Noninterference Analysis of Reversible Probabilistic Systems

Andrea Esposito

University of Urbino

Joint work with Alessandro Aldini and Marco Bernardo

Noninterference

- The notion of noninterference was first introduced by Goguen and Meseguer (1982).
- Used to reason about the way in which illegitimate information flows can occur in multi-level security systems by exploiting covert channels.
- Noninterference guarantees that low-level agents can never infer from their observations what high-level agents are doing.
- Regardless of the specific implementation, noninterference is closely tied to the notion of behavioral equivalence among processes.

- One of the most established formal definitions of equivalence employed for noninterference properties is weak bisimilarity.

Noninterference in Nondeterministic Reversible Systems

- One of the most established formal definitions of equivalence employed for noninterference properties is weak bisimilarity.
- It is not adequate to study noninterference in reversible systems.

Noninterference in Nondeterministic Reversible Systems

- One of the most established formal definitions of equivalence employed for noninterference properties is weak bisimilarity.
- It is not adequate to study noninterference in reversible systems.
- Esposito, Aldini, and Bernardo (2023) have shown that an adequate semantics is given by branching bisimilarity.
- The reason is that it has been proven to coincide with weak back-and-forth bisimilarity [DMV90].

Noninterference in Probabilistic Reversible Systems

- How to analyze noninterference in probabilistic reversible systems?
- Probabilistic noninterference has been investigated by Aldini, Bravetti, and Gorrieri (2004) in the generative-reactive model, where only a very limited form of nondeterminism is allowed.
- In their calculus, in addition to probabilistic choice, other operators such as parallel composition and hiding are decorated with a probabilistic parameter.
- This complicates the definitions of noninterference properties as they require universal quantifications over probabilistic parameters.

Noninterference in Probabilistic Reversible Systems

- We want to study noninterference for reversible systems that feature both nondeterminism and probabilities.
- A more expressive probabilistic model is the strictly alternating model introduced by Hansson e Jonsson (1990):
  - States are divided into nondeterministic ($\mathcal{S}_{\mathrm{n}}$) and probabilistic ($\mathcal{S}_{\mathrm{p}}$).
  - Transitions are divided into:
    - action transitions, from $\mathcal{S}_{\mathrm{n}}$ to $\mathcal{S}_{\mathrm{p}}$
    - probabilistic transitions, from $\mathcal{S}_{\mathrm{p}}$ to $\mathcal{S}_{\mathrm{n}}$.
- We use weak and branching bisimilarities for this model to recast a variety of noninterference properties (they are decidable in polynomial time).
- A process calculus in which to express noninterference properties, where only the probabilistic choice operator is decorated.

Probabilistic Labeled Transition Systems

Definition

A *probabilistic labeled transition system (PLTS)* is a triple $(\mathcal{S}, \mathcal{A}_\tau, \longrightarrow)$:

- $\mathcal{S} = \mathcal{S}_\mathrm{n} \cup \mathcal{S}_\mathrm{p}$ is a nonempty set of nondet. ($\mathcal{S}_\mathrm{n}$) and prob. ($\mathcal{S}_\mathrm{p}$) states with $\mathcal{S}_\mathrm{n} \cap \mathcal{S}_\mathrm{p} = \emptyset$.
- $\mathcal{A}_\tau = \mathcal{A} \cup \{\tau\}$ is a countable set of actions with $\tau \notin \mathcal{A}$ denoting the unobservable action.
- $\longrightarrow\ =\ \longrightarrow_\mathrm{a} \cup \longrightarrow_\mathrm{p}$ is a transition relation where:
  - $\longrightarrow_\mathrm{a} \subseteq \mathcal{S}_\mathrm{n} \times \mathcal{A}_\tau \times \mathcal{S}_\mathrm{p}$ is the action transition relation.
  - $\longrightarrow_\mathrm{p} \subseteq \mathcal{S}_\mathrm{p} \times \mathbb{R}_{]0,1[} \times \mathcal{S}_\mathrm{n}$ is the probabilistic transition relation where $\sum_{s \xrightarrow{p}_\mathrm{p} s'} p \in \{0, 1\}$ for all $s \in \mathcal{S}_\mathrm{p}$.

Probabilistic Bisimilarities

- Identifying nondeterministic (resp. probabilistic) states when they behave the same based on their transitions [HJ90].
- Philippou, Lee, and Sokolsky (2000) additionally allows a nondeterministic state and a probabilistic state to be identified when the latter concentrates all of its probabilistic mass in reaching the former.
- To this purpose the following function is introduced:

$$\mathit{prob}(s, s') = \begin{cases} p & \text{if } s \in \mathcal{S}_{\mathrm{p}} \wedge \sum_{s \xrightarrow{p'}_{\mathrm{p}} s'} p' = p > 0 \\ 1 & \text{if } s \in \mathcal{S}_{\mathrm{n}} \wedge s' = s \\ 0 & \text{otherwise} \end{cases}$$

- The function is then lifted to a set C of states by letting $\mathit{prob}(s, C) = \sum_{s' \in C} \mathit{prob}(s, s')$.

Weak Probabilistic Bisimilarity

- Weak bisimilarity $\approx_{\rm w}$ was introduced by Milner (1989) to abstract from the unobservable action τ.
- $\Longrightarrow$ is a finite sequence of alternating $\xrightarrow{\tau}_{\rm a}$ and $\xrightarrow{p}_{\rm p}$.
- $\stackrel{\hat{a}}{\Longrightarrow}$ is $\Longrightarrow$ if $a = \tau$, $\Longrightarrow \xrightarrow{a}_{\rm a} \Longrightarrow$ otherwise.

Definition

$s_1 \approx_{\rm p} s_2$ iff $(s_1, s_2) \in \mathcal{B}$ for some weak probabilistic bisimulation $\mathcal{B}$.
An equivalence relation $\mathcal{B}$ over $\mathcal{S}$ is a weak probabilistic bisimulation iff, whenever $(s_1, s_2) \in \mathcal{B}$, then:

- For each $s_1 \xrightarrow{a}_{\rm a} s'_1$ there exists $s_2 \stackrel{\hat{a}}{\Longrightarrow} s'_2$ with $(s'_1, s'_2) \in \mathcal{B}$.
- $prob(s_1, C) = prob(s_2, C)$ for all equivalence classes $C \in \mathcal{S}/\mathcal{B}$.

- By restricting the definition to nondeterministic states and ignoring *prob* we obtain $\approx_{\rm w}$.

Probabilistic Branching Bisimilarity

- Branching bisimilarity $\approx_{\mathrm{b}}$ was introduced by Van Glabbeek and Wejland (1996) as a refinement of weak bisimilarity.
- A probabilistic variant for the non-strictly alternating model was introduced by Andova, Georgievska, and Trčka (2012).

Definition

$s_1 \approx_{\mathrm{pb}} s_2$ iff $(s_1, s_2) \in \mathcal{B}$ for some probabilistic branching bisimulation $\mathcal{B}$.
An equivalence relation $\mathcal{B}$ over $\mathcal{S}$ is a probabilistic branching bisimulation iff, whenever $(s_1, s_2) \in \mathcal{B}$, then:

- For each $s_1 \xrightarrow{a}_{\mathrm{a}} s_1'$:
  - either $a = \tau$ and $(s_1', s_2) \in \mathcal{B}$;
  - or there exist $s_2 \Longrightarrow \bar{s}_2 \xrightarrow{a}_{\mathrm{a}} s_2'$ with $(s_1, \bar{s}_2) \in \mathcal{B}$ and $(s_1', s_2') \in \mathcal{B}$.
- *prob*$(s_1, C) =$ *prob*(s_2, C) for all equivalence classes $C \in \mathcal{S}/\mathcal{B}$.

- By restricting the definition to nondeterministic states and ignoring *prob* we obtain $\approx_{\mathrm{b}}$.

Process Language: High and Low Actions

- Two sets of actions for multi-level security systems:
  - High level actions: $\mathcal{A}_{\mathcal{H}}$.
  - Low level actions: $\mathcal{A}_{\mathcal{L}}$.
- Set of visible actions: $\mathcal{A} := \mathcal{A}_{\mathcal{H}} \cup \mathcal{A}_{\mathcal{L}}$.
- Overall set of actions: $\mathcal{A}_{\tau} := \mathcal{A} \cup \{\tau\}$.

Process Language: Nondeterministic Processes

- The overall set of process terms is $\mathbb{P} = \mathbb{P}_{\mathrm{n}} \cup \mathbb{P}_{\mathrm{p}}$.
- The set of nondeterministic process terms $\mathbb{P}_{\mathrm{n}}$ is the following where $a \in \mathcal{A}_\tau$ and $L \subseteq \mathcal{A}$:

$$
\begin{array}{rclll}
N & ::= & \underline{0} & & \text{terminated process} \\
 & | & a \,.\, P & & \text{action prefix} \\
 & | & N_1 + N_2 & & \text{nondeterministic choice} \\
 & | & N_1 \parallel_L N_2 & & \text{parallel composition} \\
 & | & N \setminus L & & \text{restriction} \\
 & | & N \,/\, L & & \text{hiding}
\end{array}
$$

- The set of probabilistic process terms $\mathbb{P}_{\mathrm{p}}$ is the following:

$$\begin{array}{rcll} P & ::= & \bigoplus_{i \in I} [p_i] N_i & \text{probabilistic choice} \\ & \mid & P_1 \parallel_L P_2 & \text{parallel composition} \\ & \mid & P \setminus L & \text{restriction} \\ & \mid & P \,/\, L & \text{hiding} \end{array}$$

Process Language: Probabilistic Processes

- The set of probabilistic process terms $\mathbb{P}_{\mathrm{p}}$ is the following:

$$
\begin{array}{rclll}
P & ::= & \bigoplus_{i \in I} [p_i] N_i & \text{probabilistic choice} \\
& | & P_1 \parallel_L P_2 & \text{parallel composition} \\
& | & P \setminus L & \text{restriction} \\
& | & P \,/\, L & \text{hiding}
\end{array}
$$

- $\bigoplus_{i \in I} [p_i]\, \_$ is the generalized probabilistic composition operator expressing a probabilistic choice among finitely many processes each with probability $p_i \in \mathbb{R}_{]0,1]}$ and such that $\sum_{i \in I} p_i = 1$.

Operational Semantic Rules: Nondeterministic Processes

- Operational semantic rule for action prefix:

$$\boxed{a \,.\, P \xrightarrow{a}_{\mathrm{a}} P}$$

- Operational semantic rules for nondeterministic choice:

$$\boxed{\frac{N_1 \xrightarrow{a}_{\mathrm{a}} P_1}{N_1 + N_2 \xrightarrow{a}_{\mathrm{a}} P_1} \quad \frac{N_2 \xrightarrow{a}_{\mathrm{a}} P_2}{N_1 + N_2 \xrightarrow{a}_{\mathrm{a}} P_2}}$$

Operational Semantic Rules: Nondeterministic Processes

- Operational semantic rules for parallel composition:

$$\frac{N_1 \xrightarrow{a}_{\mathrm{a}} P_1 \quad a \notin L}{N_1 \parallel_L N_2 \xrightarrow{a}_{\mathrm{a}} P_1 \parallel_L [1] N_2} \qquad \frac{N_2 \xrightarrow{a}_{\mathrm{a}} P_2 \quad a \notin L}{N_1 \parallel_L N_2 \xrightarrow{a}_{\mathrm{a}} [1] N_1 \parallel_L P_2}$$

- Operational semantic rule for synchronization:

$$\frac{N_1 \xrightarrow{a}_{\mathrm{a}} P_1 \quad N_2 \xrightarrow{a}_{\mathrm{a}} P_2 \quad a \in L}{N_1 \parallel_L N_2 \xrightarrow{a}_{\mathrm{a}} P_1 \parallel_L P_2}$$

- Operational semantic rules for restriction and hiding:

$$\frac{N \xrightarrow{a}_{\mathrm{a}} P \quad a \notin L}{N \setminus L \xrightarrow{a}_{\mathrm{a}} P \setminus L}$$

$$\frac{N \xrightarrow{a}_{\mathrm{a}} P \quad a \in L}{N \,/\, L \xrightarrow{\tau}_{\mathrm{a}} P \,/\, L} \qquad \frac{N \xrightarrow{a}_{\mathrm{a}} P \quad a \notin L}{N \,/\, L \xrightarrow{a}_{\mathrm{a}} P \,/\, L}$$

- Operational semantic rule for probabilistic choice:

$$\boxed{\frac{j \in I}{\bigoplus_{i \in I} [p_i] N_i \xrightarrow{p_j}_{\mathrm{p}} N_j}}$$

- Operational semantic rule for parallel composition:

$$\boxed{\frac{P_1 \xrightarrow{p_1}_{\mathrm{p}} N_1 \quad P_2 \xrightarrow{p_2}_{\mathrm{p}} N_2}{P_1 \parallel_L P_2 \xrightarrow{p_1 \cdot p_2}_{\mathrm{p}} N_1 \parallel_L N_2}}$$

- Operational semantic rules for restriction and hiding:

$$\frac{P \xrightarrow{p}_{\mathrm{p}} N}{P \setminus L \xrightarrow{p}_{\mathrm{p}} N \setminus L}$$

$$\frac{P \xrightarrow{p}_{\mathrm{p}} N}{P \,/\, L \xrightarrow{p}_{\mathrm{p}} N \,/\, L}$$

Nondeterministic Noninterference

- Whenever a group of agents at the high security level performs some actions, the effect of those actions should not be seen by any agent at the low security level.
- We recall some bisimilarity-based noninterference properties.
- Focardi and Gorrieri (2001) provided a characterization of these properties by employing weak bisimilarity in a nondeterministic process algebraic framework, resulting in a study of their features and comparisons between them.
- In [EAB23] we extended their approach to reversible systems by recasting the same properties with branching bisimilarity.
- We provide a further extension by recasting the properties with probabilistic bisimilarities.

Bisimulation-Based Properties

- The first property we examine is the *Bisimulation-based Strong Nondeterministic Non Interference* (BSNNI).
- It is satisfied by any process that behaves the same when its high-level actions are forbidden or hidden.

Definition

Let $E \in \mathbb{P}$ and $\approx$ a weak bisimilarity.
$E \in \mathrm{BSNNI}_{\approx} \iff E \setminus \mathcal{A}_{\mathcal{H}} \approx E \,/\, \mathcal{A}_{\mathcal{H}}$.

Bisimulation-Based Properties

- The first property we examine is the *Bisimulation-based Strong Nondeterministic Non Interference* (BSNNI).
- It is satisfied by any process that behaves the same when its high-level actions are forbidden or hidden.

Definition

Let $E \in \mathbb{P}$ and $\approx$ a weak bisimilarity.
$E \in \mathrm{BSNNI}_{\approx} \iff E \setminus \mathcal{A}_{\mathcal{H}} \approx E \,/\, \mathcal{A}_{\mathcal{H}}$.

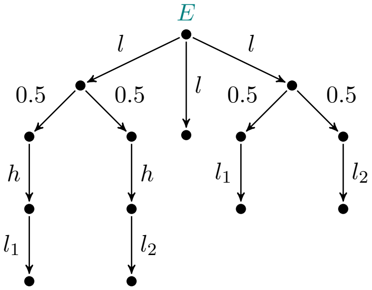

Bisimulation-Based Properties

- The first property we examine is the *Bisimulation-based Strong Nondeterministic Non Interference* (BSNNI).
- It is satisfied by any process that behaves the same when its high-level actions are forbidden or hidden.

Definition

Let $E \in \mathbb{P}$ and $\approx$ a weak bisimilarity.
$E \in \mathrm{BSNNI}_{\approx} \iff E \setminus \mathcal{A}_{\mathcal{H}} \approx E \,/\, \mathcal{A}_{\mathcal{H}}$.

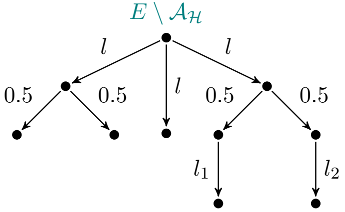

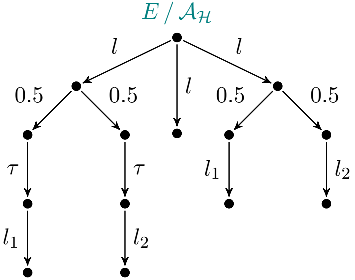

Bisimulation-Based Properties

- BSNNI is not powerful enough to capture covert channels that derive from the behavior of high-level agents interacting with the system, so other stronger properties have been studied in the literature.
- *Non Deducibility on Composition* (BNDC) requires to check the interaction between the system and every possible high-level agent.
- *Strong* BSNNI (SBSNNI) requires that at any reachable state the property BSNNI must be satisfied.
- *Persistent* BNDC (P_BNDC) requires that at any reachable state the property BNDC must be satisfied.
- *Strong BNDC* (SBNDC) requires that the low-level view of every reachable state of a system must be the same before and after the execution of every high level action.

Bisimulation-Based Properties

Definition

Let $E \in \mathbb{P}$ and $\approx$ a weak bisimilarity:

- $E \in \mathrm{BSNNI}_{\approx} \iff E \setminus \mathcal{A}_{\mathcal{H}} \approx E \,/\, \mathcal{A}_{\mathcal{H}}$.
- $E \in \mathrm{BNDC}_{\approx} \iff$ for all $F \in \mathbb{P}$ such that every $F' \in \mathit{reach}(F)$ can execute only actions in $\mathcal{A}_{\mathcal{H}}$ and for all $L \subseteq \mathcal{A}_{\mathcal{H}}$, $E \setminus \mathcal{A}_{\mathcal{H}} \approx ((E \parallel_L F) \,/\, L) \setminus \mathcal{A}_{\mathcal{H}}$.
- $E \in \mathrm{SBSNNI}_{\approx} \iff$ for all $E' \in \mathit{reach}(E)$, $E' \in \mathrm{BSNNI}_{\approx}$.
- $E \in \mathrm{P\_BNDC}_{\approx} \iff$ for all $E' \in \mathit{reach}(E)$, $E' \in \mathrm{BNDC}_{\approx}$.
- $E \in \mathrm{SBNDC}_{\approx} \iff$ for all $E' \in \mathit{reach}(E)$ for all E'' such that $E' \xrightarrow{a}_{\mathrm{a}} E''$ for some $a \in \mathcal{A}_{\mathcal{H}}$, $E' \setminus \mathcal{A}_{\mathcal{H}} \approx E'' \setminus \mathcal{A}_{\mathcal{H}}$.

Relation among Properties

- Focardi and Gorrieri showed the following taxonomy:

$$\mathrm{SBNDC}_{\approx_\mathrm{w}} \longrightarrow \mathrm{SBSNNI}_{\approx_\mathrm{w}} \longrightarrow \mathrm{BNDC}_{\approx_\mathrm{w}} \longrightarrow \mathrm{BSNNI}_{\approx_\mathrm{w}}$$

- Later on, $\mathrm{P\_BNDC}_{\approx_\mathrm{w}}$ was introduced by Focardi and Rossi (2006) and proven to be equivalent to $\mathrm{SBSNNI}_{\approx_\mathrm{w}}$.

Nondeterministic Taxonomy

- In [EAB23] branching bisimilarity has been used to recast the nonintenference properties and extend the taxonomy:

$$
\begin{tikzcd}
\mathrm{SBNDC}_{\approx_b} \arrow[r] \arrow[dr] & \mathrm{SBNDC}_{\approx_w} \arrow[d, leftrightarrow, "/" marking] \arrow[dr] & & & \\
& \begin{array}{l}\mathrm{SBSNNI}_{\approx_b}\\ \mathrm{P\_BNDC}_{\approx_b}\end{array} \arrow[r] \arrow[dr] & \begin{array}{l}\mathrm{SBSNNI}_{\approx_w}\\ \mathrm{P\_BNDC}_{\approx_w}\end{array} \arrow[d, leftrightarrow, "/" marking] \arrow[dr] & & \\
& & \mathrm{BNDC}_{\approx_b} \arrow[r] \arrow[dr] & \mathrm{BNDC}_{\approx_w} \arrow[d, leftrightarrow, "/" marking] \arrow[dr] & \\
& & & \mathrm{BSNNI}_{\approx_b} \arrow[r] & \mathrm{BSNNI}_{\approx_w}
\end{tikzcd}
$$

Preservation

- By recasting noninterference properties using $\approx_{\mathrm{p}}$ and $\approx_{\mathrm{pb}}$ we can study their features and characteristics.
- $\approx_{\mathrm{p}}$ and $\approx_{\mathrm{pb}}$ preserve all the five properties.

Theorem

Let $E_1, E_2 \in \mathbb{P}$, $\approx \in \{\approx_{\mathrm{p}}, \approx_{\mathrm{pb}}\}$, *and* $\mathcal{P} \in \{\mathrm{BSNNI}_{\approx}, \mathrm{BNDC}_{\approx}, \mathrm{SBSNNI}_{\approx}, \mathrm{P\_BNDC}_{\approx}, \mathrm{SBNDC}_{\approx}\}$.

$$\text{If } E_1 \approx E_2, \text{ then } E_1 \in \mathcal{P} \iff E_2 \in \mathcal{P}.$$

- This is very useful in automated property verification as it can be more convenient to work with a reduced system, i.e., a system equivalent to the one we are checking but with a smaller state space.

Compositionality

- The stronger properties are preserved by (most of) the operators of $\mathbb{P}$.

Theorem

Let $E, E_1, E_2 \in \mathbb{P}$, $\approx \in \{\approx_{\mathrm{p}}, \approx_{\mathrm{pb}}\}$, $\mathcal{P} \in \{\mathrm{SBSNNI}_{\approx}, \mathrm{P\_BNDC}_{\approx} \mathrm{SBNDC}_{\approx}\}$. *Then:*

1. $E \in \mathcal{P} \Longrightarrow a\,.\,E \in \mathcal{P}$ *for all* $a \in \mathcal{A}_{\mathcal{L}} \cup \{\tau\}$ *and* $E \in \mathbb{P}_{\mathrm{p}}$.
2. $E_1, E_2 \in \mathcal{P} \Longrightarrow E_1 \parallel_L E_2 \in \mathcal{P}$ *for all* $L \subseteq \mathcal{A}_{\mathcal{L}}$ *if* $\mathcal{P} \in \{\mathrm{SBSNNI}_{\approx_{\mathrm{pb}}}, \mathrm{P\_BNDC}_{\approx_{\mathrm{pb}}}\}$,
   $L \subseteq \mathcal{A}$ *if* $\mathcal{P} \in \{\mathrm{SBSNNI}_{\approx_{\mathrm{p}}}, \mathrm{P\_BNDC}_{\approx_{\mathrm{p}}}, \mathrm{SBNDC}_{\approx_{\mathrm{p}}}, \mathrm{SBNDC}_{\approx_{\mathrm{pb}}}\}$.
3. $E \in \mathcal{P} \Longrightarrow E \setminus L \in \mathcal{P}$ *for all* $L \subseteq \mathcal{A}$.
4. $E \in \mathcal{P} \Longrightarrow E \,/\, L \in \mathcal{P}$ *for all* $L \subseteq \mathcal{A}_{\mathcal{L}}$.

Extended Probabilistic Taxonomy

- Taxonomy of security properties based on weak and branching probabilistic bisimilarities:

$$
\begin{tikzcd}
\mathrm{SBNDC}_{\approx_{\mathrm{pb}}} \arrow[r] \arrow[dr] & \mathrm{SBNDC}_{\approx_{\mathrm{p}}} \arrow[dr] \arrow[d, leftrightarrow, "\times" marking] & & & \\
& \begin{array}{l}\mathrm{SBSNNI}_{\approx_{\mathrm{pb}}} \\ \mathrm{P\_BNDC}_{\approx_{\mathrm{pb}}}\end{array} \arrow[r] \arrow[dr] & \begin{array}{l}\mathrm{SBSNNI}_{\approx_{\mathrm{p}}} \\ \mathrm{P\_BNDC}_{\approx_{\mathrm{p}}}\end{array} \arrow[dr] \arrow[d, leftrightarrow, "\times" marking] & & \\
& & \mathrm{BNDC}_{\approx_{\mathrm{pb}}} \arrow[r] \arrow[dr] & \mathrm{BNDC}_{\approx_{\mathrm{p}}} \arrow[dr] \arrow[d, leftrightarrow, "\times" marking] & \\
& & & \mathrm{BSNNI}_{\approx_{\mathrm{pb}}} \arrow[r] & \mathrm{BSNNI}_{\approx_{\mathrm{p}}}
\end{tikzcd}
$$

Relating Nondeterministic and Probabilistic Taxonomies

- Given a process $E \in \mathbb{P}$, we can obtain its nondet. variant $nd(E)$.
- We replace each $\bigoplus_{i \in I} [p_i] E_i$ with $\sum_{i \in I} \tau \, . \, E_i$.

Theorem

Let $E_1, E_2 \in \mathbb{P}$. Then:

- $E_1 \approx_{\mathrm{p}} E_2 \implies nd(E_1) \approx_{\mathrm{w}} nd(E_2)$.
- $E_1 \approx_{\mathrm{pb}} E_2 \implies nd(E_1) \approx_{\mathrm{b}} nd(E_2)$.

Relating Nondeterministic and Probabilistic Taxonomies

- Given a process $E \in \mathbb{P}$, we can obtain its nondet. variant $nd(E)$.
- We replace each $\bigoplus_{i \in I} [p_i] E_i$ with $\sum_{i \in I} \tau \,.\, E_i$.

Theorem

Let $E_1, E_2 \in \mathbb{P}$. Then:

- $E_1 \approx_{\mathrm{p}} E_2 \implies nd(E_1) \approx_{\mathrm{w}} nd(E_2)$.
- $E_1 \approx_{\mathrm{pb}} E_2 \implies nd(E_1) \approx_{\mathrm{b}} nd(E_2)$.

- The inverse is not true.

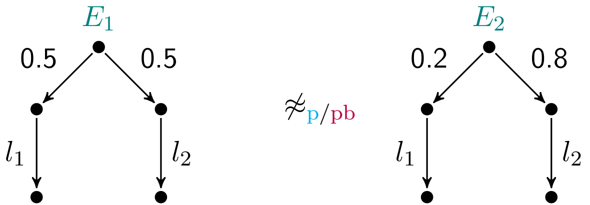

- Given a process $E \in \mathbb{P}$, we can obtain its nondet. variant $nd(E)$.
- We replace each $\bigoplus_{i \in I} [p_i] E_i$ with $\sum_{i \in I} \tau \,.\, E_i$.

Theorem

Let $E_1, E_2 \in \mathbb{P}$. Then:

- $E_1 \approx_{\mathrm{p}} E_2 \Longrightarrow nd(E_1) \approx_{\mathrm{w}} nd(E_2)$.
- $E_1 \approx_{\mathrm{pb}} E_2 \Longrightarrow nd(E_1) \approx_{\mathrm{b}} nd(E_2)$.

- The inverse is not true.

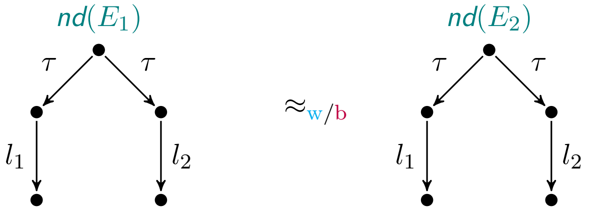

Relating Nondeterministic and Probabilistic Taxonomies

- A consequence is that if a process E is secure under a probabilistic noninterference property, then $nd(E)$ is secure under the corresponding nondeterministic property.

Corollary

Let $E \in \mathbb{P}, \approx_{\mathrm{pr}} \in \{\approx_{\mathrm{p}}, \approx_{\mathrm{pb}}\}, \approx_{\mathrm{nd}} \in \{\approx_{\mathrm{w}}, \approx_{\mathrm{b}}\}$,
$\mathcal{P}_{\mathrm{pr}} \in \{\mathrm{BSNNI}_{\approx_{\mathrm{pr}}}, \mathrm{BNDC}_{\approx_{\mathrm{pr}}}, \mathrm{SBSNNI}_{\approx_{\mathrm{pr}}}, \mathrm{P\_BNDC}_{\approx_{\mathrm{pr}}}, \mathrm{SBNDC}_{\approx_{\mathrm{pr}}}\}$,
$\mathcal{P}_{\mathrm{nd}} \in \{\mathrm{BSNNI}_{\approx_{\mathrm{nd}}}, \mathrm{BNDC}_{\approx_{\mathrm{nd}}}, \mathrm{SBSNNI}_{\approx_{\mathrm{nd}}}, \mathrm{P\_BNDC}_{\approx_{\mathrm{nd}}}, \mathrm{SBNDC}_{\approx_{\mathrm{nd}}}\}$.
Then:

$$E \in \mathcal{P}_{\mathrm{pr}} \implies nd(E) \in \mathcal{P}_{\mathrm{nd}}$$

- This means that our results further extend the nondeterministic taxonomy.

Back-and-Forth Bisimilarities

- Introduced by De Nicola, Montanari, and Vaandraager (1990).
- Back-and-forth bisimulations are defined over *computational paths* instead of states.
- This is needed to remain in an interleaving setting of concurrency.
- It preserves not only causality but also history.
- Whenever a process returns to a past state it must do it by reverting the same computational path performed in going forward.
- In the nondeterministic setting, weak back-and-forth bisimilarity is finer than weak bisimilarity, and coincides with branching bisimilarity.

- The bisimulation is defined over the set of computational paths $\mathcal{U}$ instead of the set of states $\mathcal{S}$.

Definition

$s_1 \approx_{\mathrm{pbf}} s_2$ iff $((s_1, \varepsilon), (s_2, \varepsilon)) \in \mathcal{B}$ for some weak probabilistic back-and-forth bisimulation $\mathcal{B}$.
An equivalence relation $\mathcal{B}$ over $\mathcal{U}$ is a weak probabilistic back-and-forth bisimulation iff, whenever $(\rho_1, \rho_2) \in \mathcal{B}$, then:

- For each $\rho_1 \xrightarrow{a}_{\mathrm{a}} \rho_1'$ there exists $\rho_2 \overset{\hat{a}}{\Longrightarrow} \rho_2'$ with $(\rho_1', \rho_2') \in \mathcal{B}$.
- For each $\rho_1' \xrightarrow{a}_{\mathrm{a}} \rho_1$ there exists $\rho_2' \overset{\hat{a}}{\Longrightarrow} \rho_2$ with $(\rho_1', \rho_2') \in \mathcal{B}$.
- $\mathit{prob}(\rho_1, C) = \mathit{prob}(\rho_2, C)$ for all equivalence classes $C \in \mathcal{U}/\mathcal{B}$.

Comparisons

- As in the nondeterministic case, weak probabilistic back-and-forth bisimilarity coincides with probabilistic branching bisimilarity.

Theorem

$$s_1 \approx_{\mathrm{pbf}} s_2 \text{ iff } s_1 \approx_{\mathrm{pb}} s_2.$$

- Therefore:
  - We can reason about reversible systems without resorting to a reversible calculus nor a path-based equivalence.
  - All the results for probabilistic branching-bisimulation-based properties can be extended to probabilistic reversible systems.

Case Study: Probabilistic Smart Contract

- Consider a lottery implemented in a probabilistic smart contract.
- Anyone can buy a ticket.
- When the lottery is closed, anyone can invoke another smart contract function, draw(), in which a random number x, between 1 and the amount of sold tickets, is drawn and the entire money is paid to the owner of the extracted value x.
- We will examine two vulnerabilities.
  - The first one emphasizes the need for passing from the nondeterministic noninterference to the probabilistic one.
  - The second one emphasizes the difference between $\approx_{\mathrm{p}}$ and $\approx_{\mathrm{pb}}$ when dealing with reversibility.

- In the first case the critical point is the randomization process of the function draw(), not natively available to smart contract programmer.
- A widely adopted approach consists of using the timestamp of the block including the transaction of the draw invocation as the seed for random number generation.
- A malicious participant can mine the block above and manipulate the timestamp to win the lottery.
- We consider the following transitions:
  - h which represent the interaction of a malicious miner.
  - $draw$ expressing the invocation of the draw() function.
  - add_i expressing the determination of the winner.
  - win_i expressing the notification of the winner.
- For simplicity, we consider a lottery with only two participants.

Case Study: Probabilistic Smart Contract

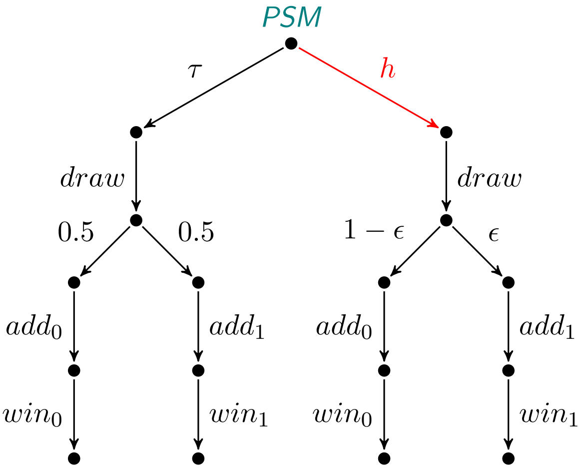

Case Study: Probabilistic Smart Contract

- The processes $PSM \setminus \mathcal{A}_{\mathcal{H}}$ and $PSM \,/\, \mathcal{A}_{\mathcal{H}}$ are not $\approx_{\mathrm{p}/\mathrm{pb}}$.

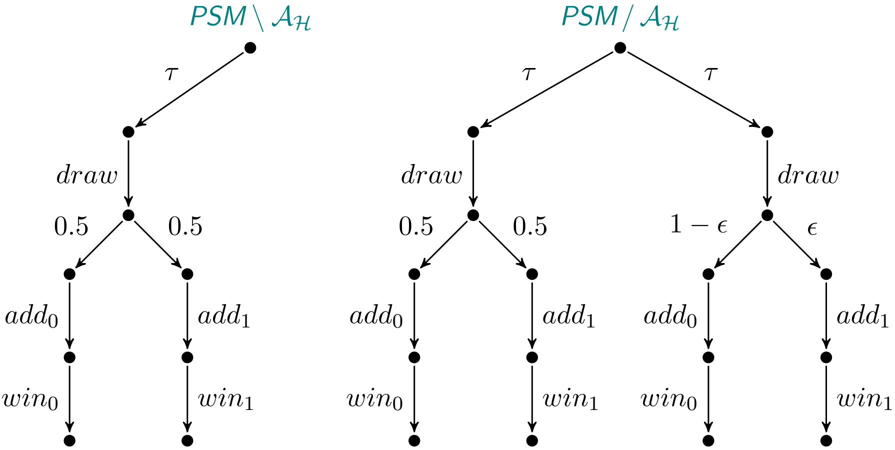

Case Study: Probabilistic Smart Contract

- But the processes $nd(PSM \setminus \mathcal{A}_{\mathcal{H}})$ and $nd(PSM \,/\, \mathcal{A}_{\mathcal{H}})$ are $\approx_{\mathrm{w}/\mathrm{b}}$.

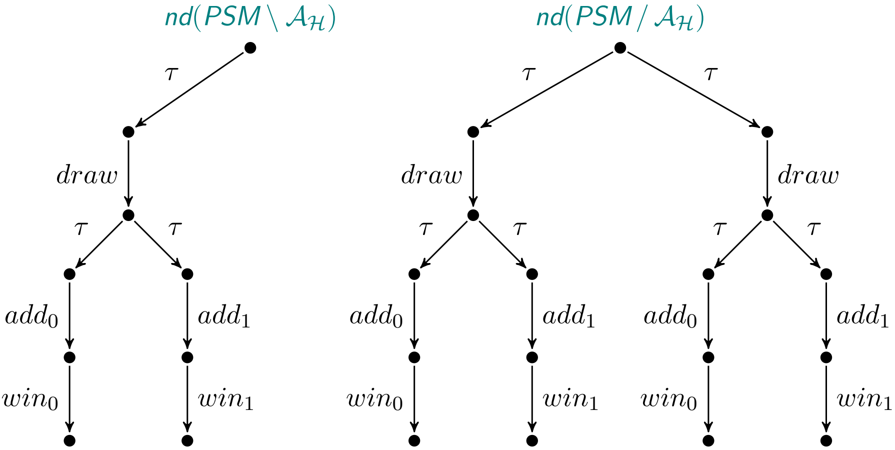

Case Study: Probabilistic Smart Contract

- In the second case the critical point is the mining procedure.
- The seed governing the probabilistic extraction cannot be manipulated.
- A malicious miner invokes the function draw() but is going to lose.
- He can force the mining failure and a rollback of the lottery.
- We add the following transitions:
  - $mine$ expressing the mining of a block, by either an honest or dishonest miner.
  - $succ$ expressing the successful termination of the mining.
  - $fail$ expressing the failed termination of the mining, it can either be forced or occur for other reasons (a wrong transaction in the block or a fork in the blockchain).

Case Study: Probabilistic Smart Contract

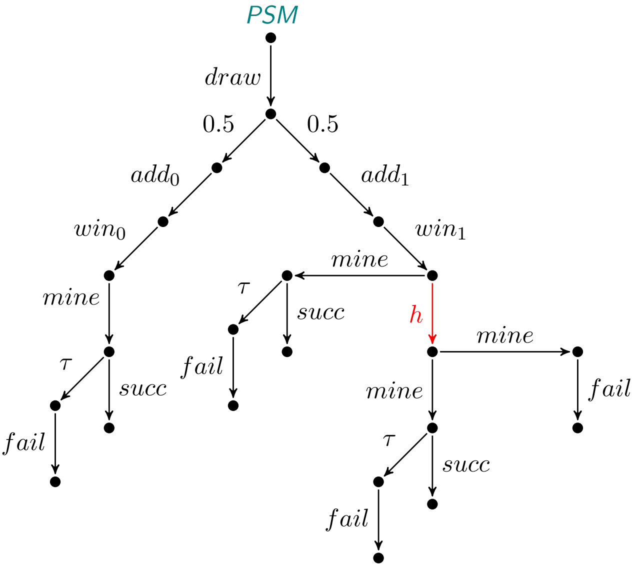

Case Study: Probabilistic Smart Contract

- The processes $nd(PSM \setminus \mathcal{A}_{\mathcal{H}})$ and $nd(PSM \,/\, \mathcal{A}_{\mathcal{H}})$ are $\approx_{\mathrm{p}}$ but not $\approx_{\mathrm{pb}}$.

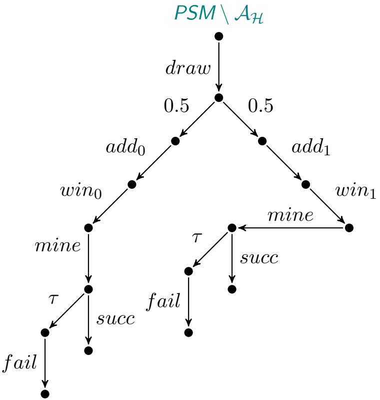

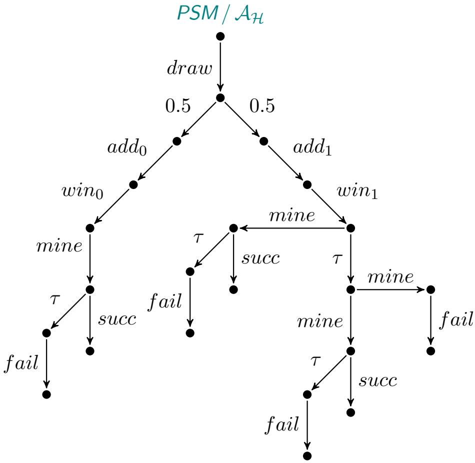

Case Study: Probabilistic Smart Contract

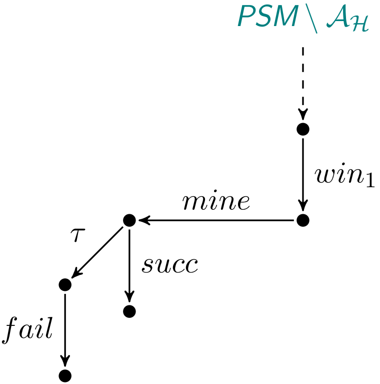

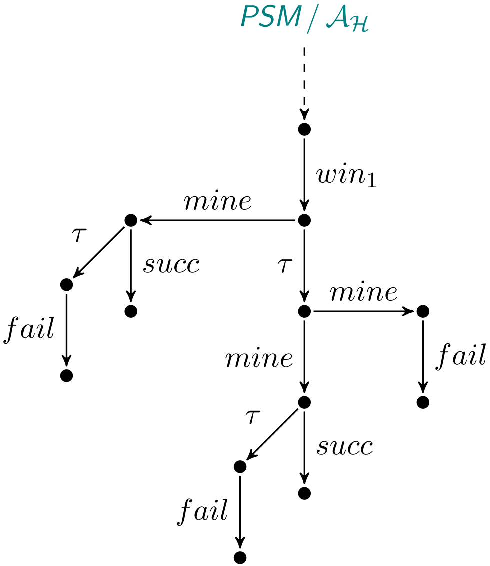

Case Study: Probabilistic Smart Contract

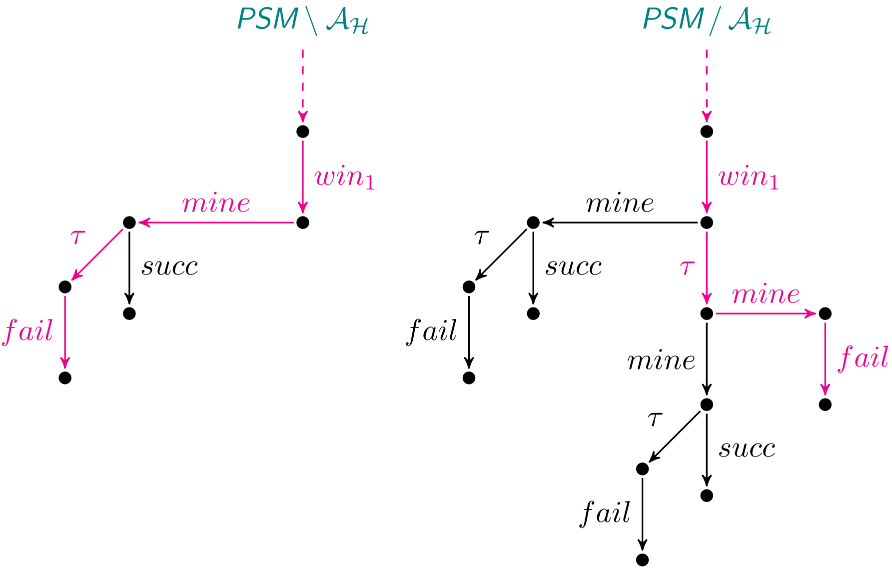

Case Study: Probabilistic Smart Contract

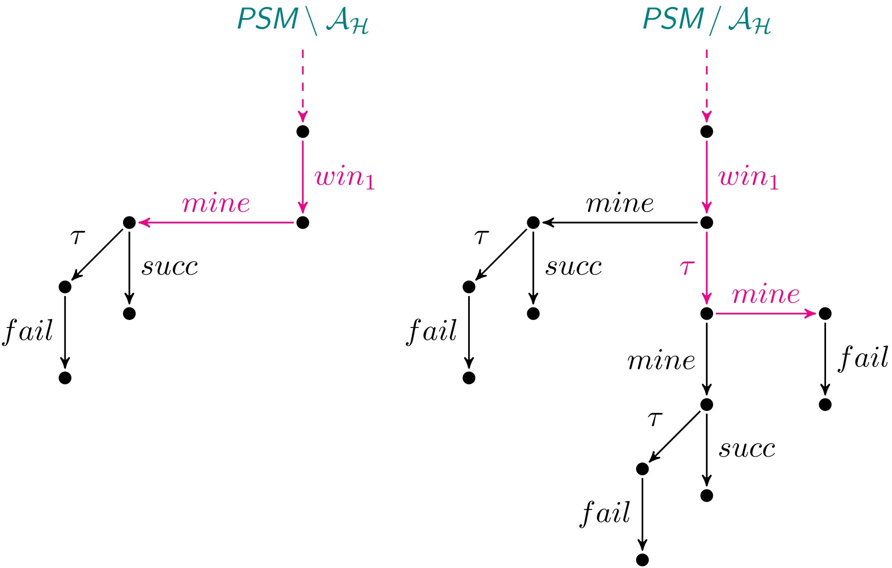

Conclusions

- We have recast a variety of noninterference properties in a probabilistic setting, studying their features and taxonomy.
- Potential covert channels arising in probabilistic reversible systems cannot be revealed by employing weak probabilistic bisimulation.
- Indeed, the higher discriminating power of probabilistic branching bisimilarity is necessary to capture information flows emerging whenever backward computations are activated.
- Since some proofs required the representation of processes as trees, we could not include recursion in our language.
- As future work we plan to find alternative proof techniques to add recursion.