A Probabilistic Property-Specific Approach to Information Flow

Danièle Beauquier[1,⋆], Marie Duflot[1,⋆], and Marius Minea[2,⋆]

[1] University Paris 12, France
{beauquier, duflot}@univ-paris12.fr
[2] Institute e-Austria Timişoara, Romania
marius@cs.utt.ro

Abstract. We study probabilistic information flow from a property-specific viewpoint. For a given property of interest, specified as set of traces, we examine whether different low-level observations imply different probabilities for the occurrence of the property. Quantifying over all properties in a given class (e.g., high-level traces, or high-level sequences separated by low-level events) we obtain different notions of information flow. We give characterizations of systems that are secure according to these definitions. We consider both properties that are expressed over whole traces and those that distinguish between past and future given a reference point. In this framework, we can express several classical definitions of possibilistic security, as well as giving a more detailed, quantitative measure of information flow.

1 Introduction

Several classical treatments of information flow exist in the literature. Trace-based approaches assume a set of observable low-level events L and a set of (not directly observable) high-level events H. The question is whether observing a certain low-level trace can give information about the occurrence of high-level events, either in a possibilistic sense (the possibility or impossibility of a certain high-level interleaving) or in a probabilistic sense, yielding quantitative information about high-level activity.

It is generally accepted that there is no single all-encompassing definition of information flow. Different notions are noninterference [5], generalized noninterference [11], noninference [14], generalized noninference and separability [13], depending on the kind of information about high-level behavior considered relevant. In these *possibilistic* approaches, information flow is prevented if the trace set corresponding to a low-level observation contains "enough" traces to make inferences about high-level behavior impossible. Indeed, there can be no information flow if all high-level behaviors of interest are possible, i.e., included in the set of traces corresponding to a low-level observation. Precisely which traces must be present depends on the individual notion of information flow.

⋆ Partially supported by ECO-NET project No 08112WJ.

V. Gorodetsky, I. Kotenko, and V. Skormin (Eds.): MMM-ACNS 2005, LNCS 3685, pp. 206–220, 2005.

Subsequently, various frameworks for information flow [13,18,10], have been developed, attempting to unify the various existing definitions. McLean's introduction of selective interleaving functions [13] provides a way to reason about the relative strength of different security properties and their preservation under composition. Zakinthinos and Lee [18] propose "perfect security" as the weakest property on trace sets which guarantees absence of information flow (in a rather informally defined sense). In contrast, Mantel [10] argues the need for variety and modularity, and provides a library of basic security predicates from which common notions of security properties can be constructed.

In the same view, that an analysis of information flow must be flexible enough to be adapted to the specific features and needs of the considered application, we propose a *parameterized* view of information flow that develops a quantitative, probabilistic approach sketched in [17]. We define information flow with respect to a *property* (a set of system traces, possibly abstracted in its low-level part) which is deemed important for the system under scrutiny. The system has information flow with respect to the given property if there exist two low-level observations for which the chosen property has different probabilities of occurrence. In this case, the quantitative, probabilistic knowledge about the given property is sensitive to the observation which can be made, and so there is information flow in the system with respect to this property.

From this starting point, we define several generic notions of information flow, corresponding to different classes of properties of interest. These include *high-level* information flow, in which properties are sets of sequences of high-level events, and *sequential* information flow, in which properties can describe not only sequences of high-level events but also how these sequences are interrupted by the low-level, following the view of [12].

In examining information flow, we consider two views on the sequence of events in a trace. In the first, a *global* view, properties are simply sets of traces (infinite sequences of events). Alternatively, in a *relativized* view, the present timepoint splits a trace into a pair: a finite sequence of past events and an infinite sequence of future events. In this way, we can express properties that link the past behavior with the future behavior of the system; we have absence of information flow if such a behavior set is equiprobable regardless of the low-level observation up to the current timepoint. For instance, a property may state that if the last event before the time point is a then the next event is a' and if the last event before the time point is b then the next event cannot be a'.

We then give characterizations of systems that are secure according to these views of information flow, describing the structure of their trace sets in terms of high/low-level events and their probabilities.

Using this framework, and choosing appropriate sets of properties, we can express several classical definitions of possibilistic security: generalized noninterference [11], noninference [14], and separability [13]. At the same time, by supporting a user-defined choice of properties, we allow a finer granularity for the definition of information flow than previous approaches. In addition, for systems

that are not secure according to one of these notions, the probabilistic approach allows us to give a quantitative measure of the appearing information flow.

An important issue when defining security properties is deciding what kinds of information flow are acceptable. In some existing definitions of information flow, such as noninference [14] or the perfect security property [18], covert channels already existent in the description of a system are allowed, such as auditing or copying low-level events on a high-level. Such definitions take a *causal* view, defining information flow as the fact that high-level behavior influences low-level behavior. Conversely, this means that viewing a string of low-level events may allow us to deduce something about the high-level events that have occurred in the past, prior to these observations.

In contrast, we take a purely *observational* view. Thus, if a low-level observation is compatible only with an interleaving of high-level events, but not with another, this constitutes information flow, regardless whether this knowledge is already present in the description (trace set) of the system. Indeed, the probability of a given interleaving of high-level events depends in this situation on the low-level observation, which corresponds to our definition of information flow.

Related Work

Work on tailoring security properties to the system under consideration originates with the string of different definitions for information flow [5,11,14,13]. Following the recognition that security is a property of trace sets rather than traces (e.g., [13]), in [18], security properties are defined uniformly by specifying a predicate that the low-level equivalent bunch of a trace has to satisfy. The approach is taken further in [10] by defining basic security predicates in terms of a restriction and a closure requirement on a trace set. The parameterization in the latter paper is given by the variants in which the basic operations of inserting and deleting high-level events in a trace (to keep their absence and presence, respectively, confidential) can be performed.

Probabilistic information flow has naturally been more difficult to treat than the possibilistic version. McLean [12] introduces the *flow model* which distinguishes mere correlation from actual causal influence. Gray [7] introduces probabilistic interference in a context of finite state machines and gives a more general information-theoretic framework, including probabilistic channel capacity [6]. Sabelfeld and Sands [16] define probabilistic noninterference in the context of schedulers for multithreaded programs, based on the concept of probabilistic bisimulation, and show compositionality properties. Lowe [9] treats quantitative information flow distinguishing probabilistic aspects from nondeterminism, which is handled from an adversarial worst-case perspective; the treatment is done in a discrete-time context, considering also the rate of information flow. A probabilistic process-algebraic approach is given in [1], focused on noninterference, generalizing the possibilistic variant and allowing formal reasoning about the amount of information flow.

All these approaches, whether possibilistic or probabilistic, treat general, system-independent notion of information flow. A framework which parameter-

izes information flow is defined in [8] by giving a definition of secrecy in multi-agent systems, using a modal logic of knowledge in a state-based model. This generalizes several existing approaches and can be extended to probabilistic security. Their parameterization stems from defining formulas (knowledge) of what must be kept secret, thus providing a fine-grained way of characterizing security requirements. Since the approach is state-based, our model appears complementary in that it can talk about both past and future evolution of the system.

Other perspectives on information flow include that of [2] which offers a variety of characterizations of non-interference, expressed in Hoare logics and CTL; however, the variety is not given by parameterization, but language aspects such as sequential vs. concurrent, or termination sensitivity. Closer to a parametric view is the approach of [4], where the parameter is an observable property (an abstraction) of the public observations of a program. Thus, the attacker is a data-flow analyzer, and can be specified in an abstract interpretation framework. Both approaches deal with much more specific systems, described in particular programming languages, and the class of expressed properties, though parameterized to some extent, is not as general.

Beyond the possibilistic approaches, [3] analyzes quantitative information flow for a simple imperative language from a semantic point of view, whereas [15] replace indistinguishability in the formalization of non-interference by similarity based on the notion of distance, in a process-algebraic setting. In comparison, we also define quantity of information flow based on the distance between the probability of a property given an observation.

Our approach to parameterization allows properties that range from the general to the entirely system-specific. Thus we can select the granularity (a particular trace set or even a single trace) with respect to which information flow is analyzed. Alternatively, quantifying over classes of such properties, we can still obtain and reason about several of the classic definitions of information flow.

Paper Outline. We first introduce the mathematical model of probabilistic event systems which we use throughout the paper. Section 3 gives property-based definitions for three classes of probabilistic information flow, and theorems characterizing systems that conform to these notions. These results are extended in Section 4 to properties which distinguish between past and future with respect to the reference point defined by the observation. Section 5 shows how some of the classic definitions of information flow can be expressed in this formalism.

2 Probabilistic Event Systems

Notations

Given a finite alphabet A, we let A^* (resp. A^ω) denote the set of finite (resp. infinite) sequences (or traces) over this alphabet. The set A^∞ is the union of A^* and A^ω. The empty sequence is denoted ϵ. Given a sub-alphabet $A' \subset A$ and a trace λ, $\lambda_{|A'}$ denotes the projection of λ onto this sub-alphabet. If λ is a finite non-empty trace, $last(\lambda)$ denotes the last letter of λ.

Let λ be a (finite or infinite) trace. We denote by $Pref(\lambda)$ the set of finite prefixes of λ. More generally, if Tr is a set of traces, $Pref(Tr) = \bigcup_{\lambda \in Tr} Pref(\lambda)$.

Let $u, v \in (A^*)^n$, $u = (x_1, x_2, \ldots, x_n), v = (y_1, y_2, \ldots, y_n)$. We denote by $u \otimes v$ the *simple interleaving* of u and v defined as $u \otimes v = x_1 y_1 x_2 y_2 \ldots x_n y_n$.

If $U, V \subset (A^*)^n$, we denote by $U \otimes V$ the set: $U \otimes V = \{u \otimes v | u \in U, v \in V\}$.

If $U, V \subset (A^*)^\omega$, the definition of $U \otimes V$ is extended in a standard way.

The interleaving of two sequences x, y, denoted by $interl(x, y)$ is the set of sequences: $\{x_1 y_1 x_2 y_2 ... x_n y_n \mid x = x_1 x_2 ... x_n, n \in \mathbb{N},\ y = y_1 y_2 ... y_n,\ x_i, y_i \in A^*\}$. This extends to sets of sequences: $interl(X, Y) = \{interl(x, y) |\ x \in X, y \in Y\}$.

Probabilistic Event System

The execution of a system is modeled by its set Tr of traces which are finite or infinite sequences of atomic events from a set E. A particular atomic event τ is distinguished which represents the halting of the system. For example, if λ is a sequence of atomic events, it is useful to distinguish between "λ has occurred but the system still executes", and "λ has occurred and the system has stopped". The latter case is modeled by the event $\lambda\tau$. To unify the presentation, it is convenient to use only infinite sequences, writing $\lambda\tau^\omega$ instead of $\lambda\tau$. Then, from now on, Tr is a set of infinite sequences which do not contain any occurrence of τ except when they are of the form $\lambda\tau^\omega$ where λ contains no occurrence of τ.

The set of atomic events E is divided into two disjoint sets, the set H of high-level atomic events and the set L of low-level ones. Depending on the situation, the stop event τ can be considered as a low-level or a high-level event. In this paper, we only consider the case when the low-level user can observe that the system has stopped, i.e., $\tau \in L$.

The set of traces Tr is equipped with a probability measure μ over the σ-algebra generated by the cylinders λE^ω, $\lambda \in E^*$, such that Tr is μ-measurable. The measure $\mu(X)$ of a measurable set X is denoted as $Pr_\mu(X)$, or shortly $Pr(X)$. Thus if we consider the infinite tree T built from Tr with edges labeled by atomic events, each edge of the tree is equipped with a non-zero probability. (We assume that every prefix of a trace in Tr has a non-zero probability).

Traditionally, an event is a measurable set in the theory of probabilities, so to avoid confusion, the atomic events of the system will be called actions.

We use the customary notation for conditional probabilities: if P and Q are two measurable events and $Pr(Q) \neq 0$, the conditional probability $Pr(P|Q)$ is $Pr(P \cap Q)/Pr(Q)$. Since we are interested only in traces of the system S we deal only with conditional probabilities relative to Tr. Thus, for each measurable event X we denote by $Pr_S(X)$ the probability $Pr(X|S)$ (assuming $Pr(S) > 0$).

Definition 1. *An event system S is a tuple (E, H, L, Tr, μ) where $E = H \cup L$, and H (resp. L) is the set of high-level (resp. low-level) actions, Tr is the set of traces of the system, and μ is a probabilistic measure on Tr.*

We assume that only low-level actions are observable on the low-level, i.e., for a trace λ the projection $\lambda_{|L}$ is observable by low-level users. More precisely, a finite prefix of $\lambda_{|L}$ is observable. Thus, from the observation of $u \in L^*$, the

low-level user who is supposed to know the entire system can construct the *bunch* $B_S(u) = \{\lambda \in Tr \mid u \text{ is a prefix of } \lambda_{|L}\}$ and possibly deduce some information about what happened or what will happen at the high-level. When there is no ambiguity, we will write $B(u)$ instead of $B_S(u)$. For every u such that $B(u)$ is non empty, $B(u)$ is supposed to be measurable and without lost of generality the measure $Pr_S(B(u))$ is supposed to be positive. A projection $u \in L^*$ such that $B(u)$ is non empty is called *possible*.

3 Global Information Flow

Depending on the level of information we are interested in, we introduce an abstraction function $\phi : L \to L' \cup \{\epsilon\}$, where L' is some set with $|L'| \leq |L|$ and express properties as sets of infinite traces on $(H \cup L')^\omega$. We extend ϕ on H as the identity, and then on E^ω in a classical way. Notice that it is possible that an infinite trace of E^ω has an image which is finite.

A property of abstraction level ϕ is a subset of $(H \cup L')^\infty$. We consider only properties P such that $\phi^{-1}(P) \cap Tr$ is a measurable subset of Tr. By abuse of notation we write $Pr_S(P) =_{df} Pr_S(\phi^{-1}(P) \cap Tr)$, and write P instead of $\phi^{-1}(P) \cap Tr$ everytime we compute probabilities, e.g., in $Pr(P \cap A)$ or $Pr(P|A)$.

Definition 2. *Given a system S, the quantity of information flow for a property P of abstraction level ϕ is the value $IF(P,S) = max_{u,v}|Pr_S(P|B(u)) - Pr_S(P|B(v))|$ for all possible $u, v \in L^*$.*

A system S is without information flow for a property P of abstraction level ϕ if $IF(P,S) = 0$.

We can also consider a "qualitative" version of this definition:

Definition 3. *A system S is without qualitative information flow for a property P of abstraction level ϕ if for every $u \in L^*$ such that $B(u)$ is non-empty, $Pr_S(P) \neq 0 \to Pr_S(P|B(u)) \neq 0$.*

Definition 4. *A system is without information flow for a given abstraction level if it is without information flow for all properties of this level.*

We will consider three abstraction functions which are of interest in an obvious way. If $L = L'$ and ϕ is identity, i.e., there is no abstraction, we will speak of *general* information flow. If L' is a singleton $\{l\}$, and $\phi(l_i) = l$ for every $l_i \in L$, a trace on $(H \cup L')^\omega$ expresses what happens on the high-level, as well as whether two high-level events have been separated by a low-level event or not (the identity of this low-level event does not matter). In this second case we speak of *sequential* information flow. Finally, if $L' = \{\tau\}$ and $\phi(\tau) = \tau$, $\phi(l_i) = \epsilon$ for every $l_i \in L \setminus \{\tau\}$, that is we are interested only in the projection on the high-level of a trace, we will speak of *high-level* information flow.

The intuition behind this hierarchy of abstractions stems from the fact that we may be interested whether an event x is followed by an event y, in other words, in the presence of the pattern xy in a system trace.

If x and y are both high-level events, this property cannot be expressed using the definition of high-level information flow, since any intervening low-level events are projected out by the abstraction. However, it can be expressed as a sequential property: $(H \cup \{l\})^* xy (H \cup \{l\})^\omega$.

If one of the events (say x) is low-level and the other one high-level, the property can no longer be expressed by a sequential property, since the identity of y is lost by abstraction to l. However, the presence of the pattern xy can still be expressed as a general property: $(H \cup L)^* xy (H \cup L)^\omega$. This motivates considering properties which preserve full information for both high- and low-level events.

Another example to motivate our framework is the following. Consider a program where variables are classified as low (observable by low level users) or high. The system consists of the set of executions of the program. A regular property like "during every time duration t (the duration is measured by the number of events and t is a fixed integer), the high level variable x is updated at least once", in other words, it is impossible that there exists a time duration t without an update of variable x can be of interest, and one can require that the system does not suffer information flow for this property.

Let $L_0 = L \setminus \{\tau\}$.

We write $\mathcal{E} = (H \cup L_0)^\omega \cup (H \cup L_0)^* \tau^\omega$ for the set of all infinite words formed by actions from H and L. This is a superset of the set of system traces: $Tr \subseteq \mathcal{E}$.

In the following, low level actions are denoted $a, b, ...$, sequences of low-level actions $u, v, ...$, sequences of high-level actions $\alpha, \beta, ...$ and traces $\lambda, \lambda', ...$

Let $S = (E, H, L, Tr, \mu)$ be a system and T be the associated probabilistic tree. We define:

$$H_n(Tr) = \{(\alpha_1, ..., \alpha_n) \in (H^*)^n | \exists a_1 ... a_n \in L \ \alpha_1 a_1 \alpha_2 a_2 ... \alpha_n a_n \in Pref(Tr)\}.$$
$$H_n^\omega(Tr) = \{(\alpha_1, ..., \alpha_n) \in (H^*)^{n-1} H^\omega | \exists a_1 ... a_{n-1} \in L \ \alpha_1 a_1 \alpha_2 a_2 ... \alpha_n \in Tr\}.$$
$$L_n(Tr) = \{(a_1, ..., a_n) \in L^n | \exists \ \alpha_1 \ldots \alpha_n \in H^* \ \alpha_1 a_1 \alpha_2 a_2 ... \alpha_n a_n \in Pref(Tr)\}.$$
$$Tr_n = \{\alpha_1 a_1 \alpha_2 a_2 \ldots \alpha_n a_n \in Pref(Tr) | \ \alpha_i \in H^*, \ a_i \in L\}.$$

We give below a characteristic property for a system S to be without sequential information flow. For this we need to introduce some technical terms related to the probabilistic tree T.

We color edges labeled by a high-level action black and edges labeled by a low-level action red. We are interested in the set of sequences of high-level actions (including the empty word) which can occur starting from a node x. To make this set of sequences more explicit we build for each such node x a "black" probabilistic tree T_x in the following way: we keep only the black edges reachable in T from x, and for each node y (including x) accessible from x by a black path, we add a node y' and an edge (y, y') labelled by ϵ and with a probability equal to the sum p of the probabilities of red edges starting from y in T. The tree T_x is a probabilistic tree which has the following meaning: the probability of a path in T_x starting from x labelled by α (without ϵ labels) is exactly the probability that the sequence of high-level actions α occurs from x; the probability of a path in T_x starting from x labelled by α and ending in a leaf is the probability that from x the sequence of actions α followed by a low-level action occurs.

A node has the color of the edge ending in this node. The root is red.

Two red nodes x and x' of T are *H-equivalent* if there exists an integer n such that the labels of the paths from the root to x and x' are respectively $\alpha_1 a_1 \alpha_2 a_2 ... \alpha_n a_n$ and $\alpha_1 b_1 \alpha_2 b_2 ... \alpha_n b_n$ where $\alpha_i \in H^*$ and $a_i, b_i \in L$.

We also need to state an equivalence property on L. Two nodes x and x' of T are *L-equivalent* if there exists an integer n such that the labels of the paths from the root to x and x' are respectively $\alpha_1 a_1 \alpha_2 a_2 ... \alpha_{n-1} a_{n-1} \alpha_n a_n$ and $\beta_1 a_1 \beta_2 a_2 ... \beta_{n-1} a_{n-1} \beta_n a_n$ where $\alpha_i, \beta_i \in H^*$ and $a_i \in L$.

A tuple (x, x', y, y') of red nodes of the tree T is *H, L-compatible* if x and x' are H-equivalent, y and y' are H-equivalent, x and y are L-equivalent and x' and y' are L-equivalent, i.e., there exist $(\alpha_1, ..., \alpha_n), (\beta_1, ..., \beta_n) \in H_n$, and $(a_1, ..., a_n), (b_1, ..., b_n) \in L_n$ such that the paths from the root to x, x', y, y' are labeled respectively by $\alpha_1 a_1 ... \alpha_n a_n$, $\alpha_1 b_1 ... \alpha_n b_n$, $\beta_1 a_1 ... \beta_n a_n$ and $\beta_1 b_1 ... \beta_n b_n$.

Let $p_1, ..., p_n$, $q_1, ..., q_n$ be the probabilities of edges labeled by $a_1, ..., a_n$ on the path from the root to x (resp. y). Let $p'_1, ..., p'_n$, $q'_1, ..., q'_n$ be the probabilities of edges labeled by $b_1, ..., b_n$ on the path from the root to x' (resp. y').

A H, L-compatible tuple (x, x', y, y') is *perfect* if for every $i = 1, ..., n$ we have $p_i/q_i = p'_i/q'_i$.

The systems we consider are supposed to satisfy:

(1) Tr is a closed subset of $\mathcal{E}$
(2) For each measurable subset X of Tr, the closure $\bar{X}$ is measurable and $Pr_S(X) = Pr_S(\bar{X})$.

We start by characterizing *sequential* information flow, where the identity of low-level events is abstracted out, and only their position in the sequence of events is preserved.

Theorem 1. *A probabilistic system S such that $Tr \not\subset H^\omega$ is without sequential information flow iff*

(1) $\forall n > 0$ $Tr_n = H_n(Tr) \otimes L_n(Tr)$.
(2) Every H, L-compatible tuple of the tree T is perfect.
(3) For every pair of H-equivalent nodes x, x' of T, the probabilistic trees T_x and $T_{x'}$ are isomorphic.
(4) For every $n > 0$ $(L_n(Tr) \neq \emptyset \rightarrow Pr_S(Tr \cap (H^*L)^{n-1}H^\omega) = 0)$.

The intuition behind this characterization is the following: we don't want the low-level traces to give any information on the interleavings with the high level. Then, if a sequential high-level trace is possible, this trace can occur whatever the trace on the low level is. Point (4) states that observing that k low-level actions have occurred doesn't give any additional information, since all traces of Tr have the same number of low-level events. Points (2) and (3) state that probabilities of certain subtrees have to be equal or in equal ratios.

We give here only a sketch of the proof.

If the system has no information flow, then we prove (1), (4) and, by contradiction, the existence of the same edges in T_x and $T_{x'}$ in (2). For the latter, we exhibit properties for which, if one edge is not in T then for some u, v,

$Pr(P|B(u)) > 0$ and $Pr(P|B(u)) = 0$. The probabilistic parts of (2) and (3) are proven by contradiction as well, assuming that there exist nodes with different ratios, considering the pair of nodes with the highest ratio and obtaining information flow for some property.

The converse is proven by considering basic cylinders for which it is possible to show that there is no information flow. Then we define measurable subsets P_n which are disjoint unions of cylinders and we prove that there is no information flow for these sets. Taking the limit of these sets we show that the absence of information flow follows for P.

Next, we characterize general information flow, which turns out to be a very strong property:

Theorem 2. *The only systems with $Tr \not\subset H^\omega$ which are without general information flow are those which have a projection on L reduced to a single trace.*

Proof. Suppose that the projection of Tr on L is a trace w. Since $Tr \not\subset H^\omega$ this trace w is different from ϵ and the finite non-empty low-level words u such that $B(u) \neq \emptyset$ are the finite prefixes of w. Moreover for such a trace u, we have $B(u) = Tr$ and in this case, the system is without general information flow.

Conversely, suppose that the projection on L of the trace set Tr contains two different traces w and w', and let u be their longest common prefix. Let $a \in L$ such that ua is a prefix of w'. Let P be the property which consists of the infinite sequences in Tr whose projection on L is equal to w. We have $Pr_S(P \mid B(u)) > 0$ and $Pr_S(P \mid B(ua)) = 0$. Therefore S has general information flow. ■

To our knowledge, there is no simple characterization of systems which are without high-level information flow. It is immediate that any system without sequential information flow is without high-level information flow, since the definition of the latter has a coarser abstraction function. Also directly from the definition, it follows that the projection of any nonempty bunch $B(u)$ onto H must be the same, otherwise, for a high-level sequence $\alpha \in H^*$ distinguishing between $B(u)$ and $B(v)$ we can take $P = \alpha H^\omega$ and we have $Pr_S(P|B(u)) \neq Prb_s(P|B(v))$, since one is zero and the other one not.

4 Relativized Information Flow

The definitions of the previous section capture information flow, but provide no specific information about the time moment of the low-level observation and the events whose occurrence are linked to it. For a more refined and relativized view, one may wish to introduce the moment of observation in the property under consideration. For example a question of interest could be: observing some partial low-level trace at the current moment, what is the probability that the potential trace satisfies some past or future or more generally some relativized property? For example, what is the probability that starting from the current time, there is still one high-level action which will occur? Or, what is the probability that at current time, an event has occurred in the past, and will never occur in the future?

In this case, properties we are interested in are called relativized properties and are defined as subsets of $(\phi(H\cup L))^*\times(\phi(H\cup L))^\omega$, where ϕ is the abstraction function. The first component represents the past, and the second one the future.

A property P is a *past property* (resp. *future property*) if $P = R\times\phi((H\cup L)^\omega)$ (resp. $P = \phi((H\cup L)^*)\times R$) where $R\subset\phi((H\cup L)^*)$ (resp. $R\subset\phi((H\cup L)^\omega)$).

We state the definition of information flow in this relativized situation.

Let $u\in L^+$ with $B(u)\neq\emptyset$. For a relativized property P we define $Pr_S(P,u)= Pr_S(\{\gamma\in Tr|\ \gamma=\gamma_1\gamma_2,\gamma_{1|_L}=u,last(\gamma_1)=last(u),(\gamma_1,\gamma_2)\in P\})/Pr_S(B(u))$.

The event $\{\gamma\in Tr|\ \gamma=\gamma_1\gamma_2,\gamma_{1|_L}=u,last(\gamma_1)=last(u),(\gamma_1,\gamma_2)\in P\}$ corresponds to the situation when the low-level user observes u and the last action which occurred is a low-level action. We assume that P is well-behaved such that this event is a measurable set for every $u\in L^+$.

We can give now a definition of relativized information flow:

Definition 5. *A system S is without relativized information flow for a relativized property P of abstraction level ϕ if for every $u,v\in L^+$ such that $B_S(u)$ and $B_S(v)$ are nonempty, $Pr_S(P,u)=Pr_S(P,v)$.*

Definition 6. *A system is without relativized information flow for a given abstraction level if it is without relativized information flow for all relativized properties of this level.*

Again, one can use different levels of abstraction depending on the type of the events whose occurrence is of interest. For instance, consider the high-level event sequence xy, and assume one wishes to express that it occurs without any low-level event intervening after the last event of the low-level observation u. This can be expressed by the sequential relative property $(H\cup\{l\})^*\times H^*xy(H\cup\{l\})^\omega$. (A sequential property is needed to express the fact that x and y are not separated by low-level events). If now one of the interesting events (say y) is low-level, we need a general relative property so the identity of y is not abstracted away. For instance, $(H\cup L)^*\times(H\cup L)^2xy(H\cup L)^\omega$ expresses that xy will occur with two intervening events after the last low-level event of the given observation.

Theorem 3. *The only systems such that $Tr\not\subset H^\omega$ which are without relativized general information flow are those which have a projection on L equal to τ^ω.*

Proof. Suppose that the projection of Tr on L is equal to τ^ω, then the only finite sequences $u\neq\epsilon$ such that $B(u)$ is non-empty are $\tau^n,\ n>0$, and in that case $Pr_S(P,\tau^n)=Pr_S(P,\tau^m)$ for all positive integers m,n for every relativized general property P. We conclude that the system S has no relativized general information flow.

Conversely, suppose that the projection of Tr on L contains a trace $w\neq\tau^\omega$. Then the first action a of w is different from τ, otherwise w would be equal to τ^ω. Consider the property $P=\{(\gamma_1 a,\gamma_2)\in E^*\times E^\omega\mid\gamma_{1|L}=\epsilon\}$. We have $Pr_S(P,a)>0$ and $Pr_S(P,\tau)=0$. Therefore S has a relativized general information flow. ■

The next theorem characterizes the systems without relativized sequential information flow. Recall that in this case the abstraction function ϕ collapses all the low-level actions into a single one, the action l.

Theorem 4. *The only systems with $Tr \not\subset H^\omega$ which are without relativized sequential information flow are those which satisfy one of the following conditions:*

(1) the projection of Tr on L is reduced to τ^ω
(2) the projection of Tr on L is a subset M of L and $Tr = U \otimes (M \times \{\epsilon\})$ where $U = \{(\alpha_1, \alpha_2) \mid \alpha_1 l \alpha_2 \in \phi(Tr)\}$ and and for every pair of H-equivalent nodes x, x' of T, of depth one, the probabilistic trees $T(x)$ and $T(x')$ are isomorphic.

Proof. If the system S satisfies condition (1) it is easy to conclude like in Theorem 3 that S is without relativized sequential information flow.

If the system S satisfies condition (2), the only finite non-empty traces $u \in L^+$ such that the bunch $B(u)$ is non-empty are actions the $a \in M$. Clearly for every relativized sequential property P, $Pr_S(P, a) = Pr_S(P, b)$ for $a, b \in M$.

Conversely, let S be a system without relativized sequential information flow. Suppose that the projection of Tr on L is not reduced to τ^ω. We have to prove that S satisfies (2). The projection of Tr on L cannot contain a trace w with more than one action and different from τ^ω. Indeed suppose that $w = abw'$, $a, b \in L$. Then Tr contains a trace $\alpha a \beta b \lambda$, where $\alpha, \beta \in H^*$, and $\lambda \in (H \cup L)^\omega$. Consider now the relativized sequential property $P = \{\alpha l\} \times \{\beta l\}(H \cup \{l\})^\omega$. We have $Pr_S(P, a) \neq 0$ and $Pr_S(P, ab) = 0$. Contradiction. So the projection of Tr on L is a subset M of L. Let us prove that $Tr = U \otimes (M \times \{\epsilon\})$. Suppose that there exists $\alpha_1 l \alpha_2 \in \phi(Tr)$ and some $a \in M$ such that $\alpha_1 a \alpha_2 \notin Tr$. Then there is information flow for the property $P = \{\alpha l\} \times (H \cup \{l\})^\omega$: $Pr_S(P, a) = 0$ and there exists $b \in M$ such that $Pr_S(P, b) \neq 0$. Proving the other conditions of (2) is straightforward, following steps of the proof of Theorem 1. ■

The absence of relativized sequential information flow is a very strong property, and as seen from the conditions in Theorem 4, very few probabilistic event systems have this property. This stems from the fact that, in expressing the property P, a trace is split into two parts, just after the occurrence of a low-level event. If it is possible to observe more or fewer low-level actions in a trace than specified in the property, there is information flow.

But it is still interesting to consider low-level traces of the same length n, and examine if they give some additional high-level information (besides the fact that n low-level events have occurred). We are then interested in a weaker notion of "no information flow" for a relativized sequential property, namely:

Definition 7. *A system S is without information flow at each fixed step for a relativized property P if $Pr_S(P, u) = Pr_S(P, v)$ for every $u, v \in L^+$ such that $|u| = |v|$ and $B(u), B(v)$ are non-empty.*

In order to characterize the systems without sequential relativized information flow at each fixed step we need to introduce a new definition. In the probabilistic tree T of the system, the *red depth* of a node is the number of red edges on the path from the root to it.

Theorem 5. *A system S such that $Tr \not\subset H^\omega$ is without sequential relativized information flow at each fixed step iff*

(1) $\forall n > 0\ Tr_n = H_n(Tr) \otimes L_n(Tr)$.
(2) $\forall n > 0$, *all nodes of red depth* n *with outgoing red edges are equivalent*
(3) For every H*-equivalent nodes* x, x' *of* $T(S)$*, the probabilistic trees* T_x *and* $T_{x'}$ *are isomorphic.*

The proof of this theorem is based on the lemma given below which links sequential relativized information flow at each fixed step with sequential relativized information flow. Then we can reuse the proof of Theorem 1.

Lemma 1. *Let* $R\mathcal{E}'$ *be a sequential property on traces where* $R \subset (H \cup l)^*$ *and* $\mathcal{E}' = (H \cup l)^\omega \cup (H \cup l)^*\tau^\omega$ *. Then, for* $P_R = \{(\gamma_1, \gamma_2) \mid \ |\gamma_{1|_L}| = n, last(\gamma_1) = l, \gamma_1\gamma_2 \in R\mathcal{E}'\}$*, for every* u *of length* n *we have* $Pr_S(P_R, u) = Pr_S(R\mathcal{E}'|B(u))$.

5 Comparison with Some Classical Security Properties

In this section we restrict ourselves to finite systems, for which $Tr \subseteq (H \cup L)^*\tau^\omega$, and we suppose that $\tau \in L$. Denote by E_0 the set $H \cup L_0$, where $L_0 = L \setminus \{\tau\}$.

We identify an element of Tr with its shortest prefix ending with the action τ. Given a trace λ and a system S, the low-level user observing $\lambda_{|L_0}\tau$ can construct the set of system traces which correspond to the same observation, the *low-level equivalent set* [18] of λ:

For $\lambda \in E_0^*\{\tau\}$, $LLES(\lambda, S) = \{\beta \in Tr \mid \lambda_{|L_0} = \beta_{|L_0}\}$.

We will show that *separability*, *noninterference* and *noninference* can be expressed in our framework and correspond to the absence of information flow for some classes of properties.

1. Noninference

Noninference is a security property which was introduced by O'Halloran [14]. It requires that every trace λ of the system admits in its low-level equivalent set its projection $\lambda_{|L_0}$. As a consequence a low-level user cannot deduce from an observation the existence of any occurrence of a high-level action:

$Noninference(S) \equiv \forall \lambda \in Tr\ \exists u \in LLES(\lambda, S)\ u \in L_0^*\tau$.

Consider the property $NonInf = L_0^*\tau \subset (H \cup L_0)^*\tau$. A trace satisfies this property iff it does not contain high-level actions. Thus this property exactly focuses on the (non) existence of a high-level activity. It turns out that *noninference* can be expressed in terms of information flow for the property $NonInf$.

Theorem 6. *For a probabilistic system* S*,* $Noninference(S)$ *holds iff* $Pr_S(NonInf) \neq 0$ *and there is no qualitative general information flow for the property* $NonInf$.

Proof. Suppose $Pr_S(NonInf) \neq 0$ and there is no qualitative general information flow for the property $NonInf$. Let λ be a trace $\in Tr$. Consider the projection $u = \lambda_{|L}$. Since $B(u)$ is non-empty, $Pr_S(NonInf) \neq 0$ and there is no qualitative general information flow for the property $NonInf$. So we have $Pr_S(NonInf, u) \neq 0$. It proves that $u \in Tr$ because $B(u) \cap NonInf = \{u\}$. Thus, $Noninference(S)$ is true.

Conversely, suppose that $Noninference(S)$ holds. Let $\lambda \in Tr$, and $u = \lambda_{|L}$. Then $Pr_S(NonInf) \neq 0$, since u is also in Tr. Suppose there exists some $v \in L^*$ such that $Pr_S(NonInf, v) = 0$ and $B(v)$ is non-empty. There exists some λ' in $B(v)$, and the projection w of λ' on L belongs to Tr and v is a prefix of w. So, $Pr_S(NonInf, w) > 0$, but $Pr_S(NonInf, v) > Pr_S(NonInf, w)$, a contradiction. No qualitative general information flow for the property $NonInf$ can occur. ■

Moreover, we can quantify the degree of noninference by measuring the maximal value of $|Pr_S(NonInf) - Pr_S(NonInf|B(u))|$ for all non-empty $B(u)$, $u \in L^*$.

2. Separability

Separability is aimed to express a complete independence between the sequences of actions at high and low level:

$Separability(S) \equiv \forall \lambda \in Tr\ \forall \lambda' \in Tr\ interl(\lambda_{|L_0}, \lambda'_{|H})\tau \in Tr$.

Again this security property can be expressed in terms of qualitative sequential information flow for some set of properties. For each $\xi_1, ..., \xi_n \in H^*$, let $Sep_{\xi_1,...,\xi_n}$ be the following predicate defined on $(H \cup \{l\})^*$:

$Sep_{\xi_1,...,\xi_n}(\lambda)$ holds iff $\lambda = \xi_1 l \xi_2 l ... \xi_p l \xi_{p+1} \xi_{p+2} ... \xi_n l$ for some $p \leq n$.

Theorem 7. *For a probabilistic system S, $Separability(S)$ holds iff for any property $Sep_{\xi_1,...,\xi_n}$, where $\xi_1...\xi_n \in Tr_{|H}$, $Pr_S(Sep_{\xi_1,...,\xi_n}) \neq 0$ and there is no qualitative sequential information flow for these properties.*

Proof. Suppose $Separability(S)$ holds. Consider the property $Sep_{\xi_1,...,\xi_n}$ for some $\xi_1, ..., \xi_n \in Tr_{|H}$. Suppose $Pr_S(Sep_{\xi_1,...,\xi_n}) = 0$. Let $v = a_1 a_2 ... a_p$ be the projection on L of some trace in Tr. If $p \geq n$ then $\xi_1 a_1 \xi_2 a_2 ... \xi_n a_n a_{n+1} ... a_p \in Tr$, and if $p < n$ then $\xi_1 a_1 \xi_2 a_2 ... \xi_p a_p \xi_{p+1} ... \xi_n \in Tr$. The two cases contradict $Pr_S(Sep_{\xi_1,...,\xi_n}) = 0$. Suppose that for some $\xi_1...\xi_n \in Tr_{|H}$, there is qualitative sequential information flow for property $Sep_{\xi_1,...,\xi_n}$. This means that $Pr_S(Sep_{\xi_1,...,\xi_n}) \neq 0$ and there exists $u \in L^+$ with $Pr_S(Sep_{\xi_1,...,\xi_n} \mid B(u)) = 0$ and $B(u)$ is non-empty.

Let $v = a_1 a_2 ... a_p$ be the projection on L of some trace in $B(u)$. If $p \geq n$ then $\xi_1 a_1 \xi_2 a_2 ..., \xi_n a_n a_{n+1} ... a_p \in Tr$ which contradicts $Pr_S(Sep_{\xi_1,...,\xi_n} \mid B(u)) = 0$. If $p < n$ then $\xi_1 a_1 \xi_2 a_2 ..., \xi_p a_p \xi_{p+1} ... \xi_n \in Tr$ which contradicts again the fact that $Pr_S(Sep_{\xi_1,...,\xi_n} \mid B(u)) = 0$.

Conversely, suppose there is no qualitative sequential information flow for any property $Sep_{\xi_1,...,\xi_n}$, where $(\xi_1, ..., \xi_n) \in H_n(Tr)$ and there exists $\lambda, \lambda' \in Tr$ and $\nu \in interl(\lambda_{|L_0}, \lambda'_{|H})\tau$ such that $\nu \notin Tr$.

The trace ν can be written $\xi_1 a_1 \xi_2 a_2 ... \xi_{n-1} a_{n-1} \xi_n \tau$, where $\xi \in H^*$, and $a_i \in L_0$. Thus $Pr_S(Sep_{\xi_1,...,\xi_n} \mid B(a_1 a_2 ... a_{n-1}\tau)) = 0$ with $B(a_1 a_2 ... a_{n-1}\tau)$ non-empty since $a_1 a_2 ... a_{n-1}\tau = \lambda'_{|H}$. Therefore $Pr_S(Sep_{\xi_1,...,\xi_n})$ must be equal to zero since there is no information flow for this property. ■

3. Noninterference

Noninterference is a security property introduced by Goguen and Meseguer [5] and generalized by McCullough [11]. It demands that a low-level user cannot infer that any sequence of high-level inputs has (not) occurred. Let $HI \subset H$ (resp. HO) is the set of high-level input (resp. output) actions. We have $HI \cap HO = \emptyset$.

$\forall \lambda \in Tr\ \forall \gamma \in interl(HI^*, \lambda_{|L_0}) \exists \delta \in LLES(\lambda, S)\ \gamma = \delta_{|L_0 \cup HI}$

For each $\mu_1, ..., \mu_n \in HI^*$ let $Noninter_{\mu_1,...,\mu_n} = interl(HO^*, \mu_1 l \mu_2 l ... \mu_n l) \times (H \cup l)^\omega$. In a similar way to Theorem 7, one can prove

Theorem 8. *For a given probabilistic system S, Noninterference(S) holds iff for each n, for each $\mu_1, ..., \mu_n \in HI^*$ $Pr_S(Noninter_{\mu_1,...,\mu_n},\ u) \neq 0$ for every $u \in L^n$ such that $B(u)$ is non-empty.*

6 Conclusion

We have studied probabilistic information flow from a point of view parameterized by user-specified properties of interest. A property is a set of system traces, possibly viewed through an abstraction function. Our definitions support a range of property classes, e.g., referring to high-level events only, or high-level sequences separated by low-level events. We also allow specifications where a distinction is made between the past and future fragments of a trace. In this way, we can define (absence of) information flow for a given property, or for an entire set of properties of a given class.

We have given theorems that characterize the structure of systems for which absence of information flow according to these notions is guaranteed: for instance, a certain isomorphism between probabilistic trees is needed for properties which can distinguish subsequences of high-level events separated by low-level ones. We have also shown how several classic notions of possibilistic information flow (noninference, noninterference and separability) can be expressed using qualitative versions of our definitions.

We believe that this property-specific fashion of characterizing information flow is useful because it can be adapted to the particularities of the system under analysis. In many cases, a mere division into high- and low-level events and a single definition of information flow policy may not be enough, whereas our approach allows for a finer granularity of reasoning depending on the property.

An issue for future research is to apply this framework in the case where systems and properties are explicitly given as Markov chains and regular languages, respectively, and to investigate the decidability of the above notions of information flow in this setting.

Acknowledgements. We are grateful to Anatol Slissenko for the numerous and fruitful discussions of the approach studied in this paper.

References

1. Aldini, A., Bravetti, M., Gorrieri, R.: A process-algebraic approach for the analysis of probabilistic noninterference. Journal of Computer Security, 12 (2004) 191–246
2. Barthe, G., D'Argenio, P.R., Rezk, T.: Secure information flow by self-composition. 17th IEEE Computer Security Foundations Workshop. IEEE Computer Society (2004) 100–114

3. Clark, D., Hunt, S., Malacaria P.: Quantified interference for a while language. Electronic Notes Theoretical Computer Science, 112 (2005) 149–166
4. Giacobazzi, R., Mastroeni, I.: Abstract non-interference: parameterizing non-interference by abstract interpretation. Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. ACM (2004) 186–197
5. Goguen, J.A., Meseguer, J.: Security policies and security models. Proc. IEEE Symp. on Security and Privacy (April 1982) 11–20
6. James W. Gray III: Toward a mathematical foundation for information flow security. Proc. 1991 IEEE Symposium on Security and Privacy. IEEE Computer Society Press (1991) 21–35
7. J.W. Gray III: Probabilistic interference. Proc. IEEE Symp. on Security and Privacy (May 1990) 170–179
8. Halpern, J.Y., O'Neill, K.R.: Secrecy in multiagent systems. Proc. IEEE Computer Security Foundations Workshop (2002)
9. Lowe, G.: Quantifying information flow. Proc. IEEE Computer Security Foundations Workshop (June 2002) 18–31
10. Mantel, H.: Possibilistic definitions of security – An assembly kit. Proc. IEEE Computer Security Foundations Workshop (July 2000) 185–199
11. McCullough, D.: Specifications for multi-level security and hook-up property. Proc. IEEE Symp. on Security and Privacy (April 1987) 161–166
12. McLean, J.: Security models and information flow. Proc. IEEE Symp. on Security and Privacy (May 1990) 180–187
13. McLean, J.: A general theory of composition for trace sets closed under selective interleaving functions. Proc. IEEE Symp. on Security and Privacy (May 1994) 79–93
14. O'Halloran, C.: A calculus of information flow. Proc. of the European Symposium on Research in Security and Privacy (ESoRiCS'90) (1990) 180–187
15. Di Pierro, A., Hankin, C., Wiklicky, H.: Approximate non-interference. Journal of Computer Security, 12 (2004) 37–82
16. Sabelfeld, A., Sands, D.: Probabilistic noninterference for multi-threaded programs. Proc. IEEE Computer Security Foundations Workshop (July 2000) 200–214
17. Slissenko, A.: On probabilistic modeling of information flow. Talk at a working seminar of LACL (2004)
18. Zakinthinos, A., Lee, E.S.: A general theory of security properties. Proc. IEEE Symp. on Security and Privacy. IEEE Computer Society Press (1997) 74–102