Probabilistic Disclosure: Maximisation vs. Minimisation*

Béatrice Bérard[1], Serge Haddad[†2], and Engel Lefaucheux[3]

1 **Sorbonne Universités, UPMC Univ. Paris 06, CNRS UMR 7606, LIP6, Paris, and LSV, ENS Paris-Saclay & CNRS & Inria, Université Paris-Saclay, France**
beatrice.berard@lip6.fr

2 **LSV, ENS Paris-Saclay & CNRS & Inria, Université Paris-Saclay, France**
serge.haddad@lsv.fr

3 **Inria, Campus Universitaire de Beaulieu, Rennes, France and LSV, ENS Paris-Saclay & CNRS & Inria, Université Paris-Saclay, France**
engel.lefaucheux@inria.fr

Abstract

We consider opacity questions where an observation function provides to an external attacker a view of the states along executions and secret executions are those visiting some state from a fixed subset. Disclosure occurs when the observer can deduce from a finite observation that the execution is secret, the ε-disclosure variant corresponding to the execution being secret with probability greater than $1-\varepsilon$. In a probabilistic and non deterministic setting, where an internal agent can choose between actions, there are two points of view, depending on the status of this agent: the successive choices can either help the attacker trying to disclose the secret, if the system has been corrupted, or they can prevent disclosure as much as possible if these choices are part of the system design. In the former situation, corresponding to a worst case, the disclosure value is the supremum over the strategies of the probability to disclose the secret (maximisation), whereas in the latter case, the disclosure is the infimum (minimisation). We address quantitative problems (comparing the optimal value with a threshold) and qualitative ones (when the threshold is zero or one) related to both forms of disclosure for a fixed or finite horizon. For all problems, we characterise their decidability status and their complexity. We discover a surprising asymmetry: on the one hand optimal strategies may be chosen among deterministic ones in maximisation problems, while it is not the case for minimisation. On the other hand, for the questions addressed here, more minimisation problems than maximisation ones are decidable.

1998 ACM Subject Classification D.2.4 Software, Program Verification

Keywords and phrases Partially observed systems, Opacity, Markov chain, Markov decision process

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2017.13

1 Introduction

Opacity. Opacity of an information system is a key security property: an external user should not, by observing an execution of a system, acquire the guarantee that it is a secret one. This property was first formalised for labelled transition systems [7], by specifying

* Full proofs can found in research report https://hal.inria.fr/hal-01618955.

† The work of S. Haddad has been supported by ERC project EQualIS (FP7-308087).

© Béatrice Bérard, Serge Haddad, and Engel Lefaucheux;
licensed under Creative Commons License CC-BY

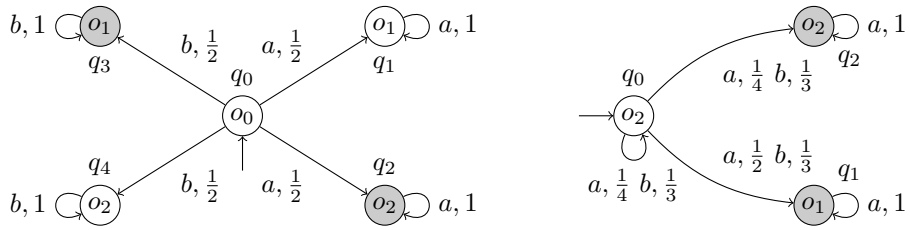

(a) $Disc = 0$ or $\frac{1}{2}$. **(b)** $Disc = \frac{2}{3}$.

Figure 1 When strategies help (or not) to disclose a secret.

a subset of secret paths and requiring that, for any secret path, there is a non-secret one with the same observation. The *disclosure set* of a system is then the set of (secret) paths violating opacity.

Opacity raises challenging research issues like (1) formal specification in various frameworks [13, 7], (2) the design of mechanisms to ensure opacity while preserving functionality and performance [2], and (3) the verification of opacity properties [14, 7].

Attacks against opacity. In order to quantify the size of the leak, various measures for the disclosure set, called probabilistic disclosure, were introduced in [17, 3, 5, 4]. For probabilistic and non deterministic systems like Markov Decision Processes (MDP), where an internal agent can choose between several actions, the disclosure can be either maximised or minimised depending on the status of the agent. In previous works, the situation considered corresponded to maximisation, where the system has been corrupted (*e.g.*, by a virus), and the internal agent cooperates with the attacker to disclose the secret. In [3], the set of secret paths is specified by some deterministic automaton and the attacker does not know which strategy is applied, hence the set of executions leaking the secret is fixed by the structure of the system and does not vary according to the strategies. This is illustrated in Figure 1a with an MDP where actions a and b are possible from state q_0. Action a (resp. b) has a uniform distribution over states q_1 and q_2 (resp. q_3 and q_4). States q_1 and q_3 produce observation o_1, while q_2 and q_4 produce o_2. The secret paths are those in $q_0q_2^\omega \cup q_0q_3^\omega$ (reaching either q_2 or q_3, in grey). If the observer is not aware of the strategy and thus has to consider all possible paths whatever the strategy, the non secret paths are $q_0q_1^\omega$ and $q_0q_4^\omega$, hence there is no disclosing path, leading to a null disclosure (as in [3]). On the other hand, if the observer is aware of the strategy, assuming that initially a is chosen produces a maximal disclosure of $\frac{1}{2}$. This example also shows that deterministic strategies are not sufficient to achieve minimisation: value 0 for the disclosure can only be obtained by randomised strategies, choosing a with probability p and b with probability $1-p$ (for $0 < p < 1$) in q_0.

In contrast, Figure 1b represents an MDP where both actions a and b have the same support in state q_0, hence choosing a or b does not change the states that can be reached. Under a strategy which always plays a, the disclosure is equal to $\frac{2}{3}$ which is the probability of reaching q_1.

Given some $\varepsilon > 0$, a path could alternatively be considered as disclosing the secret, if the measure of the set of paths with same observation that are not secret is less than ε. This notion is called ε-disclosure. For instance in Figure 1a with $\varepsilon < \frac{1}{2}$ in order to achieve a minimal ε-disclosure of 0 the strategy must select p between ε and $1-\varepsilon$. However in Figure 1b for every $\varepsilon > 0$ the ε-disclosure is equal to 1 as the probability to be in a secret state converges to 1 on every path.

Table 1 Complexity results for maximisation and minimisation of disclosure.

Disclosure	General	Limit-sure	Almost-sure
Maximisation finite horizon	undecidable	undecidable	EXPTIME-c
Minimisation finite horizon	PSPACE-hard $\leq$ Min $\leq$ EXPTIME		
Maximisation fixed horizon	PSPACE-c.	PTIME	PTIME
Minimisation fixed horizon	PSPACE-c	PSPACE-c	PSPACE-c

This figure also illustrates the drastic restriction used in [4] where no edge in an MDP can be blocked by a strategy. With this restricted power of the internal component, the authors can assume that the observer knows the strategy, which is an important requirement since the security of a system should not be based on hiding its design. The model used in [4] is in fact a restricted case of Interval Markov Decision Processes (IMDPs) so that after some transformation, the problem boils down to IMDP model checking [10]. The general decidability status of the disclosure problem is left open.

Contributions. Here we focus on several problems in MDPs under partial observation that cannot be formalised as problems for classical POMDPs (Partially Observable MDPs). The notion of disclosure is defined with respect to a fixed subset Sec of states: A (finite or infinite) path is secret if it has visited some state of Sec. Other variants of secret path specifications have been proposed with deterministic finite automata accepting finite or infinite paths. The former case can be easily translated in our setting while we believe that the latter one is debatable: a system is not really vulnerable if the attacker can only know the secret at infinite horizon!

Once a strategy is fixed, the behaviour of the system is described by a possibly infinite partially observable Markov chain, so we start in Section 2 by establishing several results on the semantical aspects of disclosure in Markov chains. In addition, we prove undecidability for the positive ε-disclosure problem (deciding if ε-disclosure is positive) within finite horizon. We then consider two different settings depending on the status of the strategies. Like in previous work, maximisation of disclosure corresponds to the internal agent cooperating with the attacker to disclose a secret. Dually, minimisation is interesting to study during the system design process, in order to optimise the choices of the internal agent to defend the system. We address various problems in these settings, for a finite horizon but also for a fixed horizon (given in unary representation), corresponding to real-time constraints requiring the number of steps to be fixed in advance. The quantitative decision problem asks whether the disclosure is above or below some threshold, while qualitative problems consider extremal values (0 or 1) of the disclosure. We prove that observation-based strategies (*i.e.*, which only depend on the sequence of observations and the current state) are dominant in both cases.

The main complexity results for decision problems are gathered in Table 1. For the maximisation objective (Section 3), we show that deterministic strategies are dominant. We answer negatively to the decidability issues left open in [4], proving that both the quantitative problem and the limit-sure problem (asking whether the supremum over all strategies is 1) are undecidable for a finite horizon. Then, we show that the almost-sure problem (asking whether there is a strategy producing a value 1 for disclosure) is EXPTIME-complete. For minimisation (Section 4), we introduce families of randomised strategies, necessary to asymptotically reach minimal disclosure, even within fixed horizon. For finite horizon, we show that the computation and decision problems belong to EXPTIME. Hence surprisingly, although the problem seems more difficult due to the necessity of randomised strategies, the

disclosure problem for minimisation is decidable whereas it is not for maximisation. Section 5 is devoted to the fixed horizon problems. For maximisation, we prove that the disclosure value can be computed in PSPACE (while its associated strategy can be computed in EXPTIME) and also establish that the corresponding decision problem is PSPACE-complete. The almost-sure and limit-sure decision problem however are easier and can be solved in PTIME. Refining the techniques to take randomised strategies into account, we obtain PSPACE-completeness of the various decision problems for minimisation.

2 Specification

We denote by $\mathbb{N}$ the set of natural numbers. For a finite alphabet Σ, we denote by Σ^* (resp. Σ^ω) the set of finite (resp. infinite) words over Σ, with $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$ and ε the empty word. The length of a word w is denoted by $|w| \in \mathbb{N} \cup \{\infty\}$ and for $n \in \mathbb{N}$, Σ^n is the set of words of length n. A word $u \in \Sigma^*$ is a prefix of $v \in \Sigma^\infty$, written $u \leq v$, if $v = uw$ for some $w \in \Sigma^\infty$. The prefix is strict if $w \neq \varepsilon$. Given a countable set Z, a distribution on Z is a mapping $\mu : Z \to [0,1]$ such that $\sum_{z \in Z} \mu(z) = 1$. The support of μ is $Supp(\mu) = \{z \in Z \mid \mu(z) > 0\}$. If $Supp(\mu) = \{z\}$ is a single element, μ is a Dirac distribution on z written $\mathbf{1}_z$. We denote by $\mathsf{Dist}(S)$ the set of distributions on S.

2.1 Opacity for Markov chains

For the purpose of opacity questions, the models are equipped with a labelling function on states, called *observation function*, describing what an external observer can see. We first define observable Markov chains (MCs for short).

▶ **Definition 1** (Markov chains)**.** An observable Markov chain (MC) over alphabet Σ is a tuple $\mathcal{M} = (S, p, \mathsf{O})$ where S is a countable set of states, $p : S \to \mathsf{Dist}(S)$ is the transition function, and $\mathsf{O} : S \to \Sigma \cup \{\varepsilon\}$ is the observation function.

We write $p(s'|s)$ instead of $p(s)(s')$ to emphasise the probability of going to state s' conditionned by being in state s. Given a distribution μ_0 on S, we denote by $\mathcal{M}(\mu_0)$ the chain with initial distribution μ_0. An infinite path of $\mathcal{M}(\mu_0)$ is a sequence of states $\rho = s_0 s_1 \ldots \in S^\omega$ such that $\mu_0(s_0) > 0$ and for each $i \geq 0$, $p(s_{i+1}|s_i) > 0$. A finite path of length n is a prefix $\rho = s_0 s_1 \ldots s_n$ of an infinite path, ending in state $\mathsf{last}(\rho) = s_n$. We denote by $\mathsf{Path}(\mathcal{M}(\mu_0))$ (resp. $\mathsf{FPath}(\mathcal{M}(\mu_0))$) the set of infinite (finite) paths of $\mathcal{M}(\mu_0)$. The observation of path $\rho = s_0 s_1 \ldots$ is the word $\mathsf{O}(\rho) = \mathsf{O}(s_0)\mathsf{O}(s_1)... \in \Sigma^\infty$. For a set R of paths, $\mathsf{O}(R) = \{\mathsf{O}(\rho) \mid \rho \in R\}$ and for a set W of observations, $\mathsf{O}^{-1}(W) = \{\rho \mid \mathsf{O}(\rho) \in W\}$. The observation function is called *non erasing* if $\mathsf{O}(S) \subseteq \Sigma$ (all states are visible).

A probability measure $\mathbf{P}_{\mathcal{M}(\mu_0)}$ is defined on $\mathsf{Path}(\mathcal{M}(\mu_0))$, where the measurable sets are generated by the cylinders $\mathsf{Cyl}(\rho)$, for $\rho \in \mathsf{FPath}(\mathcal{M}(\mu_0))$, containing the infinite paths having ρ as prefix. Then $\mathbf{P}_{\mathcal{M}(\mu_0)}$ is inductively defined by: $\mathbf{P}_{\mathcal{M}(\mu_0)}(s) = \mu_0(s)$ for $s \in S$ and for $\rho' = \rho s'$, with $\mathsf{last}(\rho) = s$, $\mathbf{P}_{\mathcal{M}(\mu_0)}(\mathsf{Cyl}(\rho')) = \mathbf{P}_{\mathcal{M}(\mu_0)}(\mathsf{Cyl}(\rho))p(s'|s)$. We sometimes write $\mathbf{P}_{\mathcal{M}(\mu_0)}(\rho)$ instead of $\mathbf{P}_{\mathcal{M}(\mu_0)}(\mathsf{Cyl}(\rho))$ for $\rho \in \mathsf{FPath}(\mathcal{M})$ and for $w \in \Sigma^*$, $\mathbf{P}_{\mathcal{M}(\mu_0)}(w)$ instead of $\mathbf{P}_{\mathcal{M}(\mu_0)}(\cup_{\rho \in \mathsf{O}^{-1}(w)} \mathsf{Cyl}(\rho))$.

We consider here the particular case where the secret is given by a subset of states $Sec \subseteq S$ of the model: a (finite of infinite) path $s_0 s_1 \ldots$ is secret if $s_i \in Sec$ for some i. We first define a probabilistic version of disclosure w.r.t. some $\varepsilon > 0$ to answer the question: Is there non-zero probability of observing some w that has probability more than $1 - \varepsilon$ of coming from a secret path?

▶ **Definition 2** (ε-Disclosure)**.** Given an MC $\mathcal{M} = (S, p, \mathsf{O})$, an initial distribution μ_0, $Sec \subseteq S$ and an observation $w \in \Sigma^*$, the proportion of secret paths with observation w is:

$$\mathsf{Psec}_{\mathcal{M}(\mu_0)}(w) = \frac{\mathbf{P}_{\mathcal{M}(\mu_0)}(\{\rho \in \mathsf{O}^{-1}(w) \mid \rho \text{ is secret}\})}{\mathbf{P}_{\mathcal{M}(\mu_0)}(w)}.$$

For $\varepsilon > 0$, w is ε-min-disclosing if $\mathsf{Psec}_{\mathcal{M}(\mu_0)}(w) > 1 - \varepsilon$ and no prefix of w satisfies this inequality. Writing $D^{\varepsilon}_{\min}$ for the set of ε-min-disclosing observations, the ε-disclosure is defined by $Disc^{\varepsilon}(\mathcal{M}(\mu_0)) = \sum_{w \in D^{\varepsilon}_{\min}} \mathbf{P}_{\mathcal{M}(\mu_0)}(w)$. The positive ε-disclosure problem consists in deciding if $Disc^{\varepsilon}(\mathcal{M}(\mu_0)) > 0$.

While being the most realistic notion of probabilistic disclosure, unfortunately the problem is already undecidable for Markov chains:

▶ **Theorem 3.** *The positive ε-disclosure problem is undecidable for MCs.*

Like in further proofs, we use a reduction from a problem on Probabilistic Automata (PA). Recall that a PA is a tuple $\mathcal{A} = (Q, q_0, \mathsf{Act}, T, F)$ where Q is a finite set of states with $q_0 \in Q$ the initial state, Act is a finite set of actions, $T : Q \times \mathsf{Act} \to \mathsf{Dist}(Q)$ is the transition function and $F \subseteq Q$ is the set of final states.

For a finite path $\rho = q_0 \xrightarrow{(} a_1) q_1 \dots \xrightarrow{(} a_n) q_n$ of $\mathcal{A}$, the word $a_1 \dots a_n \in \mathsf{Act}^*$ is called the *trace* of ρ and denoted by $tr(\rho)$. Writing $\mathsf{FPath}(w, q) = \{\rho \in \mathsf{FPath} \mid tr(\rho) = w \text{ and } \mathsf{last}(\rho) = q\}$ for $w \in \mathsf{Act}^*$ and $q \in Q$, we define $\mathbf{P}^q_{\mathcal{A}}(w) = \mathbf{P}_{\mathcal{A}}(\cup_{\rho \in \mathsf{FPath}(w,q)} \mathsf{Cyl}_\rho)$, $\mathbf{P}^F_{\mathcal{A}}(w) = \sum_{q \in F} \mathbf{P}^q_{\mathcal{A}}(w)$ and $val(\mathcal{A}) = sup_{w \in \mathsf{Act}^*} \mathbf{P}^F_{\mathcal{A}}(w)$.

Given a threshold $\theta \in [0, 1[$, we set $\mathcal{L}_{>\theta}(\mathcal{A}) = \{w \in \mathsf{Act}^* \mid \mathbf{P}^F_{\mathcal{A}}(w) > \theta\}$. The strict emptiness problem for $\mathcal{A}$, asking whether this set is empty or not, is known to be undecidable for $\theta > 0$ [16]. The value 1 problem, asking whether $val(\mathcal{A}) = 1$ is undecidable as well [11].

Sketch of Proof. Given a PA $\mathcal{A}$, we build a Markov chain $\mathcal{M}_{\mathcal{A}}$ with initial distribution μ_0 and secret Sec such that for any ε, $0 < \varepsilon < 1$, $\mathcal{L}_{>1-\varepsilon}(\mathcal{A})$ is not empty iff $Disc^{\varepsilon}(\mathcal{M}_{\mathcal{A}}(\mu_0)) > 0$. ◀

This leads us to return to the simpler case where the disclosure is the probability of the set of paths leaking the secret, *i.e.*, such that *all* paths with the same observation are secret. The ω-disclosure (corresponding to measures in [3, 5, 4]) was defined for a Markov chain $\mathcal{M} = (S, p, \mathsf{O})$ with initial distribution μ_0 by considering a measurable set of secret paths $\mathsf{SPath} \subseteq \mathsf{Path}(\mathcal{M}(\mu_0))$. Here, as mentioned above, SPath is $\mathsf{Reach}(Sec)$, the set of infinite paths visiting a state from Sec, and an infinite observation $w \in \Sigma^\omega$ discloses the secret if all paths $\rho \in \mathsf{O}^{-1}(w)$ are secret. Setting $\overline{\mathsf{SPath}} = \mathsf{Path}(\mathcal{M}(\mu_0)) \setminus \mathsf{SPath}$, we define:

▶ **Definition 4** (ω-Disclosure)**.** For an MC $\mathcal{M} = (S, p, \mathsf{O})$, an initial distribution μ_0 and a subset $Sec \subseteq S$, with $\mathsf{SPath} = \mathsf{Reach}(Sec)$, the ω-disclosure is defined by:
$Disc_\omega(\mathcal{M}(\mu_0)) = \mathbf{P}_{\mathcal{M}(\mu_0)}(\mathsf{SPath} \setminus \mathsf{O}^{-1}(\mathsf{O}(\overline{\mathsf{SPath}})))$.

To obtain measures directly related to the finite observation of a possible attacker, we assume that $\mathcal{M} = (S, p, \mathsf{O})$ is *convergent*: each infinite path ρ has an infinite observation $\mathsf{O}(\rho) \in \Sigma^\omega$. Two measures can then be defined, when considering a fixed or finite horizon. In the former case, we consider a non-erasing function O to obtain real-time observations.

▶ **Definition 5** (Disclosure of MCs)**.** Let $\mathcal{M} = (S, p, \mathsf{O})$ be an MC, μ_0 an initial distribution and $Sec \subseteq S$. A finite observation $w \in \Sigma^*$ discloses the secret if all paths $\rho \in \mathsf{O}^{-1}(w)$ are secret. It is min-disclosing if it discloses the secret and no strict prefix of w does.

***n*-disclosure:** When O is non-erasing, we denote by D_n, for $n \in \mathbb{N}$, the set of disclosing observations of length n. The n-disclosure is $Disc_n(\mathcal{M}(\mu_0)) = \sum_{w \in D_n} \mathbf{P}_{\mathcal{M}(\mu_0)}(w)$;

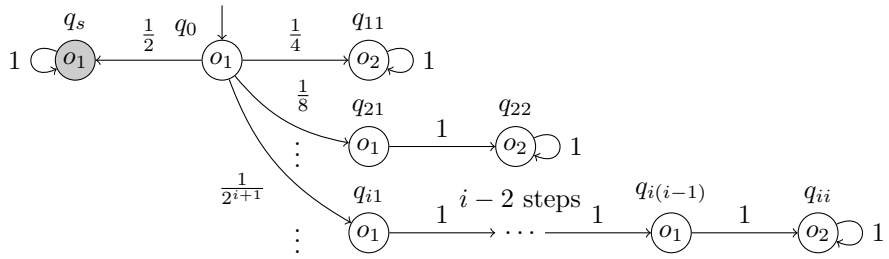

Figure 2 An infinitely branching MC with $Sec = \{q_s\}$, $Disc_\omega = \frac{1}{2}$ and $Disc = 0$.

Disclosure: Writing $D_{\min}$ for the set of min-disclosing observations, the disclosure (w.r.t. finite horizon) is defined by $Disc(\mathcal{M}(\mu_0)) = \sum_{w \in D_{\min}} \mathbf{P}_{\mathcal{M}(\mu_0)}(w)$.

Note that if D is the set of disclosing observations, and $\mathcal{V}(\mu_0) = \cup_{w \in D} \cup_{\rho \in \mathsf{O}^{-1}(w)} \mathsf{Cyl}(\rho)$ the set of paths disclosing the secret, then $Disc(\mathcal{M}(\mu_0))$ is also equal to $\mathbf{P}_{\mathcal{M}(\mu_0)}(\mathcal{V}(\mu_0))$.

▶ Remark. Without loss of generality, we can assume that once a secret state has been reached by an execution, all subsequent states remain secret. For this, a new Markov chain $\mathcal{M}' = (S', p', \mathsf{O}')$ is defined from $\mathcal{M}$ by: $S' = Sec \uplus ((S \setminus Sec) \times \{0,1\})$, where $(s,0)$ represents state s where the secret has not been visited while $(s,1)$ represents the opposite situation. The transitions are then duplicated accordingly: (1) $p'((s',i)|(s,i)) = p(s'|s)$ for all $s, s' \in S \setminus Sec$, and $i = 0,1$, (2) $p'((s',1)|s) = p(s'|s)$ for all $s \in Sec$, and $s' \in S \setminus Sec$, (3) $p'(s'|(s,i)) = p(s'|s)$ for all $s \in S \setminus Sec$, $i = 0,1$, and $s' \in Sec$, and (4) $p'(s'|s) = p(s'|s)$ for all $s, s' \in Sec$. The observation function is extended by $\mathsf{O}'((s,i)) = \mathsf{O}(s)$ for all $s \in S \setminus Sec$ and $i = 0,1$ and the new set of secrets is $Sec \uplus ((S \setminus Sec) \times \{1\})$. There is a one-to-one probability-preserving correspondence between the paths in $\mathcal{M}$ and those in $\mathcal{M}'$.

We show that disclosure and ω-disclosure may be different by consider the infinitely branching MC of Figure 2, with initial distribution $\mathbf{1}_{q_0}$, $Sec = \{q_s\}$ hence $\mathsf{SPath} = \{q_0 q_s^\omega\}$, $\mathsf{O}(\overline{\mathsf{SPath}}) = o_1^+ o_2^\omega$. Then $Disc_\omega = \frac{1}{2}$ but since no finite observation is disclosing, $Disc = 0$.

However, both notions coincide for convergent finitely branching MCs.

▶ **Lemma 6** (Comparison of Disclosure Notions). *Let $\mathcal{M} = (S, p, \mathsf{O})$ be a Markov chain, μ_0 an initial distribution and $Sec \subseteq S$. For $\mathsf{SPath} = \mathsf{Reach}(Sec)$, $Disc(\mathcal{M}(\mu_0)) \leq Disc_\omega(\mathcal{M}(\mu_0))$ and equality holds if $\mathcal{M}$ is convergent and finitely branching.*

2.2 Opacity for Markov Decision Processes

We now turn to MDPs that combine non determinism with probabilistic transitions.

▶ **Definition 7** (MDP). A Markov Decision Process (MDP) over alphabet Σ is a tuple $\mathsf{M} = (S, \mathsf{Act}, p, \mathsf{O})$ where S is a finite set of states, $\mathsf{Act} = \cup_{s \in S} A(s)$ where $A(s)$ is a finite non-empty set of actions for each state $s \in S$, $p : S \times \mathsf{Act} \to \mathsf{Dist}(S)$ is the (partial) transition function defined for (s,a) when $a \in A(s)$ and $\mathsf{O} : S \to \Sigma \cup \{\varepsilon\}$ is the observation function.

As before, we write $p(s'|s,a)$ instead of $p(s,a)(s')$. Given an initial distribution μ_0, an infinite path of M is a sequence $\rho = s_0 a_0 s_1 a_1 \ldots$ where $\mu_0(s_0) > 0$ and $p(s_{i+1}|s_i, a_i) > 0$, for $s_i \in S$, $a_i \in A(s_i)$, for all $i \geq 0$. Finite paths (ending in a state) and observation of a path are defined like for Markov chains, and we use similar notations for the various sets of paths. For decidability and complexity results, we assume that all probabilities occurring in the model (transition probabilities and initial distribution) are rationals.

Nondeterminism is resolved by strategies. Given a finite path ρ with $\mathsf{last}(\rho) = s$, a *decision rule* for ρ is a distribution on the possible actions in $A(s)$ chosen at this point. For such a decision rule δ, we write $p(s'|s,\delta) = \sum_{a\in A(s)} \delta(a)p(s'|s,a)$.

▶ **Definition 8** (Strategy). A strategy of MDP $\mathsf{M} = (S, \mathsf{Act}, p, \mathsf{O})$ with initial distribution μ_0 is a mapping $\sigma : \mathsf{FPath}(\mathsf{M}(\mu_0)) \to \mathsf{Dist}(\mathsf{Act})$ associating with ρ a decision rule $\sigma(\rho)$.

Given a strategy σ, a path $\rho = s_0a_0s_1a_1\dots$ of M is σ-compatible if for all i, $a_i \in Supp(\sigma(s_0a_0s_1a_1\dots s_i))$. A strategy σ is *deterministic* if $\sigma(\rho)$ is a Dirac distribution for each finite path ρ. In this case, we denote by $\sigma(\rho)$ the single action $a \in A(\mathsf{last}(\rho))$ such that $\sigma(\rho) = \mathbf{1}_a$. A strategy σ is *observation-based* if for any finite path ρ, $\sigma(\rho)$ only depends on the observation sequence $O(\rho)$ and the current state $\mathsf{last}(\rho)$, writing $\sigma(O(\rho), \mathsf{last}(\rho))$ for $\sigma(\rho)$.

Let σ be a strategy and ρ be a σ-compatible path. We define B^σ_ρ the *belief* of ρ w.r.t. σ about states corresponding to the last observation as follows:

$$B^\sigma_\rho = \{s \mid \exists \rho' \ \sigma\text{-compatible},\ O(\rho') = O(\rho) \wedge s = \mathsf{last}(\rho') \wedge O(s) \neq \varepsilon\}$$

A strategy σ is *belief-based* if for all ρ, $\sigma(\rho)$ only depends on its belief B^σ_ρ and its current state $\mathsf{last}(\rho)$. Observe that a belief-based strategy is observation-based since B^σ_ρ only depends on $w = \mathsf{O}(\rho)$. So we also write B^σ_w for B^σ_ρ. A strategy σ is *memoryless* if $\sigma(\rho)$ only depends on $\mathsf{last}(\rho)$ for all ρ.

A strategy σ on $\mathsf{M}(\mu_0)$ defines a (possibly infinite) Markov chain $\mathsf{M}_\sigma(\mu_0)$ with set of states $\mathsf{FPath}(\mathsf{M}_\sigma(\mu_0))$ (the finite σ-compatible paths), that can be equipped with the observation function associating $\mathsf{O}(\mathsf{last}(\rho))$ with the finite path ρ. The transition function p_σ is defined for $\rho \in \mathsf{FPath}(\mathsf{M}_\sigma(\mu_0))$ and $\rho' = \rho a s'$ by $p_\sigma(\rho'|\rho) = \sigma(\rho)(a)p(s'|s,a)$ and we denote by $\mathbf{P}_{\mathsf{M}_\sigma(\mu_0)}$ (or $\mathbf{P}_\sigma$ for short when there is no ambiguity) the associated probability measure. Writing $\mathcal{V}_\sigma(\mu_0)$ for the set of paths disclosing the secret in $\mathsf{M}_\sigma(\mu_0)$, we have $Disc(\mathsf{M}_\sigma(\mu_0)) = \mathbf{P}_{\mathsf{M}_\sigma(\mu_0)}(\mathcal{V}_\sigma(\mu_0))$.

Disclosure values for MDPs are defined according to the status of the strategies, by considering them as adversarial or cooperative with respect to the system (we only consider ε-disclosure for fixed horizon in view of the undecidability result of Theorem 3).

▶ **Definition 9** (Disclosure of an MDP). Given an MDP $\mathsf{M} = (S, \mathsf{Act}, p, \mathsf{O})$, an initial distribution μ_0 and a secret $Sec \subseteq S$, the maximal disclosure of Sec in M is $disc_{\max}(\mathsf{M}(\mu_0)) = \sup_\sigma disc(\mathsf{M}_\sigma(\mu_0))$ and the minimal disclosure is $disc_{\min}(\mathsf{M}(\mu_0)) = \inf_\sigma disc(\mathsf{M}_\sigma(\mu_0))$ for $disc \in \{Disc, Disc_n, Disc^\varepsilon_n\}$, $n \in \mathbb{N}$ and $0 < \varepsilon < 1$.

Note that the construction ensuring that once a secret state is visited, the path remains secret forever, extends naturally from Markov chains to MDPs. We consider only MDPs of this form in the sequel. We now show that for disclosure problems we can restrict strategies to observation-based ones.

▶ **Proposition 10** (Observation-based strategies). *Given an MDP, a secret and a strategy σ, there exists an observation-based strategy σ' with the same disclosure values.*

Erasing observations leads to technical and cumbersome developments. In order to avoid them in the design of procedures for the finite horizon case, we apply the preliminary transformation described in the next proposition. We precisely state the size of the transformed MDP in view of complexity results.

▶ **Proposition 11** (Avoiding erasing observations). *Given an MDP $\mathsf{M} = (S, \mathsf{Act}, p, \mathsf{O})$, an initial distribution μ_0 and a secret Sec, one can build in exponential time an MDP $\mathsf{M}' =$*

$(S', \mathsf{Act}', p', \mathsf{O}')$, an initial distribution μ'_0 and a secret Sec' where O' is non-erasing and for $disc \in \{Disc_{min}, Disc_{max}\}$ $disc(\mathsf{M}(\mu_0)) = disc(\mathsf{M}'(\mu'_0))$. In addition, the size of S', p' and μ'_0 is polynomial w.r.t. those of S', p' and μ'_0. The size of Act' is polynomial w.r.t. the size of Act and exponential w.r.t. the size of S.

We study the following problems over MDPs:

- **Computation problems.** The *value problem*: compute the disclosure and the *strategy problem*: compute an optimal strategy whenever it exists;
- **Quantitative decision problems.** The *disclosure problem*: Given M and a threshold $\theta \in [0, 1]$, is $disc(\mathsf{M}) \bowtie \theta$? with $\bowtie = \geq$ for maximisation and $\bowtie = \leq$ for minimisation, and the more demanding *strategy decision problem*: does there exist a strategy σ such that $disc(\mathsf{M}_\sigma) \bowtie \theta$?
- **Qualitative decision problems.** The *limit-sure disclosure problem*: the disclosure problem when $\theta = 1$ for maximisation and $\theta = 0$ for minimisation and the *almost-sure disclosure problem*: the strategy decision problem with the same restrictions.

For the complexity results regarding a fixed horizon n, we will assume that n is written in unary representation or bounded by a polynomial in the size of the model where the polynomial is independent of the model as done in classical studies (see for instance [15]).

3 Maximisation with finite horizon

While strategies may be randomised, this additional power is not necessary for maximisation:

▶ **Proposition 12** (Dominance of deterministic strategies)**.** *Given an MDP, a secret and an observation-based strategy σ there exists a deterministic observation-based strategy σ' with greater or equal disclosure of the secret.*

Sketch of Proof. The Lemma 1 of [8] (or alternatively [12]) does not directly give the result as, contrary to the objectives used in their paper, disclosure depends on the strategy. However, as a disclosing path for a randomised strategy is also a disclosing path for a deterministic strategy that does not introduce new paths, we can use parts of their proof to show our result. ◀

An edge can be completely blocked by some strategy, modifying the set of paths that disclose the secret. This was illustrated in Figure 1a, where choosing action a in state q_0 removes the edges to q_3 and q_4. This situation was excluded in the disclosure computation from [4], where the general problem was left open for Interval Markov Chains (IMCs). We answer negatively by proving undecidability of the disclosure problem, hence the disclosure cannot be computed in general. Undecidability also holds for limit-sure disclosure.

Writing $\mathbb{I}$ for the set of intervals in $[0, 1]$, an IMC (with observation) is a tuple $\mathsf{M} = (S, s_{init}, I, \mathsf{O})$ where S is the set of states, s_{init} is the initial state, $I : S \to \mathbb{I}^S$ associates with any state $s \in S$ a mapping from S into $\mathbb{I}$, and $\mathsf{O} : S \to \Sigma \cup \{\varepsilon\}$ is the observation function. An IMC can be transformed into an (exponentially larger) MDP where actions are the basic feasible solutions of the linear program specified by the constraints associated with intervals [18]. Thus undecidability results for IMCs also hold for MDPs.

▶ **Theorem 13** (Undecidability of maximal finite horizon disclosure)**.** *The maximal finite horizon disclosure problem is undecidable for MDPs, even when the secret is reached with probability 1 and for a non-erasing observation function.*
The maximal finite horizon disclosure problem when restricted to finite-memory strategies is also undecidable (with the same additional assumptions).

Sketch of Proof. Starting from a PA $\mathcal{A}$, we build an IMC $\mathsf{M}_{\mathcal{A}} = (S, s_0, I, \mathsf{O})$ such that there exists a word $w \in \{a, b\}^*$ with $\mathbf{P}^F_{\mathcal{A}}(w) > \frac{1}{2}$ if and only if $Disc_{\max}(\mathsf{M}_{\mathcal{A}}) > \frac{1}{4}$. While similar to the one of Theorem 3, the proof is more involved because the strategies must be taken into account. ◀

As a consequence, we obtain:

▶ **Corollary 14.** *The maximal finite horizon disclosure of an MDP cannot be computed.*

Using a reduction of the value 1 problem in PA, we also have:

▶ **Theorem 15** (Undecidability of maximal finite horizon limit-sure disclosure)**.** *The maximal finite horizon limit-sure disclosure problem is undecidable for MDPs.*

Fortunately the maximal finite-horizon almost-sure disclosure problem is decidable. The proof relies on results for partially observable MDPs (POMDPs): a POMDP is an MDP where the strategies resolving the non determinism only depend on the observation sequence and do not take the current state into account.

▶ **Theorem 16** (Decidability of maximal finite-horizon almost-sure disclosure)**.** *The maximal finite-horizon almost-sure disclosure problem in MDPs is* EXPTIME*-complete. Moreover, if the system is almost-surely disclosing, one can build a belief-based strategy with disclosure* 1*.*

Sketch of Proof. We reduce the almost-sure disclosure problem for maximisation in MDPs to almost-sure reachability in a POMDP. The POMDP we build is exponential in the size of the original MDP and the algorithm to solve almost-sure reachability is exponential in the size of the POMDP [9]. This gives an EXPTIME algorithm as those two exponentials do not stack. The hardness is obtained by a reduction from the safety problem in games with imperfect information that was shown to be EXPTIME-complete in [6]. ◀

4 Minimisation with finite horizon

Recall from the example illustrated in Figure 1a of introduction, that randomised strategies are necessary for minimisation. To address this issue we introduce *families of almost deterministic strategies* based on *ε-decision rules*, that will be used in the decision procedures.

▶ **Definition 17.** Let δ be the deterministic decision rule for state s selecting action $a \in A(s)$. Then $\delta_\varepsilon \in \mathsf{Dist}(A(s))$ is a (randomised) ε-decision rule, said to *favour* a, and defined by:
1. If $|A(s)| > 1$ then $\delta_\varepsilon(a) = 1 - \varepsilon$ and for all $b \in A(s) \setminus \{a\}$, $\delta_\varepsilon(b) = \frac{\varepsilon}{|A(s)|-1}$;
2. Else $\delta_\varepsilon(a) = 1$.

▶ **Definition 18.** Let σ be an observation-based deterministic strategy. Then $\{\sigma_\varepsilon\}_{\varepsilon>0}$ is a family of *observation-based almost deterministic strategies* defined for any state s and $w \in \Sigma^n$, an observation of length $n \in \mathbb{N}$, by: $\sigma_\varepsilon(w, s) = \sigma(w, s)_{2^{-n}\varepsilon}$.

Using Proposition 11, we assume that the observation function O is non-erasing. The complexity of the transformation does not affect the results since the complexities are all polynomial in the number of actions. To compute the minimal disclosure value, we build from an MDP M, another MDP $\mathsf{M}_{\min}$ which is a "correct abstraction" (as stated by Proposition 19) for reducing minimal disclosure problems to minimal reachability problems, by enlarging states with the maximal belief that can occur independently of the action that has been selected.

Given a set of potential current states B and a new observation o, we define the maximal set of potential next states $\mathsf{NextMax}(B, o)$ over decision rules applied to B by:

$$\mathsf{NextMax}(B, o) = \{s' \in \mathsf{O}^{-1}(o) \mid \exists s \in B\ \exists a \in A(s)\ p(s'|s, a) > 0\}$$

Observe that given a family of almost deterministic strategies $\{\sigma_\varepsilon\}$ and a path ρas of M with $\mathsf{O}(s) = o$, one has $B^{\sigma_\varepsilon}_{\rho as} = \mathsf{NextMax}(B^{\sigma_\varepsilon}_{\rho}, o)$. Then $\mathsf{M}_{\min}$ is formally defined as follows:

- $S_{\min}$, the set of states, is defined by: $S_{\min} = \{(s, B) \mid s \in B \subseteq \mathsf{O}^{-1}(\mathsf{O}(s))\}$;
- Let $(s, B) \in S_{\min}$. Then $A(s, B) = A(s)$;
- Let $(s, B), (s', B') \in S_{\min}$. If $B' = \mathsf{NextMax}(B, \mathsf{O}(s'))$ then $p((s', B')|(s, B), a) = p(s'|s, a)$ else $p((s', B')|(s, B), a) = 0$.

Given μ_0 an initial distribution over S, the associated initial distribution $\mu_{\min}$ over $S_{\min}$ is defined by $\mu_{\min}(s, Supp(\mu_0) \cap \mathsf{O}^{-1}(\mathsf{O}(s))) = \mu_0(s)$ and $\mu_{\min}(s, B) = 0$ for all other B. We define the subset $\mathsf{Avoid}(Sec) \subseteq S_{\min}$ by $\mathsf{Avoid}(Sec) = \{(s, B)|B \subseteq Sec\}$.

▶ **Proposition 19.** *The minimal disclosure value for Sec in* $\mathsf{M}(\mu_0)$ *is equal to the minimal probability to reach* $\mathsf{Avoid}(Sec)$ *in* $\mathsf{M}_{\min}(\mu_{\min})$*. Furthermore it is asymptotically reached by a family of belief-based almost deterministic strategies.*

Since minimal reachability probability in MDPs can be computed in polynomial time we immediately obtain the first part of the next theorem. We establish the second part (PSPACE-hardness) in the proof of Theorem 23.

▶ **Theorem 20.** *The minimal disclosure value of* $\mathsf{M}(\mu_0)$ *can be computed in* EXPTIME*. The associated decision problem is* PSPACE*-hard.*

We now turn to the existence of a strategy that achieves the minimal value and establish that it can be analysed without additional complexity. The main ingredient of the proof is an equation system over states of the MDP whose unique solution is the minimal reachability probability vector.

Notations. Given μ a distribution over states and $\vec{\delta}$ a vector of decision rules over states in the support of μ, we define $\mathsf{NextDist}(\mu, \vec{\delta})$ the next distribution over S when applying $\vec{\delta}$ by:

$$\mathsf{NextDist}(\mu, \vec{\delta})(s') = \sum_{s \in Supp(\mu)} \mu(s) p(s'|s, \vec{\delta}[s]) \quad \text{for any } s' \in S.$$

For a distribution μ over S and $o \in \Sigma$, we write $\mu(o)$ for $\mu(\mathsf{O}^{-1}(o))$. If $Supp(\mu) \cap \mathsf{O}^{-1}(o) \neq \emptyset$, the relative distribution μ_o over $\mathsf{O}^{-1}(o)$ is defined by: $\mu_o(s) = \frac{\mu(s)}{\mu(o)}$ for $s \in \mathsf{O}^{-1}(o)$ and $\mu_o(s) = 0$ otherwise.

We have $Disc_{\min}(\mathsf{M}(\mu)) = \sum_{o \in \Sigma} \mu(o) Disc_{\min}(\mathsf{M}(\mu_o))$. For use in the next proof, we define $disc^*(\mathsf{M}(s, B))$ as the minimal disclosure value when starting in M in state s with belief B. Given some belief B and some decision rule vector $\vec{\delta}$ over B we introduce the possible successors of B when applying $\vec{\delta}$: $\mathsf{Next}(B, \vec{\delta}) = \{s' \mid \exists s \in B\ \exists a \in Supp(\vec{\delta}[s])\ p(s'|s, a) > 0\}$ and $\mathsf{Next}(B, \vec{\delta}, o) = \mathsf{Next}(B, \vec{\delta}) \cap \mathsf{O}^{-1}(o)$.

▶ **Theorem 21.** *The existence of a strategy that achieves the minimal disclosure value can be decided in* EXPTIME*. In the positive case, this strategy can be computed in* EXPTIME*.*

Proof. The algorithm simultaneously solves the existence and the synthesis problem.
Using proposition 19, the algorithm computes for all $(s, B) \in S_{\min}$, $disc^*(\mathsf{M}(s, B))$.
Then it maintains a set Win of beliefs initially set to all beliefs from which it iteratively eliminates items and stops when no more elimination is possible.
Given $B \in Win$, it looks for a decision rule vector $\vec{\delta}$ over B such that:

- for all $o \in \mathsf{O}(\mathsf{Next}(B, \vec{\delta}))$, $\mathsf{Next}(B, \vec{\delta}, o) \in Win$;
- for all $s \in B$, $disc^*(\mathsf{M}(s, B)) = \sum_{o \in \Sigma} \sum_{s' \in \mathsf{O}^{-1}(o)} p(s'|s, \vec{\delta}[s]) disc^*(\mathsf{M}(s', \mathsf{Next}(B, \vec{\delta}, o)))$.

If such a $\vec{\delta}$ does not exist then B is eliminated from Win. Each iteration can be performed in polynomial time w.r.t. $|S_{\min}|$ and the number of iterations is at most $|S_{\min}|$. Observe that when a belief is eliminated, it should not be "reached" by a strategy that obtains the minimal disclosure value. So the elimination is sound.

When the elimination stops, the algorithm answers positively iff for all $o \in \mathsf{O}(Supp(\mu_0))$, $Supp(\mu_0) \cap \mathsf{O}^{-1}(o) \in Win$. Thus by the soundness of the elimination step, if the answer is negative there is no optimal strategy for minimal disclosure value.

If the answer is positive, let us consider the belief-based strategy σ defined by applying the decision rules obtained during the last iteration of the algorithm. On the one hand, under σ when visiting a state s with belief B such that $disc^*(\mathsf{M}(s, B)) = 0$, one never leaves such kind of pairs of states and beliefs. So the secret is never disclosed, showing that the disclosure value obtained by σ for such (s, B) is null. Under σ the disclosure value of all the other pairs of state and belief fulfill the equations of the elimination step. It is known that the single solution of this system is the vector of minimal reachability probabilities of Avoid in $\mathsf{M}_{\min}(\mu_{\min})$ (see [1] for instance) which yields the result. ◀

5 Fixed horizon problems

5.1 Maximal disclosure

In order to compute the value of the maximal disclosure within a fixed horizon, one could build the POMDP described in the proof of Theorem 16 then use pre-existing results on POMDPs. This would result in an EXPTIME algorithm, whereas we obtain in the result below an algorithm with a better complexity in PSPACE.

▶ **Theorem 22** (Computation of the maximal disclosure value within fixed-horizon). *The fixed-horizon maximal value (when the horizon n is described in unary representation) is computable in* PSPACE *and the fixed-horizon maximal disclosure problem is* PSPACE*-complete.*

Sketch of proof - value and membership. We first order the observation alphabet Σ. Then a non deterministic decision procedure operating in PSPACE orderly reads every observation sequence of length n while maintaining the sets of states that were possible after every prefix of this observation, the actions that were chosen nondeterministically in those states and values used in the computation of the disclosure. The information kept is of polynomial size and when every observation has been read, one of the values computed will be exactly the disclosure of the system at time n. We then remove the non determinism using Savitch's Theorem. In order to get the value we observe that we can compute the polynomially sized denominator of this value and then we proceed by iterations of the decision algorithm. ◀

As can be seen in the proof, the optimal strategy could be computed when solving the value problem. However the size of this strategy may be exponential due to the beliefs and thus this strategy is computable in EXPTIME.

For the hardness result, we reduce the truth of a Quantified Boolean Formula (QBF). Recall that QBFs are extension of propositional formulas where boolean variables can be quantified. Syntactically, the formulas are described by the following grammar:

$$\phi ::= \psi \mid \exists x.\phi \mid \forall x.\phi$$
$$\psi ::= x \mid \psi \wedge \psi \mid \psi \vee \psi \mid \neg\psi \mid \texttt{true}$$

A QBF is *closed* if every boolean variable is bound by a quantifier. Deciding if a closed QBF is equivalent to `true` is PSPACE-hard [19].

Sketch of proof - hardness. Given ϕ a closed QBF (w.l.o.g. in 3CNF with n variables and m clauses), we build an MDP M such that ϕ is true iff the disclosure of M is greater or equal to $\frac{1}{2^{2n}}$ in $2(n+m)+3$ steps. In fact, $\frac{1}{2^{2n}}$ is exactly the measure of paths reaching the secret in $2(n+m)+3$ steps, thus every path reaching the secret must be disclosing. Such a path discloses the secret iff a boolean variable of ϕ and its negation (x and $\neg x$ for example) do not occur in its observation.

In M, during the first $2n$ steps, an assignment will be 'given' to each boolean variable: (i) for each existentially quantified boolean variable x, the strategy chooses whether x or $\neg x$ occurs in the observation and (ii) for each universally quantified boolean variable y, by a random choice with probability $\frac{1}{2}$. During the last $2m$ steps, the strategy must trigger a boolean variable in every clause of ϕ so that if a clause is not satisfied by the current assignment, then a boolean variable will be observed as both true and false during the path. Thus the observation would not disclose the secret. ◀

The existence of an optimal strategy here implies that the limit-sure and the almost-sure problem are equivalent. Moreover, the secret being revealed with probability 1 in a given number of steps implies that every path reaches the secret in this number of steps. Therefore the almost-sure problem can be seen as a reachability problem in an MDP which can be solved in polynomial time.

The proof of hardness can be adapted for ε-disclosure, but the algorithm for membership can not be directly applied. The ε-disclosure could however be computed by minimising an exponential system of equations, resulting in an exponential time algorithm.

5.2 Minimal disclosure

The proofs of the two first assertions of the next theorem are similar to the proof of Theorem 22. However in order to get the same complexity for the last assertion, we establish that when a randomised decision rule must be selected in the optimal strategy, it can always be uniformly distributed over its support.

▶ **Theorem 23** (Minimal disclosure within fixed horizon)**.** *The fixed horizon minimal value is computable in* PSPACE*. The fixed horizon minimal disclosure problem is* PSPACE*-complete. In addition, the strategy decision problem is also decidable in* PSPACE*.*

The above proof implies PSPACE-completeness for the limit-sure and almost-sure problem for minimisation. The remark on ε-disclosure of the previous subsection holds again here.

6 Conclusion

We revisit the problems of disclosure for MDPs by (1) taking into account general actions contrary to previous work and (2) considering both maximisation and minimisation problems. We almost fully characterise the decidability and complexity of those problems establishing an asymmetry between minimisation and maximisation problems: the former ones being easier although they require families of randomised strategies for reaching the optimal value.

There remains a complexity gap (PSPACE versus EXPTIME) for the finite-horizon minimisation problem that we want to fill. From a qualitative point of view, observe that disclosure is a hyperproperty as its truth value is defined relatively to a set of paths. Thus we plan to

address such kinds of properties in a restricted setting in order to get other decidability results. Another direction would be to strengthen the requirement for approximate ε-disclosure to regain decidability within finite horizon.

References

1 C. Baier and J.-P. Katoen. *Principles of model checking*. MIT Press, 2008.

2 Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2001. doi:10.1007/3-540-44647-8_1.

3 Béatrice Bérard, Krishnendu Chatterjee, and Nathalie Sznajder. Probabilistic opacity for markov decision processes. *Inf. Process. Lett.*, 115(1):52–59, 2015. doi:10.1016/j.ipl.2014.09.001.

4 Béatrice Bérard, Olga Kouchnarenko, John Mullins, and Mathieu Sassolas. Preserving opacity on interval markov chains under simulation. In Christos G. Cassandras, Alessandro Giua, and Zhiwu Li, editors, *13th International Workshop on Discrete Event Systems, WODES 2016, Xi'an, China, May 30 - June 1, 2016*, pages 319–324. IEEE, 2016. doi:10.1109/WODES.2016.7497866.

5 Béatrice Bérard, John Mullins, and Mathieu Sassolas. Quantifying opacity. *Mathematical Structures in Computer Science*, 25(2):361–403, 2015. doi:10.1017/S0960129513000637.

6 Dietmar Berwanger and Laurent Doyen. On the power of imperfect information. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2008, December 9-11, 2008, Bangalore, India*, volume 2 of *LIPIcs*, pages 73–82. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008. doi:10.4230/LIPIcs.FSTTCS.2008.1742.

7 Jeremy Bryans, Maciej Koutny, Laurent Mazaré, and Peter Y. A. Ryan. Opacity generalised to transition systems. *Int. J. Inf. Sec.*, 7(6):421–435, 2008. doi:10.1007/s10207-008-0058-x.

8 Krishnendu Chatterjee, Laurent Doyen, Hugo Gimbert, and Thomas A. Henzinger. Randomness for free. In Petr Hlinený and Antonín Kucera, editors, *Mathematical Foundations of Computer Science 2010, 35th International Symposium, MFCS 2010, Brno, Czech Republic, August 23-27, 2010. Proceedings*, volume 6281 of *Lecture Notes in Computer Science*, pages 246–257. Springer, 2010. doi:10.1007/978-3-642-15155-2_23.

9 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Qualitative analysis of partially-observable markov decision processes. In Petr Hlinený and Antonín Kucera, editors, *Mathematical Foundations of Computer Science 2010, 35th International Symposium, MFCS 2010, Brno, Czech Republic, August 23-27, 2010. Proceedings*, volume 6281 of *Lecture Notes in Computer Science*, pages 258–269. Springer, 2010. doi:10.1007/978-3-642-15155-2_24.

10 Krishnendu Chatterjee, Koushik Sen, and Thomas A. Henzinger. Model-checking omega-regular properties of interval markov chains. In Roberto M. Amadio, editor, *Foundations of Software Science and Computational Structures, 11th International Conference, FOSSACS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29 - April 6, 2008. Proceedings*, volume 4962 of *Lecture Notes in Computer Science*, pages 302–317. Springer, 2008. doi:10.1007/978-3-540-78499-9_22.

11 Hugo Gimbert and Youssouf Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Automata, Languages and*

Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part II, volume 6199 of *Lecture Notes in Computer Science*, pages 527–538. Springer, 2010. doi:10.1007/978-3-642-14162-1_44.

12 Jean Goubault-Larrecq and Roberto Segala. Random measurable selections. In Franck van Breugel, Elham Kashefi, Catuscia Palamidessi, and Jan Rutten, editors, *Horizons of the Mind. A Tribute to Prakash Panangaden - Essays Dedicated to Prakash Panangaden on the Occasion of His 60th Birthday*, volume 8464 of *Lecture Notes in Computer Science*, pages 343–362. Springer, 2014. doi:10.1007/978-3-319-06880-0_18.

13 D. J. D. Hughes and V. Shmatikov. Information Hiding, Anonymity and Privacy: a Modular Approach. *Journal of Computer Security*, 12(1):3–36, 2004.

14 Laurent Mazaré. Decidability of opacity with non-atomic keys. In Theodosis Dimitrakos and Fabio Martinelli, editors, *Formal Aspects in Security and Trust: Second IFIP TC1 WG1.7 Workshop on Formal Aspects in Security and Trust (FAST), an event of the 18th IFIP World Computer Congress, August 22-27, 2004, Toulouse, France*, volume 173 of *IFIP*, pages 71–84. Springer, 2004. doi:10.1007/0-387-24098-5_6.

15 Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of markov decision processes. *Math. Oper. Res.*, 12(3):441–450, 1987. doi:10.1287/moor.12.3.441.

16 A. Paz. *Introduction to Probabilistic Automata*. Academic Press, 1971.

17 Anooshiravan Saboori and Christoforos N. Hadjicostis. Current-state opacity formulations in probabilistic finite automata. *IEEE Trans. Automat. Contr.*, 59(1):120–133, 2014. doi:10.1109/TAC.2013.2279914.

18 Koushik Sen, Mahesh Viswanathan, and Gul Agha. Model-checking markov chains in the presence of uncertainties. In Holger Hermanns and Jens Palsberg, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 12th International Conference, TACAS 2006 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2006, Vienna, Austria, March 25 - April 2, 2006, Proceedings*, volume 3920 of *Lecture Notes in Computer Science*, pages 394–410. Springer, 2006. doi:10.1007/11691372_26.

19 M. Sipser. *Introduction to the theory of computation*. Thomson Course Technology, 2006.