COMPUTING SCIENCE

Towards Quantitative Analysis of Opacity

Jeremy Bryans, Maciej Koutny and Chunyan Mu

TECHNICAL REPORT SERIES

No. CS-TR-1304 November 2011

TECHNICAL REPORT SERIES

No. CS-TR-1304 November, 2011

Towards Quantitative Analysis of Opacity

J. Bryans, M. Koutny and C. Mu

Abstract

Opacity is a general approach for describing and unifying security properties expressed as predicates. A predicate is opaque if an observer of the system is unable to determine the satisfaction of the predicate in a given run of the system. The meaning of opacity is straightforward when considering the standard (qualitative) operational semantics, but there are a number of possible interpretations in a context where quantitative information about system evolutions is available. We propose four variants of quantitative opacity defined for probabilistic labelled transition systems, with each variant capturing a different aspect of quantifying the opacity of a predicate. Moreover, we present results showing how these four properties can be checked or approximated for specific classes of probabilistic labelled transition systems, observation functions, and system predicates.

© 2011 Newcastle University.
Printed and published by Newcastle University,
Computing Science, Claremont Tower, Claremont Road,
Newcastle upon Tyne, NE1 7RU, England.

Bibliographical details

BRYANS, J., KOUTNY, M., MU, C.

Towards Quantitative Analysis of Opacity
[By] J. Bryans, M. Koutny, C. Mu
Newcastle upon Tyne: Newcastle University: Computing Science, 2011.

(Newcastle University, Computing Science, Technical Report Series, No. CS-TR-1304)

Added entries

NEWCASTLE UNIVERSITY
Computing Science. Technical Report Series. CS-TR-1304

Abstract

Opacity is a general approach for describing and unifying security properties expressed as predicates. A predicate is opaque if an observer of the system is unable to determine the satisfaction of the predicate in a given run of the system. The meaning of opacity is straightforward when considering the standard (qualitative) operational semantics, but there are a number of possible interpretations in a context where quantitative information about system evolutions is available. We propose four variants of quantitative opacity defined for probabilistic labelled transition systems, with each variant capturing a different aspect of quantifying the opacity of a predicate. Moreover, we present results showing how these four properties can be checked or approximated for specific classes of probabilistic labelled transition systems, observation functions, and system predicates.

About the authors

Jeremy received his BSc in Mathematics and Computer Science from Reading University in 1993, and his PhD in 1997, also from Reading University. He has worked in a number of university departments, including Royal Holloway, Kent and Stirling, and has been at Newcastle since December 2002. His research is in the security of information within large computer-based systems. A particular area of current interest is access control the development and maintenance of access control policies within dynamic coalitions. In the past at Newcastle he has worked on including DIRC (the Interdisciplinary Research Collaboration on Dependability) and GOLD (Grid Oriented Lifecycle Development) He is currently employed on the User Friendly Grid Security project and TrAmS (Trustworthy Ambient Systems). He is part of the RESIST network, and a member of RESIST's working group on Verification.

Maciej Koutny obtained his MSc (1982) and PhD (1984) from the Warsaw University of Technology. In 1985 he joined the then Computing Laboratory of the University of Newcastle upon Tyne to work as a Research Associate. In 1986 he became a Lecturer in Computing Science at Newcastle, and in 1994 was promoted to an established Readership at Newcastle. In 2000 he became a Professor of Computing Science.

Chunyan Mu received her PhD from King's College London, and is currently a temporary lecture of Computing Science at Newcastle University. Her research interests include: language-based security, programming languages, information flow security.

Suggested keywords

PROBABILISTIC OPACITY
PROBABILISTIC LABELLED TRANSITION SYSTEMS
OBSERVATIONS

Towards Quantitative Analysis of Opacity

Jeremy Bryans, Maciej Koutny and Chunyan Mu

School of Computing Science, Newcastle University,
Newcastle upon Tyne, NE1 7RU, U.K.

Abstract. Opacity is a general approach for describing and unifying security properties expressed as predicates. A predicate is opaque if an observer of the system is unable to determine the satisfaction of the predicate in a given run of the system. The meaning of opacity is straightforward when considering the standard (qualitative) operational semantics, but there are a number of possible interpretations in a context where quantitative information about system evolutions is available. We propose four variants of quantitative opacity defined for probabilistic labelled transition systems, with each variant capturing a different aspect of quantifying the opacity of a predicate. Moreover, we present results showing how these four properties can be checked or approximated for specific classes of probabilistic labelled transition systems, observation functions, and system predicates.

keywords: Probabilistic opacity, Probabilistic labelled transition systems, Observations

1 Introduction

Opacity has been shown to be a promising technique for describing and unifying security properties [6]. For a given observer of a system (or adversary), a predicate capturing a system property is opaque if the observer will never be able to determine the truth of that predicate.

The definition of [6] is based on a qualitative operational semantics. In it, observation functions are used in order to give fine-grained control over the capabilities of an observer. Through such observations, an observer can establish certain properties of the system. Informally, an observer cannot establish the predicate (and hence the predicate is opaque) if for any run of the system in which the predicate is true, there is a run for which the predicate is false, and the two runs are observationally equivalent under the defined observation function. However, in the case where the probability of the first run is significantly higher than the probability of the second, the observer (although not able to be certain) may have good reason to believe that the predicate (although opaque) is none the less true. This paper presents the results of our initial investigations into this probabilistic case.

The contribution of this paper is to show how the work referenced above extends to the more general case when the information given about the system

is qualitative. We therefore consider the general theory of probabilistic opacity in the context of *probabilistic labelled transition systems* which allows us to reason about the quantitative properties of systems. Based on the probabilistic model of opacity, we introduce four alternative definitions of probabilistic opacity, and investigate the efficiency with which they can be verified or approximated. We relate the definitions to the existing work on qualitative opacity. The obtained results can be used in a quantified information flow analysis of a system.

This paper is organised as follows. In Section 2 we recall some definitions from the literature in particular relating to probability distributions, and in Section 3 we give the definition of probabilistic labelled transition systems and prove a property which is then needed to estimate the efficiency of our approximations of probabilistic opacity. Section 4 contains our main contribution, i.e., the definitions of four variants of opacity together with an investigation of their basic properties. Section 5 contains a brief comparison with other work, and in Section 6 we present our concluding remarks.

2 Preliminaries

We use the standard mathematical notation. In particular, ϵ denotes the empty sequence, $|\lambda|$ denotes the length of a finite sequence λ, and λ^k denotes the concatenation of k copies of λ.

A *probability distribution* on a countable set X is a function $f : X \to [0,1]$ such that $\sum_{x\in X} f(x) = 1$. To measure difference between probability distributions on the same set, we will use Jensen-Shannon divergence [11] which is related to information-theoretical functionals, such as Kullback-Leibler distance (the relative entropy). It therefore shares some of their properties as well as their intuitive interpretation, and measures the difference in information bits. Unlike the Kullback-Leibler distance, it is symmetric, always well-defined and bounded by 1.

Let $P = \{p_x\}_{x\in X}$ and $P' = \{p'_x\}_{x\in X}$ be two probability distributions on a countable set X with associated weights w and w', respectively ($0 \leq w, w' \leq 1$ and $w + w' = 1$). Then the weighted Jensen-Shannon divergence between P and P' is given by:

$$D_{JS}(w{\cdot}P, w'{\cdot}P') = \mathcal{H}\left(\{w \cdot p_x + w' \cdot p'_x\}_{x\in X}\right) - w{\cdot}\mathcal{H}\left(\{p_x\}_{x\in X}\right) - w'{\cdot}\mathcal{H}\left(\{p'_x\}_{x\in X}\right)$$

where $\mathcal{H}(\{q_x\}_{x\in X}) = -\sum_{x\in X} q_x \log_2 q_x$ denotes Shannon entropy [16] (note that if $q_x = 0$ then $q_x \log_2 q_x$ is taken to be 0 which is justified by $\lim_{q\to 0^+} q \log_2 q = 0$).

If, in the above formula, we denote by d_x the 'contribution' made by a single element $x \in X$, then:

$$D_{JS}(w \cdot P, w' \cdot P') = \sum_{x\in X} d_x \ ,$$

where:

$$d_x = -(w{\cdot}p_x + w'{\cdot}p'_x){\cdot}\log_2(w{\cdot}p_x + w'{\cdot}p'_x) + w{\cdot}p_x{\cdot}\log_2 p_x + w'{\cdot}p'_x{\cdot}\log_2 p'_x \ . \quad (1)$$

An individual contribution is minimal ($d_x = 0$) if $p_x = p'_x$, i.e., when P and P' do not diverge at x. It is maximal if one of the probabilities is 0, which gives $d_x = -w \cdot p_x \cdot \log_2 w$ or $d_x = -w' \cdot p'_x \cdot \log_2 w'$, and so:

$$d_x \leq -w \cdot p_x \cdot \log_2 w - w' \cdot p'_x \cdot \log_2 w' \leq c \cdot (p_x + p'_x) \ ,$$

where $c > 0$ is a constant depending on w and w'. As a consequence, if we take $Y \subset X$ then the contribution d_Y of the elements of Y to the overall divergence satisfies:

$$d_Y \leq c \cdot (P(Y) + P'(Y)) \ . \tag{2}$$

3 Probabilistic Labelled Transition Systems

In order to consider probabilistic behaviour and quantitative analysis of opacity, we use probabilistic labelled transition systems which adapts the well-known model introduced in [10].

A *labelled transition system* is a tuple:

$$LTS = (S, L, \Delta, s_0) \ ,$$

where S is a countable set of states, L is a countable set of labels, $\Delta \subseteq S \times L \times S$ is the transition set, and $s_0 \in S$ is the initial state. We consider deterministic labelled transition systems,[1] and so for any transitions $(s, l, s'), (s, l, s'') \in \Delta$, it is the case that $s' = s''$. For every state $s \in S$, we will denote by Δ_s the set of all transitions outgoing from s, i.e., $\Delta_s = \{(s', l, s'') \in \Delta \mid s = s'\}$.

A *run* of *LTS* is a finite sequence of labels $\lambda = l_1 \dots l_n$ $(n \geq 0)$[2] such that there are states $s_1, \dots, s_n$ satisfying (s_{i-1}, l_i, s_i), for $i = 1, \dots, n$. We will denote the state s_n by s_λ and call it *reachable*. Note that s_λ is well-defined since *LTS* is deterministic. Moreover, $s_\epsilon = s_0$, where ϵ denotes the empty run. The set of all runs of *LTS* will be denoted by $runs(LTS)$.

Probabilistic labelled transition systems are labelled transition systems with probability distributions attached to all the states.

Definition 1. *A* probabilistic labelled transition system *is a tuple:*

$$PLTS = (S, L, \Delta, s_0, \mu) \ ,$$

such that $LTS = (S, L, \Delta, s_0)$ *is a labelled transition system and* $\mu : S \cup \Delta \to [0, 1]$ *is a mapping satisfying the following:*

(i) for every $s \in S$, μ *restricted to* $\{s\} \cup \Delta_s$ *is a probability distribution:*

$$\mu(s) + \sum_{d \in \Delta_s} \mu(d) = 1 \ ,$$

and $\inf\{\mu(s) \mid s \in S \wedge \mu(s) \neq 0\} > 0$.

[1] Nondeterminism is introduced in the next section, through the notion of an observation function.

[2] If $n = 0$ then $\lambda = \epsilon$.

(ii) there is an integer $k \geq 1$ such that there is no sequence of transitions in Δ:

$$(s, l_1, s_2), (s_2, l_2, s_3), \ldots, (s_m, l_m, s_{m+1})$$

such that $\mu(s_2) = \mu(s_3) = \ldots = \mu(s_m) = 0$ and $m > k$.

The set of runs of PLTS, denoted runs(PLTS), is the same as that of the underlying labelled transition system. Other notations are also inherited.

Definition 1(i) ensures that for every state s, the probability $\mu(s)$ of remaining in that state together with the probabilities associated with all transitions out of that state form a probability distribution. We also require that non-empty probabilities $\mu(s)$ cannot be arbitrarily small (similarly as in [10] it was assumed that non-empty probabilities $\mu(s, l, s')$ cannot be arbitrarily small). Note that this is always the case if there are finitely many states.

We extend the mapping μ to each run $\lambda = l_1 \ldots l_k$ of *PLTS*, in the following way. Let $s_1, \ldots, s_k$ be states such that $(s_{i-1}, l_i, s_i) \in \Delta$, for $i = 1, \ldots, k$. Then:

$$\mu(\lambda) = \mu(s_k) \cdot \prod_{i=1}^{k} \mu(s_{i-1}, l_i, s_i) \; .$$

Note that $\mu(\lambda)$ is well-defined as the underlying labelled transition system is deterministic. We also denote:

$$\widetilde{\mu}(\lambda) = \prod_{i=1}^{k} \mu(s_{i-1}, l_i, s_i)$$

(i.e., $\mu(\lambda) = \mu(s_k) \cdot \widetilde{\mu}(\lambda)$) and, for every set of runs $\Lambda \subseteq \mathit{runs}(\mathit{PLTS})$:

$$\mu(\Lambda) = \sum_{\lambda \in \Lambda} \mu(\lambda) \quad \text{and} \quad \widetilde{\mu}(\Lambda) = \sum_{\lambda \in \Lambda} \widetilde{\mu}(\lambda) \; .$$

Proposition 1. *$\mu(\mathit{runs}(\mathit{PLTS})) \leq 1$.*

Proof. Follows from the fact that, for every $i \geq 0$:

$$\mu(\{\lambda \in \mathit{runs}(\mathit{PLTS}) \mid |\lambda| \leq i\}) + \widetilde{\mu}(\{\lambda \in \mathit{runs}(\mathit{PLTS}) \mid |\lambda| = i + 1\}) = 1 \; .$$

The above can be shown by a straightforward induction on i, using Definition 1(i). □

The formalisation of a probabilistic labelled transition system is tailored to reflect our understanding of observation of a computing system. In a nutshell, we treat $\prod_{i=1}^{k} \mu(s_{i-1}, l_i, s_i)$ in the standard way as the probability of executing a sequence of transitions making up the run λ. In addition to that, the factor $\mu(s_k)$ gives the probability that the observation is concluded after the state s_k has been reached. For instance, it may be the probability that the process terminates after reaching s_k, similarly as it was done in [2]. It therefore follows that Definition 1(ii) captures our intuition that the system cannot be indefinitely 'unobserved' (i.e., probability of a conclusive observation cannot be zero forever).

Example 1. Consider the following probabilistic labelled transition system:

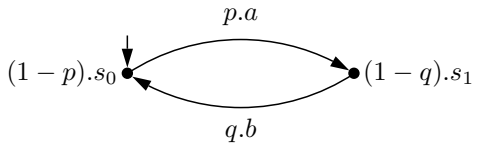

where $0 \leq p, q \leq 1$ and the notation $x.y$ indicates that $x = \mu(y)$. According to Definition 1(ii), we must have $p \cdot q \neq 1$. We can then show that μ defines a probability distribution, as follows:

$$\begin{aligned}
\mu(runs(PLTS)) &= \sum_{k=0}^{\infty} \mu((ab)^k) + \sum_{k=0}^{\infty} \mu(a(ba)^k) \\
&= \sum_{k=0}^{\infty} (p \cdot q)^k \cdot (1-p) + \sum_{k=0}^{\infty} p \cdot (q \cdot p)^k \cdot (1-q) \\
&= (1-p) \cdot \sum_{k=0}^{\infty} (p \cdot q)^k + p \cdot (1-q) \cdot \sum_{k=0}^{\infty} (q \cdot p)^k \\
&= (1-p) \cdot \frac{1}{1-p \cdot q} + p \cdot (1-q) \cdot \frac{1}{1-p \cdot q} \\
&= 1 \ .
\end{aligned}$$

Note that if in the above example we assumed that $p = q = 1$, and hence $\mu(s_0) = \mu(s_1) = 0$, then we would have $\mu(runs(PLTS)) = 0$, and so μ would not be a probability distribution on the runs of $PLTS$. To avoid this, we introduced condition (ii) in Definition 1 which rules out this case. Note also that the condition captured through Definition 1(ii) is easy to verify by checking that in the graph of $PLTS$ there are no cycles passing only through reachable states s satisfying $\mu(s) = 0$.

Since the set of runs is in general infinite, we will be approximating various quantities defined on the basis of the set of runs, by looking only at runs up to certain length. We therefore define, for every $m \geq 0$:

$$runs_m(PLTS) = \{\lambda \in runs(PLTS) \mid |\lambda| \leq m\} \ .$$

The next result is crucial for the soundness of such approximations.

Proposition 2. *There is an integer $\kappa \geq 1$ and a real number $0 \leq \delta < 1$ such that, for every $i \geq 0$:*

$$\mu(runs_{\kappa \cdot i}(PLTS)) \geq 1 - \delta^i \ .$$

Proof. In what follows, for every state $s \in S$, we denote by $PLTS_s$ the probabilistic labelled transition system obtained from $PLTS$ by setting the initial state to s.

In the first part of the proof, we assume that *PLTS* satisfies the following two properties:

(i) $\mu(s) > 0$, for all $s \in S \setminus \{s_0\}$.
(ii) If $\mu(s_0) = 0$ then there is no transition $(s, l, s') \in \Delta$ such that $s' = s_0$.

We also define:

$$\delta = \begin{cases} \sup\{1 - \mu(s) \mid s \in S\} & \text{if } \mu(s_0) > 0 \\ \sup\{1 - \mu(s) \mid s \in S \setminus \{s_0\}\} & \text{otherwise .} \end{cases}$$

Note that $0 \leq \delta < 1$ is well-defined by Definition 1(i). Proceeding by induction on $i \geq 0$, will now show that, for every $i \geq 0$ and every $s \in S$:

$$\mu(runs_i(PLTS_s)) \geq \begin{cases} 1 - \delta^i & \text{if } s = s_0 \text{ and } \mu(s_0) = 0 \\ 1 - \delta^{i+1} & \text{otherwise .} \end{cases} \tag{3}$$

In the base case:

$$\mu(runs_0(PLTS_s)) = \mu(\epsilon) = \mu(s) \geq \begin{cases} 1 - 1 = 1 - \delta^0 & \text{if } s = s_0 \text{ and } \mu(s_0) = 0 \\ 1 - \delta = 1 - \delta^1 & \text{otherwise .} \end{cases}$$

In the induction step, we assume that the thesis holds for i, and proceed as follows:

$$\begin{aligned} \mu(runs_{i+1}(PLTS_s)) &= \mu(s) \; + \sum_{(s,l_j,s_j) \in \Delta_s} \mu(s, l_j, s_j) \cdot \mu(runs_i(PLTS_{s_j})) \\ &= 1 - \mu(\Delta_s) \; + \sum_{(s,l_j,s_j) \in \Delta_s} \mu(s, l_j, s_j) \cdot \mu(runs_i(PLTS_{s_j})) \; . \end{aligned}$$

By the induction hypothesis, we obtain the following (note that $s_j \neq s_0$, for every $(s, l_j, s_j) \in \Delta_s$):

$$\mu(runs_{i+1}(PLTS_s)) \geq 1 - \mu(\Delta_s) \; + \sum_{(s,l_j,s_j) \in \Delta_s} \mu(s, l_j, s_j) \cdot (1 - \delta^{i+1}) = 1 - \delta^{i+1} \cdot \mu(\Delta_s) \; .$$

Now, if $s = s_0$ and $\mu(s_0) = 0$, then $\mu(\Delta_s) = 1$ and we get

$$\mu(runs_{i+1}(PLTS_s)) \geq 1 - \delta^{i+1} \; ;$$

otherwise, $\delta \geq \mu(\Delta_s)$ and we obtain:

$$\mu(runs_{i+1}(PLTS_s)) \geq 1 - \delta^{i+2} \; .$$

Hence (3) holds.

In the second part of the proof, we transform *PLTS* into a probabilistic labelled transition system $PLTS'$ satisfying (i) and (ii) above, in the following way:

– For every $s \in S$, we set $\mu'(s) = \mu(s)$.
– We create a fresh initial state s'_0 with $\mu'(s'_0) = \mu(s_0)$ and, for every sequence of transitions of $PLTS$ of the form:

$$(s_0, l_1, s_1), (s_1, l_2, s_2) \dots (s_{k-1}, l_k, s_k)$$

such that $0 = \mu(s_1) = \ldots = \mu(s_{k-1}) \neq \mu(s_k)$, we introduce a transition $(s'_0, l_1 l_2 \dots l_k, s_k)$ and set:

$$\mu'(s'_0, l_1 l_2 \dots l_k, s_k) = \mu(s_0, l_1, s_1) \cdot \mu(s_1, l_2, s_2) \cdot \ldots \cdot \mu(s_{k-1}, l_k, s_k) \ .$$

– For every sequence of transitions of $PLTS$ of the form:

$$(s_1, l_1, s_2), (s_2, l_2, s_3) \dots (s_k, l_k, s_{k+1})$$

such that $\mu(s_1) \neq 0 = \mu(s_2) = \ldots = \mu(s_k) \neq \mu(s_{k+1})$, we introduce a transition $(s_1, l_1 l_2 \dots l_k, s_{k+1})$ and set:

$$\mu'(s_1, l_1 l_2 \dots l_k, s_{k+1}) = \mu(s_1, l_1, s_2) \cdot \mu(s_2, l_2, s_3) \cdot \ldots \cdot \mu(s_k, l_k, s_{k+1}) \ .$$

Note that by Definition 1(ii), there is the largest k as above, denoted by k_{max}. We then delete all the states $s \in S$ with $\mu(s) = 0$ together with the adjacent arcs. Thanks to Definition 1(ii), $PLTS'$ is a well-defined probabilistic labelled transition system whose labels are finite sequences of labels from $PLTS$,

$$runs(PLTS') = \{\lambda \in runs(PLTS) \mid \mu(\lambda) > 0\} \ ,$$

and $\mu(\lambda) = \mu'(\lambda)$, for all $\lambda \in runs(PLTS')$. Moreover, we can apply (3) to $PLTS'$ and conclude that, for every $i \geq$:

$$\mu(runs_i(PLTS')) \geq 1 - \delta^i \ .$$

Thus, by:

$$runs_i(PLTS') \subseteq \{\lambda \in runs_{(k_{max}+1)\cdot i}(PLTS) \mid \mu(\lambda) > 0\} \ ,$$

we have that $\mu(runs_{\kappa \cdot i}(PLTS)) \geq 1 - \delta^i$ for $\kappa = k_{max} + 1$. □

In other words, we know how far to 'unfold' the transition system to approximate with arbitrary accuracy 'almost all' the runs (in probabilistic terms).

As a corollary of our previous results, μ always defines a probability distribution for the set of runs of the probabilistic labelled transition system.

Theorem 1. $\mu(runs(PLTS)) = 1$.

Proof. Follows directly from Propositions 1 and 2. □

4 Probabilistic Opacity

In this section, we introduce concepts relating to the definitions of probabilistic opacity, and prove our main results.

In what follows, $PLTS = (S, L, \Delta, S_0, \mu)$ is a probabilistic labelled transition system, and Obs is a set of elements called *observables*. To model the different capabilities for observing the system modelled by $PLTS$, we use an *observation function*:

$$obs : runs(PLTS) \to Obs^* \ .$$

We will, in particular, use the *static* observation function obs for which there is a map $obs' : L \to Obs \cup \{\epsilon\}$ such that, for every run $\lambda = l_1 \dots l_n$ of $PLTS$:

$$obs(\lambda) = obs'(l_1)obs'(l_2)\dots obs'(l_n) \ .$$

Consider now an observation function obs. We are interested in whether an observer (or attacker) can establish a *property* ϕ (a predicate over system runs) for a run of $PLTS$ having only access to the result of the observation function. We will identify ϕ with its characteristic set, i.e., the set of all those runs for which it holds. Now, given an observed execution of the system, we would want to find out whether the fact that the underlying run belongs to ϕ can be deduced by the observer. We will, in particular, be interested in the *final* opacity predicates, ϕ_Z, where $Z \subseteq S$, defined as the set of all the runs λ of $PLTS$ satisfying $s_\lambda \in Z$. Intuitively, this means that we are interested in finding out whether an observed run of the system represented by $PLTS$ ended in one of secret (or sensitive) states belonging to Z. (Note that we are not interested in establishing whether the underlying run does not belong to ϕ; to do this, we would consider the property $runs(PLTS) \setminus \phi$.)

We will now introduce a series of notions relating to the idea of opacity, recalling first its standard non-probabilistic (or qualitative) definition. In what follows, obs is an observation function for $PLTS$, ϕ is a subset of $runs(PLTS)$, and $\overline{\phi} = runs(PLTS) \setminus \phi$. Intuitively, ϕ captures a property which we want to declare opaque.

4.1 Qualitative (non-probabilistic) opacity of [6]

When no probabilistic information is supplied, or if one is simply not interested in probabilistic aspects of the system, we say that ϕ is *opaque* w.r.t. obs if, for every run $\lambda \in \phi$, there is another run $\lambda' \notin \phi$ such that $obs(\lambda) = obs(\lambda')$, i.e., λ' *covers* λ. In other words, all runs in ϕ are covered by runs in $\overline{\phi}$:

$$obs(\phi) \subseteq obs(\overline{\phi}) \ . \tag{4}$$

Different variants of qualitative opacity have been discussed in, for example, [6].

4.2 Quantitative (probabilistic) opacity

What it means to deduce (or satisfactorily cover) a property expressed as ϕ in the probabilistic case may mean different things, depending on what is relevant or important from the point of view of a real application. In particular, one may consider different ways of quantifying the degree to which runs contained in ϕ are covered by the runs in $\overline{\phi}$ (c.f. the inclusion (4)), leading to different variants of quantitative opacity.

π-opacity. A straightforward approach to defining probabilistic opacity might be to require that the likelihood of ever witnessing an uncovered run of ϕ is negligible. That is, we say that ϕ is *π-opaque* w.r.t. *obs* if the probability of having a run in ϕ which is not covered by a run in $\overline{\phi}$ is zero:

$$\mu(\phi \setminus obs^{-1}(obs(\overline{\phi}))) = 0\ . \tag{5}$$

Example 2. Consider the following probabilistic labelled transition system:

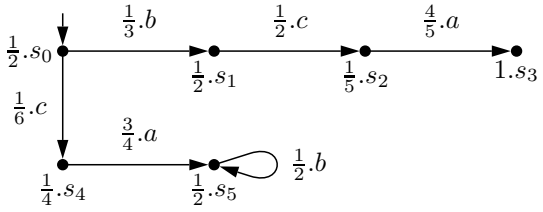

where $obs(a) = a$, $obs(b) = \epsilon$ and $obs(c) = \epsilon$, as well as the property $\phi = \phi_{\{s_2,s_3\}}$. We then have:

$$\phi = \{bc, bca\} \ \text{ and } \ \overline{\phi} = \{\epsilon, b, c, ca, cab, cabb, \ldots\}\ .$$

In this case, we have $obs(runs(PLTS)) = obs(\phi) = obs(\overline{\phi}) = \{\epsilon, a\}$ and so, obviously, π-opacity is satisfied:

$$\mu(\phi \setminus obs^{-1}(obs(\overline{\phi}))) = \mu(\varnothing) = 0.$$

Checking π-opacity may be straightforward, as shown by the next result and its proof.

Theorem 2. *For a finite PLTS and a static observation function obs, it is decidable whether ϕ_Z (where $Z \subset S$) is π-opaque.*

Proof. We first observe that in this case $\mu(\phi_Z \setminus obs^{-1}(obs(\overline{\phi}_Z))) = 0$ is equivalent to $obs(L) \subseteq obs(L')$, where L is the regular language obtained from *PLTS* by changing each label l to $obs(l)$ and setting as the final states all those $s \in Z$ for which $\mu(s) > 0$; and L' is the regular language obtained from *PLTS* by changing each label l to $obs(l)$ and setting $S \setminus Z$ as the final states. Since the inclusion of two regular languages is decidable, the result follows. □

π_ξ-opacity. One could argue that the probabilistic opacity captured by (5) is too demanding, and one might require only that the probability of witnessing an uncovered run of ϕ is low. To capture this, we say that ϕ is π_ξ-*opaque* w.r.t. *obs* if $0 \leq \xi \leq 1$ is the probability of having a run in ϕ which is not covered:

$$\mu(\phi \setminus obs^{-1}(obs(\overline{\phi}))) = \xi \ . \tag{6}$$

One would then declare ϕ opaque if ξ was sufficiently small number. Note that π_0-opacity coincides with π-opacity.

In practice, knowing the value of ξ with high accuracy (see Theorem 3) would allow a designer or verifier to compare it with a given required opacity level, ξ_{req}. The system represented by *PLTS* would then satisfy the opacity w.r.t. ϕ if $\xi \leq \xi_{req}$. Similar comment applies to other opacity notions introduced in the rest of this paper.

Example 3. Consider the following probabilistic labelled transition system:

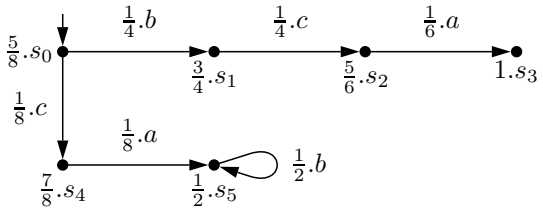

where $obs(a) = a$, $obs(b) = \epsilon$ and $obs(c) = c$, as well as the property $\phi = \phi_{\{s_2,s_3,s_5\}}$. We then have:

$$\phi = \{bc, bca, ca, cab, cabb, \ldots\} \ \text{ and } \ \overline{\phi} = \{\epsilon, b, c\} \ .$$

Hence:

$$\phi \setminus obs^{-1}(obs(\overline{\phi})) = \{bca, ca, cab, cabb, \ldots\} \ .$$

and we obtain:

$$\begin{aligned} \mu(\phi \setminus obs^{-1}(obs(\overline{\phi}))) &= \frac{1}{4} \cdot \frac{1}{4} \cdot \frac{1}{6} + \frac{1}{8} \cdot \frac{1}{8} \cdot \frac{1}{2} + \frac{1}{8} \cdot \frac{1}{8} \cdot \frac{1}{2} \cdot \frac{1}{2} + \ldots \\ &= \frac{5}{192} \approx 0.026. \end{aligned}$$

The property $\phi_{\{s_2,s_3,s_5\}}$ is therefore $\pi_{0.026}$-opaque.

Although determining the π_ξ-opacity may in general be difficult, in several important cases it is possible to approximate the value of ξ with a desired accuracy.

The next result requires that the observation function is such that one does not have to wait for 'too long' in order to find a run in $\overline{\phi}$ covering $\lambda \in \phi$. More

precisely, we say that *obs* is N-*efficient* for $PLTS$ and ϕ, if N is a positive integer such that, for every run $\lambda \in \phi$ which is covered by runs in $\overline{\phi}$, there exists a run $\lambda' \in \overline{\phi}$ covering λ and satisfying $|\lambda'| \leq N \cdot |\lambda|$. Note that being efficient is not too demanding a requirement; in particular, each static observation function *obs* is N-efficient provided that $PLTS$ is finite and $\phi = \phi_Z$ for some $Z \subset S$ (N can then be taken to be the number of states of $PLTS$). Below, $\phi_k = \phi \cap runs_k(PLTS)$ and $\overline{\phi}_k = \overline{\phi} \cap runs_k(PLTS)$, for every $k \geq 0$.

Theorem 3. *If obs is N-efficient for $PLTS$ and ϕ, then there is an integer $\zeta \geq 1$ and a real number $0 \leq \eta < 1$ such that, for every $i \geq 0$:*

$$|\mu(\phi \setminus obs^{-1}(obs(\overline{\phi}))) - \mu(\phi_{\zeta \cdot i} \setminus obs^{-1}(obs(\overline{\phi}_{\zeta \cdot i})))| \leq \eta^i \ .$$

Proof. By Proposition 2, there exists a positive integer κ and a real number $0 \leq \delta < 1$ such that $\mu(runs_{\kappa \cdot i}(PLTS)) \geq 1 - \delta^i$, for every $i \geq 0$. In other words, for every $i \geq 0$:

$$\mu(runs(PLTS) \setminus runs_{\kappa \cdot i}(PLTS)) \leq \delta^i \ . \tag{7}$$

Let us now take any $i \geq 0$, and consider:

$$x = \mu(\phi \setminus obs^{-1}(obs(\overline{\phi})))$$
$$y = \mu(\phi_{N \cdot \kappa \cdot i} \setminus obs^{-1}(obs(\overline{\phi}_{N \cdot \kappa \cdot i}))) \ .$$

We then observe that, by *obs* being N-efficient, we have:

$$\begin{aligned} y &= \mu(\phi_{\kappa \cdot i} \setminus obs^{-1}(obs(\overline{\phi}_{N \cdot \kappa \cdot i}))) + \mu((\phi_{N \cdot \kappa \cdot i} \setminus \phi_{\kappa \cdot i}) \setminus obs^{-1}(obs(\overline{\phi}_{N \cdot \kappa \cdot i}))) \\ &= \mu(\phi_{\kappa \cdot i} \setminus obs^{-1}(obs(\overline{\phi}))) + \mu((\phi_{N \cdot \kappa \cdot i} \setminus \phi_{\kappa \cdot i}) \setminus obs^{-1}(obs(\overline{\phi}_{N \cdot \kappa \cdot i}))) \ . \end{aligned}$$

We therefore obtain:

$$x - y = \mu((\phi \setminus \phi_{\kappa \cdot i}) \setminus obs^{-1}(obs(\overline{\phi}))) - \mu((\phi_{N \cdot \kappa \cdot i} \setminus \phi_{\kappa \cdot i}) \setminus obs^{-1}(obs(\overline{\phi}_{N \cdot \kappa \cdot i})) \ .$$

Now, since

$$(\phi \setminus \phi_{\kappa \cdot i}) \setminus obs^{-1}(obs(\overline{\phi})) \subseteq runs(PLTS) \setminus runs_{\kappa \cdot i}(PLTS)$$
$$(\phi_{N \cdot \kappa \cdot i} \setminus \phi_{\kappa \cdot i}) \setminus obs^{-1}(obs(\overline{\phi}_{N \cdot \kappa \cdot i}) \subseteq runs(PLTS) \setminus runs_{\kappa \cdot i}(PLTS)$$

and $x, y \geq 0$, we obtain from (7) that $|x - y| \leq \delta^i$.
Hence the result holds with $\zeta = N \cdot \kappa$ and $\xi = \delta$. □

In other words, finite unfoldings of a probabilistic labelled transition system can approximate the probability of the uncovered runs in ϕ with a desired precision, providing a natural way of estimating π_ξ-opacity.

$\overline{\pi}_\gamma$-opacity. Let us consider the set ϕ^{cov} of runs of $\overline{\phi}$ which cover at least one run in ϕ, i.e., $\phi^{cov} = \overline{\phi} \cap obs^{-1}(obs(\phi))$. The first two notions of quantitative opacity retained the flavour of the original (qualitative) opacity. In particular, so far we have accepted that a set of runs ϕ with non-zero occurrence probability, $\mu(\phi) > 0$, can be covered by a set of runs with occurrence probability much lower than that of ϕ, $\mu(\phi^{cov}) \ll \mu(\phi)$, or indeed with a negligible chance of ever occurring, $\mu(\phi^{cov}) = 0$. That is, we were basically demanding very low occurrence probability of totally uncovered runs in ϕ. In our next definition, we remedy this by intuitively requiring that each run in ϕ is covered by γ runs, where $\gamma > 0$ would normally be expected to be (much) greater than 1. More precisely, for every $\gamma \geq 0$, we say that ϕ is $\overline{\pi}_\gamma$-*opaque* if:

$$\mu(\phi) > 0 \quad \text{and} \quad \frac{\mu(\phi^{cov})}{\mu(\phi)} = \gamma \,, \tag{8}$$

or, slightly more generally (as we do not have to assume that $\mu(\phi) > 0$), if the following holds:

$$\mu(\phi^{cov}) - \gamma \cdot \mu(\phi) = 0 \,. \tag{9}$$

In combination with π_ξ-opacity for small ξ, $\overline{\pi}_\gamma$-opacity for large γ would clearly increase our confidence in declaring ϕ opaque.

Example 4. Consider the following probabilistic labelled transition system:

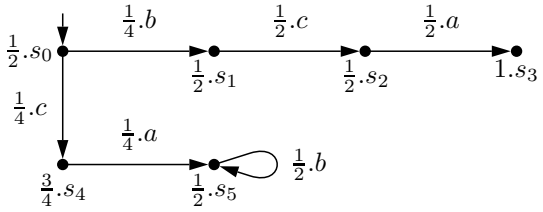

where $obs(a) = \epsilon$, $obs(b) = b$ and $obs(c) = c$, as well as the property $\phi = \phi_{\{s_3,s_5\}}$. We then have:

$$\phi = \{bca, ca, cab, cabb, \ldots\} \;\text{ and }\; \phi^{cov} = \{bc, c\} \,.$$

and so we obtain:

$$\mu(\phi) = \frac{1}{4}\cdot\frac{1}{2}\cdot\frac{1}{2} + \frac{1}{4}\cdot\frac{1}{4}\cdot\frac{1}{2} + \frac{1}{4}\cdot\frac{1}{4}\cdot\frac{1}{2}\cdot\frac{1}{2} + \frac{1}{4}\cdot\frac{1}{4}\cdot\frac{1}{2}\cdot\frac{1}{2} + \ldots = \frac{1}{8}$$

$$\mu(\phi^{cov}) = \frac{1}{4}\cdot\frac{1}{2}\cdot\frac{1}{2} + \frac{1}{4}\cdot\frac{3}{4} = \frac{1}{4}$$

which leads to:

$$\frac{\mu(\phi^{cov})}{\mu(\phi)} = 2.$$

The property $\phi_{\{s_3,s_5\}}$ is therefore $\overline{\pi}_2$-opaque.

As before, we will now investigate how one could approximate $\overline{\pi}_\gamma$-opacity. Below we assume that *obs* is *inversely M-efficient* for *PLTS* and ϕ, by which we mean that M is a positive integer such that, for every run $\lambda \in \phi^{cov}$, there exists a run $\lambda' \in \phi$ covered by λ and satisfying $|\lambda'| \leq M \cdot |\lambda|$. Again, being inversely efficient is not too demanding a requirement; in particular, each static observation function *obs* is inversely M-efficient provided that *PLTS* is finite and $\phi = \phi_Z$ for some $Z \subset S$ (M can then be taken to be the number of states of *PLTS*).

Theorem 4. *Let ϕ be $\overline{\pi}_\gamma$-opaque. If obs is inversely M-efficient for PLTS and ϕ, then there is an integer $\rho \geq 1$ and a real number $0 \leq \nu < 1$ such that, for every $i \geq 0$:*

$$|(\mu(\phi^{cov}) - \gamma \cdot \mu(\phi)) - (\mu(\phi^{cov}_{\rho \cdot i}) - \gamma \cdot \mu(\phi_{\rho \cdot i}))| \leq (1+\gamma) \cdot \nu^i \ .$$

Proof. By Proposition 2, there exists a positive integer κ and a real number $0 \leq \delta < 1$ such that, for every $i \geq 0$:

$$\mu(runs(PLTS) \setminus runs_{\kappa \cdot i}(PLTS)) \leq \delta^i \ . \tag{10}$$

Let us now take any $i \geq 0$, and consider:

$$x = \mu(\phi^{cov}) - \mu(\phi^{cov}_{M \cdot \kappa \cdot i})$$
$$y = \mu(\phi) - \mu(\phi_{M \cdot \kappa \cdot i}) \ .$$

We then observe that, by *obs* being inversely M-efficient, we have:

$$\begin{aligned}
x &= \mu(\overline{\phi} \cap obs^{-1}(obs(\phi))) \\
&\quad - (\mu(\overline{\phi}_{\kappa \cdot i} \cap obs^{-1}(obs(\phi_{M \cdot \kappa \cdot i}))) + \mu((\overline{\phi}_{M \cdot \kappa \cdot i} \setminus \overline{\phi}_{\kappa \cdot i}) \cap obs^{-1}(obs(\phi_{M \cdot \kappa \cdot i})))) \\
&= \mu(\overline{\phi} \cap obs^{-1}(obs(\phi))) \\
&\quad - (\mu(\overline{\phi}_{\kappa \cdot i} \cap obs^{-1}(obs(\phi))) + \mu((\overline{\phi}_{M \cdot \kappa \cdot i} \setminus \overline{\phi}_{\kappa \cdot i}) \cap obs^{-1}(obs(\phi_{M \cdot \kappa \cdot i})))) \\
&= \mu((\overline{\phi} \setminus \overline{\phi}_{\kappa \cdot i}) \cap obs^{-1}(obs(\phi))) - \mu((\overline{\phi}_{M \cdot \kappa \cdot i} \setminus \overline{\phi}_{\kappa \cdot i}) \cap obs^{-1}(obs(\phi_{M \cdot \kappa \cdot i}))) \\
y &= \mu(\phi \setminus \phi_{M \cdot \kappa \cdot i}) \ \leq \ \mu(\phi \setminus \phi_{\kappa \cdot i}) \ .
\end{aligned}$$

We therefore obtain from (10) that $|x| \leq \delta^i$ and $y \leq \delta^i$. Consequently, we obtain that

$$|x - \gamma \cdot y| \leq (\gamma + 1) \cdot \delta^i \ .$$

Hence the result holds with $\rho = M \cdot \kappa$ and $\nu = \delta$. □

In practice, one could approximate γ using the inequalities $|x| \leq \delta^i$ and $y \leq \delta^i$ from the above proof.

$\widetilde{\pi}_\psi$-opacity. The above notions of defining probabilistic opacity may still find it difficult to distinguish between subtle differences in which *obs* acts upon ϕ and $\overline{\phi}$. A possible way to assess such differences could be, e.g., to look at the probability distributions induced by $obs(\phi)$ and $obs(\phi^{cov})$ and conclude that they are rather similar.

In our last attempt at a notion of quantitative opacity, we define $\widetilde{\pi}_\psi$-opacity which uses Jensen-Shannon divergence as a way to measure the differences in which *obs* acts upon ϕ and ϕ^{cov}. Below we assume that $\mu(\phi) > 0$ and $\mu(\phi^{cov}) > 0$.

Since $runs(PLTS)$ with μ is a probabilistic space, $obs(runs(PLTS))$ can also be turned into a probabilistic space by defining

$$\pi(o) = \mu(obs^{-1}(o) \cap runs(PLTS)) \ ,$$

for every $o \in \mathcal{O} = obs(runs(PLTS))$. Moreover, any subset Λ of $runs(PLTS)$ with $\mu(\Lambda) \geq 0$ gives rise to a probability distribution Π_Λ on $\mathcal{O}$. More precisely, for every $o \in \mathcal{O}$:

$$\Pi_\Lambda(o) = \frac{\mu(obs^{-1}(o) \cap \Lambda)}{\mu(\Lambda)}.$$

Then, for a property ϕ, we can define Π_ϕ and $\Pi_{\phi^{cov}}$ and say that ϕ is $\widetilde{\pi}_\psi$-*opaque* if $0 \leq \psi \leq 1$ is their weighted Jensen-Shannon divergence:

$$D_{JS}(w \cdot \Pi_\phi, w' \cdot \Pi_{\phi^{cov}}) = \psi \ ,$$

where $w = \frac{\mu(\phi)}{\mu(\phi)+\mu(\phi^{cov})}$ and $w' = \frac{\mu(\phi^{cov})}{\mu(\phi)+\mu(\phi^{cov})}$.

Example 5. Consider the following probabilistic labelled transition system:

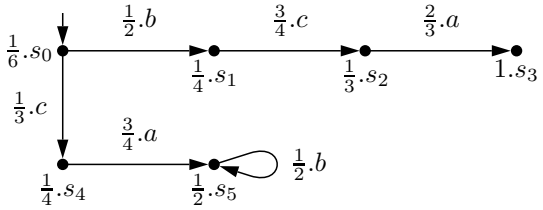

where $obs(a) = a$, $obs(b) = \epsilon$ and $obs(c) = c$, as well as the property $\phi = \phi_{\{s_2,s_3\}}$. We then have:

$$\phi = \{bc, bca\} \ \text{ and } \ \phi^{cov} = \{c, ca, cab, cabb, cabbb, \ldots\} \ .$$

and so:

$$\mu(\phi) = \frac{1}{2} \cdot \frac{3}{4} \cdot \frac{1}{3} + \frac{1}{2} \cdot \frac{3}{4} \cdot \frac{2}{3} \cdot 1 = \frac{3}{8}$$

$$\mu(\phi^{cov}) = \frac{1}{3} \cdot \frac{1}{4} + \frac{1}{3} \cdot \frac{3}{4} \cdot \frac{1}{2} + \frac{1}{3} \cdot \frac{3}{4} \cdot \frac{1}{2} \cdot \frac{1}{2} + \cdots = \frac{1}{3}$$

In this case, $\mathcal{O} = \{\epsilon, c, ca\}$, and we obtain:

$$\Pi_\phi = \left\{\epsilon \mapsto 0,\ c \mapsto \frac{\frac{1}{8}}{\frac{3}{8}},\ ca \mapsto \frac{\frac{1}{4}}{\frac{3}{8}}\right\} = \left\{\epsilon \mapsto 0,\ c \mapsto \frac{1}{3},\ ca \mapsto \frac{2}{3}\right\}$$

$$\Pi_{\phi^{cov}} = \left\{\epsilon \mapsto 0,\ c \mapsto \frac{\frac{1}{12}}{\frac{1}{3}},\ ca \mapsto \frac{\frac{1}{4}}{\frac{1}{3}}\right\} = \left\{\epsilon \mapsto 0,\ c \mapsto \frac{1}{4},\ ca \mapsto \frac{3}{4}\right\}$$

We calculate the weights of Π_ϕ and $\Pi_{\phi^{cov}}$ as:

$$w_{\Pi_\phi} = \frac{\frac{3}{8}}{\frac{3}{8}+\frac{1}{3}} = \frac{9}{17} \text{ and } w_{\Pi_{\phi^{cov}}} = \frac{\frac{1}{3}}{\frac{3}{8}+\frac{1}{3}} = \frac{8}{17}$$

and finally calculate:

$$\begin{aligned} D_{JS}(w_{\Pi_\phi} \cdot \Pi_\phi, w_{\Pi_{\phi^{cov}}} \cdot \Pi_{\phi^{cov}}) &= \mathcal{H}\left(0\ ,\ \frac{9}{17}\cdot\frac{1}{3} + \frac{8}{17}\cdot\frac{1}{4}\ ,\ \frac{9}{17}\cdot\frac{2}{3} + \frac{8}{17}\cdot\frac{3}{4}\right) \\ &\quad - \left(\frac{9}{17}\cdot\mathcal{H}\left(0, \frac{1}{3}, \frac{2}{3}\right) + \frac{8}{17}\cdot\mathcal{H}\left(0, \frac{1}{4}, \frac{3}{4}\right)\right) \\ &\approx 0.006\ . \end{aligned}$$

The property $\phi_{\{s_2,s_3\}}$ is therefore $\widetilde{\pi}_{0.006}$-opaque.

Similarly as in the previous cases, it may be possible to approximate the value of ψ in $\widetilde{\pi}_\psi$-opacity with a desired accuracy, using finite unfoldings of the probabilistic transition system. Below we assume that *obs* is K, L-*bounded* for *PLTS* and ϕ, by which we mean that K and L are positive integers such that:

- for every observation $o \in \mathcal{O}$ and $\lambda \in \phi \cup \phi^{cov}$, if $obs(\lambda) = o$ then $|\lambda| \leq K \cdot |o|$.
- for every run $\lambda \in \phi \cup \phi^{cov}$, $|obs(\lambda)| \leq L \cdot |\lambda|$.

Note that each static observation function *obs* is $K, 1$-bounded provided that *PLTS* is finite and *obs* does not induce ϵ-loops in the part of *PLTS* which is covered by the runs in $\phi \cup \phi^{cov}$ (K can then be taken to be the length of the longest ϵ-path in such a part of *PLTS* plus 1).

In the next result, we attempt to approximate the value of:

$$D_{JS}\left(w \cdot \left\{\frac{\mu(obs^{-1}(o) \cap \phi)}{\mu(\phi)}\right\}_{o\in\mathcal{O}}\ ,\ w' \cdot \left\{\frac{\mu(obs^{-1}(o) \cap \phi^{cov})}{\mu(\phi^{cov})}\right\}_{o\in\mathcal{O}}\right)\ .$$

To simplify the discussion, we assume that we are given the values of w, w', $\mu(\phi)$ and $\mu(\phi^{cov})$ (note that we can calculate them with a desired accuracy using Proposition 2).

Below $\mathcal{O}_k$ denotes $\{o \in \mathcal{O} \mid |o| \leq k\}$, for every $k \geq 0$. Moreover, for every $o \in \mathcal{O}$ and $m \geq 0$:

$$d_o^m = -(w \cdot p + w' \cdot p') \cdot \log_2(w \cdot p + w' \cdot p') + w \cdot p \cdot \log_2 p + w' \cdot p' \cdot \log_2 p'\ ,$$

where:

$$p = \frac{\mu(obs^{-1}(o) \cap \phi_m)}{\mu(\phi)} \quad \text{and} \quad p' = \frac{\mu(obs^{-1}(o) \cap \phi_m^{cov})}{\mu(\phi^{cov})} \ .$$

Theorem 5. *Let obs be K, L-bounded for PLTS and ϕ, and κ and $0 \leq \delta < 1$ be as in Proposition 2. Then there is $\alpha > 0$ such that, for every $i \geq 0$:*

$$0 \quad \leq \quad D_{JS}(w \cdot \Pi_\phi, w' \cdot \Pi_{\phi^{cov}}) - \sum_{o \in \mathcal{O}_{\kappa \cdot L \cdot i}} d_o^{\kappa \cdot L \cdot K \cdot i} \quad \leq \quad \alpha \cdot \delta^i \ . \tag{11}$$

Proof. Let d_o be the individual contribution of each $o \in \mathcal{O}$ to $D_{JS}(w \cdot \Pi_\phi, w' \cdot \Pi_{\phi^{cov}})$ as defined in (1). By the first part of K, L-boundedness, we obtain $d_o = d_o^{\kappa \cdot L \cdot K \cdot i}$, for every $o \in \mathcal{O}_{\kappa \cdot L \cdot i}$. This and (2) yields:

$$\begin{aligned} 0 &\leq D_{JS}(w \cdot \Pi_\phi, w' \cdot \Pi_{\phi^{cov}}) - \sum_{o \in \mathcal{O}_{\kappa \cdot L \cdot i}} d_o^{\kappa \cdot L \cdot K \cdot i} \\ &= \sum_{o \in \mathcal{O} \setminus \mathcal{O}_{\kappa \cdot L \cdot i}} d_o \quad \leq \quad c \cdot (\Pi_\phi(\mathcal{O} \setminus \mathcal{O}_{\kappa \cdot L \cdot i}) + \Pi_{\phi^{cov}}(\mathcal{O} \setminus \mathcal{O}_{\kappa \cdot L \cdot i})) \ . \end{aligned}$$

Now, by the second part of K, L-boundedness, we have that:

$$obs^{-1}(\mathcal{O} \setminus \mathcal{O}_{\kappa \cdot L \cdot i}) \subseteq runs_{\kappa \cdot i}(PLTS) \ .$$

Hence, by Proposition 2, we obtain

$$\Pi_\phi(\mathcal{O} \setminus \mathcal{O}_{\kappa \cdot L \cdot i}) \leq \frac{\delta^i}{\mu(\phi)} \ \text{ and } \ \Pi_{\phi^{cov}}(\mathcal{O} \setminus \mathcal{O}_{\kappa \cdot L \cdot i}) \leq \frac{\delta^i}{\mu(\phi^{cov})} \ .$$

As a result, (11) holds with $\alpha = c \cdot \left(\frac{1}{\mu(\phi)} + \frac{1}{\mu(\phi^{cov})}\right)$. □

5 Related work

Opacity and related concepts were first studied and related to information flow properties in a qualitative context in [6, 7, 5]. In the probabilistic context, opacity has been studied in [9, 2]. [9] studied the notion of opacity in the probabilistic computational world. There opacity was based on the probabilities of observer's pre-beliefs on the truth of the predicate. The work in [2] presents a quantitative information leakage analysis concerning probabilistic opacity, and there is a clear relationship between that work and the work in this paper. Indeed, although the setting in [2] is based on finite probabilistic automata, our probabilistic labelled transition system could be viewed as a generalisation of the fully probabilistic automaton (FPA) considered there. Note, however, that the automata in [2] are always finite and the notion of opacity is symmetric, while our system model allows infinite state spaces and we consider asymmetric opacity. Our work can also be related to quantitative analysis for secure information flow, including

[1, 3, 4, 8, 12–15]. Most of these works relate to the property of non-interference from the security literature, and they focus on flow measurement.

Opacity has already provided a promising technique for describing and unifying more general security properties. This paper has extended the notion of opacity to the model of probabilistic labelled transition systems. The results presented allow one to investigate and represent concepts from the literature on secure flow analysis.

6 Conclusions and Further Work

We have presented a formal model for the description of probabilistic opacity based on probabilistic labelled transition systems. We extend and generalise the notion of qualitative opacity and show how it applies to probabilistic and quantitative systems. We have investigated four alternative definitions of probabilistic opacity and given initial efficiency and approximation results. We believe that these results are promising and merit further consideration.

There is a clear link between the work presented here and the work on quantified information flow within the security community. Information flow security aims to ensure that information propagates throughout the execution environment without security violations such that only controlled secure information is deducible from observations of the system. The information we require to be confidential can be described as a predicate which we require to be opaque. By studying opacity in a quantified context we can relax the strict qualitative security policies, and tolerate a low probability that a quantitatively 'small' amount of secure information is leaked. We therefore believe our general model and results will be useful for quantified flow analysis in the security community.

More generally, we believe our work can provide a framework for the measurement of system security, by quantifying the opacity of key predicates with respect to the system. In future work, we plan to develop and extend the initial results presented here, as well as investigate and establish links between our work and the other work in the security community on the measurement of quantified information flow.

References

1. A. Aldini & A. Di Pierro (2004): *A Quantitative Approach to Noninterference for Probabilistic Systems.*
2. B. Bérard, J. Mullins & M. Sassolas (2010): *Quantifying Opacity.* In: *QEST*, pp. 263–272.
3. M. Boreale, F. Pampaloni & M. Paolini (2011): *Asymptotic Information Leakage under One-Try Attacks.* In: *FOSSACS*, pp. 396–410.
4. M. Boreale, F. Pampaloni & M. Paolini (2011): *Quantitative Information Flow, with a View.* In: *ESORICS*, pp. 588–606.
5. J. Bryans, M. Koutny, L. Mazaré & P. Y. A. Ryan (2005): *Opacity Generalised to Transition Systems.* In: *FAST*, pp. 81–95.

6. J. Bryans, M. Koutny, L. Mazaré & P. Y. A. Ryan (2008): *Opacity Generalised to Transition Systems. Int. J. Inf. Sec.* 7(6), pp. 421–435.
7. J. Bryans, M. Koutny & P. Y. A. Ryan (2004): *Modelling Dynamic Opacity Using Petri Nets with Silent Actions.* In: *FAST*, pp. 159–172.
8. K. Chatzikokolakis, T. Chothia & A. Guha (2010): *Statistical Measurement of Information Leakage.* In: *TACAS*, pp. 390–404.
9. Y. Lakhnech & L. Mazaré (2005): *Probabilistic Opacity for a Passive Adversary and its Application to Chaum's Voting Scheme.* Technical Report 4, Verimag.
10. K. G. Larsen & A. Skou (1989): *Bisimulation Through Probabilistic Testing (Preliminary Report).* In: *POPL*, ACM, New York, NY, USA, pp. 344–352.
11. J. Lin (1991): *Divergence Measures Based on the Shannon Entropy. IEEE Transactions on Information Theory* 37, pp. 145–151.
12. G. Lowe (2004): *Defining Information Flow Quantity. Journal of Computer Security* 12(3-4), pp. 619–653.
13. A. McIver, L. Meinicke & C. Morgan (2011): *Hidden-Markov Program Algebra with Iteration. CoRR* abs/1102.0333.
14. A. Di Pierro, C. Hankin & H. Wiklicky (2002): *Approximate Non-Interference.* In: *CSFW*, pp. 3–17.
15. A. Di Pierro, C. Hankin & H. Wiklicky (2005): *Quantitative Static Analysis of Distributed Systems. J. Funct. Program.* 15(5), pp. 703–749.
16. C. E. Shannon (1948): *A Mathematical Theory of Communication. SIGMOBILE Mob. Comput. Commun. Rev.* 5(1), pp. 3–55.