Interval Labelled Markov Chains

### 4.1 Definition and semantics

In general, it is hard to know exactly the probabilities of all the transitions in the system. To model this uncertainty, we introduce Interval LMCs, inspired by Interval Markov Chains and Labelled Markov Chains, where transitions are associated with intervals rather than precise probabilities. We denote by $\mathcal{I}$ the set of closed intervals $[u, v]$ included in $[0, 1]$, i.e. with $0 \leq u \leq v \leq 1$.

**Definition 3.** *An* Interval Labelled Markov Chains (ILMC) *on finite alphabet $\Sigma$ is a tuple $\mathcal{A} = (S, \delta, \mu_0)$ with $S$ the set of states, $\delta : S \times \Sigma \times S \to \mathcal{I}$ is the transition function associating an interval to each transition.*

Given a state $s \in S$, we say that a distribution $\mu : \Sigma \times S \longrightarrow [0, 1]$ is an assignment for $s$ if $\mu(a, s') \in \delta(s, a, s')$ for all $a \in \Sigma, s' \in S$. For $\delta : S \times \Sigma \times S \to \mathcal{I}$, we denote $\delta^-, \delta^+ : S \times \Sigma \times S \to [0, 1]$ such that $\delta(s, a, s') = [\delta^-(s, a, s'), \delta^+(s, a, s')]$ for all $s, a, s'$. Note that an assignment for state $s \in S$ only exists iff $\sum_{s' \in S, a \in \Sigma} \delta^-(s, a, s') \leq 1 \leq \sum_{s' \in S, a \in \Sigma} \delta^+(s, a, s')$. The ILMC is *well-formed* if there exists at least one assignment for each state. In the following, we will consider that all ILMCs are well-formed. For instance, the ILMC in Figure 3 is well-formed: For state $q_3$, the following $\mu_1, \mu_2$ are assignments: $\mu_1(a, q_4) = \frac{1}{2}$ and $\mu_1(a, q_5) = \frac{1}{2}$, but also e.g. $\mu_2(a, q_4) = 0$ and $\mu_2(a, q_5) = 1$.

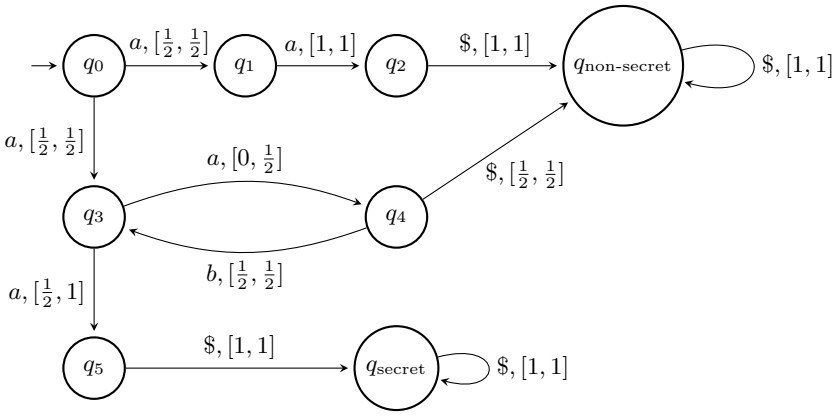


Fig. 3: An ILMC with 8 states.

Two semantics exist for ILMCs, namely *Uncertain Markov Chain(UMC)* and *Interval Markov Decision Process(IMDP)*: the difference is whether the assignment is fixed for each state, or whether they can depend upon the path reaching the state. Let $\mathcal{A} = (S, \delta, \mu_0)$ be an ILMC.

The *UMC semantics* of $\mathcal{A}$ is the set $\text{UMC}(\mathcal{A})$ of LMCs $A = (S, M, \mu_0)$ s.t. for all $s \in S$, we have $M(s, a, s') \in \delta(s, a, s')$.

The *IMDP semantics* of $\mathcal{A}$ is the set $\text{IMDP}(\mathcal{A})$ of (infinite-state) LMCs $A = (S \times (\Sigma \times S)^*, M, \mu_0)$ where states are finite paths in $\mathcal{A}$, and such that for all $\rho$ reaching $s \in S$, for all $s' \in S$, we have $M(\rho, a, \rho\, a\, s') \in \delta(s, a, s')$ and $M(\rho, a, \rho') = 0$ if $\rho'$ is not of the form $\rho\, a\, s$. The name IMDP comes from the fact that probabilities depend upon the history, similarly as schedulers in MDPs (see appendix or [2]).

We require that for all $A \in \text{UMC}(\mathcal{A})$ (resp. $A \in \text{IMDP}(\mathcal{A})$), for all states $s$ of $A$, there is probability 1 to reach $\{q_{\text{secret}}, q_{\text{non-secret}}\}$ from $s$.

*Example 2.* Consider the ILMC $\mathcal{A}$ in Figure 3. Let $A \in \text{UMC}(\mathcal{A})$, and let $p = M_A(q_3, a, q_4)$. There are two cases: either $p > 0$ or $p = 0$. If $p = 0$, then the only stopping observation is $w = aa\$$, and we have $d(w) = \frac{1}{2}$. Otherwise, $p > 0$, and the set of stopping observation is $a(ab)^*a\$$. For all $w \in a(ab)^*a\$$, we have $\mathcal{P}_{\text{non-secret}}(w) > 0$, and thus $d(w) > 0$.

Now, consider $A \in \text{IMDP}(\mathcal{A})$, such that the first time $q_3$ is seen, transition $(q_3, a, q_4)$ is assigned probability $\frac{1}{2}$ (formally, $M_A(q_0aq_3, a, q_0aq_3aq_4) = \frac{1}{2}$ ), and every other time, $(q_3, a, q_4)$ is assigned 0. We have $\mathcal{P}_{\text{secret}}(aaba\$) > 0$ but $\mathcal{P}^A_{\text{non-secret}}(aaba\$) = 0$. Thus, $d^A(aaba\$) = 0$.

It is easy to see that both semantics encompass any LMC $A = (S, M, \mu_0)$: it suffices to set $\delta(t) = [M(t), M(t)]$ for each transition $t$. We obtain $\text{UMC}(\mathcal{A}) = \text{IMDP}(\mathcal{A}) = \{A\}$. Thus, the complexity of computing the opacity degree for ILMCs is at least as hard as for LMCs. In particular, it is undecidable to compute the worst-case opacity degree of an ILMC, and it even cannot be approximated.

Hence we turn to the qualitative questions on opacity degree: is it the case that $d_A(w) = 1$ for all $w$? is it the case that $d_A(w) > 0$ for all $w$? For both semantics, we have two ways to look at the ILMC: either we have an uncertain model of the system, and we ask the question for all LMCs $A$ in the semantics. The second way to look at ILMCs is as a specification, and the actual implementation, chosen by the designer, could be any LMC $A$ in the semantics. In this case, the question becomes: does there exists an $A$ in the semantics ensuring a given property.

Concerning the average opacity degree, standard algorithms [11] computing the optimal probability to reach a set of states in interval Markov chains yield:

**Proposition 3.** *Computing* $\min_{A \in IMDP(\mathcal{A})} d_{avg}(A) = \min_{A \in UMC(\mathcal{A})} d_{avg}(A)$ *and* $\max_{A \in IMDP(\mathcal{A})} d_{avg}(A) = \max_{A \in UMC(\mathcal{A})} d_{avg}(A)$ *can be done in PTIME.*

In the following, we will focus on computing $\inf_{A \in \text{ LMC}(\mathcal{A})} d_\theta(A)$ and $\sup_{A \in \text{ LMC}(\mathcal{A})} d_\theta(A)$, considering $d_\theta$ instead of the average opacity degree $d_{avg}$.

### 4.2 Uniform reachability set for ILMCs

In this subsection we give a construction which will be used in Theorems 2, 4, 6 to handle ILMCs. Let $\mathcal{A}$ be an ILMC. Let $F \subseteq S$ be a subset of states. We want to compute the set $X(F) = \{s \mid \forall A \in \text{ LMC}(\mathcal{A}),$ there exists a path from s to F$\}$ of states from which $F$ can be reached in all $A \in \text{LMC}(\mathcal{A})$, for LMC $\in \{$UMC,IMDP$\}$.

**Lemma 1.** *Given a set $F$, one can compute $X(F)$ in PTIME.*

In particular, $X(F)$ does not depend upon the choice of the semantics.

*Proof.* We initialize $X_0 = F$. We then compute inductively $X^{i+1}$ from $X^i$:

- We look for a state $s \in S \setminus X^i$ such that, either there is a transition $t$ from $s$ to $X^i$ with $\delta^-(t) > 0$, or the sum of $\delta^+(t')$ over transition from $s$ to $S \setminus X_i$ is $p < 1$. In both cases, we set $X^{i+1} = X^i \cup \{s\}$.
- If we cannot find such a state $s$, then we return $X(F) = X^i$.
- The process ends after at most $|S|$ iterations (the construction can be done in PTIME, but not in NLOGSPACE).

At the end of the algorithm, we have for all $s \in X(F)$, for all $A \in \text{LMC}(\mathcal{A})$, there is a path from $s$ to $F$ in $A$. On the other hand, there is a $A \in \text{LMC}(\mathcal{A})$ such that for all $s \notin X(F)$, there is no path from $s$ to $F$.

Let $i$ with $X^i = X(F)$. For the second part, it is easy to build an LMC $A = (S, M, \mu_0) \in \text{UMC}(\mathcal{A})$ (and thus also in IMDP$(\mathcal{A})$) such that for all $s \in S \setminus X$, for all $t \in X$, $M(s, a, t) = 0$. By construction, we can set all transitions from $s$ to $X^i$ to 0, and we have the sum of $\delta^+(t')$ over transition from $s$ to $S \setminus X_i = S \setminus X$ is $p \geq 1$. In particular, we can set an assignment from $s$ staying in $S \setminus X$ and respecting the requirement of $\mathcal{A}$.

We prove the first part by induction of the minimal $i$ such that $s \in X^i$. First, the assertion is trivial for $s \in X^0 = F$. Inductively assuming it is true for all $s \in X^i$, for any $A \in \text{LMC}(\mathcal{A})$, taking $s \in X^{i+1} \setminus X^i$, either there is a transition

$(s, a, s')$ with $s' \in X^i$ and $\delta^-(s, a, s') > 0$ and $M(s, a, s') \geq \delta^-(s, a, s') > 0$ and we have the proof. Otherwise, we have the sum of $\delta^+(t')$ over transition from $s$ to $S \setminus X_i = S \setminus X$ is $p < 1$, which means that the assignment in $A$ must have a $s' \in X^i$ with $M(s, a, s') > 0$. □

## 5 Approximating the weighted $\theta$-threshold degree

As explained in Section 3, weighted $\theta$-threshold degrees allows us to quantify opacity in a more tractable way than worst-case degree which is non approximable. Further, considering several values of $\theta$ allows us to understand precisely the amount of runs in a particular opacity degree range.

We show that while one cannot compute exactly the weighted $\theta$-threshold degree for $0 < \theta < 1$, one can approximate its value as close as desired. We can approximate the weighted $\theta$-threshold degree very efficiently with high confidence, and we can also approximate it with total certainty, although not as efficiently.

To do so, we first prove that considering runs of size at most $k$, for $k$ not so large, suffices to approximate the weighted $\theta$-threshold degree, because almost all of these runs are stopped.

**Lemma 2.** *Let $A$ be an LMC, and let $\varepsilon > 0$. Then there exists $k$ such that $\mathcal{P}(\Sigma^{>k} \cap L(A)) \leq \varepsilon$.*

*Proof.* By hypothesis, $\forall s \in S, \exists \rho_s, s^-(\rho_s) = s, s^+(\rho_s) \in \{q_{\text{secret}}, q_{\text{non-secret}}\}$. Let $n_s < \infty$ be the length of $\rho_s$ and $p_s = \mathcal{P}(\rho_s) < \infty$ be the probability of $\rho_s$. Let $n = \max_{s \in S} n_s$, and $p = \min_{s \in S} p_s$. We use max and not sup as there is only a finite number of states in $S$. Hence $n < \infty$ is finite and $p > 0$. Hence the probability not to stop after $n$ steps is at most $1 - p$, whatever the initial state. These non-stopping paths end in some $s \in S$, and hence after $2n$ steps, there is probability at most $(1-p)^2$ to be non-stopping. By a trivial induction, we obtain probability $(1-p)^{k'}$ to be non-stopping after $k' \cdot n$ steps. As $0 < 1 - p < 1$, we have $(1-p)^{k'} \underset{k' \to +\infty}{\longrightarrow} 0$. Hence there exists $k'$ such that $(1-p)^{k'} < \varepsilon$, and we deduce that with $k = k'n$, $\mathcal{P}(\Sigma^{>k} \cap L(A)) \leq \varepsilon$. □

This lemma allows us to approximate accurately the weighted degree by considering the finite set of runs of size at most $k$.

**Theorem 1.** *Input: an LMC $A = (S, M, \mu_0)$, $0 < \theta < 1$ and $\lambda \in [0, 1]$,*

1. *Knowing whether $d_\theta(A) > \lambda$ is undecidable. However,*
2. *Given an error bound $\varepsilon > 0$, and a high confidence value $\delta < 1$, one can compute a value $\gamma$ such that $P(|d_\theta(A) - \gamma| \leq \varepsilon) \geq \delta$ in polynomial time $O(|A| \frac{log(1/(1-\delta))}{\varepsilon^2})$.*
3. *Given an error bound $\varepsilon > 0$, it is possible to compute a value $\gamma$ such that $|d_\theta(A) - \lambda| \leq \varepsilon$ in PSPACE in $|A|$ and $|\varepsilon|$.*

*Proof.* First, point 1) can be reduced with the undecidability in Prop. 1: Given $\alpha > 0$, for all observation $w$ of $A$, we have that $d_A(w) \geq \alpha$ iff $d_\alpha(A) = 0$. Hence $d_\alpha(A) = 0$ is undecidable, as well as its complementary $d_\alpha(A) > 0$ (i.e. taking $\theta = \alpha$ and $\lambda = 0$).

For point 2), we approximate the average of the probability distribution of $d_\theta$ values on words conditionally to a length bound on these words. We also prove that this average is close to the average of the probability distribution of $d_\theta$ values on all words. This is inspired by a similar result from [23].

Let $Y_{\leq k} : (\Sigma^{\leq k} \cap L(A)) \to [0,1]$ be the random variable associated to the values of $d_\theta$ of words of length smaller than $k$ (the $k$ of lemma 2) and $Y$ be the one associated to the value of $d_\theta$ for all words. We have $d_\theta(A) = \mathbb{E}(Y)$. Next, we will show that $\mathbb{E}(Y)$ is close to $\mathbb{E}(Y_{\leq k})$ and that we can approximate $Y_{\leq k}$. The interest of $Y_{\leq k}$ is that its samples will have a polynomial length.

We have that $\mathbb{E}(Y) = (\sum_{w \in \Sigma^{\leq k} \cap L(A)} \mathcal{P}^A(w) d_\theta(w))/(\mathcal{P}(\Sigma^{\leq k} \cap L(A)))$ and $\mathbb{E}(Y_{\leq k}) = \sum_{w \in \Sigma^{\leq k} \cap L(A)} \mathcal{P}^A(w) d_\theta(w) + \sum_{w \in \Sigma^{>k} \cap L(A)} \mathcal{P}^A(w) d_\theta(w)$. Then,

$$\begin{aligned}
|\mathbb{E}(Y) - \mathbb{E}(Y_{\leq k})| \leq \quad & |\sum_{w \in \Sigma^{\leq k} \cap L(A)} \mathcal{P}^A(w) d_\theta(w)| (1 - 1/(\mathcal{P}(\Sigma^{\leq k} \cap L(A)))) \\
& + |\sum_{w \in \Sigma^{>k} \cap L(A)} \mathcal{P}^A(w) d_\theta(w)| \\
\leq \quad & (1 - 1/(\mathcal{P}(\Sigma^{\leq k} \cap L(A))) + \varepsilon \\
\leq \quad & 3\varepsilon
\end{aligned}$$

Let $X_1, \ldots, X_n$ be independent random variables following the same law as $Y_{\leq k}$ and $X = \sum_{i=1}^n X_i$. Then, for all $i$, we have that $0 \leq X \leq 1$ and we denote $\mu = \mathbb{E}(X) = n \cdot \mathbb{E}(Y_{\leq k})$ and $\gamma = X/n$. Applying Chernoff's inequality, we obtain that for all $\tau > 0$

$$P(X \geq \mu(1+\tau)) \leq e^{\frac{-2\tau^2\mu^2}{n}}$$

$$P(X \geq \mu(1-\tau)) \leq e^{\frac{-\tau^2\mu^2}{n}}$$

Thus, $P(\gamma \geq \mathbb{E}(Y_{\leq k}) + \mu\tau/n) \leq e^{\frac{-2\tau^2\mu^2}{n}}$ and by denoting $\varepsilon = \mu\tau/n$, we obtain $P(\gamma \geq \mathbb{E}(Y_{\leq k}) + \varepsilon) \leq e^{-2n\varepsilon^2}$. Similarly, we obtain that $P(\gamma \geq \mathbb{E}(Y_{\leq k}) - \varepsilon) \leq e^{-n\varepsilon^2}$.

So, the probability that the $\gamma$ diverges from $\mathbb{E}(Y_{\leq k})$ by more than $\varepsilon$ is smaller than $2e^{-n\varepsilon^2}$. In order to have confidence $\delta$, we can take $n \geq \log(\frac{2}{(1-\delta)})\varepsilon^{-2}$.

By combining these two results, we obtain that with $n \geq \log(\frac{2}{(1-\delta)})\varepsilon^{-2}$ samples, $\mathcal{P}(|\gamma - \mathbb{E}(Y)| \leq 4\varepsilon) \geq \delta$.

For point 3), we sketch the proof: we now explain how to approximate $d_\theta(A)$ using Lemma 2, inspired by [13]. We have $d_\theta(A) = \sum_w \mathcal{P}(w) \cdot d_\theta(w)$. We apply Lemma 2 above to obtain a $k$. We have $d_\theta = \sum_{w \in (L(A) \cap \Sigma^{\leq k})} \mathcal{P}(w) \cdot d_\theta(w) + \sum_{w \in (L(A) \cap \Sigma^{>k})} \mathcal{P}(w) \cdot d_\theta(w)$. We can compute in EXPTIME the first term, and

ignore the second term, as it is smaller than $\varepsilon$. While we cannot compute the first term in PSPACE, we can approximate it in PSPACE using Ladner's result [15]. This is sufficient to conclude (see appendix). □

We now show that we can approximate in the same way $\min_{A\in\mathrm{LMC}(\mathcal{A})}(d_\theta(A))$ and $\max_{A\in\mathrm{LMC}(\mathcal{A})}(d_\theta(A))$ for LMC $\in \{\text{UMC,IMDP}\}$. That is, one can approximate the degree for all 4 cases of ILMCs. We proceed in the same way as for LMCs: we prove that there exists a constant $k$, uniform over all LMCs in $\mathrm{UMC}(\mathcal{A})$ and $\mathrm{IMDP}(\mathcal{A})$, after which almost all runs are stopped, using Lemma 1.

**Lemma 3.** *Let $\mathcal{A}$ be an ILMC. Let $\varepsilon > 0$. Then there exists $K$ such that $\forall A \in LMC(\mathcal{A}), \mathcal{P}(\Sigma^{>K} \cap L(A)) \leq \varepsilon$, for LMC $\in \{UMC,IMDP\}$.*

*Proof.* We use Lemma 1 with $F = \{q_{\text{secret}}, q_{\text{non-secret}}\}$, and obtain the set $X(F)$ of states from which $F$ can be reached in every LMC $A \in \mathrm{LMC}(\mathcal{A})$. By definition of ILMCs, we should have $X(F) = S$. During the inductive construction of Lemma 1, we compute for each state $s \in X^i$ a probability $p_s > 0$ such that for all $A \in \mathrm{LMC}(\mathcal{A})$, we have that the probability to reach $\{q_{\text{secret}}, q_{\text{non-secret}}\}$ from $s$ is at least $p_s$. We initialize $p_s = 1$ for $s \in X^0 = \{q_s, q_{ns}\}$. We then compute inductively $p_s$ in $s \in X^{i+1}$ from $p_s$ for $s \in X^i$: When a state $s$ is added to $X^{i+1}$, either we have a transition $t$ from $s$ to $X^i$ with $\delta^-(t) > 0$, and we set $p_s = \delta^-(t) \min_{x\in X^i} p_x$, or the sum of $\delta^+(t')$ over transition from $s$ to $S \setminus X^i$ is $p < 1$, and we set $p_s = (1-p)\min_{x\in X^i} p_x$.

Hence for all $s \in S$, we have $p_s > 0$. In particular, for all $A$, there is probability at least $\min_{s\in S} p_s > 0$ to terminate in at most $|S|$ steps. From there, we conclude similarly as in the proof of Lemma 2. □

We can now state an approximability result for ILMCs. Notice that we do not state the case with high confidence as we do not see how to improve the complexity better than with full certainty in the case of ILMCs:

**Theorem 2.** *Given an ILMC $\mathcal{A} = (S, \delta, \mu_0)$, and a constant $\varepsilon > 0$, one can compute in EXPTIME a number $\lambda_{min} \in [0,1]$ such that $|\min_{A\in LMC(\mathcal{A})}(d_\theta(A)) - \lambda_{min}| \leq \varepsilon$. Also, one can compute in EXPSPACE a number $\lambda_{max} \in [0,1]$ such that $|\max_{A\in LMC(\mathcal{A})}(d_\theta(A)) - \lambda_{max}| \leq \varepsilon$, for LMC $\in \{UMC,IMDP\}$.*

*Proof (Sketch of.).* We prove Theorem 2 using Lemma 3. Let $\mathrm{contrib}^A(w) = \max(0, \theta\mathcal{P}^A_{\text{stop}}(w) - \mathcal{P}^A_{\text{non-secret}}(w))$. Then $\sum_{w\in\Sigma^{\leq K}} \mathrm{contrib}^A(w) \leq d_\theta(A) \leq \sum_{w\in\Sigma^{\leq K}} \mathrm{contrib}^A(w) + \varepsilon$. We can compute $\min_{A\in\mathrm{LMC}(\mathcal{A})}(\sum_{w\in\Sigma^{\leq K}} \mathrm{contrib}^A(w))$ using linear optimization with an exponential number of variables. We can compute $\max_{A\in\mathrm{LMC}(\mathcal{A})}(\sum_{w\in\Sigma^{\leq K}} \mathrm{contrib}^A(w))$ using the first order theory of reals (see appendix).

## 6 Qualitative worst-case opacity degree

Our results concerning qualitative worst-case opacity are summed-up in the table of Figure 4. We prove decidability for all the variant considered, and exhibit their

| Model | $d_A(w) > 0$ for all $w \in \Sigma^*$ | $d_A(w) = 1$ for all $w \in \Sigma^*$ |
|---|---|---|
| LMC $A$ | PSPACE-Complete (Thm. 3) | LOGSPACE (Thm. 5) |
| $\forall A \in$ UMC($\mathcal{A}$) | PSPACE-Complete (Thm. 3) | LOGSPACE (Thm. 5) |
| $\forall A \in$ IMDP($\mathcal{A}$) | PSPACE-Complete (Thm. 3) | LOGSPACE (Thm. 5) |
| $\exists A \in$ UMC($\mathcal{A}$) | PSPACE-Complete (Thm. 4) | P-Complete (Thm. 6) |
| $\exists A \in$ IMDP($\mathcal{A}$) | EXPTIME-Complete (Thm. 4) | P-Complete (Thm. 6) |

Fig. 4: Exact complexity for qualitative Worst-case Opacity. Notice that the decidability of cases handled by Theorem 4 can be derived from [5].

exact complexities. Interestingly, checking whether opacity holds no matter the system following an interval LMCs ($\forall A \in$ UMC($\mathcal{A}$) or IMDP($\mathcal{A}$)) is not more complex than checking it when the LMC is entirely known. Further, it is usually slightly more complex (EXPTIME-complete instead of PSPACE-complete, or P-complete instead of LOGSPACE) when trying to synthesize one opaque LMC satisfying an ILMC specification ($\exists A \in$ UMC($\mathcal{A}$) or IMDP($\mathcal{A}$)), except for one case where it is not harder ($\exists A \in UMC(\mathcal{A}), d_A(w) > 0$ for all $w \in \Sigma^*$). We provide proofs here for the most interesting results, in particular Theorem 4 showing this discrepancy between the UMC and the IMDP semantics. The other proofs can be found in the appendix.

**No run fully reveals whether it is secret** Consider now opacity with the meaning that no run fully reveals whether it is secret. This question is harder than a plain reachability question, as we wonder whether for all observation $w$, $\mathcal{P}_{\text{secret}}(w) > 0$ implies $\mathcal{P}_{\text{non-secret}}(w) > 0$. For that, we need to follow the set of states reachable with the same observation. Let $w \in \Sigma^*$. The belief associated with $w$ is $B(w) = \{s \mid s_0 \rightarrow^w s\}$. The deterministic belief automaton associated with $A$ is $\mathcal{B} = (2^S, T, B_0)$, where $B_0 = \{q \mid \mu_0(q) > 0\}$ and $(S, a, S') \in T$ iff $S' = \{q' \mid \exists q \in S, M(q, a, q') > 0\}$. The question whether there exists an observation $w$ with $\mathcal{P}_{\text{non-secret}}(w) = 0$ and $\mathcal{P}_{\text{secret}}(w) > 0$ (i.e. a run revealing the secret) is equivalent with the existence of a belief $B$ of $\mathcal{B}$ reachable from $B_0$ containing $q_{\text{secret}}$ but not $q_{\text{non-secret}}$.

**Theorem 3.** *Given one LMC $A$, testing whether $d_A(w) > 0$ for all $w \in \Sigma^*$ is PSPACE-complete. For an ILMC $\mathcal{A}$, testing $d_A(w) > 0$ for all $w \in \Sigma^*$ and for all $A \in$ IMDP($\mathcal{A}$) or $A \in$ UMC($\mathcal{A}$) is also PSPACE-complete.*

*Proof (Sketch of.).* For an LMC, checking for the existence of a belief containing $q_{\text{secret}}$ but not $q_{\text{non-secret}}$ can be performed on the fly in PSPACE. We obtain the hardness by reduction from the language equality of two automata.

For ILMCs under the UMC semantics, we can guess the non-zero transitions such that one can reach $q_{\text{secret}}$ but not $q_{\text{non-secret}}$. For ILMC under the IMDP semantics, this is not possible as there are too many possible LMCs in IMDP($\mathcal{A}$). However, as the choice of transition can be different between different path, one

can guess directly a path in any LMC of IMDP($\mathcal{A}$) that reaches $q_{\text{secret}}$ but not $q_{\text{non-secret}}$ in PSPACE. □

When considering the synthesis problem, it becomes harder for ILMCs viewed as a specification with the IMDP semantics, namely EXPTIME-complete rather than PSPACE-complete. Notice that decidability (but not the exact complexity, nor the fact that complexity is different between UMC and IMDP semantics) is known from [5]. Their Corollary 11 states that it is decidable whether there exists a MDP for which no observation has an opacity degree 0 (qualitative opacity degree). It could be tempting to believe that our result is related with the EXPTIME-completeness result of Theorem 15 from [5]. The hardness result in our proof is actually much stronger: first, the synthesis in [5] is done based on the observation, but not on the exact run, while the LMC we synthesized from an ILMC knows the exact run. Although we seem to be in a much less complex setting, we show that the complexity of both problems is actually the same. Another point is that to obtain decidability, [5] needs to restrict to finite memory controllers, which is not necessary in our setting. Notice that without finite memory, opacity is undecidable in the setting of [5] (point 2) of Theorem 13 from [5]).

**Theorem 4.** *Given an ILMC $\mathcal{A}$, testing whether there exists $A \in$ UMC($\mathcal{A}$) such that $d_A(w) > 0$ for all $w \in \Sigma^*$ is PSPACE-complete, while testing whether there exists $A \in$ IMDP($\mathcal{A}$) such that $d_A(w) > 0$ for all $w \in \Sigma^*$ is EXPTIME-complete.*

*Proof.* Under the UMC semantics, the proof is similar as in Theorem 3: we can guess the support of transitions of an LMC $A \in \text{UMC}(\mathcal{A})$ in PSPACE, and check in PSPACE whether $d_A(w) > 0$ for all $w \in \Sigma^*$. PSPACE-hardness comes from the result on LMCs in Theorem 3.

We now turn to the IMDP semantics: We prove that this problem is EXPTIME-complete. We first show that it is in EXPTIME. We consider the non-deterministic belief automaton $\mathcal{B}'$ associated with ILMC $\mathcal{A}$, representing all possible choices of synthesis. It differs from $\mathcal{B}$ on the (non-deterministic) transition relation $\Delta'$:

- For $B, B' \in 2^S, a \in \Sigma$, we have $(B, a, B') \in \Delta'$ if and only if there exists a $\sigma : S \times \Sigma \mapsto Dist(S)$ with $\sigma(s, a, s') \in \delta(s, a, s')$ for all $s, a, s'$ such that $B' = \{s' \in S \mid \exists s \in B, \sigma(s, a, s') > 0\}$ (notice that as this is the ILMC semantics, the next step needs not use the same scheduler $\sigma$).

We want to know whether it is possible to avoid belief $\{q_{\text{secret}}\}$. For that, we use Lemma 1 on the non-deterministic belief automaton $\mathcal{B}'$ (which is of exponential size), using $F = \{q_{\text{secret}}\}$. It allows to compute $X(F)$ in EXPTIME, the set of beliefs such that $F$ is unavoidable. If $B_0 \in X(F)$, then we know that there is no $A \in \text{IMDP}(\mathcal{A})$ such that $d_A(w) > 0$ for all $w \in \Sigma^*$. Otherwise, there is an $A \in \text{IMDP}(\mathcal{A})$ such that $d_A(w) > 0$ for all $w \in \Sigma^*$, choosing an LMC following the choices avoiding $F = \{q_{\text{secret}}\}$.

Last, we prove that this problem is EXPTIME-hard, by reducing the problem of sure winning in a two player game over a polynomial space alternating Turing

machine (ATM), which is an ALT-PSPACE by definition, which is equal to EXP-TIME [9]. An alternating Turing machine is a tuple $\mathcal{M} = (S, q_0, g, \Sigma_i, \Sigma_t, \Delta, F)$ where :-

- $S$ is a finite set of control states;
- $q_0 \in S$ is the initial control state;
- $g : S \to \{\vee, \wedge\}$;
- $\Sigma_i = \{0, 1\}$ is the input alphabet;
- $\Sigma_t = \{0, 1, 2\}$ is the tape alphabet and 2 is the blank symbol;
- $\Delta \subseteq S \times \Sigma_t \times S \times \Sigma_t \times \{-1, 1\}$ where -1(1) denotes left and right direction respectively is a transition relation;
- $F \subseteq S$ is the set of accepting states

Without loss of generality, we make the hypothesis that the initial control state of the machine is a $\vee$-state and that transitions link $\vee$-state to $\wedge$-state and vice versa. Player 1 selects the transition from $\vee-$state to $\wedge-$state while player 2 selects the transition from $\wedge-$state to $\vee-$state. The main problem is to encode the tape: encoding it explicitly would result an exponential number of states. Instead, the ILMC asks the observation to tell him the symbol. For each observation which does not correspond to the tape symbol, the ILMC as a chance to catch it, sending to state $q_{\text{non-secret}}$. Hence, it only needs to handle the correct observation. Given an ATM with with a tape of size $k$, we define the ILMC $\mathcal{A} = (S', \delta, \mu_0)$ over alphabet $\Delta \uplus \Sigma \uplus \{\$\}$ by the following steps:-

- $S'$ contains the following two types of states: the first type is of the form $\{(i, q, a, l), (j, b)\} \in S'$ where $1 \leq i \leq n$ denotes the position in the input, $q \in S$ denotes the current state control state, $a \in \Sigma_t$ denotes the tape symbol at position i, $l \in \{\vee, \wedge\}$ denotes whether player 1 or player 2 has to select next move, $1 \leq j \leq n, b \in \Sigma_t$ which is an information stored meaning that the $j$ symbol of the tape is a $b$.
  The second type is of the form $\{(i, q, l), (j, b)\} \in S'$ where all variables have the same meaning, but we don't have letter $a \in \Sigma_t$,
- $S'$ also contains special states $q_{\text{secret}}, q_{\text{non-secret}}, q_{cheat}$,
- For all $i \leq k$, we have $\mu_0(\{(1, q_0, 2, \vee), (i, 2)\} = 1/k$, i.e. we start with an empty input tape made of 2's.

The transition function is defined as follows:

- For $q \in F$ final, we have $\delta(\{(i, q, a, l), (j, b)\}, \$, q_{\text{secret}}) = [1, 1]$ for all $i, a, l, j, b$.
- From $q_{cheat}$, we have $\delta(q_{cheat}, a, q_{cheat}) = [0.3, 0.3]$ for all $a \in \{0, 1, 2\}$, and $\delta(q_{cheat}, \$, q_{\text{non-secret}}) = [0.1, 0.1]$.
- For all state of the form $s = \{(i, q, op), (j, b)\}$ with $op \in \{\vee, \wedge\}$, we have $\delta(s, a, q_{cheat}) = [1/3, 1/3]$ for $j = i$ but $a \neq b$, and $\delta(s, a, \{(i, q, a, op), (j, b)\}) = [1/3, 1/3]$ in any other case (remember there are only 3 possible tape symbols $a = 0, 1, 2$),

– For all state of the form $s = \{(i, q, a, \vee), (j, b)\}$, for all transition $t = (q, a, p, c, k) \in \Delta$ of the Turing machine, we have both $\delta(s, t, \{(i + k, p, \wedge), (j, b)\}) = [0, 1]$ (keep the previous information $(j, b)$), and $\delta(s, t, \{(i+k, p, \wedge), (i, c)\}) = [0, 1]$ (keep the new information that symbol $c$ is at the $i$-th position of the tape). Player $\vee$ can choose a particular transition $t$ by choosing an $A$ where $\delta(s, t', \{(i + k, p, \wedge), (j, b)\}) = 0$ for all $t' \neq t$, and it will allow the two transition labeled by $t$ with probability $1/2$ each,
– For all state of the form $s = \{(i, q, a, \wedge), (j, b)\}$, for all transition $t = (q, a, p, c, k) \in \Delta$ of the Turing machine, we have $\delta(s, t, \{(i + k, p, \wedge), (j, b)\}) = [\frac{1}{2n}, \frac{1}{2n}]$ and $\delta(s, t, \{(i + k, p, \wedge), (i, c)\}) = [\frac{1}{2n}, \frac{1}{2n}]$, where $n$ is the number of transitions from $s$. Player $\vee$ needs to accommodate all possible choices for $t$.

Notice that observations $w$ with $P(w) > 0$ are of the form $w = \tau_1 a_1 \tau_2 a_2 \cdots$, where transition $\tau_i, i$ odd chosen by $\vee$ then $\tau_i, i$ even chosen by $\wedge$ alternates with observations $a_i$ describing the current tape symbol.

Assume that there is a strategy $\sigma$ for player $\vee$ such for all choice from nodes $\wedge$, we can avoid reaching final states $F$ in the alternating Turing machine. Let $A \in \text{IMDP}(\mathcal{A})$ be the LMC associated with this strategy. Let $w = \tau_1 a_1 \tau_2 a_2 \cdots$ an observation with $P_A(w) > 0$. Consider $\tau_1 \tau_2 \cdots$ following the strategy $\sigma$ avoiding $F$. Consider $b_1, \ldots$ the sequence of symbol on the tape at the position of the head when we follow $\tau_1 \tau_2 \cdots$. Assume that $a_n = b_n$ for all $n$: by definition, this observation $w$ will avoid $F$, and thus $P_s(w) = 0$ and $d_A(w) = 1 > 0$. Now, consider the first symbol $a_n \neq b_n$, and let $i$ the head position after $\tau_1, \ldots, \tau_n$: there is probability $p > 0$ to reach state $\{(i, q, op), (b_n, i)\}$ after $\tau_1 a_1 \tau_2 a_2 \cdots \tau_n$, and thus when $a_n$ is observed, we have probability $> 0$ to reach $q_{cheat}$. From there, there will be probability $> 0$ to reach $q_{\text{non-secret}}$ no matter the observation, and in particular $d_A(w) > 0$. Hence for all observation $w$, we have $d_A(w) > 0$.

On the other hand, assume that there is a strategy $\tau$ from $\wedge$ to reach $F$ (the game is determined, so such a strategy exists if there is no strategy for $\vee$). Take any $A \in \text{ILMC}(\mathcal{A})$, and any maximal observation $w = \tau_1 a_1 \tau_2 a_2 \cdots$ corresponding to strategy $\tau$ and $A$, with $a_i$ following the sequence $\tau_1 \tau_2, \cdots$. For this observation, we have probability 0 to reach $q_{cheat}$ and hence probability 0 to reach $q_{\text{non-secret}}$. However, as $\tau_1 \tau_2, \cdots$ eventually reaches $F$, observation $w$ eventually reaches $q_{\text{secret}}$, and $d_A(w) = 0$. That is, for all $A \in \text{ILMC}(\mathcal{A})$, we have $d_A(w) = 0$ for some $w \in \Sigma^*$. □

This is interesting since it shows that in this case, dealing with the IMDP semantics is harder than dealing with the UMC semantics. This contrasts with model checking, where model checking the UMC semantics is harder than model checking the IMDP semantics [11]. Notice that the decidability of these cases have been shown for MDPs in [5].

**Total opacity** The situation is quite similar when asking when the system is totally opaque: $d_A(w) = 1$ for all $w$. However, the complexity does not change when we consider the ILMC as a specification under the UMC semantics.

**Theorem 5.** *Given an LMC $A$, testing whether $d_A(w) = 1$ for all $w \in \Sigma^*$ can be done in NLOGSPACE. For an ILMC $\mathcal{A}$, checking whether $d_A(w) = 1$ for all $w \in \Sigma^*$ and for all $A \in$ LMC($\mathcal{A}$) is LOGSPACE, for LMC $\in \{$UMC, IMDP$\}$.*

*Proof.* $d_A(w) = 1$ for all $w \in \Sigma^*$ is equivalent with $q_{\text{secret}}$ not being reachable. For an LMC $A$, we thus just need to use a standard graph reachability algorithm, which is NLOGSPACE. For ILMCs, there is no difference between the UMC and the IMDP semantics as we need to handle a reachability objective (positional strategies are sufficient to win reachability objectives in MDPs [2]).

Now, when we view an ILMC as describing an uncertain system, it suffices to consider an LMC $A_{\max} = (S, M, \mu_0)$ such that for all transitions $t$ of $\mathcal{A}$ with $\delta^+(t) > 0$, $M(t) > 0$. Then $d_A(w) = 1$ for all $w \in \Sigma^*, A \in \text{LMC}(\mathcal{A})$ iff $d_{(A_{\max})}(w) = 1$ for all $w \in \Sigma^*$, which can be done in NLOGSPACE.

However, if one considers an ILMC as a specification to choose from, the complexity becomes slightly higher:

**Theorem 6.** *Given an ILMC $\mathcal{A}$, testing whether there exists $A \in$ UMC($\mathcal{A}$) or $A \in$ IMDP($\mathcal{A}$) such that $d_A(w) = 1$ for all $w \in \Sigma^*$ is PTIME-complete.*

*Proof.* We use Lemma 1 on $F = \{q_{\text{secret}}\}$ to compute the set $X(F)$ of states which can reach $q_{\text{secret}}$ in all LMCs $A \in \text{LMC}(\mathcal{A})$. We let Init$=\{s \mid \mu_0(s) > 0\}$. We have that there exists $A \in \text{LMC}(\mathcal{A})$ with $d_A(w) = 1$ for all $w \in \Sigma^*$ iff $\text{Init} \cap X(F) = \emptyset$. We prove the PTIME-completeness by reduction with safety in MDP (is there a strategy avoiding a set of state), which is PTIME-complete [7], [14].

## 7 Conclusion

Opacity degrees allow one to understand how much information a system can disclose through observable events: we show that weighted degrees are easier to compute than worst-case degrees, as they can be $\varepsilon$-approximated with high confidence in polynomial time, while worst case degrees cannot be approximated.

To handle model uncertainties, we defined and studied *Interval Labelled* Markov Chains, for the first time as far as we know. We provided methods to handle them in a complexity close to but sometimes slightly higher than the complexity for LMCs. We did not find an exponential gap, except in one case (approximating the weighted degree for an ILMC seen as a specification). In this partially observable setting, there can be differences between the UMC and the IMDP semantics: our most interesting result in the respect of ILMCs is to show that checking whether one can design an LMC not totally revealing and respecting an ILMC specification is EXPTIME-complete for the IMDP semantics, while it is PSPACE-complete for the UMC semantics.

## References


1. S. Akshay, Hugo Bazille, Eric Fabre, Blaise Genest. *Classification among hidden Markov models*. FSTTCS'19, 29:1-29:14, 2019.
2. Christel Baier, Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
3. Béatrice Bérard, Serge Haddad, Engel Lefaucheux. *Probabilistic Disclosure: Maximisation vs. Minimisation*. FSTTCS'17, pp.13:1-13:14, 2017.
4. Béatrice Bérard, John Mullins, Mathieu Sassolas. *Quantifying opacity*. Mathematical Structures in Computer Science 25(2): pp.361-403, 2015.
5. Béatrice Bérard, Krishnendu Chatterjee, Nathalie Sznajder. *Probabilistic opacity for Markov decision processes*. Information Processing Letters, pp.52-59, 2015.
6. Andrea Bianco, Luca de Alfaro. Model checking of probabilistic and nondeterministic systems. In *FSTTCS'95*, pp.499-513, 1995.
7. C. Beeri. On the membership problem for functional and multivalued dependencies in relational databases. *ACM Transactions on Database Systems*, pp.5:241-259, 1980.
8. Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger. Qualitative Analysis of Partially-observable Markov Decision Processes, *International Symposium on Mathematical Foundations of Computer Science*, pp.258-269, 2010.
9. Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Jean-Francois Raskin. Algorithm for Omega-Regular Games with Imperfect Information, *Logical Methods in Computer Science* 3(3), pp.287-302, 2007.
10. Kostas Chatzikokolakis, Catuscia Palamidessi. A Framework to Analyze Probabilistic Protocols and its Application to the Partial Secrets Exchange. *Theoretical Computer Science*, 389(3):512-527, 2007.
11. Taolue Chen, Tingting Han, Marta Z. Kwiatkowska. On the complexity of model checking interval-valued discrete time Markov chains. *Information Processing Letters* 113(7): pp.210-216, 2013.
12. Jérémy Dubreil, Philippe Darondeau, Hervé Marchand. Supervisory Control for Opacity. *IEEE Transactions on Automatic Control* 55(5): pp.1089-1100, 2010.
13. Stefan Kiefer: On Computing the Total Variation Distance of Hidden Markov Models. *ICALP'18*, pp.130:1-130:13, 2018.
14. N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences* pp.22:384-406, 1981.
15. R. E. Ladner. Polynomial space counting problems. *SIAM Journal on Computing*, 18(6), pp.1087-1097, 1989.
16. Omid Madani, Steve Hanks, Anne Condon. On the undecidability of probabilistic planning and related stochastic optimization problems, *Artificial Intelligence* 147(1-2), pp.5-34, 2003.
17. Laurent Mazaré. Decidability of opacity with non-atomic keys. *FAST'04*, pp.71-84, 2005.
18. Martin L. Puterman. Markov Decision Processes: Discrete Stochastic Dynamic Programming, Wiley & Sons, 1994.
19. Anooshiravan Saboori, Christoforos N. Hadjicostis. Notions of security and opacity in discrete event systems. *CDC'07*, pp.5056-5061, IEEE, 2007.
20. Anooshiravan Saboori, Christoforos N. Hadjicostis. Opacity verification in stochastic discrete event systems. *CDC'10*, pp.6759-6764, IEEE, 2010.
21. A. Saboori and C. N. Hadjicostis. Verification of $K$-step opacity and analysis of its complexity," *IEEE Transactions on Automation Science and Engineering*, 8(3), pp.549-559, 2011.

22. Anooshiravan Saboori, Christoforos N. Hadjicostis. Current state opacity formulations in Probabilistic finite Automata. *Transactions on Automatic Control* 59(1), pp.120-133, IEEE, 2014.
23. Arnab Bhattacharyya, Sutanu Gayen, Kuldeep S. Meel, N. V. Vinodchandran. Efficient Distance Approximation for Structured High-Dimensional Distributions via Learning. Available on Arxiv: https://arxiv.org/abs/2002.05378.

# Appendix

### Definitions of MDPs and PFAs

We introduce here *Markov Decision Processes (MDPs)* [2], to explain the relation with IMDP semantics, and because they are used in some lower bounds.

An MDP $A$ on finite alphabet $\Sigma$ is a tuple $\mathcal{A} = (S, (M_a)_{a\in\Sigma}, \mu_0)$ with $S$ a set of states, $\mu_0$ an initial distribution, $M_a : S \times S \to [0, 1]$ for each $a \in \Sigma$, such that for all $a$, for all $s$, $\sum_{s'} M_a(s, s') = 1$.

MDPs are fully observable: the semantic of an MDP is the set of schedulers $\sigma : S \times (\Sigma \times S)^* \to Dist(\Sigma)$ telling which distribution over letters is chosen according to the path from the initial to the current state.

For comparison, the IMDP semantics of an ILMC can be described by MDP with either an infinite number of actions (alphabet), or with the same alphabet, but some restrictions on the Distribution being chosen.

While MDPs are fully observable (we know the sequence of states visited), *Rabin's Probabilistic Finite Automata (PFAs)* is the same model as MDPs, but interpreted as having no information at all about the states. That is, the semantics is given as the set of schedulers $\sigma : \mathbb{N} \to Dist(\Sigma)$, where the choice of actions can only differ based on the number of actions played so far. Given a distribution $\sigma_1$ over letters, we associate the Markov Chain $M_{\sigma_1} = \sum_{a\in\Sigma} \sigma_1(a)M_a$. We associate to a scheduler $\sigma$ the sequence of Markov chain $M^\sigma_1, \ldots,$ with $M^\sigma_i = M_{\sigma(i)} \ldots M_{\sigma(1)}$. An important question is to know given a scheduler $\sigma$ and a number $n$ of steps what is the probability to be in $F$ from $\mu_0$ following $n$ steps of scheduler $\sigma$, that is $P^n(\sigma) = \sum_{s\in S,t\in F} \mu_0(s)M^\sigma_i[t]$.

### Proofs in Section 3

**Proposition 1.** *It is undecidable to know, given $\alpha \in (0, 1)$ and an LMC $A$, whether for all observation $w$, $d_A(w) \geq \alpha$ [3]. Further, one cannot even approximate $d_{worst}(A) = \inf_{w\in\Sigma^*} d_A(w)$.*

*Proof.* We want to know whether there exists an observation $w$ such that $d(w) < \alpha$. We prove that this problem is undecidable by reducing emptiness problem in a Rabin's probabilistic finite automaton to our problem.

Emptiness Problem :- Given a PFA $H = (Q, \Sigma, s_0, (M_a)_{a\in\Sigma}, F)$, and a constant $\alpha$ with $0 < \alpha < 1$, the emptiness problem asks whether there exists a finite number of distributions $\sigma_1, \ldots, \sigma_n$ a number $n \in \mathbb{N}$ and a scheduler $\sigma : \mathbb{N} \to Distrib(\Sigma)$ such that $\mathcal{P}^n(\sigma) \geq \alpha$. If this exists, then there is a pure strategy $\sigma$, which can be written as a word $w$ of size $n$ being the sequnence of choices in $\Sigma$ being made. In that case, we will write $\mathcal{P}^H(w) = \mathcal{P}^n_\sigma(w)$. It is thus also undecidable to know whether there exists $w$ such that $\mathcal{P}^H(w) \geq \theta$.

Let $H = (S, \Sigma, \mu_0, (M_a)_{a\in\Sigma}, F)$ be a PFA. We can assume without loss of generality that it is complete. We build LMC $A = (S', M, \mu_0)$ over $\Sigma' = \Sigma \uplus \{\$\}$ as follows:

- We have $S' = S \uplus \{q_s\} \uplus \{q_{ns}\}$,

– For $s, s' \in S, a \in \Sigma$, we define $M(s, a, s') = M_{a[s,s']}/(|\Sigma| + 1)$,
– For $s \in F$, we define $M(s, \$, q_{ns}) = 1/|\Sigma| + 1)$,
– For $s \in S' \setminus F$, we define $M(s, \$, q_s) = 1/(|\Sigma| + 1)$,
– For $M(q_s, \$, q_s) = 1, M(q_{ns}, \$, q_{ns}) = 1$.

We have for all $s \in S', \sum_{a\in\Sigma', s'\in S'} M(s, a, s') = 1$, and there is a path to $q_{\text{secret}}, q_{\text{non-secret}}$ (playing just one letter). We represent the probability of a word $w$ (seen as a sequence of actions chosen) with respect to PFA $H$ as $P^H(w)$ and the probability of $w$ (seen as an observation) with respect to LMC $A$ as $P^A(w)$.

Let $w \in \Sigma^*$. We have easily $\mathcal{P}^A_{\text{non-secret}}(w\$) = P^H(w)/(|\Sigma| + 1)^{|w|+1}$ and $\mathcal{P}^A_{\text{secret}}(w\$) = (1 - P^H(w))/(|\Sigma| + 1)^{|w|+1}$. Hence, $P^A(w\$) = 1/(|\Sigma| + 1)^{|w|+1}$.

That is, $d(w\$) = \mathcal{P}^A_{\text{non-secret}}(w\$)/P^A(w\$) = (P^H(w)\cdot(1/(|\Sigma|+1)^{|w|+1}))/1/(|\Sigma|+1)^{|w|+1} = P^H(w)$. In particular, $d(w\$) < \alpha$ iff $P^H(w) < \alpha$. Hence, we obtain the undecidability of $d_{\text{worse}}(A) \geq \alpha$.Similar result can be found in [22].

Further, we can show that it is unapproximable by using the following undecidable problem on PFAs [16]: Given $0 < \epsilon < 1$ and a PFA $H$ for which one of the two cases hold :
(1) there is some strategy in the PFA such that the probability to accept is lower than $\epsilon$, or
(2) for all strategy, the probability to accept is at least $1 - \epsilon$.
Decide whether case (1) holds.

Given a PFA $H = (S, \Sigma, s_0, (M_a)_{a\in\Sigma}, F)$ we construct the same LMC $A$ as above. As derived above, $\forall w \in \Sigma^*, P^H(w) = d_s(w\$)$ so either
(1) $\exists w, d(w\$) < \epsilon$, or (2) $\forall w, d(w\$) \geq 1 - \epsilon$

We cannot decide which case holds else we could decide which case hold for hte PFA, a contradiction. In particular, if we could approximate $d_{worst}(A)$, we could approximate $min_{w\in\Sigma^*} P^H(w) = d_{worst}(A)$, which cannot be. □

**Proofs in Section 5**

We start by proving the proof of point 3) of Theorem 1 (recalled below), which was only sketched before.

**Theorem 1.** *Input: an LMC $A = (S, M, \mu_0)$, $0 < \theta < 1$ and $\lambda \in [0, 1]$,*

1. *Knowing whether $d_\theta(A) > \lambda$ is undecidable. However,*
2. *Given an error bound $\varepsilon > 0$, and a high confidence value $\delta < 1$, one can compute a value $\gamma$ such that $P(|d_\theta(A) - \gamma| \leq \varepsilon) \geq \delta$ in polynomial time $O(|A|^{\frac{log(1/(1-\delta))}{\varepsilon^2}})$.*
3. *Given an error bound $\varepsilon > 0$, it is possible to compute a value $\gamma$ such that $|d_\theta(A) - \lambda| \leq \varepsilon$ in PSPACE in $|A|$ and $|\varepsilon|$.*

*Proof.* We have $d_\theta(A) = \sum_w \mathcal{P}_{\text{stop}}(w) \cdot d_\theta(w)$. We apply Lemma 1 and obtain a number $k$ of steps after which almost all runs are stopped. We have $d_\theta = \sum_{w\in(L(A)\cap\Sigma^{\leq k})} \mathcal{P}_{\text{stop}}(w) \cdot d_\theta(w) + \sum_{w\in(L(A)\cap\Sigma^{>k})} \mathcal{P}_{\text{stop}}(w) \cdot d_\theta(w)$. Consider the two terms separately:

We can compute the first term $\lambda = \sum_{w \in (L(A) \cap \Sigma^{\leq k})} \mathcal{P}_{\text{stop}}(w) \cdot d_\theta(w)$ by calculating the sum over all the possible words satisfying these conditions: we can do this in EXPTIME as there are at most an exponential number of such runs wrt the size of the LMC. The second term is trivially non-negative. Also, by choice of $k$ and because $0 \leq d_\theta(w) \leq 1$, we have that $0 \leq \sum_{w \in (L(A) \cap \Sigma^{>k})} \mathcal{P}_{\text{stop}}(w) \cdot d_\theta(w) \leq \varepsilon$. We thus have $0 \leq d_\theta(A) - \lambda \leq \varepsilon$.

We now turn to the PSPACE complexity result: First, notice that $\mathcal{P}_{\text{stop}}(w)$ cannot be written with polynomially many bits for $w \in \Sigma^k$. The trick is to approximate $\mathcal{P}_{\text{stop}}(w)$ using floating-point arithmetic with small relative error: $\widetilde{\mathcal{P}_{\text{stop}}}(w) \in [\mathcal{P}_{\text{stop}}(w)(1-\alpha), \mathcal{P}_{\text{stop}}(w)(1+\alpha)]$ for small $\alpha > 0$. We compute similarly an approximation $\widetilde{\mathcal{P}_{\text{non-secret}}}(w)$ of $\mathcal{P}_{\text{non-secret}}(w)$, for all $w \in \Sigma^{\leq k}$. We now show that from there, we can compute a good approximation $\widetilde{d}(w)$ of $d(w)$, and then a good approximation $\widetilde{d_\theta(A)} = \sum \widetilde{\mathcal{P}_{\text{stop}}}(w) \cdot max(0, \theta - \widetilde{d}(w))$.

We define $\widetilde{d}(w) = \frac{\widetilde{\mathcal{P}_{\text{non-secret}}(w)}}{\widetilde{\mathcal{P}_{\text{stop}}}(w)} \in [\frac{\mathcal{P}_{\text{non-secret}}(w)(1-\alpha)}{\mathcal{P}_{\text{stop}}(w)(1+\alpha)}, \frac{\mathcal{P}_{\text{non-secret}}(w)(1+\alpha)}{\mathcal{P}_{\text{stop}}(w)(1-\alpha)}]$.

Hence $\theta - \widetilde{d}(w) \in [\theta - \frac{\mathcal{P}_{\text{non-secret}}(w)(1+\alpha)}{\mathcal{P}_{\text{stop}}(w)(1-\alpha)}, \theta - \frac{\mathcal{P}_{\text{non-secret}}(w)(1-\alpha)}{\mathcal{P}_{\text{stop}}(w)(1+\alpha)}]$

Let us compare $\max(0, \theta - d(w))$ and $\max(0, \theta - \widetilde{d(w)})$. There are three cases:

- if $\theta < d(w)\frac{(1-\alpha)}{(1+\alpha)}$, then $max(0, \theta - d(w)) = 0 = max(0, \theta - \widetilde{d}(w))$.
- If $\theta > d(w)\frac{(1+\alpha)}{(1-\alpha)}$, then $|d_\theta(w) - \widetilde{d}_\theta(w)| \leq \theta - d(w) - (1-\alpha)(\theta - d(w)\frac{1+\alpha}{1-\alpha})$.
  Hence, $|d_\theta(w) - \widetilde{d}_\theta(w)| \leq (\theta + d(w)) \cdot \alpha \leq 2\alpha$.
- Otherwise, $d(w)\frac{(1+\alpha)}{(1-\alpha)} \leq \theta \leq d(w)\frac{(1-\alpha)}{(1+\alpha)}$. Thus, $|d_\theta(w) - \widetilde{d}_\theta(w)| \leq |(1+\alpha) \cdot (d(w)\frac{1+\alpha}{1-\alpha} - d(w)\frac{1-\alpha}{1+\alpha})| \leq d(w)\frac{4\alpha}{1-\alpha} \leq 8\alpha$.

So given an $\varepsilon > 0, 0 < \theta < 1$ we select $\alpha$ such that $8\alpha \leq \varepsilon$, which implies that $|d_\theta(A) - \widetilde{d}_\theta(A)| \leq \sum_w \mathcal{P}(w) \cdot |d_\theta(w) - \widetilde{d}_\theta(w)| \leq \varepsilon \sum_w \mathcal{P}(w) \leq \varepsilon$. We know by Ladner's result [15] on counting that we can perform $\lambda = \widetilde{d}_\theta(A)$ in polynomial space (sum over polynomial number of bits), which makes the overall problem solvable in PSPACE. □

Let $\mathcal{A}$ be an ILMC. Let $F$ a subset of states. We define $X(F) = \{s \mid \forall A \in \text{LMC}(\mathcal{A})$, there exists a path from $s$ to $F\}$ of states from which $F$ can be reached in all $A \in \text{LMC}(\mathcal{A})$, for LMC $\in \{$UMC,IMDP$\}$.

We now turn to the case of ILMCs.

**Theorem 2.** *Given an ILMC $\mathcal{A} = (S, \delta, \mu_0)$, and a constant $\varepsilon > 0$, we can compute in EXPTIME a number $\lambda_{min} \in [0,1]$ such that $|\min_{A \in LMC(\mathcal{A})}(d_\theta(A)) - \lambda_{min}| \leq \varepsilon$. Also, we can compute in EXPSPACE a number $\lambda_{max} \in [0,1]$ such that $|\max_{A \in LMC(\mathcal{A})}(d_\theta(A)) - \lambda_{max}| \leq \varepsilon$, for LMC $\in \{$UMC,IMDP$\}$.*

*Proof.* The idea is to consider $\sum_{w\in\Sigma^{\leq K}} \text{contrib}^A(w) \leq d_\theta(A) \leq \sum_{w\in\Sigma^{\leq K}} \text{contrib}^A(w) + \varepsilon$, where $\text{contrib}^A(w) = \theta\mathcal{P}^A_{\text{stop}}(w) - \mathcal{P}^A_{\text{non-secret}}(w)$ or 0 if this is negative.

First, remark that $d_\theta(A) = \sum_w(\theta\mathcal{P}_{\text{stop}}(w) - \min(\mathcal{P}_{\text{non-secret}}(w), \theta\mathcal{P}(w))$. For all $w$, the contribution from $w$ is thus either $\text{contrib}(w) = \theta\mathcal{P}_{\text{stop}}(w) - \mathcal{P}_{\text{non-secret}}(w)$ or 0, and it is smaller than $\mathcal{P}_{\text{stop}}(w)$.

We have $\sum_{w\in\Sigma^{\leq K}} \text{contrib}(w) \leq d_\theta(A) \leq \sum_{w\in\Sigma^{\leq K}} \text{contrib}(w) + \varepsilon$. It thus suffices to compute $\lambda_{min} = \min_{A\in\text{IMDP}(\mathcal{A})} \sum_{w\in\Sigma^{\leq K}} \text{contrib}(w)$ to obtain a desired number $\varepsilon$-close to the minimal value.

For that, we use the following linear optimization program:

- variables: $\text{contrib}(w)$ for all $w \in \Sigma^{\leq K}$ and $\mu(s,n)$ for all $s \in S$ and $n \leq K$.
- constraints: for all $s, n$, $\mu(s,n)$ should be a valid assignment ie with $\mu((s,n),a,(s',n+1)) \in \delta(s,a,s')$ and $\sum_{a,s'} \mu((s,n),a,(s',n+1)) = 1$. We denote by $A$ the associated LMC.
- further, for all $w \in \Sigma^{\leq K}$, we need to have $\text{contrib}(w) \geq 0$ and $\text{contrib}(w) \geq \theta \sum_{\text{path } \rho \text{ labelled } w} \mathcal{P}^A_{\text{stop}}(\rho) - \sum_{\text{path } \rho \text{ to non-secret labelled } w} \mathcal{P}^A_{\text{stop}}(\rho)$
- Objective: minimize $\sum_{w\in\Sigma^{\leq K}} \text{contrib}(w)$.

Notice that when the minimum is reached, we have either $\text{contrib}(w) = \theta \sum_{\text{path } \rho \text{ labelled } w} \mathcal{P}^A_{\text{stop}}(\rho) - \sum_{\text{path } \rho \text{ to non-secret labelled } w} \mathcal{P}^A_{\text{stop}}(\rho)$ or 0. This gives us an EXPTIME algorithm.

This is not true for max. Instead of using linear programming, we rely on the existential theory of real to compute $\lambda_{max} = \max_{A\in\text{IMDP}(\mathcal{A})} \sum_{w\in\Sigma^{\leq K}} \text{contrib}(w)$. This gives us an EXPSPACE algorithm. □

## Proofs in Section 6

### Proofs for no run fully reveals whether it is secret

**Theorem 3.** *Given one LMC $A$, testing whether $d_A(w) > 0$ for all $w \in \Sigma^*$ is PSPACE-complete. For an ILMC $\mathcal{A}$, testing $d_A(w) > 0$ for all $w \in \Sigma^*$ and for all $A \in$ IMDP($\mathcal{A}$) or $A \in$ UMC($\mathcal{A}$) is also PSPACE-complete.*

*Proof.* We start by considering an LMC $A = (S, M, \mu_0)$.

First we will prove that the problem is in PSPACE. The problem to know whether $d_A(w) > 0$ for all $w \in \Sigma^*$ is equivalent with there does not exist an observation $w$ with $\mathcal{P}_{\text{non-secret}}(w) = 0$ and $\mathcal{P}_{\text{secret}}(w) > 0$. The algorithm is classical, see e.g. [21]: it suffices to consider the deterministic belief automaton associated with $A$, namely $\mathcal{B} = (2^S, T, B_0)$, where $B_0 = \{q \mid \mu_0(q) > 0\}$ and $(S, a, S') \in T$ iff $S' = \{q' \mid \exists q \in S, M(q,a,q') > 0\}$. We have that there exists $w$ with $\mathcal{P}_{\text{non-secret}}(w) = 0$ and $\mathcal{P}_{\text{secret}}(w) > 0$ iff one can reach a belief $B$ containing $q_{\text{secret}}$ but not $q_{\text{non-secret}}$ from $B_0$. This is a reachability question in a graph with an exponential number of nodes, i.e. it can be done in PSPACE.

We now prove that this problem is PSPACE-Hard. For that, we reduce the PSPACE-Hard Problem of inclusion of language of two automaton, similar to

the proof of theorem 14 in [1]. Given two automata for $A_i = (S_i, \Sigma, s_0^i, \Delta^i, F_i)$ for $i = 1, 2$, we know that checking whether $L(A_1) \subseteq L(A_2)$ is a PSPACE-Hard problem. We construct an LMC $A = (S_1 \uplus S_2 \uplus \{q_{\text{secret}}\} \uplus \{q_{\text{non-secret}}\}, M, \mu_0)$ over alphabet $\Sigma \uplus \{\$\}$ with:

- $\mu_0(s_0^1) = 0.5, \mu_0(s_0^2) = 0.5$
- For all $a \in \Sigma$, $M(s, a, s') > 0$ iff $(s, a, s') \in \Delta_1$ or $(s, a, s') \in \Delta_2$
- $M(s, \$, q_{\text{secret}}) > 0$ iff $s \in F_1$
- $M(s, \$, q_{\text{non-secret}}) > 0$ iff $s \in F_2$
- $M(q_{\text{secret}}, \$, q_{\text{secret}}) = 1, M(q_{\text{non-secret}}, \$, q_{\text{non-secret}}) = 1$

Now we claim that $d_A(w) > 0$ for all $w \in \Sigma^*$ iff $L(B1) \subseteq L(B2)$. Indeed, there is some $w$ such that $d_A(w) = 0$, iff $w \in L(A_1)$ but $w \notin L(A_2)$, i.e. iff $L(B1) \not\subseteq L(B2)$. Hence the problem is PSPACE-Complete.

We now consider an ILMC $\mathcal{A}$, viewed as an uncertain model. We first consider the UMC semantics of $\mathcal{A}$. As shown above, the support of the transitions is important, but the actual $> 0$ probability of each $> 0$ transition does not matter. We can thus guess in PSPACE the support of the transitions of an UMC $A$, check that $A \in \text{UMC}(\mathcal{A})$, and then use the above PSPACE algorithm to check whether $d_{\text{worst}}(A) = 0$. We thus obtain a PSPACE algorithm. The PSPACE-hardness follows from the hardness on LMCs.

We now consider the IMDP semantics of $\mathcal{A}$: The PSPACE-hardness follows from the PSPACE-hardness for LMCs. We now prove that one can check whether there exists $A \in \text{IMDP}(\mathcal{A})$ and $w \in \Sigma^*$ with $d_A(w) = 0$, which is the opposite of our problem. We conclude as co-PSPACE=PSPACE by Immerman-Szelepcsenyi theorem [14].

As transitions do not only depend upon the action, but also on the whole history, we cannot encode a choice of support of transition with a polynomial number of bits. Hence we cannot reason as for the UMC semantics. Instead, we will guess on the fly the transitions. Thus, 2 states with different history may not have the same set of transitions, but that is exactly what is allowed in the IMDP semantics. To represent that, we build a non-deterministic belief automaton $\mathcal{B}' = (2^S, \Delta, B_0)$ on alphabet $\Sigma$ where $2^S$ represents the set of states, $B_0$ the initial state, and $\Delta$ the set of transitions corresponding to the ILMC $\mathcal{A} = (S, \Sigma, \mu_0, \delta)$ as follows :-

- $B_0 = \{s \in S | \mu_0(s) > 0\}$
- For $B, B' \in 2^S, a \in \Sigma$, we have $(B, a, B') \in \Delta'$ if and only if there exists a $\sigma : S \times \Sigma \mapsto Dist(S)$ with $\sigma(s, a, s') \in \delta(s, a, s')$ for all $s, a, s'$ such that $B' = \{s' \in S \mid \exists s \in B, \sigma(s, a, s') > 0\}$.

In Fig. 5, we give the belief automaton $\mathcal{B}$ corresponding to the ILMC of Fig. 3. The two non-deterministic transitions are reading $a$ from either node $\{q_1, q_3\}$ or $\{q_3\}$, corresponding to $M(q_3, a, q_4) = 0$ and $M(q_3, a, q_4) > 0$.

We have that there exists $A \in \text{IMDP}(\mathcal{A})$ and $w \in \Sigma^*$ with $d_A(w) = 0$ iff there is a path from $B_0$ to $\{q_{\text{secret}}\}$ in $\mathcal{B}$. This can be tested in PSPACE, as the non-deterministic belief automaton has an exponential number of states.

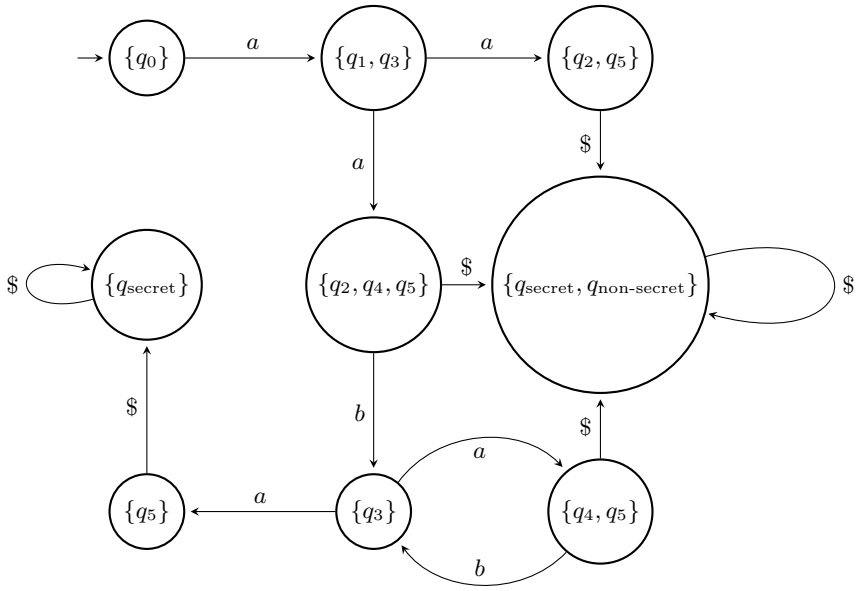


Fig. 5: Non-deterministic belief automaton for the ILMC in figure 3

Indeed, assume that there exists an LMC $A \in \text{IMDP}(\mathcal{A})$ such that $\exists w, d_A(w) = 0$. That is, the path labeled by $w$ in the deterministic belief automaton associated with $A$ from $B_0$ leads to $\{q_{\text{secret}}\}$. By construction, we will find such a path in the non-deterministic belief automaton associated with the ILMC $\mathcal{A}$.

On the other hand, if there is a path $\rho$ to $\{q_{\text{secret}}\}$ in the non-deterministic belief automaton associated with the ILMC $\mathcal{A}$, then one can build an LMC $A \in \text{IMDP}(\mathcal{A})$ such that $d_A(w) = 0$: it suffices to choose any sequence of assignments according to $\rho$. □

**Proofs for total opacity**

**Theorem 5.** *Given an LMC $A$, testing whether $d_A(w) = 1$ for all $w \in \Sigma^*$ can be done in NLOGSPACE. For an ILMC $\mathcal{A}$, checking whether $d_A(w) = 1$ for all $w \in \Sigma^*$ and for all $A \in LMC(\mathcal{A})$ is LOGSPACE, for LMC $\in \{UMC, IMDP\}$.*

*Proof.* Let $A$ be an LMC. $d_A(w) = 1$ for all $w$ is equivalent to the non-existence of $w$ such that $\mathcal{P}_s^A(w\$) > 0$, that is state $q_{\text{secret}}$ is not reachable. This can be checked in NLOGSPACE.

For ILMCs, there is no difference between the UMC and the IMDP semantics as we need to handle a reachability objective (positional strategies are sufficient to win reachability objectives in MDPs [2]). Without restriction, we can thus consider the UMC semantics. We rely on the following lemmas.

**Lemma 5.** *Let $\mathcal{A} = (S, \delta, \mu_0)$ be an ILMC and $A_1 = (S, M_1, \mu_0), A_2 = (S, M_2, \mu_0) \in UMC(\mathcal{A})$. Let $A = (S, M, \mu_0)$ such that for all $s, s' \in S, a \in \Sigma$, $M(s, a, s') = (M_1(s, a, s') + M_2(s, a, s'))/2$. Then $A \in UMC(\mathcal{A})$.*

*Proof.* For all $s, s' \in S$ and $a \in \Sigma$, $M(s, a, s') \in [M_1(s, a, s'), M_2(s, a, s')]$ (or $[M_2(s, a, s'), M_1(s, a, s')]$) and thus $M(s, a, s') \in \delta(s, a, s')$. □

Notice that the edges of $A$ are the union of the edges of $A_1$ and $A_2$.

**Lemma 6.** *Let $\mathcal{A} = (S, \delta, \mu_0)$ be an ILMC under the UMC semantic such that $UMC(\mathcal{A})$ is non-empty. There exists an LMC $A \in UMC(\mathcal{A})$ such that for all $s, s' \in S, a \in \Sigma$, $\exists A' = (S, M', \mu_0) \in UMC(\mathcal{A})$ and $M'(s, a, s') > 0 \Leftrightarrow M(s, a, s') > 0$.*

That is, the support of $A$ is the maximal wrt inclusion of edges.

*Proof.* Let us suppose there is no such $A$. Let $A_1, A_2 \in UMC(\mathcal{A})$ such that their supports are maximal wrt inclusion of edges and different. By lemma 5, there is an $A$ such that the support of $A$ the union of those of $A_1$ and $A_2$. Thus, there were not maximal, hence a contradiction. □

Thus, if there exists an LMC such that $q_{\text{secret}}$ is reachable, then $q_{\text{secret}}$ is reachable in $A$ containing all the possible edges. Hence, the problem of deciding if $d_{\text{worst}}(A) = 1$ is equivalent to decide if $q_{\text{secret}}$ is reachable in that $A$, hence the complexity.

□

**Theorem 6.** *Given an ILMC $\mathcal{A}$, testing whether there exists $A \in UMC(\mathcal{A})$ or $A \in IMDP(\mathcal{A})$ such that $d_A(w) = 1$ for all $w \in \Sigma^*$ is PTIME-complete.*

*Proof.* As for Theorem 5, there is no difference between the UMC and the IMDP semantics for this case. In order to prove the PTIME completeness, we first give an algorithm that answers the problem for the UMC semantics (and thus for the IMDP semantics), and then show that the total opacity problem for UMCs is P-hard.

We use Lemma 1 with $F = \{q_{\text{secret}}\}$ to compute the set $X(F)$ of states which can reach $q_{\text{secret}}$ in all LMCs $A \in \text{LMC}(\mathcal{A})$. We let Init$=\{s \mid \mu_0(s) > 0\}$. There exists an $A \in \text{LMC}(\mathcal{A})$ and some $w$ with $\mathcal{P}_{\text{secret}}(w) > 0$ iff Init $\cap X(F) \neq \emptyset$, by definition of $X(F)$. By contraposition, there exists $A \in \text{LMC}(\mathcal{A})$ with $d_A(w) = 1$ for all $w \in \Sigma^*$ iff Init $\cap X(F) = \emptyset$.

*Hardness:* Now we prove that this problem is PTIME-hard:

**Lemma 7.** *Given an ILMC $\mathcal{A}$, testing whether there exists $A \in UMC(\mathcal{A})$ such that $d_A(w) = 1$ for all $w \in \Sigma^*$ is PTIME-hard.*

*Proof.* The proof is a reduction from the reachability problem over MDPs, which is known to be P-complete [7, 14]. Given a complete MDP $\mathcal{M} = (S, \Sigma, \mu_0, \Delta)$, a target $T \subseteq S$ we construct the ILMC $B = (Q, \Sigma \uplus \{\$\}, \mu_0, \delta)$ as follows:

- $Q = S \cup \{(s, a) \in S \times \Sigma\} \cup \{q_s, q_{ns}\}$,
- $\forall s \in T, \delta(s, \$, q_s) = [1, 1]$,
- $\forall s \in S \setminus T, \delta(s, \$, q_{ns}) = [0.1, 0.1]$,
- $\forall s \in S \setminus T, a \in \Sigma,\ \delta(s, \varepsilon, (s, a))) = [0, 0.9]$,
- $\forall s, s' \in S, a \in \Sigma,\ \delta((s, a), a, s') = [\Delta(s, a, s'), \Delta(s, a, s')]$,
- $\delta(q_s, \$, q_s) = [1, 1]$ and $\delta(q_{ns}, \$, q_{ns}) = [1, 1]$,
- for any other $s, s'$ and $a \in \Sigma$ $\delta(s, a, s') = [0, 0]$.

Remember that if there is a strategy $\sigma'$ such that $\mathcal{P}_{\sigma'}(Reach(T)) = 0$, then there is a memoryless strategy $\sigma$ such that $\mathcal{P}_{\sigma}(Reach(T)) = 0$. Let $\sigma$ be such a memoryless strategy. Then let be the LMC $A \in UMC(B)$ with $\forall s$, $\delta(s, \varepsilon, (s, a))) = 0.9$ if $a = \sigma(s)$ and 0 else. It is easy to see that in $A$, $q_s$ is not reachable as $T$ is not either, and thus $d_A(w) = 1$ for all $w \in \Sigma^*$.

Conversely, let us suppose there exists $A' \in UMC(B)$ such that $d_{A'}(w) = 1$ for all $w \in \Sigma^*$. Then, there exists $A \in UMC(B)$ such that $d_A(w) = 1$ for all $w \in \Sigma^*$ with for all $s, a$, $\delta(s, \varepsilon, (s, a)))$ equals either 0 or 0.9: removing edges will not enable reachability of $q_s$. Then, let $\sigma$ be such that $a = \sigma(s)$ iff $\delta(s, \varepsilon, (s, a))) = 0.9$. Then $\mathcal{P}_{\sigma}(Reach(T)) = 0$. □