# Secure information flow for concurrent processes

Jan Jürjens

LFCS, Division of Informatics, University of Edinburgh, GB *

**Abstract.** Information flow security is that aspect of computer security concerned with how confidential information is allowed to flow through a computer system. This is especially subtle when considering processes that are executed concurrently. We consider the notion of *Probabilistic Noninterference* (PNI) proposed in the literature to ensure secure information flow in concurrent processes. In the setting of a model of *probabilistic dataflow*, we provide a number of important results towards simplified verification that suggest relevance in the interaction of probabilistic processes outside this particular framework:
PNI is shown to be compositional by casting it into a rely-guarantee framework, where the proof yields a more general Inductive Compositionality Principle. We deliver a considerably simplified criterion equivalent to PNI by "factoring out" the probabilistic behaviour of the environment. We show that the simpler nonprobabilistic notion of *Nondeducibility-on-Strategies* proposed in the literature is an instantiation of PNI, allowing us to extend our results to it.

## 1 Introduction

### 1.1 Motivation

Information security of computer systems has received increased attention in recent years. A topical source of concern is the security of the Internet, which is becoming the world's largest public electronic marketplace. In addition to existing threats e. g. from hacker attacks, there are issues arising from the increasing use of mobile code in languages such as Java. If the Internet should provide the platform for commercial transactions, it is vital that sensitive information (like a credit card number or a cryptographic key) is not leaked out from a business site (e. g. by hackers) or from a user's computer (e. g. by malicious applets). Another example are multi-application smart cards (possibly containing sensitive medical data), especially when considering post-issuance code downloading.

In this paper we consider *confidentiality*, i. e. the prevention of unauthorised access to information, in a model of concurrent processes. More precisely, we examine the situation of "multilevel-security", where e. g. programs $H$ and $L$ (that could be *Trojan Horses*) running on an operating system $M$ have access

* `http://www.jurjens.de/jan`. Supported by the Studienstiftung des deutschen Volkes and the Division of Informatics. – In: C. Palamidessi (ed.), CONCUR 2000, LNCS 1877, Springer, 2000. Online version (30/01/01).

to a shared resource. $H$ (for high) is assumed to have access to confidential information that $L$ (for low) should not obtain. Here ensuring confidentiality means preventing information flow from $H$ through $M$ to $L$. To minimise the verification necessary one would prefer the security of the system composed of $H$, $M$ and $L$ to only depend on suitable requirements on $M$.

Leakage of information can happen in rather subtle ways via *covert channels* (discovered by Butler Lampson in 1973). An example for a means of sending a bit of information is the increased use of processor power by $H$ that could be detected by $L$.

Since eliminating covert channels in existing systems can be expensive, mathematical models have been proposed to detect them in the early phases of system development. This research offers the possibility to mechanically verify formal specifications with respect to security requirements. Since security issues can be very subtle on the one hand, and very costly when neglected on the other, it seems a realistic area for practical use of formal methods.

Recently there has been work on applying the approach to secure information flow considered in this paper to safety-critical systems (e. g. with the goal to avoid interference of possibly faulty off-the-shelf components with safety-relevant ones [DS99]). This suggests usefulness of the results given in this paper outside security of concurrent systems.

### 1.2 Previous Work

Early approaches to ensure confidentiality in multi-level secure systems (most well-known the model proposed by Bell and LaPadula in 1975) made use of access control, i. e. restricting the operations the respective users could perform. However, these models did not treat covert channels. Therefore the notion of *Noninterference* [GM82] was proposed to reason about secure information flow in deterministic systems. It formalises the idea that absence of information flow from $H$ to $L$ means that $L$ cannot distinguish different behaviours of $H$.

There are several approaches to generalize Noninterference to nondeterministic systems. [WJ90] showed that earlier models were either too weak or too strong and introduced the notion of *Nondeducibility-on-Strategies* (NOS). It did not take into account probabilistic aspects, and thus considered systems to be secure that are insecure when allowing an attacker to draw conclusions from statistical inference (for an example cf. section 5).

This motivated probabilistic notions inspired by Shannon's information theory, most importantly the *Applied Flow Model* (AFM, [McL90, Gra92]) and *Probabilistic Noninterference* (PNI, [Gra92]). [Gra92] showed that (in the above situation) if $M$ satisfies PNI then the communication channel from $H$ to $L$ has capacity zero.

To enable modular design of and reasoning about secure systems, special attention has been paid to compositionality of security properties, e. g. in [McL96]. For an excellent overview cf. [McL94].

Additionally there has been a significant amount of work on process algebras for information flow including [Ros95, FG97, Low99, RG99, RS99], cf. Section 6.

### 1.3 Present work

The main contribution of this paper is the proof of several important properties of Probabilistic Noninterference with the ultimate goal to make it a feasible property for mechanical verification. These results give insights into the interaction between probabilistic processes outside the scope of the specific notions considered.

For comparison, we note some weaknesses of the Applied Flow Model.

We then consider PNI in the setting of a model of *probabilistic dataflow* and prove various results towards a feasible verification of PNI, including equivalent definitions that are easier to verify and compositionality.

The first result shows that in the definition of PNI one only has to quantify over deterministic environments. This greatly reduces the amount of checking to be done when verifying PNI. Further research into this observation should provide similar simplifications for other properties of interacting probabilistic processes.

We show compositionality of PNI, solving an open problem posed in [Gra92]. Rely-guarantee specifications aim to permit compositionality without additional verification effort. We prove compositionality using a rely-guarantee formulation. This facilitates future research in providing weaker conditions on components with secure composition. The proof uses an Inductive Compositionality principle to be used in more general situations.

Finally we show that the nonprobabilistic notion of Nondeducibility-on-Strategies proposed in the literature is an instantiation of PNI, such that our results on PNI specialise to NOS.

Ultimate goal of this research is the mechanical verification of secure information flow. In a companion paper [Jür00] we express PNI in the framework of discrete Markov chains used for probabilistic modelchecking [dAKN$^+$00] by making use of a new equivalent coinductive definition in a state machine setting, and provide further simplifications of checking PNI.

Due to space limitations we can only sketch the proofs; the complete proofs will appear in an extended version.

## 2 Probabilistic models of Secure Information Flow

### 2.1 System Model

Before considering AFM and PNI we define the underlying trace-based model [Gra92] similar to dataflow models.

We consider concurrently executing probabilistic processes interacting by transmitting sequences of data values over unidirectional FIFO communication

channels. Communication is asynchronous in the sense that transmission of a value cannot be prevented by the receiver. Technically, this means that processes are input-total (i. e. at each state the receiving of any input induces a state change) which is not a real restriction since one can model a "deadlocked" process by one that never outputs anything (arguably, this is the only property of a deadlocked process its environment can observe anyhow).

Processes can take two roles: that of a *system* whose security wrt. information flow should be examined and that of an *environment* that may try to exploit possible holes in the system's security.

The interface of a process is given by disjoint finite sets of input resp. output ports $I$ resp. $O$ (possibly empty, with union $C := I \cup O$). We write $\mathcal{E}_O$ for the set of functions (*output events*) $o : O \to \mathcal{D}_o$ and $\mathcal{E}_I$ for the *input events* $i : I \to \mathcal{D}_i$ (where $\mathcal{D}_o, \mathcal{D}_i$ are finite data sets containing the possible data values on the channels $o, i$, including an element $\varepsilon$ representing absence of a value). For a system $M$ with $I$ (resp. $O$) the set of input (resp. output) ports we write $\mathcal{E}^M := (\mathcal{E}_O \times \mathcal{E}_I)^*$ (set of *system histories*, i. e. even-length alternating sequences of outputs and inputs), and for an environment $E$ (with $I, O$ as above) we write $\mathcal{E}^E := \mathcal{E}_I \times (\mathcal{E}_O \times \mathcal{E}_I)^*$ (set of *environment histories*, odd-length alternating sequences starting with input). Intuitively, the difference is that the system starts the interaction with the environment by outputting a value, and after that both sides take turns (the fact that the system starts simplifies the composition of systems defined later). This is a slight strengthening of the capabilities of the environment from [Gra92] necessary to ensure compositionality (cf. section 4). All other results in this paper do not depend on this.

A *system* is a triple $M = (I, O, P)$ with sets of input (resp. output) ports $I$ (resp. $O$) and $P : \mathcal{E}_O \times \mathcal{E}^M \to [0, 1]$ such that for each system history $w \in \mathcal{E}^M$ we have $\sum_{o \in \mathcal{E}_O} P(o, w) = 1$.

An *environment* is a triple $E = (I, O, P)$ with $I, O$ as above and $P : \mathcal{E}_O \times \mathcal{E}^E \to [0, 1]$ such that for each system history $w \in \mathcal{E}^E$ we have $\sum_{o \in \mathcal{E}_O} P(o, w) = 1$.

We write $P(o \mid w)$ for $P(o, w)$. $P(o|w)$ is the probability that after the history $w$ the next output of the process is $o$. Say that a history has length $t$ if it contains $t$ outputs.

In the following we consider systems $M$ with sets of input resp. output channels divided disjointly into low and high input resp. output channels: $I = I_l \cup I_h$, $O = O_l \cup O_h$. We write the high input values as a tuple $a$, and similarly the low input, high output and low output values as $b$, $x$ and $y$, resp. .[1]

We assume the environment to consist of two independent components $H$ and $L$ such that $M$ is connected via the high channels to $H$ and via the low channels to $L$. The output from $L$ and the low output from $M$ are additionally given as input to $H$, but $L$ cannot directly observe the values on the high channels (*secure environment criterion*). Thus $M = (I, O, P_M((x, y) \mid w))$ ($w$ ranging over $\mathcal{E}^M$), $L = (O_l, I_l, P_L(b \mid w \restriction_l))$ ($w \restriction_l$ the restriction of $w \in \mathcal{E}^{H,L}$ to the low channels)

[1] Thus in the picture below a line represents a *tuple* of channels.

and $H = (O \cup I_l, I_h, P_H(a \mid w))$ ($w$ in $\mathcal{E}^{H,L}$).

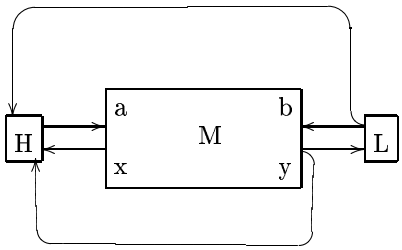


### 2.2 Applied Flow Model

We give the simpler verification condition [Gra92] equivalent to AFM.

**Definition 1.** *A system $M$ satisfies the* Applied Flow Model *(AFM) if for all traces $w, w'$ that are identical wrt. the low input and output, and for all low outputs $y$,*

$$\sum_x P_M((x,y) \mid w) = \sum_x P_M((x,y) \mid w').$$

This means that the probability of the next low output of $M$ depends only on the *low* part of the previous input-/output-history.

AFM is relatively simple and implies PNI, but it is too strict: there are visibly secure systems (satisfying PNI), that do not satisfy AFM [Gra92]. Worse, below we show AFM to contradict the usual assumptions on the model.

The example $S$ for a system that satisfies PNI but not AFM [Gra92] consists of a random number generator $R$ whose output is directly passed on to the high output of $S$, and simultaneously is sent to a delay component $D$ in $S$. $D$ outputs any value received after one time step, and its output is connected to the low output of $S$. $S$ is secure: since $H$ has no way of influencing the low output of $S$, it cannot send information to $L$.

However, if we drop the one-step delay, i. e. if the low output receives the data at the same time as the high output, then we obtain a system $S'$ that does satisfy AFM ! This is unrealistic since by definition of the Applied Flow Model $L$ has state (i. e. memory), so in the case of $S'$, $L$ has all information it has in $S$, and even one step earlier.

The following obvious weakening of AFM still fails to capture all secure processes:

**Definition 2.** *A process $M$ fulfils* weak AFM *if for all possible traces $w$ and $w'$ that are identical wrt. the low input and output* and the high output*, and all low-level output events $y$,*

$$\sum_x P_M((x,y) \mid w) = \sum_x P_M((x,y) \mid w').$$

Weak AFM correctly considers the example above to be secure. But the following secure process (that satisfies PNI) does not satisfy weak AFM. Let $\mathcal{D} = \{0, 1\}$. $M$ starts by outputting $x_1 = 0$ resp. $x_1 = 1$ each with probability $1/2$ to $H$, and $y_1 = \varepsilon$ to $L$. Upon input of values $a_1$ (resp. $b_1$) from $H$ (resp. $L$) $M$ then outputs $x_2 = \varepsilon$ to $H$ and $y_2$ to $L$ depending on $x_1$:

$x_1 = 0$: $y_2 := a_1$ with probability $1/4$ and $\neg a_1$ with probability $3/4$
$x_1 = 1$: $y_2 := \neg a_1$ with probability $1/4$ and $a_1$ with probability $3/4$

$M$ is secure, because independently of $H$, $L$ receives 0 and 1 each with probability $1/2$, so there is no flow from $H$ to $L$. But $M$ does not satisfy weak AFM: $P(a_1 \mid (a_1, x_1 = 0)) = 1/4 \neq 3/4 = P(a_1 \mid (\neg a_1, x_1 = 0))$.

### 2.3 Probabilistic Noninterference

For a system $M$ and environment $H, L$ we define the probability $P^M_{H,L}(w)$ of a system history $w \in \mathcal{E}^M$ inductively by $P^M_{H,L}(\varepsilon) = 1$ and

$$P^M_{H,L}(w.(x,y).(a,b)) = P^M_{H,L}(w) \cdot P_M((x,y) \mid w) \cdot P_H(a \mid w.(x,y)) \cdot P_L(b \mid w\restriction_l .y)$$

where . denotes concatenation (we leave out the superscript $M$ when possible).

$P^M_{H,L}(w)$ gives the probability that the execution of the system $M$ in the environment $H, L$ results in $w$.

The probability $P^l_{H,L}(w_l)$ of a low system history $w_l$ is

$$P^l_{H,L}(w_l) = \sum_{w : w\restriction_l = w_l} P_{H,L}(w)$$

**Definition 3.** *A system $M$ satisfies* Probabilistic Noninterference *if for all $H, H', L$ and for all low system histories $w_l$*

$$P_{H,L}(w_l) = P_{H',L}(w_l).$$

This means that the probability of a low trace arising from the interaction between $M$ with high and low processes only depends on the low process.

## 3 Deterministic environments are sufficient

**Definition 4.** *A probabilistic process $(I, O, P)$ is called* deterministic *if 0 and 1 are the only occurring probabilities, i. e. for all histories $w$ and all output events $o$ we have $P(o \mid w) \in \{0, 1\}$.*[2]

We now show that in fact in the definition of PNI it is sufficient to consider 'deterministic $H, H', L$:

[2] Thus exactly one output may occur at each point; in the following we specify deterministic processes by indicating these values.

**Theorem 1 (Determinism).** *A system $M$ satisfies Probabilistic Noninterference iff for all deterministic $H, H', L$ and for all low system histories $w_l$ we have*

$$P_{H,L}(w_l) = P_{H',L}(w_l).$$

The reason is that each of $M, H, L$ can only observe the values coming from the other components, and not directly the probability that certain values are sent. Therefore allowing $H$ and $L$ to send values with varying probabilities does not increase their capability to communicate through $M$ since by construction probabilities are propagated only indirectly.

This theorem simplifies verification of PNI considerably and allows significantly simpler proofs of earlier results on PNI. Further results on simplified verification are given in [Jür00].

*Proof.* (of the Theorem)

Call a process $(I, O, P)$ *$n$-deterministic* if its first $n$ outputs are deterministically determined by the previous history, i. e. for all histories $w$ of length $< n$ and all output events $o$ we have $P(o \mid w) \in \{0, 1\}$.

Fix $M$ and assume that for all deterministic $H, H', L$ and all low system histories $w^l$, $P_{H,L}(w^l) = P_{H',L}(w^l)$. Given $w^l = (y_1, b_1, \ldots, y_t, b_t)$ we must show $P_{H,L}(w^l) = P_{H',L}(w^l)$ for all $H, H', L$. Define $w^l_s = (y_1, b_1, \ldots, y_s, b_s)$ for $s \leq t$. We show inductively for $n = t, \ldots, 0$ that $P_{H,L}(w^l) = P_{H',L}(w^l)$ for all $n$-deterministic $H, H', L$.

This is clear by assumption for $n = t$ since for any $t$-deterministic process one can find a deterministic one with the same behaviour at the first $t$ outputs. Supposing the statement for $t \geq n > 0$ we prove it for $s = n - 1$: Suppose we are given $s$-deterministic $H, H', L$. Thus for each $r \leq s$ there exists $a(x_1, \ldots, x_r)$ such that $P_H(a(x_1, \ldots, x_r) \mid w(x_1, \ldots, x_{r-1}).(x_r, y_r)) = 1$ for $w(x_1, \ldots, x_{r-1}) = ((x_1, y_1), (a(x_1), b_1,), \ldots, (x_{r-1}, y_{r-1}), (a(x_1, \ldots, x_{r-1}), b_{r-1}))$. Similarly we have $a'(x_1, \ldots, x_r)$ for $H'$. By assumption on $L$ we may assume $P_L(b_r \mid w^l_{r-1}.y_r) = 1$ for $r \leq s$. Thus

$$\begin{aligned}
&P_{H,L}(w^l) \\
&= \sum_{x_1,\ldots,x_s} P_M((x_1, y_1) \mid \varepsilon) \cdot \ldots \cdot P_M((x_s, y_s) \mid w(x_1, \ldots, x_{s-1})) \\
&\cdot \sum_{a_{s+1}, x_{s+1}} P_M((x_{s+1}, y_{s+1}) \mid w(x_1, \ldots, x_s)) \\
&\cdot P_H(a_{s+1} \mid w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1})) \cdot P_L(b_{s+1} \mid w^l_s.y_{s+1}) \\
&\cdot P_{H,L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a_{s+1}, b_{s+1}))
\end{aligned}$$

with

$$P_{H,L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | v) := \sum_{a_{s+2}, x_{s+2}, \ldots, a_t, x_t} P_{H,L}(v.(x_{s+2}, y_{s+2}). \ldots .(a_t, b_t)) / P_{H,L}(v)$$

To derive a contradiction suppose $P_{H,L}(w^l) \neq P_{H',L}(w^l)$, so wlog. $P_{H,L}(w^l) > P_{H',L}(w^l)$. Construct $\tilde{L}$ from $L$ by defining $P_{\tilde{L}}(b_{s+1} \mid w_s^l.y_{s+1}) = 1$ (and leaving the probabilities for time steps $\neq s+1$ unchanged). Then $P_{H,\tilde{L}}(w^l) > P_{H',\tilde{L}}(w^l)$ since $P_{H,\tilde{L}}(w^l) = P_{H,L}(w^l)/P_{\tilde{L}}(b_{s+1} \mid w_s^l)$ (and similarly for $H'$). For all $x_1, \ldots, x_{s+1}, a_{s+1}$,

$$\begin{aligned}
&P_{H,\tilde{L}}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a_{s+1}, b_{s+1})) \\
&= P_{H,L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a_{s+1}, b_{s+1}))
\end{aligned}$$

since $L$ and $\tilde{L}$ differ only at the $s+1$st output.

For each $(x_1, \ldots, x_{s+1})$ choose $a(x_1, \ldots, x_{s+1})$ such that

$$\begin{aligned}
&\sum_{x_{s+1}} P_M((x_{s+1}, y_{s+1}) \mid w(x_1, \ldots, x_s)) \\
&\cdot P_{H,L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a(x_1, \ldots, x_{s+1}), b_{s+1}))
\end{aligned}$$

is maximal (this is possible since $\mathcal{D}$ is finite). Construct $\tilde{H}$ from $H$ by setting $P_H(a(x_1, \ldots, x_{s+1}) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1})) := 1$.

Similarly, for each $(x_1, \ldots, x_{s+1})$ choose $a'(x_1, \ldots, x_{s+1})$ such that

$$\begin{aligned}
&\sum_{x_{s+1}} P_M((x_{s+1}, y_{s+1}) \mid w(x_1, \ldots, x_s)) \\
&\cdot P_{H',L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a'(x_1, \ldots, x_{s+1}), b_{s+1}))
\end{aligned}$$

is minimal and construct $\tilde{H}'$ from $H'$ by setting

$$P_{H'}(a'(x_1, \ldots, x_{s+1}) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1})) := 1.$$

Then $P_{\tilde{H},\tilde{L}}(w^l) \geq P_{H,\tilde{L}}(w^l)$ because for every $x_1, \ldots, x_s, a_{s+1}$,

$$\begin{aligned}
&\sum_{x_{s+1}} P_M((x_{s+1}, y_{s+1}) \mid w(x_1, \ldots, x_s)) \\
&\cdot P_{H,L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a(x_1, \ldots, x_{s+1}), b_{s+1})) \\
&\geq \sum_{x_{s+1}} P_M((x_{s+1}, y_{s+1}) \mid w(x_1, \ldots, x_s)) \\
&\cdot P_{H,L}((b_{s+2}, y_{s+2}), \ldots, (b_t, y_t) | w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}).(a_{s+1}, b_{s+1}))
\end{aligned}$$

by definition of $a(x_1, \ldots, x_{s+1})$ and since

$$\begin{aligned}
&\sum_{a_{s+1}} P_H(a_{s+1} \mid w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1})) = 1 \\
&= P_H(a(x_1, \ldots, x_{s+1}) \mid w(x_1, \ldots, x_s).(x_{s+1}, y_{s+1}))
\end{aligned}$$

Similarly $P_{\tilde{H}',\tilde{L}}(w^l) \leq P_{H,\tilde{L}}(w^l)$. Thus $P_{\tilde{H},\tilde{L}}(w^l) > P_{\tilde{H}',\tilde{L}}(w^l)$ contradicting the inductive assumption. □

## 4 Compositionality

**Definition 5.** *The* composition $M \otimes N$ *of two systems* $M, N$ *with* $O_M \cap O_N = I_M \cap I_N = \emptyset$ *(which can always be achieved by renaming) is defined as follows:*

- $O_{M\otimes N} = O_M \cup O_N$
- $I_{M\otimes N} = I_M \cup I_N \setminus O_{M\otimes N}$
- $P^{M\otimes N}(o \otimes p | v \otimes w) = P^M(o|v) \cdot P^N(p|w)$ *for* compatible *histories* $v, w$.

*Here* $e \otimes f := (e \cup f)$ *for all output events* $e : O_M \to \mathcal{D}$, $f : O_N :\to \mathcal{D}$ *where the union of two functions is the union of the underlying relations (which by disjointness of the channel sets is a function). For input events* $e : I_M \to \mathcal{D}$, $f : I_N :\to \mathcal{D}$, $e \otimes f := (e \cup f)\restriction_{I_{M\otimes N}}$. *The tensor of histories is defined elementwise:* $(v_1, \ldots, v_t) \otimes (w_1, \ldots, w_t) = (v_1 \otimes w_1, \ldots, v_t \otimes w_t)$. *Histories* $v, w$ *of the same length are* compatible *if any value on an output channel appears at the input channel connected to it (if any) at the next time step.*

Note that the above composition makes the model essentially a probabilistic version of I/O-systems [Jon87]. As usual in models of dataflow, noncommunicating parallel composition and feedback are instantiations of the general composition.

In the following we always require that systems are only composed in a way such that low (resp. high) output channels are connected to low (resp. high) input channels.

The following example[3] of a system $M$ shows why the strengthening of the environment is needed to ensure compositionality: $M$ has two input channels $A_1, A_2$ (both high), a high output channel $X$ and a low output channel $Y$. $M$ is supposed to encrypt the high input on $A_1$ with the key provided on $A_2$ and then with an internally randomly generated one-time key pad and send the result out on $Y$. The random number should also be output on $X$, but so that the high input on $A_i$ cannot depend on it.

In the model in [Gra92], the system and the environment output values in parallel (not in turn) each independently of the input received from the other party at the same time. Then the following system meets the above description: At time 1, $M$ inputs binary values $a_1, a_2$ on $D_1, D_2$, resp. , and outputs a random binary value $x$ on $X$. At time 2, $M$ outputs $y := a_1 \oplus a_2 \oplus x$ on $Y$ (where $\oplus$ is exclusive or). Then $M$ is considered secure since $a_1 \oplus a_2$ does not depend on $x$ and so $y$ is 0 or 1 each with probability 1/2 independent of the high process $H$ (assuming a true random generator).

But the system derived from $M$ by feedback of $X$ into $A_2$ is not secure: it takes a value on $A_1$ at time 1 and outputs the same value on $Y$ at time 2.

The model considered in this paper is derived from the one in [Gra92] by allowing the output of the environment to depend on input received at the same time (this is the *worst case assumption* that the environment could have a substantially faster reaction). Then $M$ is correctly considered insecure.

[3] A variation of an example in [Sha58].

**Theorem 2 (Compositionality).** *If $M$ and $N$ satisfy PNI, so does $M \otimes N$.*

The proof of this theorem uses a rather general principle in a rely-guarantee setting inspired by [AL93] which is given below.

Their result on composing rely-guarantee specifications does not apply directly. Already in the nonprobabilistic case the information flow properties are sets of sets of traces [McL96], and not just sets of traces, as considered in [AL93]. The earlier approaches to compositionality of security properties including [McL96] only consider the nonprobabilistic case and so do not apply here.

The general setting is the following: $\mathcal{S}$ is a set of "systems" and $\mathcal{E}$ of "environments" such that $\mathcal{S} \subseteq \mathcal{E}$ (any system can be the environment of another system). $\otimes$ is a binary composition on $\mathcal{E}$ restricting to $\mathcal{S}$. For each $t \in \mathbb{N}_{\geq 0}$ there are properties $\phi_t$ on environments (i. e. $\phi_t \subseteq \mathcal{E}$) and $\sigma_t$ on systems relative to environments ($\sigma_t \subseteq \mathcal{S} \times \mathcal{E}$) with $\phi_0 = \mathcal{E}$ and $\sigma_0 = \mathcal{S} \times \mathcal{E}$. If $(M, E) \in \sigma_t$ say "$M$ guarantees $\sigma_t$ when placed into $E$". If $M$ guarantees $\sigma_t$ when placed into any $E$ satisfying $\phi_s$ we say "$M$ guarantees $\sigma_t$ assuming $\phi_s$". Write $\phi$ for $\bigwedge_t \phi_t$ and $\sigma$ for $\bigwedge_t \sigma_t$.

We also suppose that the rely-guarantee condition "guarantee $\sigma$ assuming $\phi$" can be obtained using the "approximations" "guarantee $\sigma_t$ assuming $\phi_t$": For any system $M$ and any $t$, if $M$ guarantees $\sigma_{t+1}$ assuming $\phi_{t+1}$ then $M$ guarantees $\sigma_t$ assuming $\phi_t$. Also $M$ guarantees $\sigma$ assuming $\phi$ iff for all $t$, $M$ guarantees $\sigma_t$ assuming $\phi_t$.

**Principle 1 (Inductive Compositionality)** *Suppose that for each $t$:*

**a)** *For any $E$, if $E$ satisfies $\phi_{t+1}$ and $M$ guarantees $\sigma_t$ assuming $\phi_t$ then $E \otimes M$ satisfies $\phi_{t+1}$ provided "$\sigma_t$ for $\phi_t$" is compositional (i. e. provided for each $M_1, M_2$, if each $M_i$ guarantees $\sigma_s$ assuming $\phi_s$, then $M_1 \otimes M_2$ guarantees $\sigma_s$ assuming $\phi_s$).*

**b)** *For each $E, M_1, M_2$: $M_1 \otimes M_2$ guarantees $\sigma_t$ in the environment $E$ if*
- *$M_1$ guarantees $\sigma_t$ in $E \otimes M_2$ and*
- *$M_2$ guarantees $\sigma_t$ in $E \otimes M_1$.*

*Then the following holds:*

*If $M_1$ and $M_2$ guarantee $\sigma$ assuming $\phi$ then $M_1 \otimes M_2$ guarantees $\sigma$ assuming $\phi$.*

The proof of the principle has to be omitted

*Proof.* (of the Theorem) Space permits only to give an idea how to express PNI using this principle. The idea is to weaken the "secure environment criterion" by requiring the environment itself only to fulfil (essentially) PNI, but only up to a specified length of traces. The circularity is broken since the definition of "PNI up to trace length $t$" on a system only requires assuming "PNI up to trace length $t-1$" on the environment, because of delay in the communication. To formulate PNI, the elements of $\mathcal{E}$ have to be finite sets of environments $H, L$ considered above. □

## 5 Probabilistic vs. Possibilistic

The Determinism theorem allows us to relate PNI to the nonprobabilistic property *Nondeducibility-on-Strategies* [WJ90].

The model used here is derived from the probabilistic one by disregarding probabilities. The behaviour of a process $(I, O, \mathcal{H})$ is given by its (prefix-closed) set of histories $\mathcal{H}$.

Even though the nondeterminism in this model is technically the same as the usual ("don't care") nondeterminism, the interpretation is different: In the "don't care" case one usually considers properties of traces independent of the existence of other traces. Here the nondeterminism is used to model specific security mechanisms like encryption where the relevant properties do not just concern an executed trace, but have to consider the traces that *could* have been executed instead. Noninterference properties do not concern traces but sets of traces [McL96] and one cannot use the usual refinement by reverse trace set inclusion. We refer to such a use of nondeterminism as *possibilistic nondeterminism.*

We define NOS: A *strategy* for $H$ is a function $\pi$ that assigns to every environment history $w \in \mathcal{E}_E$ a high input event $a = \pi(w)$. Thus a strategy defines a deterministic high system $H$ and vc. vs..

A history $(x_1, y_1), (a_1, b_1,), \ldots, (a_t, b_t)$ is *compatible* with a strategy $\pi$ if for each $s \leq t$, $\pi((x_1, y_1), (a_1, b_1,), \ldots, (x_s, y_s)) = a_s$, i. e. $\pi$ behaves according to the history. Thus a history $w$ of $M$ is compatible with $\pi$ iff it is an element of the set of possible traces of the execution of $M$ composed with the $H$ corresponding to $\pi$ and any deterministic $L$ behaving according to $w \restriction_l$ (i. e. an $L$ that outputs $b_s$ when the previous history is $y_1, b_1, \ldots, y_s$). Call a low history $w_l$ *compatible* with $\pi$ if there is a history $w$ with $w \restriction_l = w_l$ such that $w$ is compatible with $\pi$, and $w_l$ is *possible* if there is a strategy with which it is compatible.

**Definition 6.** *A system $M$ satisfies* Nondeducibility-on-Strategies *if for any strategy $\pi$ and any possible low history $w_l$ there is a history $w$ compatible with $\pi$ such that $w \restriction_l = w_l$ for its low view $w \restriction_l$.*

Intuitively, $L$ cannot observe which strategy $H$ follows.

One can "embed" possibilistic nondeterminism into probabilistic nondeterminism in a natural way as follows: For any possibilistic system $M$ construct a probabilistic system $\hat{M}$ by assigning to all outputs possible at a given point the same (nonzero) probabilities: if for a possible history $w$ the next possible outputs from $M$ are $o_1, \ldots, o_n$, define $P^M(o_i|w) = 1/n$ for each $i$.

Conversely, from any probabilistic system $M$ one can derive a possibilistic system $|M|$ by "forgetting" the probabilities: each history is possible in $|M|$ if it has nonzero probability in $M$.

Note that while $|\hat{M}| = M$ always holds, $\widehat{|M|}$ is the system derived from $M$ by "levelling out" the nonzero probabilities at a given state (i. e. assigning them uniform probabilities) and thus in general not equal to $M$.

We say that a possibilistic system $M$ satisfies PNI if $\hat{M}$ does. Then the Determinism theorem implies that for possibilistic systems, PNI reduces to NOS:

**Theorem 3. a)** *A possibilistic system $M$ satisfies PNI iff it satisfies NOS.*

**b)** *If a probabilistic system $M$ satisfies PNI then its "nonprobabilistic quotient" $|M|$ satisfies PNI (and thus NOS).*

Part a) means that, if the system $M$ under consideration is itself nonprobabilistic, then NOS is fully sufficient, even if the systems $H$ and $L$ that $M$ is connected with are probabilistic. On the other hand, this result allows to pass over to the possibilistic case whenever the probabilistic one is too complex to verify. Also, the above results obtained on PNI thus specialise to NOS.

Note that the implication in b) is not reversible: Let the system $M$ receive a high value $a$ and, with probability $1/2$, output it to $L$, while also with probability $1/2$ it outputs a random value to $L$. $|M|$ satisfies NOS (and the other nonprobabilistic information flow notions), but PNI correctly considers $M$ to be insecure, because there is a channel from $H$ to $L$ (even if it is not *noiseless*).

The implication "$M$ PNI $\Rightarrow$ $|M|$ NOS" has been observed without giving a proof in [Dut99].

*Proof.* (of the Theorem) Let $M$ be a possibilistic system.

a) $\Rightarrow$: Suppose $M$ satisfies PNI and we are given a strategy $\pi$ and a possible low history $w_l = (b_1, y_1), \ldots, (b_t, y_t)$. We need to show that there is a history $w$ compatible with $\pi$ such that $w \upharpoonright_l = w_l$ for its low view $w \upharpoonright_l$. Since $w_l$ is possible by assumption, there must be a history $w'$ and a strategy $\pi'$ such that $w'$ is compatible with $\pi'$ and $w \upharpoonright_l = w_l$. Let $H$ (resp. $H'$) be defined by $\pi$ (resp. $\pi'$) and let $L$ output $b_1, \ldots, b_t$ regardless of the $y_i$. By definition of PNI and of the translation $\hat{M}$ to the probabilistic setting $w_l$ is possible in an execution of $M$ composed with $H$ and $L$ (i. e. the corresponding trace set contains a trace $w$ with $w \upharpoonright_l = w_l$) iff it is possible when composed with $H'$ and $L$. But the first part of the equivalence holds by assumption on $\pi'$, and the second is what we needed to show.

$\Leftarrow$: Suppose $M$ satisfies NOS. By the Determinism Theorem and the definition of $\hat{M}$ it is sufficient to show that given deterministic systems $H, H', L$, each low history $w_l$ is possible in an execution of $M$ composed with $H$ and $L$ iff it is possible when composed with $H'$ and $L$. By symmetry, it is sufficient to show for each $H, H', L, w_l$, if $w_l$ is possible with $H, L$, then also with $H', L$. Suppose it is possible with $H, L$. This implies that $w_l$ is possible and that $L$ behaves according to $w_l$. Then by definition of NOS, for the strategy corresponding to $H'$ there exists a history $w$ compatible with $\pi'$ such that $w \upharpoonright_l = w_l$. But this means that $w_l$ is possible with $H', L$.

b) Since $\widehat{|M|}$ is the system derived from $M$ by at each state assigning equal probabilities to the output events with nonzero probability in $M$, the proof only requires the observation that if two probabilities are equal (such as in the definition of PNI) then they are in particular both nonzero or both zero. $\square$

## 6 Related Work

Goal of our work is the verification of specifications of secure information flow, including notions that allow detection of covert channels arising from statistical inference.

In recent years there have been reformulations of nonprobabilistic information flow notions in process algebras which provide mechanical support. Thus one could aim to achieve the goal by trying to extend these frameworks with probabilistic aspects. We consider the approaches to information flow using CSP and CCS, process algebras which have been very successful in other areas of concurrency (e. g. deadlock-detection) and security (e. g. cryptographic protocols).

Since the natural translation of noninterference into CSP is not easily addressable by its model checker and to avoid the "refinement paradox" [RWW94, Ros95] introduces the approach of "noninterference through determinism" which however is too restrictive according to [Low99] who points out that the other definitions up to that date for process algebras were also too weak or too strong. He argues that standard models of CSP are not sufficient to treat information flow since they do not make enough distinctions, and offers a solution using a non-standard model of CSP. However, he concludes that "the model is moderately complicated, but I suspect that this is inevitable". He also points out weaknesses of existing probabilistic process algebras.

For the CCS approach, an extensive comparison [FG94] brought forward a reformulation of NOS called BNDC as the preferred notion [FG97]. [Foc96] shows that it is under suitable conditions equivalent to the one in [Ros95], which extends the critique cited above to this approach.

The bi-directional "handshake" communication common to both process algebras requires care, since one wants to allow low behaviour to interfere with high behaviour, but not the other way [RG99].

According to [RS99] there is no consensus yet which notion of noninterference is "correct", thus it seems profitable to examine various notions, including the non-process-algebraic ones presented here.

## 7 Conclusion and Further Work

After demonstrating weaknesses of the alternative notion AFM we have given several important results on Probabilistic Noninterference towards practical verification. This includes the result that in PNI one needs only quantify over deterministic environments, the proof of compositionality of PNI (an open problem from [Gra92]) - using a new Inductive Compositionality Principle in a rely-guarantee framework - and the proof that the nonprobabilistic notion Nondeducibility-on-Strategies proposed in the literature is an instantiation of PNI.

In a companion paper [Jür00] we express PNI in the framework of discrete Markov chains used for probabilistic modelchecking [dAKN+00] by making use of a new equivalent coinductive definition in a state machine setting, and we provide further results towards simplification of PNI. Further work needs to be

done wrt. efficiency and practicality of using PNI. Also we have been considering a notion of *secure refinement* inspired by [AHKV98, Abr99].

Current work includes extending the model considered here with primitives for symmetric encryption and giving a computational interpretation following the approach of [AR00].

In another direction, since sometimes covert channels cannot be avoided completely, one should consider relaxing the conditions towards accepting small influences of high behaviour on low probabilities, possibly along the lines of [DGJP99].

## 8 Acknowledgements

Many thanks for interesting comments go to Samson Abramsky, Bruno Dutertre (also for making [Dut99] available), Riccardo Focardi, Dusko Pavlović, and furthermore to Peter Ryan for providing the opportunity to attend the workshop on Information Flow (London, Dec. 1999). Helpful comments by several anonymous referees are acknowledged.

Part of this research was done during a visit to Kestrel Institute (Palo Alto, CA) in January 2000 whose hospitality is gratefully acknowledged.

## References

[Abr99] S. Abramsky. A note on reactive refinement, 1999. Manuscript.

[AHKV98] Rajeev Alur, Thomas A. Henzinger, Orna Kupferman, and Moshe Vardi. Alternating refinement relations. In *CONCUR'98*, volume 1466 of *LNCS*, pages 163–178. Springer, 1998.

[AL93] M. Abadi and Leslie Lamport. Composing specifications. *ACM Transactions on Programming Languages and Systems 15*, 1:73–132, January 1993.

[AR00] M. Abadi and P. Rogaway. Reconciling two views of cryptography (invited lecture). In *TCS 2000 (IFIP conference)*, Japan, August 2000.

[dAKN+00] L. de Alfaro, M. Kwiatkowska, G. Norman, D. Parker, and R. Segala. Symbolic model checking of concurrent probabilistic processes using MTBDDs and the Kronecker representation. In *TACAS'2000*, LNCS, January 2000.

[DGJP99] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panagaden. Metrics for labeled markov systems. In *CONCUR*, 1999.

[DS99] Bruno Dutertre and Victoria Stavridou. A model of noninterference for integrating mixed-criticality software components. In *DCCA-7, Seventh IFIP International Working Conference on Dependable Computing for Critical Applications*, San Jose, CA, January 1999.

[Dut99] B. Dutertre. State of the art in secure noninterference, 1999. Manuscript.

[FG94] R. Focardi and R. Gorrieri. A classification of security properties for process algebra. *J. Computer Security, 3(1):5-33*, 1994.

[FG97] R. Focardi and R. Gorrieri. The compositional security checker: A tool for the verification of information flow security properties. *IEEE Transaction of Software Engineering*, 23(9), September 1997.

[Foc96] R. Focardi. Comparing two information flow security properties. In *Proceeding of the 9th IEEE Computer Security Foundation Workshop*, pages 116–122, June 1996.

[GM82] J. Goguen and J. Meseguer. Security policies and security models. In *Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society, 1982.

[Gra92] J. W. Gray. Toward A Mathematical Foundation for Information Flow Security. *Journal of Computer Security*, 3–4(1):255–294, 1992.

[Jon87] B. Jonsson. *Compositional Verification of Distributed Systems.* PhD thesis, Department of Computer Systems, Uppsala University, 1987. Tech. Rep. DoCS 87/09.

[Jür00] J. Jürjens. Verification of probabilistic secure information flow, 2000. submitted.

[Low99] Gavin Lowe. Defining information flow. Technical report, Department of Mathematics and Computer Science Technical Report 1999/3, University of Leicester, 1999.

[McL90] J. McLean. Security Models and Information Flow. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, pages 180–187, Oakland, CA, May 1990.

[McL94] J. McLean. Security models. In John Marciniak, editor, *Encyclopedia of Software Engineering.* Wiley & Sons, Inc., 1994.

[McL96] J. McLean. A general theory of composition for a class of "possibilistic" properties. *IEEE Transactions on Software Engineering*, 22(1):53–67, 1996.

[RG99] A. Roscoe and M. Goldsmith. What is intransitive noninterference ? In *IEEE Computer Security Foundations Workshop*, 1999.

[Ros95] A.W. Roscoe. CSP and determinism in security modelling. In *IEEE Symposium on Security and Privacy*, 1995.

[RS99] P. Ryan and S. Schneider. Process algebra and non-interference. In *IEEE Computer Security Foundations Workshop*, 1999.

[RWW94] A. Roscoe, J. Woodcock, and L. Wulf. Non-interference through determinism. In *ESORICS 94*, volume 875 of *LNCS*. Springer, 1994.

[Sha58] C. Shannon. Channels with side information at the transmitter. *IBM Journal of Research and Development*, 2:289–293, 1958.

[WJ90] J. T. Wittbold and D. M. Johnson. Information Flow in Nondeterministic Systems. In *IEEE Symposium on Security and Privacy*, pages 144–161, 1990.