



HAL
open science

Probabilistic Opacity in Refinement-Based Modeling

Béatrice Bérard, Olga Kouchnarenko, John Mullins, Mathieu Sassolas

► **To cite this version:**

Béatrice Bérard, Olga Kouchnarenko, John Mullins, Mathieu Sassolas. Probabilistic Opacity in Refinement-Based Modeling. 2015. hal-01347699

HAL Id: hal-01347699

<https://hal.sorbonne-universite.fr/hal-01347699v1>

Preprint submitted on 21 Jul 2016

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

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

Probabilistic Opacity in Refinement-Based Modeling^{*}

Béatrice Bérard^{1,2} and Olga Kouchnarenko³ and John Mullins^{4,**} and
Mathieu Sassolas⁵

¹ Sorbonne Université, UPMC Univ Paris 06, UMR 7606, LIP6, Paris, France

² CNRS, UMR 7606, F-75005 Paris, France

³ Université de Franche-Comté, FEMTO-ST, CNRS UMR 6174, Inria/NGE,
Besançon, France

⁴ Dept. of Computer & Software Eng., École Polytechnique de Montréal
Montreal (Quebec), Canada, H3C 3A7

⁵ Université Paris-Est, LACL, Créteil, France

Abstract. Given a probabilistic transition system (PTS) \mathcal{A} partially observed by an attacker, and an ω -regular predicate φ over the traces of \mathcal{A} , measuring the disclosure of the secret φ in \mathcal{A} means computing the probability that an attacker who observes a run of \mathcal{A} can ascertain that its trace belongs to φ . In the context of refinement, we consider specifications given as Interval-valued Discrete Time Markov Chains (IDTMCs), which are underspecified Markov chains where probabilities on edges are only required to belong to intervals. Scheduling an IDTMC \mathcal{S} produces a concrete implementation as a PTS and we define the worst case disclosure of secret φ in \mathcal{S} as the maximal disclosure of φ over all PTSs thus produced. We compute this value for a subclass of IDTMCs and we prove that refinement can only improve the opacity of implementations.

1 Introduction

Context and motivation. When modeling complex systems, a top-down approach based on the refinement of models allows to gradually specify various system requirements. These refinements are designed to preserve some behavioral properties, like safety, reachability, and liveness under some conditions.

Security requirements, which are not behavioral ones [1], may not fare well under refinement, unless tailored specially to do so, as in [2].

^{*} Partially supported by a grant from Coopération France-Québec, Service Coopération et Action Culturelle 2012/26/SCAC (French Government), the NSERC Discovery Individual grant No. 13321 (Government of Canada), the FQRNT Team grant No. 167440 (Quebec's Government) and the CFQCU France-Quebec Cooperative grant No. 167671 (Quebec's Government).

^{**} This research has been partially done while this author was visiting the LIP6, Université Pierre & Marie Curie.

Several known security properties such as noninference or anonymity can be encoded in the framework of *opacity* [3,4,2]. In this context, an external observer tries to discover whether a predicate (given as an ω -regular set) holds by partially observing the system through a projection of its actions. A system is opaque if the attacker fails to discover this information. In the possibilistic setting, a violation of opacity captures the existence of at least one perfect leak.

In probabilistic models like Discrete Time Markov Chains (DTMCs), naturally random events such as faults, message transmission failure, actions of an external environment, etc. can be taken into account. Opacity was extended in the probabilistic setting [5,6] to provide a *measure* of the set of runs *disclosing* information on the truth value of the predicate. As a result, opacity increases when the disclosure – which is the focus of this paper – decreases.

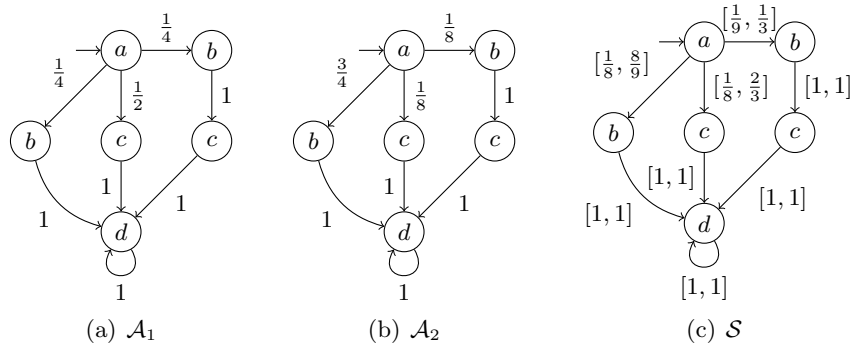


Fig. 1. Probabilistic systems \mathcal{A}_1 or \mathcal{A}_2 implementing underspecified system \mathcal{S} .

Consider for instance the two systems in Fig. 1(a)-(b), which are DTMCs with the addition of labels on states (indicated inside). We assume that the occurrence of b must be kept secret and that all labels except b are observable. In this case, the only runs *disclosing* the secret are those observed by ad^ω , since every such run betrays the occurrence of b . The probability of disclosure is $1/4$ in \mathcal{A}_1 while it is $3/4$ in \mathcal{A}_2 , hence \mathcal{A}_1 is more secure than \mathcal{A}_2 . Our aim is to establish sufficient conditions on systems like \mathcal{A}_1 and \mathcal{A}_2 , that can be compared, for one of them to be more secure than the other.

In the process of system modeling, it is common practice to use *underspecified* models as first steps of specification. A first approach is to consider *sub-stochastic* models where transition probabilities need not sum up to 1. In this framework, the notions of satisfaction and simulation were extensively studied in [7]. The second approach is to introduce non-determinism in the model which makes possible the description of some choices of the environment [8,9,10,11,6]. These models have also been studied in relation to the refinement process [8]. For example, both systems of Fig. 1(a)-(b) could have been derived from a single underspecified system \mathcal{S} with the same structure but imprecise probabilities, like the one in Fig. 1(c). A particular case of such models is the Interval-valued DTMCs (IDTMCs) where the transition probabilities are underspecified by giving only bounds in the form of intervals, as done in Fig. 1(c). The refinement consists for example in restricting these intervals. Implementation, on the other hand, reduces the bounds to a single point, thus providing a Discrete Time Markov Chain.

Scheduling is an effective way to obtain an implementation: at each step, the scheduler provides a distribution satisfying the bounds, producing a (possibly infinite) DTMC on-the-fly. As such, a scheduler can represent a strategy of an agent inside the system. In the case of opacity, this scheduler is adversarial in the sense that it tries to disclose as much information as possible to the passive observer.

Contribution. In this work, we investigate the effect of refinement of probabilistic systems on opacity. Disclosure, as any security measure, is defined on IDTMCs as the worst case scenario for all its implementations, although not every implementation is obtained through scheduling.

In order to evaluate the disclosure of an IDTMC under realistic assumptions, we study the supremum of the disclosure of all scheduled implementations: this measures the information obtained by the passive observer when the system is controlled by the smartest scheduler in coalition with the observer. We show how to compute this value for a subclass of IDTMCs, namely IDTMCs where no transition can be completely blocked by the scheduler. The computation is based on techniques proposed in [10]. Note that the previous approach in [6] has already used schedulers to evaluate disclosure, although in the context of (fully specified) Markov Decision Processes.

We then show that a refinement of an IDTMC can only improve the opacity of all implementations obtained by scheduling. This can be viewed as an extension of the work in [2] to the probabilistic setting. The main

difficulty of this result comes from the restriction of the implementations to those obtained by scheduling.

Organization of the paper. In Section 2 we present the underlying models for specification and implementation, with the satisfaction relation and the simulation relation. We define probabilistic disclosure in this context and show how to compute it (for a restricted case) in Section 3. Finally, we prove monotonicity of opacity under simulation in Section 4. Due to lack of space, several proofs are provided in a separate appendix.

2 Finite automata and probabilistic models for refinement

In this section, we present the models used in this work: finite automata and the probabilistic models for specifications and implementations, as well as the satisfaction and simulation relations.

The set of natural numbers is denoted by \mathbb{N} . The composition of relations \mathcal{R}_2 and \mathcal{R}_1 is defined by $\mathcal{R}_2 \circ \mathcal{R}_1 = \{(x, z) \mid \exists y, (x, y) \in \mathcal{R}_1 \wedge (y, z) \in \mathcal{R}_2\}$. Given 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.

Given a countable set Z , a discrete distribution is a mapping $\mu : Z \rightarrow [0, 1]$ such that $\sum_{z \in Z} \mu(z) = 1$. The support of μ is $\text{supp}(\mu) = \{z \in Z \mid \mu(z) > 0\}$. The set of all discrete distributions on Z is denoted by $\mathcal{D}ist(Z)$. When dealing with a distribution on domain $Z_1 \times Z_2$ (a *joint* distribution), it is convenient to write $\mu(Y_1, Y_2) = \sum_{y_1 \in Y_1, y_2 \in Y_2} \mu(y_1, y_2)$ for $Y_1 \subseteq Z_1$ and $Y_2 \subseteq Z_2$. Moreover, we abusively write $\mu(y_1, Y_2) = \mu(\{y_1\}, Y_2)$ and $\mu(Y_1, y_2) = \mu(Y_1, \{y_2\})$.

2.1 Models

Definition 2.1 (Deterministic Parity Automaton). *A deterministic parity automaton (DPA) is a tuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$, where Q is a finite set of states, Σ is an input alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is a transition function, $q_0 \in Q$ is an initial state, and F is a mapping from Q to a finite set of colors $\{1, \dots, k\}$.*

A run of \mathcal{A} on a word $w = a_1 a_2 \dots \in \Sigma^\omega$ is an infinite sequence $\rho = q_0 q_1 \dots \in Q^\omega$ such that for all $i \geq 0$, $q_{i+1} = \delta(q_i, a_{i+1})$. For such a run ρ , we define $\text{Inf}(\rho)$ as the set of states appearing infinitely often in the sequence. The run is accepting if $\min\{F(q) \mid q \in \text{Inf}(\rho)\}$ is even.

In this case, the corresponding word is accepted by \mathcal{A} and $\mathcal{L}(\mathcal{A})$ is the subset of Σ^ω of words accepted by \mathcal{A} . A subset K of Σ^ω is ω -regular if there is an automaton \mathcal{A} such that $K = \mathcal{L}(\mathcal{A})$.

For the probabilistic context, we consider two models: Probabilistic Transition Systems (PTSs) for implementations, and Interval-valued Discrete Time Markov chains (IDTMCs) for specifications.

Probabilistic transition systems are a particular case of the probabilistic processes of [8], with the restriction of a countable set of states. They are classical Discrete Time Markov Chains (DTMCs), with the addition of state labeling.

Definition 2.2 (Probabilistic Transition System). *Let Σ be an alphabet. A probabilistic transition system (PTS) over Σ is a 4-tuple $\mathcal{A} = \langle Q, q_{init}, \Delta, L \rangle$ where*

- Q is a countable set of states, with $q_{init} \in Q$ the initial state,
- $\Delta : Q \rightarrow \mathcal{D}ist(Q)$ is a mapping associating with any state $q \in Q$ a distribution $\Delta(q)$ over Q , with finite support,
- $L : Q \rightarrow \Sigma$ is the labeling function on states.

A (finite or infinite) run of \mathcal{A} starting from state $q \in Q$ is a sequence of states $\rho = q_0q_1q_2 \dots$ such that $q_0 = q$ and for each i , $0 \leq i < |\rho|$, $\Delta(q_i)(q_{i+1}) > 0$. When the run is finite $\rho = q_0q_1 \dots q_n$, we note $q_n = \text{lst}(\rho)$. The trace of ρ is the word $\text{tr}(\rho) = L(q_0)L(q_1) \dots \in \Sigma^\omega$. We denote by $Runs_q(\mathcal{A})$ the set of infinite runs starting from q and we set $Runs(\mathcal{A}) = Runs_{q_{init}}(\mathcal{A})$, and $Tr(\mathcal{A}) = \{\text{tr}(\rho) \mid \rho \in Runs(\mathcal{A})\}$, the set of traces of \mathcal{A} . We also define $FRuns_q(\mathcal{A})$ the set of finite runs starting from q , and similarly $FRuns(\mathcal{A}) = FRuns_{q_{init}}(\mathcal{A})$ and $FTr(\mathcal{A}) = \{\text{tr}(\rho) \mid \rho \in FRuns(\mathcal{A})\}$, the subset of Σ^* of finite traces of \mathcal{A} .

Recall [12] that a probability measure $\mathbf{P}_{\mathcal{A}}$ can be defined on $Runs(\mathcal{A})$: measurable sets are generated by *cones*, where the cone C_ρ associated with a finite run $\rho = q_0q_1 \dots q_n$ is the subset of infinite runs in $Runs(\mathcal{A})$ having ρ as prefix. The probability of C_ρ is $\mathbf{P}_{\mathcal{A}}(C_\rho) = \prod_{i=0}^{n-1} \Delta(q_i)(q_{i+1})$. The cone of a word $w \in \Sigma^*$ is defined by $C_w = \bigcup_{\rho \in \text{tr}^{-1}(w)} C_\rho$. For an ω -regular language K , we denote by $\mathbf{P}_{\mathcal{A}}(K)$ the probability of K in \mathcal{A} . This probability is well defined (since ω -regular sets are measurable [13]) and can be computed given a DPA for K [14].

The specifications we consider here are given by Interval-valued Discrete Time Markov chains, introduced in [8] and further investigated in [10,11] from a verification point of view. We denote by \mathcal{I} the set of intervals in $[0, 1]$.

Definition 2.3 (Interval-valued Discrete Time Markov Chains).

An Interval-valued Discrete Time Markov Chains (IDTMC) is a 4-tuple $\mathcal{S} = (S, s_{init}, T, \lambda)$ where

- S is a finite set of states, with $s_{init} \in S$ the initial state,
- $T : S \rightarrow (S \rightarrow \mathcal{I})$ associates with any state $s \in S$ a mapping $T(s)$ from S into \mathcal{I} ,
- $\lambda : S \rightarrow \Sigma$ is the labeling function.

By extension, $f \in T(s)$ will denote any function $f : S \rightarrow [0, 1]$ such that for all $s' \in S$, $f(s') \in T(s)(s')$ and $\sum_{s' \in S} f(s') = 1$.

Several semantic interpretations for IDTMCs have been considered in [8,10,11]. The simplest one is Uncertain Markov Chains, which corresponds to first choosing all distributions for the states, with probabilities belonging to the specified intervals. In this case, the resulting model is a PTS, with the same structure as the specification.

A richer semantics consists in choosing the distribution at each step, as in a Markov Decision Process, hence the name Interval Markov Decision Process (IMDP). Note that an IMDP can be translated into an exponentially larger MDP [10], but the transformed MDP is ill-suited to model refinement.

Finally, the most general semantics is directly given by the satisfaction relation from [8] recalled later.

We first describe the IMDP semantics. A run of \mathcal{S} starting from a state s is a sequence $s \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \dots$ where $s_i \in S$ and each μ_i is a distribution over S such that $\forall s \in S, \mu_i(s) \in T(s_{i-1})(s)$. As before, we denote by $Runs_s(\mathcal{S})$ the set of runs starting from s , we set $Runs(\mathcal{S}) = Runs_{s_{init}}(\mathcal{S})$, $FRuns(\mathcal{S})$ is the set of finite runs of \mathcal{S} starting from s_{init} , and for a run $\rho = s \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \dots s_{n-1} \xrightarrow{\mu_n} s_n$ in $FRuns(\mathcal{S})$ we define $lst(\rho) = s_n$.

To associate a probability measure with the runs, it is necessary to resolve the non determinism by a scheduler that chooses a distribution at each step. More precisely:

Definition 2.4 (Scheduler). A scheduler A for an IDTMC specification $\mathcal{S} = (S, s_{init}, T, \lambda)$, is a mapping $A : FRuns(\mathcal{S}) \rightarrow Dist(S)$ such that for each run ρ with $s = lst(\rho)$, $A(\rho)(s') \in T(s)(s')$.

We denote by $Sched(\mathcal{S})$ the set of schedulers for \mathcal{S} . Like for Markov Decision Processes, scheduling \mathcal{S} with A produces a PTS denoted by $\mathcal{S}(A)$ where states are finite runs: $Q \subseteq FRuns(\mathcal{S})$, the initial state is the run containing only the initial state of \mathcal{S} : $q_{init} = s_{init}$, and for $\rho \in Q$,

$L(\rho) = \lambda(\text{lst}(\rho))$ and $\Delta(\rho)(\rho') = A(\rho)(s')$ for $\rho' = \rho \xrightarrow{A(\rho)} s'$. We note $\underline{\text{sat}}(\mathcal{S}) = \{\mathcal{S}(A) \mid A \in \text{Sched}(\mathcal{S})\}$.

Note that the Uncertain Markov Chains semantics corresponds to the particular case of memoryless schedulers: there is a mapping $B : S \rightarrow \text{Dist}(S)$ such that $A(\rho) = B(\text{lst}(\rho))$. In this case the set of states of the PTS $\mathcal{S}(A)$ is $Q = S$ with $\Delta(s)(s') = B(s)(s') \in T(s)(s')$.

2.2 Satisfaction relation

We now turn to the general satisfaction relation defined in [8].

Definition 2.5 (Satisfaction relation). *A PTS $\mathcal{A} = \langle Q, q_{\text{init}}, \Delta, L \rangle$ satisfies IDTMC $\mathcal{S} = (S, s_{\text{init}}, T, \lambda)$, written $\mathcal{A} \models \mathcal{S}$, if there exists a relation $\mathcal{R} \subseteq Q \times S$ such that $q_{\text{init}} \mathcal{R} s_{\text{init}}$ and if $q \mathcal{R} s$ then*

1. $L(q) = \lambda(s)$,
2. *there exists a joint distribution $\delta \in \text{Dist}(Q \times S)$ such that*
 - (a) $\delta(q', S) = \Delta(q)(q')$ for all $q' \in Q$,
 - (b) $\delta(Q, s') \in T(s)(s')$ for all $s' \in S$,
 - (c) $q' \mathcal{R} s'$ whenever $\delta(q', s') > 0$.

We write $\text{sat}(\mathcal{S}) = \{\mathcal{A} \mid \mathcal{A} \models \mathcal{S}\}$ for the set of all PTSs that satisfy specification \mathcal{S} ; these are called *implementations* of \mathcal{S} .

As a first result, we show that scheduling an IDTMC specification is a particular case of implementation:

Proposition 2.6. *Let \mathcal{S} be an IDTMC specification. For each scheduler A of \mathcal{S} , we have: $\mathcal{S}(A) \models \mathcal{S}$, hence $\underline{\text{sat}}(\mathcal{S}) \subseteq \text{sat}(\mathcal{S})$.*

Proof. The relation $\mathcal{R} \subseteq Q \times S$ is defined by $\mathcal{R} = \{(\rho, s) \mid \text{lst}(\rho) = s\}$. We prove that the relation \mathcal{R} is a satisfaction relation by defining $\delta_{(\rho, s)}$ over $Q \times S$ as follows:

$$\delta_{(\rho, s)}(\rho', s') = \begin{cases} A(\rho)(s') & \text{if } \rho' = \rho \xrightarrow{A(\rho)} s', \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The first condition results from the definition of the labeling and conditions 2 and 3 come from the fact that the joint distribution $\delta_{(\rho, s)}$ is diagonal in this case.

- 2 (a) $\delta(\rho', S) = A(\rho)(s') = \Delta(\rho)(\rho')$ with $\rho' = \rho \xrightarrow{A(\rho)} s'$ for all $\rho' \in Q$.
- (b) $\delta(Q, s') = A(\rho)(s') \in T(s)(s')$ with $s = \text{lst}(\rho)$ for all $s' \in S$.

3. $\rho' \mathcal{R} s'$ whenever $\delta(\rho', s') > 0$ since $s' = \text{lst}(\rho')$ by definition of \mathcal{R} . \square

Indeed, for any scheduler A , $\mathcal{S}(A)$ is a kind of *unfolding* of \mathcal{S} , which restricts the structure of $\mathcal{S}(A)$: at each step, the scheduler chooses a valid distribution among successor states. Hence not every implementation can be mapped to a scheduler. Said otherwise, not all implementations can be put in relation with \mathcal{S} with a satisfaction relation where the joint distributions δ are diagonal. This means that the inclusion $\underline{\text{sat}}(\mathcal{S}) \subseteq \text{sat}(\mathcal{S})$ is strict.




Fig. 2. A specification with an implementation that is not the result of scheduling.

For example, consider the specification \mathcal{S}_0 of Fig. 2(a). There is a single possible scheduler for this specification: the one that picks in q_0 probability $\frac{1}{2}$ to go to either q_1 or q_2 (\mathcal{A}_1 in Fig. 2(b)). However, the PTS \mathcal{A}_2 of Fig. 2(c) is also an implementation of this specification ($\mathcal{A}_2 \models \mathcal{S}_0$) where r_2 is split between q_1 and q_2 . The corresponding matrix is

$$\delta(q_0, r_0) = \begin{matrix} & r_0 & r_1 & r_2 & r_3 \\ \begin{matrix} q_0 \\ q_1 \\ q_2 \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{6} & 0 \\ 0 & 0 & \frac{1}{6} & \frac{1}{3} \end{pmatrix} \end{matrix}$$

Finally, the refinement relation between specifications is simply defined as inclusion of the satisfaction sets:

Definition 2.7 (Refinement). For two IDTMC specifications \mathcal{S}_1 and \mathcal{S}_2 , \mathcal{S}_1 refines \mathcal{S}_2 if $\text{sat}(\mathcal{S}_1) \subseteq \text{sat}(\mathcal{S}_2)$.

2.3 Simulation relation

The notion of simulation relation between probabilistic specifications was introduced in [8], where it is proved to be a sufficient condition for refinement: \mathcal{S}_2 simulates \mathcal{S}_1 implies that all implementations of \mathcal{S}_1 are implementations of \mathcal{S}_2 . This notion is adapted to our setting in Definition 2.8 below. We then investigate the effect of simulation when implementations are considered with respect to schedulers only. Therefore we study the relations between $\underline{sat}(\mathcal{S}_1)$ and $\underline{sat}(\mathcal{S}_2)$ whenever \mathcal{S}_2 simulates \mathcal{S}_1 .

Definition 2.8 (Simulation relation). For $\mathcal{S}_1 = (S_1, s_{1,init}, T_1, \lambda_1)$ and $\mathcal{S}_2 = (S_2, s_{2,init}, T_2, \lambda_2)$ two IDTMC specifications, \mathcal{S}_2 simulates \mathcal{S}_1 if there exists a relation $\mathcal{R} \subseteq S_1 \times S_2$ such that $s_{1,init} \mathcal{R} s_{2,init}$ and if $s_1 \mathcal{R} s_2$ then:

1. $\lambda_1(s_1) = \lambda_2(s_2)$,
2. there exists a function $\delta : S_1 \rightarrow \text{Dist}(S_2)$ such that for all $f \in T_1(s_1)$ and $s'_2 \in S_2$,

$$\left(\sum_{s'_1 \in S_1} f(s'_1) \cdot \delta(s'_1)(s'_2) \right) \in T_2(s_2)(s'_2), \quad (2)$$

3. $s'_1 \mathcal{R} s'_2$ whenever $\delta(s'_1)(s'_2) > 0$.




Fig. 3. A simulation of \mathcal{S}_1 by \mathcal{S}_2 .

In Figure 3, dashed lines illustrate the simulation relation \mathcal{R} of \mathcal{S}_1 by \mathcal{S}_2 labeled with $\delta(q_i)(r_j)$ since for Condition (2), we may in this case uniformly use the function:

$$\delta(q_i)(r_j) = \begin{cases} 1 & \text{if } (q_i, q_j) \in \mathcal{R} \\ 0 & \text{otherwise} \end{cases}$$

Note that there is no simulation relation of \mathcal{S}_2 by \mathcal{S}_1 . Indeed, let $f \in T_2(r_0)$ defined as $f(r_1) = 1$. The only way to distribute f over \mathcal{S}_1 in order to satisfy Eq. 2 is to distribute $\frac{2}{3}$ of f to q_1 and $\frac{1}{3}$ to q_2 . Hence, it forces to set $\delta(r_1)(q_1)$ to $\frac{2}{3}$ and $\delta(r_1)(q_2)$ to $\frac{1}{3}$ but this choice for δ is not uniform for any $f \in T_2(r_0)$ as it does not satisfy Eq. 2 for $f(r_1) = \frac{1}{2}$ for instance.

Let us notice that in the case of PTSs, Definition 2.8 is still valid but intervals reduce to points and Equation 2 becomes

$$\sum_{s'_1 \in \mathcal{S}_1} (\Delta_1(s_1)(s'_1) \cdot \delta(s'_1)(s'_2)) = \Delta_2(s_2)(s'_2). \quad (3)$$

2.4 Simulation vs satisfaction

Remark that not only simulation and satisfaction are defined over different kind of models, they express different requirements. Indeed, simulation is symmetrical in the models used but asymmetrical in its semantics since one model must contain the behavior of the other. In contrast, satisfaction concerns different models but the behavior of the PTS must exactly comply to the IDTMC specification: it can neither add behaviors nor introduce new ones.

Nevertheless, these notions coincide when the simulated IDTMC is in fact a PTS:

Proposition 2.9. *Let \mathcal{A} be a PTS satisfying an IDTMC specification \mathcal{S} . Then \mathcal{S} simulates \mathcal{A} when seen as an IDTMC where intervals are reduced to a point.*

Proof (Sketch). The core of the proof relies on building a function δ_{sim} for the simulation relation, from function δ_{sat} given by the satisfaction relation. For any pair of states $q \in Q$ of PTS \mathcal{A} and $s \in S$ of IDTMC \mathcal{S} , define $\delta_{sim}(q, s) = \frac{\delta_{sat}(q, s)}{\delta_{sat}(q, S)}$. One can show that this δ_{sim} satisfies Condition (2) of Definition 2.8.

Also, as noted before, simulation implies inclusion of the set of implementations:

Proposition 2.10 ([8]). *\mathcal{S}_2 simulates \mathcal{S}_1 implies that $sat(\mathcal{S}_1) \subseteq sat(\mathcal{S}_2)$.*

3 Opacity

The original definition of opacity was given in [4] for (non probabilistic) transition systems, with respect to some observation function \mathcal{O} and some predicate φ (the secret) on the runs of the system.

3.1 Opacity for probabilistic models

We consider an ω -regular set $\varphi \subseteq \Sigma^\omega$ and we say that a run ρ of a PTS \mathcal{A} satisfies φ if its trace belongs to φ . We consider also an *observation function* $\mathcal{O} : \Sigma^\omega \rightarrow \Sigma_{ob}^\infty$, where Σ_{ob} is a subset of Σ representing its visible part. If π_{ob} is the standard projection from Σ^∞ onto Σ_{ob}^∞ mapping any element $a \in \Sigma \setminus \Sigma_{ob}$ to ε , the observation function is defined on a word w by $\mathcal{O}(w) = \pi_{ob}(w)$. Alternatively, we could consider a subset P_{ob} of a set of observable atomic propositions P . Then, by setting $P_u = P \setminus P_{ob}$, $\Sigma_{ob} = 2^{P_{ob}}$ and $\Sigma_u = 2^{P_u}$, Σ could be viewed as the product $\Sigma_{ob} \times \Sigma_u$. The observation function \mathcal{O} would be then defined on a word w by $\mathcal{O}(w) = \pi_1(w)$, where π_1 is the projection from $(\Sigma_{ob} \times \Sigma_u)^\infty$ onto Σ_{ob}^∞ mapping each element on its first component.

The observation class of a word w is $[w]_{\mathcal{O}} = \mathcal{O}^{-1}(\mathcal{O}(w))$: for the observer, each word in this class is undistinguishable from w . Then φ is *opaque* with respect to \mathcal{A} and \mathcal{O} if each time a word satisfies φ , another word in the same observation class does not. We denote by $\mathcal{V}(\mathcal{A}, \mathcal{O}, \varphi)$ the set of words violating this condition:

$$\mathcal{V}(\mathcal{A}, \mathcal{O}, \varphi) = (\text{Tr}(\mathcal{A}) \cap \varphi) \setminus (\mathcal{O}^{-1}(\mathcal{O}(\text{Tr}(\mathcal{A}) \setminus \varphi))).$$

Definition 3.1 (Opacity). *Let \mathcal{A} be a PTS, with observation function \mathcal{O} . A predicate φ is opaque in \mathcal{A} for \mathcal{O} if $\mathcal{V}(\mathcal{A}, \mathcal{O}, \varphi) = \emptyset$.*

Equivalently, φ is opaque if for any w satisfying φ , $[w]_{\mathcal{O}} \not\subseteq \varphi$. Variants of opacity have been defined, with other observation functions and predicates, or by requiring symmetry: the predicate φ is *symmetrically opaque* in \mathcal{A} for \mathcal{O} if both φ and $\Sigma^\omega \setminus \varphi$ are opaque.

The notion of probabilistic opacity [15] extends this boolean property by defining a measure of the set of runs violating opacity:

Definition 3.2 (Probabilistic Disclosure). *Let \mathcal{A} be a PTS, with observation function \mathcal{O} and ω -regular predicate φ . The probabilistic disclosure of φ in \mathcal{A} for \mathcal{O} is $\text{Disc}(\mathcal{A}, \mathcal{O}, \varphi) = \mathbf{P}_{\mathcal{A}}(\mathcal{V}(\mathcal{A}, \mathcal{O}, \varphi))$.*

For instance, recall systems \mathcal{A}_1 and \mathcal{A}_2 of Fig. 1. The secret predicate in this case is the set of runs $\varphi_b = ab\Sigma^\omega$ and the observation function is the projection $\pi_{\{a,c,d\}}$ onto $\{a,c,d\}^\omega$. This predicate is not opaque since the run abd^ω discloses the occurrence of b . This is measured by the disclosure: $Disc(\mathcal{A}_1, \pi_{a,c,d}, \varphi_b) = \mathbf{P}_{\mathcal{A}_1}(abd^\omega) = \frac{1}{4}$ and $Disc(\mathcal{A}_2, \pi_{a,c,d}, \varphi_b) = \mathbf{P}_{\mathcal{A}_2}(abd^\omega) = \frac{3}{4}$.

Note that dealing with symmetric opacity would simply require to add both measures for φ and $\Sigma^\omega \setminus \varphi$. Also remark that disclosure only measures probabilities of the observer being *sure* that the run is in the secret. For example, one can model anonymity of an agent α initiating some protocol by defining φ_α as the set of all runs initiated by α . Anonymity of α is then equivalent to opacity of φ_α . In the case where anonymity is not guaranteed, disclosure provides a measure of the threat. In the case where anonymity holds, this measure will be 0 and does not give any insight on the “strength” of anonymity. Other notions measuring this strength were proposed in [16,17] and quantitative opacity for partial disclosure of the secret have also been defined in [5], although they are not linear hence do not fare well under standard optimization techniques.

In order to compare systems with respect to opacity, we consider two PTSs \mathcal{A}_1 and \mathcal{A}_2 over the same alphabet Σ , with a given observation function \mathcal{O} . We say that \mathcal{A}_1 is more opaque than \mathcal{A}_2 if $Disc(\mathcal{A}_1, \mathcal{O}, \varphi) \leq Disc(\mathcal{A}_2, \mathcal{O}, \varphi)$.

We lift the definition of probabilistic disclosure from PTSs to IDTMC specifications as follows:

$$Disc(\mathcal{S}, \mathcal{O}, \varphi) = \sup_{\mathcal{A} \in sat(\mathcal{S})} Disc(\mathcal{A}, \mathcal{O}, \varphi)$$

From the results in [8] mentioned in Proposition 2.10, it is easy to see that if \mathcal{S}_2 simulates \mathcal{S}_1 , then $Disc(\mathcal{S}_1, \mathcal{O}, \varphi) \leq Disc(\mathcal{S}_2, \mathcal{O}, \varphi)$.

To obtain a similar result for the notion of satisfaction restricted to schedulers, we define the restriction of the disclosure to scheduled implementations:

$$\underline{Disc}(\mathcal{S}, \mathcal{O}, \varphi) = \sup_{\mathcal{A} \in \underline{sat}(\mathcal{S})} Disc(\mathcal{A}, \mathcal{O}, \varphi) = \sup_{\mathcal{A} \in Sched(\mathcal{S})} Disc(\mathcal{S}(\mathcal{A}), \mathcal{O}, \varphi)$$

Note that this notion differs from the similar one in [6] for Markov Decision Processes. The notion presented here is finer since the set of runs measured by the disclosure depends on the scheduled implementation. In [6], the set of runs of the disclosure is defined on the (unscheduled) MDP, and its probability is optimized afterwards. This would not be

consistent in IDTMCs, since two scheduled implementations can have different sets of edges with non-null probability, as explained below.

3.2 Computing the probabilistic disclosure of a specification

When the interval of an edge of an IDTMC is non-punctual and closed on the left by 0, then this edge may be present or not in an implementation. Otherwise said, in the case of schedulers, this action can be completely blocked. In keeping with the terminology of [8], we call these edges *modal* edges, and IDTMCs that contain such edges are also called modal IDTMCs.

Definition 3.3 (Modal edge). *An edge $T(s)(s')$ in IDTMC \mathcal{S} is modal if there exists a scheduler A such that in $\mathcal{S}(A)$, for any run ρ with $\text{lst}(\rho) = s$, $\Delta(\rho)(\rho \xrightarrow{A(\rho)} s') = 0$.*

From a modeling point of view, modal edges add a lot of flexibility for refinement. This however means that the range of potential implementation is larger and so it will be harder to obtain meaningful properties. Therefore such edges are desirable in an early modeling phase but less so in the latest refinements. In the context of opacity, removing an edge drastically changes the disclosure, since it can remove ambiguities.

We illustrate the importance of modal edges when computing the disclosure of a specification. For example, consider the modal IDTMC \mathcal{S}_m of Fig. 4(a), with an observation function that observes a and b and a secret being the presence of c . An implementation of \mathcal{S}_m that blocks the




Fig. 4. The influence of modal transitions on disclosure.

direct edge from a to b (Fig. 4(c)) has a disclosure of 1, since the secret is guaranteed to be part of the only possible run. On the other hand, in the non-modal version of the IDTMC (Fig. 4(b)), such implementations are banned and only implementations that retain a small probability to avoid c are allowed. In these implementations, the disclosure is 0, since every run is observed as ab^ω and it is possible that c did not occur.

In the case of non-modal IDTMCs, disclosure can be computed. The proof relies on a translation from IMDP to MDP and synchronization with DPA; it can be found in Appendix A.

Theorem 3.4. *Computing the value of disclosure for an IDTMC specification \mathcal{S} without modal edges can be done in $2EXPTIME$.*

Remarks on modal edges. When a scheduler is faced with the choice to include or exclude a modal edge, it can produce several versions of PTSs, say \mathcal{A}_1 and \mathcal{A}_2 , with $Tr(\mathcal{A}_1) \neq Tr(\mathcal{A}_2)$, hence $\mathcal{V}(\mathcal{A}_1, \mathcal{O}, \varphi) \neq \mathcal{V}(\mathcal{A}_2, \mathcal{O}, \varphi)$. In addition, these choices may be history dependent, as in the example of Fig. 5, with $\varphi = a\Sigma^\omega$ and only letters c and d being observed. Intuitively, a way for the scheduler to always disclose the presence of an initial a is to always follow an a by the same letter, say a c . However, this choice must be made after the first letter has been seen. Moreover, leaving the possibility of a run $ad\cdots$ to occur means that run $ac\cdots$ does not disclose φ . As a result, the scheduler should also take into account φ and the observation function before committing to a choice with respect to modal edges. So far, the general case of modal IDTMCs remains open.




Fig. 5. IDTMC where the choice on modal edge requires history.

4 Monotonicity of opacity under simulation

This section is devoted to the proof of the following result:

Theorem 4.1. *Let \mathcal{S}_1 and \mathcal{S}_2 be IDTMC specifications such that \mathcal{S}_2 simulates \mathcal{S}_1 . Then $\underline{Disc}(\mathcal{S}_1, \mathcal{O}, \varphi) \leq \underline{Disc}(\mathcal{S}_2, \mathcal{O}, \varphi)$.*

Since scheduling is a restrictive way to derive implementations from a specification, it is not the case in general that $\underline{sat}(\mathcal{S}_1) \subseteq \underline{sat}(\mathcal{S}_2)$: although any scheduling $\mathcal{S}_1(A_1)$ of \mathcal{S}_1 with A_1 is an implementation of \mathcal{S}_2 (by Propositions 2.6 and 2.10), this implementation may not be a scheduling (as illustrated by the example of Fig. 2).

Instead, the proof builds a scheduler A_2 for \mathcal{S}_2 that produces an implementation $\mathcal{S}_2(A_2)$ that *simulates* $\mathcal{S}_1(A_1)$ (Theorem 4.2, illustrated in Fig. 6). Then, this simulation is shown to ensure that the probabilities of (cones of) finite words coincide (Propositions 4.3 and 4.4). Disclosure being a measurable event, coincidence of probabilities on cones ensures coincidence of probabilities for the disclosure.




Fig. 6. The result of Theorem 4.2. Relation $sat_1 \circ \mathcal{R}$ always exists but might not be a scheduling.

Notations. Given two specifications \mathcal{S}_1 and \mathcal{S}_2 such that \mathcal{S}_2 simulates \mathcal{S}_1 through relation \mathcal{R} , we define the relation \sim on $FRuns(\mathcal{S}_1) \times FRuns(\mathcal{S}_2)$ by: $\rho_1 \sim \rho_2$ if $|\rho_1| = |\rho_2|$ and at any intermediate step i , the corresponding states satisfy $s_{1,i} \mathcal{R} s_{2,i}$.

Let A_1 and A_2 be two schedulers of \mathcal{S}_1 and \mathcal{S}_2 , respectively. We set $\mathcal{A}_1 = \mathcal{S}_1(A_1)$ and $\mathcal{A}_2 = \mathcal{S}_2(A_2)$, with respective sets of states Q_1 and Q_2 . For $\rho_2 \in Q_2$, we set $sim(\rho_2) = \{\rho_1 \in Q_1 \mid \rho_1 \sim \rho_2\}$. We now define a measure μ_{ρ_2} over $sim(\rho_2)$ by $\mu_{\rho_2}(\rho_1) = \frac{\mathbf{P}_{\mathcal{A}_1}(\rho_1)}{\mathbf{P}_{\mathcal{A}_1}(sim(\rho_2))}$ (where the probability of finite run ρ is abusively written instead of the probability of its cone C_ρ).

We first show how to build a scheduler for \mathcal{S}_2 that simulates the scheduling of \mathcal{S}_1 . The proof of the theorem below is given in Appendix B.

Theorem 4.2. *Let \mathcal{S}_1 and \mathcal{S}_2 be IDTMC specifications such that \mathcal{S}_2 simulates \mathcal{S}_1 . Then for any $A_1 \in Sched(\mathcal{S}_1)$ there exists $A_2 \in Sched(\mathcal{S}_2)$ such that $\mathcal{S}_2(A_2)$ simulates $\mathcal{S}_1(A_1)$.*

Now we show that simulation between two PTSs is sufficient to compare their disclosure. Namely, we show that the probabilities of cones of words are equal in both systems. Note that although this property is well known to hold for paths, it needs to be lifted to words in order to compare disclosure.

We start by considering the sets of traces globally; although it is folklore that simulation implies trace inclusion, we provide a proof for completeness sake.

Proposition 4.3. *Let \mathcal{A}_1 and \mathcal{A}_2 be PTSs such that \mathcal{A}_2 simulates \mathcal{A}_1 . Then $Tr(\mathcal{A}_1) = Tr(\mathcal{A}_2)$.*

Proof. We prove the proposition by induction on a strengthened statement. Namely, we claim that for every finite run in \mathcal{A}_1 there exists a similar run in \mathcal{A}_2 . Since an infinite run is the limit of the sequence of its finite prefixes, this claim is sufficient to prove the proposition. Assume by induction that the proposition holds for every word of length n . Let $w \in FTr(\mathcal{A}_1)$ of length $n+1$. We write $w = w_0a$ for some $a \in \Sigma$. Consider a run of \mathcal{A}_1 that produces w . It is of the form $\rho_0s'_1$ where $\lambda(s'_1) = a$; let $s_1 = \text{lst}(\rho_0)$. Let ρ'_0 be a run in \mathcal{A}_2 , similar to ρ_0 , and $s_2 = \text{lst}(\rho'_0)$. By definition of simulation, there exists a function δ such that for any state s'_2 of \mathcal{A}_2 ,

$$\Delta_2(s_2)(s'_2) = \sum_{\sigma_1 \in S_1} \Delta_1(s_1)(\sigma_1) \cdot \delta(\sigma_1)(s'_2).$$

Moreover, whenever $\delta(\sigma_1)(s'_2) > 0$, $\lambda(s'_1) = \lambda(s'_2)$. Since $\delta(s'_1)$ is a distribution over S_2 , $\delta(s'_1)(s'_2) > 0$ for at least one state s'_2 . Hence $\rho'_0s'_2$ is similar to ρ , which shows in particular that $w \in FTr(\mathcal{A}_2)$. \square

We additionally show that probabilities coincide:

Proposition 4.4. *Let \mathcal{A}_1 and \mathcal{A}_2 be PTSs such that \mathcal{A}_2 simulates \mathcal{A}_1 . Then for all $w \in \Sigma^*$, $\mathbf{P}_{\mathcal{A}_1}(C_w) = \mathbf{P}_{\mathcal{A}_2}(C_w)$.*

Since a given word may be produced by several paths, their probabilities should be considered altogether. Hence the proof of the above proposition is not immediate; it is quite technical and can be found in Appendix C.

Existing properties about simulation for PTSs can be retrieved as consequences of the above result. They were for example obtained as a particular case of sub-stochastic simulation in [7]. Although not necessary to prove the main theorem, these results illustrate how constraining simulation between PTSs is.

Recall that a probabilistic bisimulation [8] is a bisimulation that preserves transition probabilities, *i.e.* a bisimulation relation \mathcal{R} on states such that for any equivalence class R of \mathcal{R} , and any two related states $s\mathcal{R}s'$, $\Delta(s)(R) = \Delta(s')(R)$.

Corollary 4.5 ([7]). *Let \mathcal{A}_1 and \mathcal{A}_2 be PTSs such that \mathcal{A}_2 simulates \mathcal{A}_1 . Then there exists a probabilistic bisimulation over the union of both PTSs.*

Corollary 4.6 ([7]). *Let \mathcal{A}_1 and \mathcal{A}_2 be PTSs such that \mathcal{A}_2 simulates \mathcal{A}_1 . Then \mathcal{A}_1 also simulates \mathcal{A}_2 .*

We are now ready to prove Theorem 4.1.

Proof. Let $\mathcal{A}_1 \in \text{sat}(\mathcal{S}_1)$. By Theorem 4.2 there exists $\mathcal{A}_2 \in \text{sat}(\mathcal{S}_2)$ that simulates \mathcal{A}_1 . By Proposition 4.4, $\mathbf{P}_{\mathcal{A}_1}(C_w) = \mathbf{P}_{\mathcal{A}_2}(C_w)$ for every word $w \in \text{FTr}(\mathcal{A}_1)$. Hence, for any ω -regular (hence measurable) language \mathcal{L} , one has $\mathbf{P}_{\mathcal{A}_1}(\mathcal{L}) = \mathbf{P}_{\mathcal{A}_2}(\mathcal{L})$. It is in particular the case for $\mathcal{V}(\mathcal{A}_1, \mathcal{O}, \varphi) = \mathcal{V}(\mathcal{A}_2, \mathcal{O}, \varphi)$. Therefore, $\underline{\text{Disc}}(\mathcal{A}_1, \mathcal{O}, \varphi) = \underline{\text{Disc}}(\mathcal{A}_2, \mathcal{O}, \varphi)$. Consequently, the theorem holds. \square

5 Conclusion

This work investigates how refinement of probabilistic models impacts the security – modeled as opacity. We provide a procedure to compute the worst-case opacity for a subclass of IDTMCs, and we show that opacity is monotonic with respect to simulation when implementations are obtained through scheduling.

Directions for future work include computing disclosure for IDTMCs with modal edges. In addition, while we considered here only the *worst case scenario*, it would be interesting to handle both the worst and the best case, thus providing bounds on the disclosure of all possible implementations.

References

1. Clarkson, M.R., Schneider, F.B.: Hyperproperties. *Journal of Computer Security* **18**(6) (September 2010) 1157–1210
2. Alur, R., Černý, P., Zdancewic, S.: Preserving secrecy under refinement. In: Proc. of the 33rd Intl. Colloquium on Automata, Languages and Programming (ICALP'06). Volume 4052 of LNCS., Springer (2006) 107–118
3. Mazaré, L.: Decidability of opacity with non-atomic keys. In: Proc. 2nd Workshop on Formal Aspects in Security and Trust (FAST'04). Volume 173 of Intl. Federation for Information Processing., Springer (2005) 71–84

4. Bryans, J.W., Koutny, M., Mazaré, L., Ryan, P.Y.A.: Opacity generalised to transition systems. *Intl. Jour. of Information Security* **7**(6) (2008) 421–435
5. Bérard, B., Mullins, J., Sassolas, M.: Quantifying opacity. *Mathematical Structures in Computer Science* **25**(2) (2015) 361–403
6. Bérard, B., Chatterjee, K., Sznajder, N.: Probabilistic opacity for Markov decision processes. *Inf. Process. Lett.* **115**(1) (2015) 52–59
7. Baier, C., Katoen, J.P., Hermanns, H., Wolf, V.: Comparative branching-time semantics for Markov chains. *Information and Computation* **200** (2005) 149–214
8. Jonsson, B., Larsen, K.G.: Specification and refinement of probabilistic processes. In: *LICS*, IEEE Computer Society (1991) 266–277
9. Segala, R.: Modeling and Verification of Randomized Distributed Real-Time Systems. PhD thesis, MIT, Department of Electrical Engineering and Computer Science (1995)
10. Chatterjee, K., Henzinger, T., Sen, K.: Model-checking omega-regular properties of interval markov chains. In Amadio, R.M., ed.: *Foundations of Software Science and Computation Structure (FoSSaCS) 2008*. (2008) 302–317
11. Benedikt, M., Lenhardt, R., Worrell, J.: LTL model checking of interval markov chains. In: *TACAS*. Volume 7795 of LNCS., Springer (2013) 32–46
12. Billingsley, P.: *Probability and Measure*. 3rd edn. Wiley, New York, NY (1995)
13. Vardi, M.Y.: Automatic verification of probabilistic concurrent finite-state programs. In: *Proceedings of 26th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society (1985) 327–338
14. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. *Journal of the ACM* **42**(4) (1995) 857–907
15. Bérard, B., Mullins, J., Sassolas, M.: Quantifying opacity. In Ciardo, G., Segala, R., eds.: *Proceedings of the 7th International Conference on Quantitative Evaluation of Systems (QEST'10)*, IEEE Computer Society (September 2010) 263–272
16. Chaum, D.: The dining cryptographers problem: unconditional sender and recipient untraceability. *Journal of Cryptology* **1** (1988) 65–75
17. Bhargava, M., Palamidessi, C.: Probabilistic anonymity. In Abadi, M., de Alfaro, L., eds.: *Proceedings of the 16th International Conference on Concurrency Theory (CONCUR'05)*. Volume 3653 of LNCS. (2005) 171–185
18. Piterman, N.: From nondeterministic Büchi and Streett automata to deterministic parity automata. *Logical Methods in Computer Science* **3**(3) (2007)

A Computing the probabilistic disclosure of a specification

A.1 Detecting modal edges.

The detection of modal edges is the first step toward computation of the disclosure of an IDTMC.

Proposition A.1. *The set of modal edges can be computed in time polynomial in the number of edges.*

Proof. The decision procedure for each edge is as follows:

- if an edge is not weighted by an interval containing 0, it is not modal;
- otherwise, compute the sum of maximal interval values of all other edges stemming from the same state;
 - if this sum is > 1 , the edge is modal;
 - if this sum is < 1 , the edge is not modal;
 - otherwise (the sum is $= 1$), the edge is modal if, and only if, all intervals of other outgoing edges are closed on the right. \square

Note that the procedure does not rely entirely on the syntactic criterion of an interval closed on 0: it is sufficient but may lead to false positives. For example, consider a state with two outgoing edges $e_1, [\frac{1}{4}; \frac{2}{3}]$ and $e_2, [0; 1]$. The e_2 edge is not actually modal since any probability distribution satisfying the specification can give at most $\frac{2}{3}$ to e_1 , hence must at least give weight $\frac{1}{3}$ to e_2 . This is avoided by the pre-computation of the least possible probability that can be put on an edge.

A.2 IDTMC without modal edges.

We now prove Theorem 3.4.

Note that intervals may be closed or open on any non-zero bound, which is not the case of IDTMCs in [10] where all intervals are closed. Hence although our procedure uses some of the ideas of Chatterjee *et al* [10], it is adapted to deal with open intervals.

First remark that there exists a regular language K such that for any scheduler A , $Tr(\mathcal{S}(A)) = K$. This is only true because \mathcal{S} is assumed non-modal. Let \mathcal{A}_K be a PTS such that $Tr(\mathcal{A}_K) = K$; it can be chosen of size $|\mathcal{S}|$. By the definition of disclosure, if the secret φ is ω -regular and the observation function \mathcal{O} is a projection (see Section 3), then finding the supremum of the disclosure means finding the maximal probability to reach an ω -regular set of runs, namely $\mathcal{V}(\mathcal{A}_K, \mathcal{O}, \varphi)$.

Then open intervals can be safely ignored when trying to optimize the probability of $\mathcal{V}(\mathcal{A}_K, \mathcal{O}, \varphi)$, *i.e.* they are treated as closed ones. Indeed, if the optimal scheduler uses a value x which is the bound of an open interval, then one can build a family of schedulers using value $x \pm \frac{1}{2^n}$ for the n th scheduler. The limit probability of reaching $\mathcal{V}(\mathcal{A}_K, \mathcal{O}, \varphi)$ is therefore the one computed when using exact value x . Remark that using closed intervals may introduce intervals containing 0, although it is of no concern since the observation classes are already defined and may not change, only their probability may change. Said otherwise, this does not mean that we are computing disclosure of the closed version, since it is only a probability. On the example of Fig. 4(b), it means trying to compute the maximal probability of the empty set, which is indeed zero.

The procedure is hence as follows. Starting from a DPA \mathcal{A}_φ for φ , a DPA $\mathcal{A}_\mathcal{V}$ for $\mathcal{V}(\mathcal{A}_K, \mathcal{O}, \varphi)$ can be built, with size exponential in the size of \mathcal{S} and \mathcal{A}_φ (and with a number k of colors polynomial in the size of \mathcal{A} and \mathcal{A}_φ). This construction relies on intersections and complementations of DPA, with a determinization step that brings the exponential blowup [18].

The construction of [10] yields a memoryless scheduler, although it is memoryless on the product, and hence is finite-memory on the original IDTMC. The procedure of [10] is in EXPTIME with respect to the size of its input, hence computation of disclosure is doubly exponential: $2^{2^{O(|\mathcal{A}| \times |\mathcal{A}_\varphi|)}}$.

B Proof of Theorem 4.2

Let $\mathcal{S}_1 = (S_1, s_{init,1}, T_1, \lambda_1)$ and $\mathcal{S}_2 = (S_2, s_{init,2}, T_2, \lambda_2)$ be interval-based specifications such that \mathcal{S}_2 simulates \mathcal{S}_1 with \mathcal{R} . Let sat_1 be the satisfaction relation related to A_1 and $\mathcal{A}_1 = \mathcal{S}_1(A_1) = (Q_1, q_{init,1}, \Delta_1, L_1)$. Then we show that there exists $A_2 \in Sched(\mathcal{S}_2)$ and a simulation relation \mathcal{R}' such that $sat_1 \circ \mathcal{R} = \mathcal{R}' \circ sat_2$ where sat_2 is the satisfaction relation related to A_2 . Then we define the scheduler A_2 by:

$$A_2(\rho_2)(s'_2) = \sum_{\rho_1 \in sim(\rho_2)} \mu_{\rho_2}(\rho_1) \sum_{s'_1 \in S_1} A_1(\rho_1)(s'_1) \cdot \delta(s'_1)(s'_2)$$

for any $s'_2 \in S_2$, writing $\mathcal{A}_2 = \mathcal{S}_2(A_2) = (Q_2, q_{init,2}, \Delta_2, L_2)$. Now the relation \mathcal{R}' can be defined as \sim that relates runs that are similar “step by step”, as defined above. To see that the conditions are satisfied, let ρ_1 and ρ_2 be runs in Q_1 and Q_2 respectively. Then the distribution δ' is obtained by:

$$\delta'(\rho_1)(\rho_2) = \begin{cases} \mu_{\rho_2}(\rho_1) \delta(\text{lst}(\rho_1))(\text{lst}(\rho_2)) & \text{if } \rho_1 \sim \rho_2, \\ 0 & \text{otherwise.} \end{cases}$$

Since \mathcal{A}_1 and \mathcal{A}_2 are PTSs, we just need to show that Equation (3) holds.

$$\begin{aligned} \Delta_2(\rho_2)(\rho'_2) &= A_2(\rho_2)(s'_2) \\ &= \sum_{\rho_1 \in sim(\rho_2)} \mu_{\rho_2}(\rho_1) \sum_{s'_1 \in S_1} A_1(\rho_1)(s'_1) \cdot \delta(s'_1)(s'_2) \\ &\quad \text{and since } \rho'_2 = \rho_2 \cdot s'_2 \\ &= \sum_{\rho_1 \in sim(\rho_2)} \sum_{s'_1 \in S_1} A_1(\rho_1)(s'_1) \cdot \delta'(\rho_1 \cdot s'_1)(\rho'_2) \\ &= \sum_{\rho'_1 \in Q_1} A_1(\rho_1)(s'_1) \cdot \delta'(\rho'_1)(\rho'_2) \\ &\quad \text{by defining } \rho'_1 = \rho_1 \cdot s'_1 \text{ and remarking that} \\ &\quad \delta' = 0 \text{ if its arguments are not similar runs} \\ \Delta_2(\rho_2)(\rho'_2) &= \sum_{\rho'_1 \in Q_1} \Delta_1(\rho_1)(\rho'_1) \cdot \delta'(\rho'_1)(\rho'_2). \end{aligned} \quad \square$$

C Proof of Proposition 4.4

Assume by induction that the proposition holds for every word of length n . Let $w \in FTr(\mathcal{A}_1) = FTr(\mathcal{A}_2)$ (recall Proposition 4.3) of length $n + 1$ with $w = w_0a$ for some $a \in \Sigma$. A run ρ' of \mathcal{A}_2 that produces w can be assumed to be of the form $\rho' = \rho'_0 s'_2$ with $\text{tr}(\rho'_0) = w_0$ and $\lambda(s'_2) = a$. Then $\mathbf{P}_{\mathcal{A}_2}(C_{\rho'}) = \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0})\Delta_2(s_2)(s'_2)$ where $s_2 = \text{lst}(\rho'_0)$ and hence

$$\begin{aligned} \mathbf{P}_{\mathcal{A}_2}(C_w) &= \sum_{\substack{\rho' \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho')=w}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'}) \\ &= \sum_{s_2, s'_2 \in S_2} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0)=w_0, \\ \text{lst}(\rho'_0)=s_2, \text{lst}(\rho')=s'_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0})\Delta_2(s_2)(s'_2) \end{aligned}$$

Now let \mathcal{A}_1 s.t. \mathcal{A}_2 simulates \mathcal{A}_1 then, as

$$\sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0)=s_1}} \mu_{\rho'_0}(\rho_0) = \sum_{s_1 \in S_1} \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0)=s_1}} \mu_{\rho'_0}(\rho_0) = 1$$

we get that

$$\begin{aligned} \mathbf{P}_{\mathcal{A}_2}(C_w) &= \sum_{s_2, s'_2 \in S_2} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0)=w_0, \\ \text{lst}(\rho'_0)=s_2, \text{lst}(\rho')=s'_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0})\Delta_2(s_2)(s'_2) \cdot \\ &\quad \sum_{s_1 \in S_1} \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0)=s_1}} \mu_{\rho'_0}(\rho_0) \end{aligned}$$

$$\mathbf{P}_{\mathcal{A}_2}(C_w) = \sum_{\substack{s_1 \in S_1 \\ s_2, s'_2 \in S_2}} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0, \\ \text{lst}(\rho'_0) = s_2, \text{lst}(\rho') = s'_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \Delta_2(s_2)(s'_2) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mu_{\rho'_0}(\rho_0)$$

As the terms are null if it is not the case that $s_1 \mathcal{R} s_2$, we have:

$$= \sum_{\substack{s_1 \in S_1 \\ s_2, s'_2 \in S_2}} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0, \\ \text{lst}(\rho'_0) = s_2, \text{lst}(\rho') = s'_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \sum_{s'_1 \in S_1} \Delta_1(s_1)(s'_1) \cdot \delta_{s_1, s_2}(s'_1)(s'_2) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mu_{\rho'_0}(\rho_0)$$

And as the terms are null if $\lambda(s'_1) \neq \lambda(s'_2) = \lambda(\text{lst}(\rho))$, we have:

$$= \sum_{\substack{s_1 \in S_1 \\ s_2, s'_2 \in S_2}} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0, \\ \text{lst}(\rho'_0) = s_2, \text{lst}(\rho') = s'_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \cdot \delta_{s_1, s_2}(s'_1)(s'_2) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mu_{\rho'_0}(\rho_0)$$

$$= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \sum_{\substack{s_2 \in S_2 \\ s'_2 \in S_2}} \delta_{s_1, s_2}(s'_1)(s'_2) \cdot$$

$$\left(\begin{array}{cc} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0, \\ \text{lst}(\rho'_0) = s_2, \text{lst}(\rho') = s'_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \Delta_1(s_1)(s'_1) & \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mu_{\rho'_0}(\rho_0) \end{array} \right)$$

And since $\sum_{s'_2 \in S_2} \delta_{s_1, s_2}(s'_1)(s'_2) = 1$, we have:

$$= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \sum_{s_2 \in S_2} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0, \\ \text{lst}(\rho'_0) = s_2}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \Delta_1(s_1)(s'_1) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mu_{\rho'_0}(\rho_0)$$

$$\mathbf{P}_{\mathcal{A}_2}(C_w) = \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \Delta_1(s_1)(s'_1) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mu_{\rho'_0}(\rho_0)$$

$$\begin{aligned}
\mathbf{P}_{\mathcal{A}_2}(C_w) &= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0}} \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \mu_{\rho'_0}(\rho_0) \\
&= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0}} \\
&\quad \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \rho_0 \sim \rho'_0 \\ \text{lst}(\rho_0) = s_1}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0}) \frac{\mathbf{P}_{\mathcal{A}_1}(C_{\rho_0})}{\mathbf{P}_{\mathcal{A}_1}(\text{sim}(\rho'_0))} \\
&= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \text{tr}(\rho_0) = w_0 \\ \text{lst}(\rho_0) = s_1}} \frac{\mathbf{P}_{\mathcal{A}_1}(C_{\rho_0})}{\mathbf{P}_{\mathcal{A}_1}(\text{sim}(\rho'_0))} \cdot \\
&\quad \sum_{\substack{\rho'_0 \in FRuns(\mathcal{A}_2) \\ \text{tr}(\rho'_0) = w_0}} \mathbf{P}_{\mathcal{A}_2}(C_{\rho'_0})
\end{aligned}$$

$$= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \text{tr}(\rho_0) = w_0 \\ \text{lst}(\rho_0) = s_1}} \frac{\mathbf{P}_{\mathcal{A}_1}(C_{\rho_0})}{\mathbf{P}_{\mathcal{A}_1}(\text{sim}(\rho'_0))} \cdot \mathbf{P}_{\mathcal{A}_2}(C_{w_0})$$

and by induction hypothesis, $\mathbf{P}_{\mathcal{A}_2}(C_{w_0}) = \mathbf{P}_{\mathcal{A}_1}(C_{w_0})$:

$$= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \text{tr}(\rho_0) = w_0 \\ \text{lst}(\rho_0) = s_1}} \frac{\mathbf{P}_{\mathcal{A}_1}(C_{\rho_0})}{\mathbf{P}_{\mathcal{A}_1}(\text{sim}(\rho'_0))} \cdot \mathbf{P}_{\mathcal{A}_1}(C_{w_0})$$

and since $\mathbf{P}_{\mathcal{A}_1}(\text{sim}(\rho'_0)) = \mathbf{P}_{\mathcal{A}_1}(C_{w_0})$:

$$\begin{aligned}
&= \sum_{s_1 \in S_1} \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \text{tr}(\rho_0) = w_0 \\ \text{lst}(\rho_0) = s_1}} \mathbf{P}_{\mathcal{A}_1}(C_{\rho_0}) \\
&= \sum_{s_1 \in S_1} \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \text{tr}(\rho_0) = w_0 \\ \text{lst}(\rho_0) = s_1}} \mathbf{P}_{\mathcal{A}_1}(C_{\rho_0}) \sum_{\substack{s'_1 \in S_1 \\ \lambda(s_1) = \lambda(\text{lst}(\rho))}} \Delta_1(s_1)(s'_1) \\
&= \sum_{\substack{\rho_0 \in FRuns(\mathcal{A}_1) \\ \text{tr}(\rho_0) = w_0}} \mathbf{P}_{\mathcal{A}_1}(C_{\rho_0}) \Delta_1(\text{lst}(\rho_0))(\text{lst}(\rho))
\end{aligned}$$

$$\mathbf{P}_{\mathcal{A}_2}(C_w) = \mathbf{P}_{\mathcal{A}_1}(C_w)$$

□