

Making Random Choices Invisible to the Scheduler^{*}

Konstantinos Chatzikokolakis Catuscia Palamidessi
INRIA and LIX, École Polytechnique, Palaiseau, France
{kostas,catuscia}@lix.polytechnique.fr

Abstract. When dealing with process calculi and automata which express both nondeterministic and probabilistic behavior, it is customary to introduce the notion of scheduler to resolve the nondeterminism. It has been observed that for certain applications, notably those in security, the scheduler needs to be restricted so not to reveal the outcome of the protocol’s random choices, or otherwise the model of adversary would be too strong even for “obviously correct” protocols. We propose a process-algebraic framework in which the control on the scheduler can be specified in syntactic terms, and we show how to apply it to solve the problem mentioned above. We also consider the definition of (probabilistic) may and must preorders, and we show that they are precongruences with respect to the restricted schedulers. Furthermore, we show that all the operators of the language, except replication, distribute over probabilistic summation, which is a useful property for verification.

1 Introduction

Security protocols, in particular those for anonymity and fair exchange, often use randomization to achieve their targets. Since they usually involve more than one agent, they also give rise to concurrent and interactive activities that can be best modeled by nondeterminism. Thus it is convenient to specify them using a formalism which is able to represent both *probabilistic* and *nondeterministic* behavior. Formalisms of this kind have been explored in both Automata Theory [1–5] and in Process Algebra [6–11]. See also [12, 13] for comparative and more inclusive overviews.

Due to the presence of nondeterminism, in such formalisms it is not possible to define the probability of events in *absolute* terms. We need first to decide how each nondeterministic choice during the execution will be resolved. This decision function is called *scheduler*. Once the scheduler is fixed, the behavior of the system (*relatively* to the given scheduler) becomes fully probabilistic and a probability measure can be defined following standard techniques.

It has been observed by several researchers that in security the notion of scheduler needs to be restricted, or otherwise any secret choice of the protocol could be revealed by making the choice of the scheduler depend on it. This issue

^{*} This work has been partially supported by the INRIA DREI Équipe Associée PRINTEMPS and by the INRIA ARC project ProNoBiS.

was for instance one of the main topics of discussion at the panel of CSFW 2006. We illustrate it here with an example on anonymity. We use the standard CCS notation, plus a construct of probabilistic choice $P +_p Q$ representing a process that evolves into P with probability p and into Q with probability $1 - p$.

The following system Sys consists of one receiver R and two senders S, T which communicate via private channels a, b respectively. Which of the two senders is successful is decided probabilistically by R . After reception, R sends a signal ok .

$$R \triangleq a.\overline{ok}.0 +_{0.5} b.\overline{ok}.0 \quad S \triangleq \bar{a}.0 \quad T \triangleq \bar{b}.0 \quad Sys \triangleq (\nu a)(\nu b)(R \mid S \mid T)$$

The signal ok is not private, but since it is the same in both cases, in principle an external observer should not be able to infer from it the identity of the sender (S or T). So the system should be anonymous. However, consider a team of two attackers A and B defined as

$$A \triangleq ok.\bar{s}.0 \quad B \triangleq ok.\bar{t}.0$$

and consider the parallel composition $Sys \mid A \mid B$. We have that, under certain schedulers, the system is no longer anonymous. More precisely, a scheduler could leak the identity of the sender via the channels s, t by forcing R to synchronize with A on ok if R has chosen the first alternative, and with B otherwise. This is because in general a scheduler can see the whole history of the computation, in particular the random choices, even those which are supposed to be private. Note that the visibility of the synchronization channels to the scheduler is not crucial for this example: we would have the same problem, for instance, if S, T were both defined as $\bar{a}.0$, R as $a.\overline{ok}.0$, and Sys as $(\nu a)((S +_{0.5} T) \mid R)$.

The above example demonstrates that, with the standard definition of scheduler, it is not possible to represent a truly private random choice (or a truly private nondeterministic choice, for the matter) with the current probabilistic process calculi. This is a clear shortcoming when we want to use these formalisms for the specification and verification of security protocols.

There is another issue related to verification: a private choice has certain algebraic properties that would be useful in proving equivalences between processes. In fact, if the outcome of a choice remains private, then it should not matter at which point of the execution the process makes such choice, until it actually uses it. Consider for instance A and B defined as follows

$$A \triangleq a(x).([x = 0]\overline{ok} +_{0.5} [x = 1]\overline{ok}) \quad B \triangleq a(x).[x = 0]\overline{ok} +_{0.5} a(x).[x = 1]\overline{ok}$$

Process A receives a value and then decides randomly whether it will accept the value 0 or 1. Process B does exactly the same thing except that the choice is performed before the reception of the value. If the random choices in A and B are private, intuitively we should have that A and B are equivalent ($A \approx B$). This is because it should not matter whether the choice is done before or after receiving




Fig. 1. Execution trees for $A \mid C$ and $B \mid C$

a message, as long as the outcome of the choice is completely invisible to any other process or observer. However, consider the parallel context $C = \bar{a}0 \mid \bar{a}1$. Under any scheduler A has probability at most $1/2$ to perform \overline{ok} . With B , on the other hand, the scheduler can choose between $\bar{a}0$ and $\bar{a}1$ based on the outcome of the probabilistic choice, thus making the maximum probability of \overline{ok} equal to 1. The execution trees of $A \mid C$ and $B \mid C$ are shown in Figure 1.

In general when $+_p$ represents a private choice we would like to have

$$C[P +_p Q] \approx C[\tau.P] +_p C[\tau.Q] \quad (1)$$

for all processes P, Q and all contexts C *not containing replication (or recursion)*. In the case of replication the above cannot hold since $!(P +_p Q)$ makes available each time the choice between P and Q , while $!(\tau.P) +_p !(\tau.Q)$ chooses once and for all which of the two (P or Q) should be replicated. Similarly for recursion. The reason why we need a τ is explained in Section 5.

The algebraic property (1) expresses in an abstract way the privacy of the probabilistic choice. Moreover, this property is also useful for the verification of security properties. The interested reader can find in [14] an example of application to a fair exchange protocol. In principle (1) should be useful for any kind of verification in the process algebra style.

We propose a process-algebraic approach to the problem of hiding the outcome of random choices. Our framework is based on a calculus obtained by adding to CCS an internal probabilistic choice construct¹. This calculus, to which we refer as CCS_p , is a variant of the one studied in [11], the main differences being that we use replication instead than recursion, and we lift some restrictions that were imposed in [11] to obtain a complete axiomatization. The semantics of CCS_p is given in terms of Segala's *simple probabilistic automata* [4, 7].

In order to limit the power of the scheduler, we extend CCS_p with terms representing explicitly the notion of scheduler. The latter interact with the original processes via a labeling system. This will allow to specify at the syntactic level (by a suitable labeling) which choices should be visible to schedulers, and which ones should not.

¹ We actually consider a variant of CCS where recursion is replaced by replication. The two languages are not equivalent, but we believe that the issues regarding the differences between replication and recursion are orthogonal to the topics investigated in this paper.

1.1 Contribution

The main contributions of this paper are:

- A process calculus CCS_σ in which the scheduler is represented as a process, and whose power can therefore be controlled at the syntactic level.
- An application of CCS_σ to an extended anonymity example (the Dining Cryptographers Protocol, DCP). We also briefly outline how to extend CCS_σ so to allow the definition of private nondeterministic choice, and we apply it to the DCP with nondeterministic master. To our knowledge this is the first formal treatment of the scheduling problem in DCP and the first formalization of a nondeterministic master for the (probabilistic) DCP.
- The adaptation of the standard notions of probabilistic testing preorders to CCS_σ , and the “sanity check” that they are still precongruences with respect to all the operators except the nondeterministic sum. For the latter we have the problem that P and $\tau.P$ are must equivalent, but $Q + P$ and $Q + \tau.P$ are not. This is typical for the $\text{CCS} +$: usually it does not preserve weak equivalences.
- The proof that, under suitable conditions on the labelings of C , $\tau.P$ and $\tau.Q$, CCS_σ satisfies the property expressed by (1), where \approx is probabilistic testing equivalence.

1.2 Related work

The works that are most closely related to ours are [15–17]. The authors of [15, 16] consider probabilistic automata and introduce a restriction on the scheduler to the purpose of making them suitable to applications in security protocols. Their approach is based on dividing the actions of each component of the system in equivalence classes (*tasks*). The order of execution of different tasks is decided in advance by a so-called *task scheduler*. The remaining nondeterminism within a task is resolved by a second scheduler, which models the standard *adversarial scheduler* of the cryptographic community. This second entity has limited knowledge about the other components: it sees only the information that they communicate during execution.

Reference [17] defines a notion of admissible scheduler by introducing an equivalence relation on the nodes of the execution tree, and requiring that an admissible scheduler maps two equivalent nodes into bisimilar steps. Both our paper and [17] have developed, independently, the solution to the problem of the scheduler in the Dining Cryptographers as an example of application to security.

Another work along these lines is [18], which uses partitions on the state-space to obtain partial-information schedulers. However [18] considers a synchronous parallel composition, so the setting is rather different from [15–17] and ours.

Our approach is in a sense *dual* to the above ones. Instead of defining a restriction on the class of schedulers, we provide a way to specify that a choice is transparent to the schedulers. We achieve this by introducing labels in process terms, used to represent both the nodes of the execution tree and the next action

or step to be scheduled. We make two nodes indistinguishable to schedulers, and hence the choice between them private, by associating to them the same label. Furthermore, in contrast with [15, 16], our “equivalence classes” (schedulable actions with the same label) can change dynamically, because the same action can be associated to different labels during the execution. However we don’t know at the moment whether this difference determines a separation in the expressive power.

1.3 Plan of the paper

In the next section we briefly recall some basic notions. In Section 3 we define a preliminary version of the language CCS_σ and of the corresponding notion of scheduler. In Section 4 we compare our notion of scheduler with the more standard “semantic” notion, and we improve the definition of CCS_σ so to retrieve the full expressive power of the semantic schedulers. In Section 5 we study the probabilistic testing preorders, their compositionality properties, and the conditions under which (1) holds. Section 6 presents an application to security. Section 7 concludes.

2 Preliminaries

In this section we briefly recall some preliminary notions about the simple probabilistic automata and CCS_p .

2.1 Simple probabilistic automata [4, 7]

A *discrete probability measure* over a set X is a function $\mu : 2^X \mapsto [0, 1]$ such that $\mu(X) = 1$ and $\mu(\cup_i X_i) = \sum_i \mu(X_i)$ where X_i is a countable family of pairwise disjoint subsets of X . The set of all discrete probability measures over X will be denoted by $\text{Disc}(X)$. We will denote by $\delta(x), x \in X$ (called the *Dirac measure* on x) the probability measure that assigns probability 1 to $\{x\}$. We will also denote by $\sum_i [p_i] \mu_i$ the probability measure obtained as a convex sum of the measures μ_i .

A *simple probabilistic automaton*² is a tuple (S, q, A, \mathcal{D}) where S is a set of states, $q \in S$ is the *initial state*, A is a set of actions and $\mathcal{D} \subseteq S \times A \times \text{Disc}(S)$ is a *transition relation*. Intuitively, if $(s, a, \mu) \in \mathcal{D}$ then there is a transition from the state s performing the action a and leading to a distribution μ over the states of the automaton. The idea is that the choice of transition among the available ones in \mathcal{D} is performed nondeterministically, and the choice of the target state among the ones allowed by μ (i.e. those states q such that $\mu(q) > 0$) is performed probabilistically.

² For simplicity in the following we will refer to a simple probabilistic automaton as *probabilistic automaton*. Note however that simple probabilistic automata are a subset of the probabilistic automata defined in [4, 5].

A probabilistic automaton M is *fully probabilistic* if from each state of M there is at most one transition available. An execution α of a probabilistic automaton is a (possibly infinite) sequence $s_0 a_1 s_1 a_2 s_2 \dots$ of alternating states and actions, such that $q = s_0$, and for each i $(s_i, a_{i+1}, \mu_i) \in \mathcal{D}$ and $\mu_i(s_{i+1}) > 0$ hold. We will use $lstate(\alpha)$ to denote the last state of a finite execution α , and $exec^*(M)$ and $exec(M)$ to represent the set of all the finite and of all the executions of M , respectively.

A *scheduler* of a probabilistic automaton $M = (S, q, A, \mathcal{D})$ is a function

$$\zeta : exec^*(M) \mapsto \mathcal{D}$$

such that $\zeta(\alpha) = (s, a, \mu) \in \mathcal{D}$ implies that $s = lstate(\alpha)$. The idea is that a scheduler selects a transition among the ones available in \mathcal{D} and it can base his decision on the history of the execution. The *execution tree* of M relative to the scheduler ζ , denoted by $etree(M, \zeta)$, is a fully probabilistic automaton $M' = (S', q', A', \mathcal{D}')$ such that $S' \subseteq exec(M)$, $q' = q$, $A' = A$, and $(\alpha, a, \mu') \in \mathcal{D}'$ if and only if $\zeta(\alpha) = (lstate(\alpha), a, \mu)$ for some μ and $\mu'(\alpha a s) = \mu(s)$. Intuitively, $etree(M, \zeta)$ is produced by unfolding the executions of M and resolving all deterministic choices using ζ . Note that $etree(M, \zeta)$ is a simple³ and fully probabilistic automaton.

2.2 CCS with internal probabilistic choice

Let a range over a countable set of *channel names*. The syntax of CCS_p is the following:

$\alpha ::= a \mid \bar{a} \mid \tau$	prefixes
$P, Q ::=$	processes
$\alpha.P$	prefix
$P \mid Q$	parallel
$P + Q$	nondeterministic choice
$\sum_i p_i P_i$	internal probabilistic choice
$(\nu a)P$	restriction
$!P$	replication
0	nil

We will also use the notation $P_1 +_p P_2$ to represent a binary sum $\sum_i p_i P_i$ with $p_1 = p$ and $p_2 = 1 - p$.

The semantics of a CCS_p term is a probabilistic automaton defined inductively on the basis of the syntax according to the rules in Figure 2. We write $s \xrightarrow{a} \mu$ when (s, a, μ) is a transition of the probabilistic automaton. We also denote by $\mu \mid Q$ the measure μ' such that $\mu'(P \mid Q) = \mu(P)$ for all processes P and $\mu'(R) = 0$ if R is not of the form $P \mid Q$. Similarly $(\nu a)\mu = \mu'$ such that $\mu'((\nu a)P) = \mu(P)$.

³ This is true because we do not consider probabilistic schedulers. If we considered such schedulers then the execution tree would no longer be a simple automaton.

$$\begin{array}{ll}
\text{ACT} & \frac{}{\alpha.P \xrightarrow{\alpha} \delta(P)} \\
\text{RES} & \frac{P \xrightarrow{\alpha} \mu \quad \alpha \neq a, \bar{a}}{(\nu a)P \xrightarrow{\alpha} (\nu a)\mu} \\
\text{SUM1} & \frac{P \xrightarrow{\alpha} \mu}{P + Q \xrightarrow{\alpha} \mu} \\
\text{PAR1} & \frac{P \xrightarrow{\alpha} \mu}{P \mid Q \xrightarrow{\alpha} \mu \mid Q} \\
\text{COM} & \frac{P \xrightarrow{a} \delta(P') \quad Q \xrightarrow{\bar{a}} \delta(Q')}{P \mid Q \xrightarrow{\tau} \delta(P' \mid Q')} \\
\text{PROB} & \frac{}{\sum_i p_i P_i \xrightarrow{\tau} \sum_i [p_i] \delta(P_i)} \\
\text{REP1} & \frac{P \xrightarrow{\alpha} \mu}{!P \xrightarrow{\alpha} \mu \mid !P} \\
\text{REP2} & \frac{P \xrightarrow{a} \delta(P_1) \quad P \xrightarrow{\bar{a}} \delta(P_2)}{!P \xrightarrow{\tau} \delta(P_1 \mid P_2 \mid !P)}
\end{array}$$

Fig. 2. The semantics of CCS_p . SUM1 and PAR1 have corresponding right rules SUM2 and PAR2, omitted for simplicity.

A transition of the form $P \xrightarrow{a} \delta(P')$, i.e. a transition having for target a Dirac measure, corresponds to a transition of a non-probabilistic automaton (a standard labeled transition system). Thus, all the rules of CCS_p imitate the ones of CCS except from PROB. The latter models the internal probabilistic choice: a silent τ transition is available from the sum to a measure containing all of its operands, with the corresponding probabilities.

Note that in the produced probabilistic automaton, all transitions to non-Dirac measures are silent. This is similar to the *alternating model* [2], however our case is more general because the silent and non-silent transitions are not necessarily alternated. On the other hand, with respect to the simple probabilistic automata the fact that the probabilistic transitions are silent looks as a restriction. However, it has been proved by Bandini and Segala [7] that the simple probabilistic automata and the alternating model are essentially equivalent, so, being in the middle, our model is equivalent as well.

3 A variant of CCS with explicit scheduler

In this section we present a variant of CCS in which the scheduler is explicit, in the sense that it has a specific syntax and its behavior is defined by the operational semantics of the calculus. We will refer to this calculus as CCS_σ . Processes in CCS_σ contain labels that allow us to refer to a particular sub-process. A scheduler also behaves like a process, using however a different and much simpler syntax, and its purpose is to guide the execution of the main process using the labels that the latter provides. A *complete process* is a process running in parallel with a scheduler, and we will formally describe their interaction by defining an operational semantics for complete processes.

$I ::= 0I \mid 1I \mid \epsilon$	label indexes	$S, T ::=$	scheduler
$L ::= l^I$	labels	$L.S$	schedule single action
$P, Q ::=$	processes	$ (L, L).S$	synchronization
$L:\alpha.P$	prefix	$ \mathbf{if} L$	label test
$ P \mid Q$	parallel	$\mathbf{then} S$	
$ P + Q$	nondeterm. choice	$\mathbf{else} S$	
$ L:\sum_i p_i P_i$	internal prob. choice	$ 0$	nil
$ (\nu a)P$	restriction	$CP ::= P \parallel S$	complete process
$!P$	replication		
$ L:0$	nil		

Fig. 3. The syntax of the core CCS_σ

3.1 Syntax

Let a range over a countable set of *channel names* and l over a countable set of *atomic labels*. The syntax of CCS_σ , shown in Figure 3, is the same as the one of CCS_p except for the presence of labels. These are used to select the subprocess which “performs” a transition. Since only the operators with an initial rule can originate a transition, we only need to assign labels to the prefix and to the probabilistic sum. For reasons explained later, we also put labels on 0, even though this is not required for scheduling transitions. We use labels of the form l^s where l is an atomic label and the index s is a finite string of 0 and 1, possibly empty⁴. Indexes are used to avoid multiple copies of the same label in case of replication, which occurs dynamically due to the bang operator. As explained in the semantics, each time a process is replicated we relabel it using appropriate indexes.

A scheduler selects a sub-process for execution on the basis of its label, so we use $l.S$ to represent a scheduler that selects the process with label l and continues as S . In the case of synchronization we need to select two processes simultaneously, hence we need a scheduler of the form $(l_1, l_2).S$. Using **if-then-else** the scheduler can test whether a label is available in the process (in the top-level) and act accordingly. A complete process is a process put in parallel with a scheduler, for example $l_1:a.l_2:b \parallel l_1.l_2$. Note that for processes with an infinite execution path we need schedulers of infinite length.

3.2 Semantics

The operational semantics of the CCS_σ -calculus is given in terms of probabilistic automata defined inductively on the basis of the syntax, according to the rules shown in Figure 4.

ACT is the basic communication rule. In order for $l:\alpha.P$ to perform α , the scheduler should select this process for execution, so the scheduler needs to be

⁴ For simplicity we will write l for l^ϵ .

$$\begin{array}{l}
\text{ACT} \frac{}{l:\alpha.P \parallel l.S \xrightarrow{\alpha} \delta(P \parallel S)} \qquad \text{RES} \frac{P \parallel S \xrightarrow{\alpha} \mu \quad \alpha \neq a, \bar{a}}{(\nu a)P \parallel S \xrightarrow{\alpha} (\nu a)\mu} \\
\text{SUM1} \frac{P \parallel S \xrightarrow{\alpha} \mu}{P + Q \parallel S \xrightarrow{\alpha} \mu} \qquad \text{PAR1} \frac{P \parallel S \xrightarrow{\alpha} \mu}{P \mid Q \parallel S \xrightarrow{\alpha} \mu \mid Q} \\
\text{COM} \frac{P \parallel l_1 \xrightarrow{a} \delta(P' \parallel 0) \quad Q \parallel l_2 \xrightarrow{\bar{a}} \delta(Q' \parallel 0)}{P \mid Q \parallel (l_1, l_2).S \xrightarrow{\tau} \delta(P' \mid Q' \parallel S)} \\
\text{REP1} \frac{P \parallel S \xrightarrow{\alpha} \mu}{!P \parallel S \xrightarrow{\alpha} \rho_0(\mu) \mid \rho_1(!P)} \qquad \text{PROB} \frac{}{l:\sum_i p_i P_i \parallel l.S \xrightarrow{\tau} \sum_i [p_i] \delta(P_i \parallel S)} \\
\text{REP2} \frac{P \parallel l_1 \xrightarrow{a} \delta(P_1 \parallel 0) \quad P \parallel l_2 \xrightarrow{\bar{a}} \delta(P_2 \parallel 0)}{!P \parallel (l_1, l_2).S \xrightarrow{\tau} \delta(\rho_0(P_1) \mid \rho_{10}(P_2) \mid \rho_{11}(!P) \parallel S)} \\
\text{IF1} \frac{l \in tl(P) \quad P \parallel S_1 \xrightarrow{\alpha} \mu}{P \parallel \text{if } l \text{ then } S_1 \text{ else } S_2 \xrightarrow{\alpha} \mu} \qquad \text{IF2} \frac{l \notin tl(P) \quad P \parallel S_2 \xrightarrow{\alpha} \mu}{P \parallel \text{if } l \text{ then } S_1 \text{ else } S_2 \xrightarrow{\alpha} \mu}
\end{array}$$

Fig. 4. The semantics of CCS_σ . SUM1 and PAR1 have corresponding right rules SUM2 and PAR2, omitted for simplicity.

of the form $l.S$. After the execution the complete process will continue as $P \parallel S$. The RES rule models restriction on channel a : communication on this channel is not allowed by the restricted process. Similarly to the Section 2.2, we denote by $(\nu a)\mu$ the measure μ' such that $\mu'((\nu a)P \parallel S) = \mu(P \parallel S)$ for all processes P and $\mu'(R \parallel S) = 0$ if R is not of the form $(\nu a)P$. SUM1 models nondeterministic choice. If $P \parallel S$ can perform a transition to μ , which means that S selects one of the labels of P , then $P + Q \parallel S$ will perform the same transition, i.e. the branch P of the choice will be selected and Q will be discarded. For example

$$l_1:a.P + l_2:b.Q \parallel l_1.S \xrightarrow{a} \delta(P \parallel S)$$

Note that the operands of the sum do not have labels, the labels belong to the subprocesses of P and Q . In the case of nested choices, the scheduler must go deep and select the label of a prefix, thus resolving all the choices at once.

PAR1 has a similar behavior for parallel composition. The scheduler selects P to perform a transition on the basis of the label. The difference is that in this case Q is not discarded; it remains in the continuation. $\mu \mid Q$ denotes the measure μ' such that $\mu'(P \mid Q \parallel S) = \mu(P \parallel S)$. COM models synchronization. If $P \parallel l_1$ can perform the action a and $Q \parallel l_2$ can perform \bar{a} , then $(l_1, l_2).S$, scheduling both l_1 and l_2 at the same time, can synchronize the two. PROB models internal probabilistic choice. Note that the scheduler cannot affect the outcome of the choice, it can only schedule the choice as a whole (that's why a probabilistic sum has a label) and the process will move to a measure containing all the operands with corresponding probabilities.

REP1 and REP2 model replication. The rules are the same as in CCS_p , with the addition of a re-labeling operator ρ_k . The reason for this is that we want to avoid ending up with multiple copies of the same label as the result of replication, since this would create ambiguities in scheduling as explained in Section 3.3. $\rho_k(P)$ replaces all labels l^s inside P with l^{sk} , and it is defined as

$$\rho_k(l^s : \alpha.P) = l^{sk} : \alpha.\rho_k(P)$$

and homomorphically on the other operators (for instance $\rho_k(P \mid Q) = \rho_k(P) \mid \rho_k(Q)$). We also denote by $\rho_k(\mu)$ the measure μ' such that $\mu'(\rho_k(P) \parallel S) = \mu(P \parallel S)$. Note that we relabel only the resulting process, not the continuation of the scheduler: there is no need for relabeling the scheduler since we are free to choose the continuation as we please.

Finally **if-then-else** allows the scheduler to adjust its behaviour based on the labels that are available in P . $tl(P)$ gives the set of top-level labels of P and is defined as $tl(l : \alpha P) = tl(l : \sum p_i P_i) = tl(l : 0) = \{l\}$ and as the union of the top-level labels of all sub-processes for the other operators. Then **if l then S_1 else S_2** behaves like S_1 if l is available in P and as S_2 otherwise. This is needed when P is the outcome of a probabilistic choice, as discussed in Section 4.

3.3 Deterministic labelings

The idea in CCS_σ is that a *syntactic* scheduler will be able to completely resolve the nondeterminism of the process, without needing to rely on a *semantic* scheduler at the level of the automaton. This means that the execution of a process in parallel with a scheduler should be fully probabilistic. To achieve this we will impose a condition on the labels that we can use in CCS_σ processes. A *labeling* is an assignment of labels to the prefixes, the probabilistic sums and the nils of a process. We will require all labelings to be *deterministic* in the following sense.

Definition 1. *A labeling of a process P is deterministic iff for all schedulers S there is only one transition rule $P \parallel S \xrightarrow{\alpha} \mu$ that can be applied and the labelings of all processes P' such that $\mu(P') > 0$ are also deterministic.*

A labeling is *linear* iff all labels are pairwise distinct. We can show that linear labelings are preserved by transitions, which leads to the following proposition.

Proposition 1. *A linear labeling is deterministic.*

There are labelings that are deterministic without being linear. In fact, such labelings will be the means by which we hide information from the scheduler. However, the property of being deterministic is crucial since it implies that the scheduler will resolve all the nondeterminism of the process.

Proposition 2. *Let P be a CCS_σ process with a deterministic labeling. Then for all schedulers S , the automaton produced by $P \parallel S$ is fully probabilistic.*

4 Expressiveness of the syntactic scheduler

CCS_σ with deterministic labelings allows us to separate probabilities from non-determinism in a straightforward way: a process in parallel with a scheduler behaves in a fully probabilistic way and the nondeterminism arises from the fact that we can have many different schedulers. We may now ask the question: how powerful are the syntactic schedulers wrt the semantic ones, i.e. those defined directly over the automaton?

Let P be a CCS_p process and P_σ be the CCS_σ process obtained from P by applying a linear labeling. We say that the semantic scheduler ζ of P is equivalent to the syntactic scheduler S of P_σ , written $\zeta \sim_P S$, iff the automata $\text{etree}(P, \zeta)$ and $P_\sigma \parallel S$ are probabilistically bisimilar in the sense of [5].

A scheduler S is *non-blocking* for a process P if it always schedules some transitions, except when P itself is blocked. Let $\text{Sem}(P)$ be the set of the semantic schedulers for the process P and $\text{Syn}(P_\sigma)$ be the set of the non-blocking syntactic schedulers for process P_σ . Then we can show that for all semantic schedulers of P we can create a equivalent syntactic one for P_σ .

Proposition 3. *Let P be a CCS process and let P_σ be a CCS_σ process obtained by adding a linear labeling to P . Then $\forall \zeta \in \text{Sem}(P) \exists S \in \text{Syn}(P_\sigma) : \zeta \sim_P S$.*

To obtain this result the label test of the scheduler is crucial, in the case P performs a probabilistic choice. The scheduler uses the test to find out the result of the probabilistic choice and adapt its behaviour accordingly (as the semantic scheduler is allowed to do). For example let $P = l:(l_1:a +_p l_2 : b) \mid (l_3:c + l_4:d)$. For this process, the scheduler $l.(\mathbf{if} \ l_1 \ \mathbf{then} \ l_3.l_1 \ \mathbf{else} \ l_4.l_2)$ first performs the probabilistic choice. If the result is $l_1:a$ it performs c, a , otherwise it performs d, b . This is also the reason we need labels for 0, in case it is one of the operands of the probabilistic choice.

One would expect to obtain also the inverse of Proposition 3, showing the same expressive power for the two kinds of schedulers. We believe that this is indeed true, but it is technically more difficult to state. The reason is that the simple translation we did from CCS_p processes to CCS_σ , namely adding a linear labeling, might introduce choices that are not present in the original process. For example let $P = (a +_p a) \mid (c + d)$ and $P_\sigma = l:(l_1:a +_p l_2 : a) \mid (l_3:c + l_4:d)$. In P the choice $a +_p a$ is not a real choice, it can only do an τ transition and go to a with probability 1. But in P_σ we make the two outcomes distinct due to the labeling. So the syntactic scheduler $l.(\mathbf{if} \ l_1 \ \mathbf{then} \ l_3.l_1 \ \mathbf{else} \ l_4.l_2)$ has no semantic counterpart simply because P_σ has more choices than P , but this is an artifact of the translation. A more precise translation that would establish the exact equivalence of schedulers is left as future work.

4.1 Using non-linear labelings

Up to now we are using only linear labelings which, as we saw, give us the whole power of semantic schedulers. However, we can construct non-linear labelings that are still deterministic, that is there is still only one transition possible at

any time even though we have multiple occurrences of the same label. There are various cases of useful non-linear labelings.

Proposition 4. *Let P, Q be CCS_σ processes with deterministic labelings (not necessarily disjoint). The following labelings are all deterministic:*

$$l:(P +_p Q) \tag{2}$$

$$l_1:a.P + l_2:b.Q \tag{3}$$

$$(\nu a)(\nu b)(l_1:a.P + l_1:b.Q \mid l_2:\bar{a}) \tag{4}$$

Consider the case where P and Q in the above proposition share the same labels. In (2) the scheduler cannot select an action inside P, Q , it must select the choice itself. After the choice, only one of P, Q will be available so there will be no ambiguity in selecting transitions. The case (3) is similar but with nondeterministic choice. Now the guarding prefixes must have different labels, since the scheduler should be able to resolve the choice, however after the choice only one of P, Q will be available. Hence, again, the multiple copies of the labels do not constitute a problem. In (4) we allow the same label on the guarding prefixes of a nondeterministic choice. This is because the guarding channels a, b are restricted and only one of the corresponding output actions is available (\bar{a}). As a consequence, there is no ambiguity in selecting transitions. A scheduler (l_1, l_2) can only perform a synchronization on a , even though l_1 appears twice.

However, using multiple copies of a label limits the power of the scheduler, since the labels provide information about the outcome of a probabilistic choice (and allow the scheduler to choose different strategies through the use of the scheduler choice). In fact, this is exactly the technique we will use to archive the goals described in Section 1. Consider for example the process:

$$l:(l_1:\bar{a}.R_1 +_p l_1:\bar{a}.R_2) \mid l_2:a.P \mid l_3:a.Q \tag{5}$$

From Proposition 4(2) this labeling is deterministic. However, since both branches of the probabilistic sum have the same label l_1 , the scheduler cannot resolve the choice between P and Q based on the outcome of the choice. There is still non-determinism: the scheduler $l.(l_1, l_2)$ will select P and the scheduler $l.(l_1, l_3)$ will select Q . However this selection will be independent from the outcome of the probabilistic choice.

Note that we did not impose any direct restrictions on the schedulers, we still consider all possible syntactic schedulers for the process (5) above. However, having the same label twice limits the power of the syntactic schedulers with respect to the semantic ones. This approach has the advantage that the restrictions are limited to the choices with the same label. We already know that having pairwise distinct labels gives the full power of the semantic scheduler. So the restriction is local to the place where we, intentionally, put the same labels.

5 Testing relations for CCS_σ processes

Testing relations [19] are a method of comparing processes by considering their interaction with the environment. A *test* is a process running in parallel with the

one being tested and which can perform a distinguished action ω that represents success. Two processes are testing equivalent if they can pass the same tests. This idea is very useful for the analysis of security protocols, as suggested in [20], since a test can be seen as an adversary who interferes with a communication agent and declares ω if an attack is successful. Then two processes are testing equivalent if they are vulnerable to the same attacks.

In the probabilistic setting we take the approach of [13] which considers the exact probability of passing a test (in contrast to [10] which considers only the ability to pass a test with probability non-zero (may-testing) or one (must-testing)). This approach leads to the definition of two preorders \sqsubseteq_{may} and $\sqsubseteq_{\text{must}}$. $P \sqsubseteq_{\text{may}} Q$ means that if P can pass O then Q can also pass O with the same probability. $P \sqsubseteq_{\text{must}} Q$ means that if P always passes O with at least some probability then Q always passes O with at least the same probability.

A labeling of a process is *fresh* (with respect to a set \mathcal{P} of processes) if it is linear and its labels do not appear in any other process in \mathcal{P} . A test O is a CCS_σ process with a fresh labeling, containing the distinguished action ω . Let $\text{Test}_{\mathcal{P}}$ denote the set of all tests with respect to \mathcal{P} and let $(\nu)P$ denote the restriction on all channels of P , thus allowing only τ actions. We define $p_\omega(P, S, O)$ to be the probability of the set of executions of the fully probabilistic automaton $(\nu)(P \mid O) \parallel S$ that contain ω . Note that this set can be produced as a countable union of disjoint cones so its probability is well-defined.

Definition 2. *Let P, Q be CCS_σ processes. We define must and may testing preorders as follows:*

$$\begin{aligned} P \sqsubseteq_{\text{may}} Q & \text{ iff } \forall O \forall S_P \exists S_Q : p_\omega(P, S_P, O) \leq p_\omega(Q, S_Q, O) \\ P \sqsubseteq_{\text{must}} Q & \text{ iff } \forall O \forall S_Q \exists S_P : p_\omega(P, S_P, O) \leq p_\omega(Q, S_Q, O) \end{aligned}$$

where O ranges over $\text{Test}_{P, Q}$ and S_X ranges over $\text{Syn}((\nu)(X \mid O))$.

We also define $\approx_{\text{may}}, \approx_{\text{must}}$ to be the equivalences induced by $\sqsubseteq_{\text{may}}, \sqsubseteq_{\text{must}}$ respectively.

A context C is a process with a hole. A preorder \sqsubseteq is a precongruence if $P \sqsubseteq Q$ implies $C[P] \sqsubseteq C[Q]$ for all contexts C . May and must testing are precongruences if we restrict to contexts with fresh labelings and without occurrences of $+$. This result is essentially an adaptation to our framework of the analogous precongruence property in [3].

Proposition 5. *Let P, Q be CCS_σ processes such that $P \sqsubseteq_{\text{may}} Q$ and let C be a context with a fresh labeling and in which $+$ does not occur. Then $C[P] \sqsubseteq_{\text{may}} C[Q]$. Similarly for $\sqsubseteq_{\text{must}}$.*

This also implies that $\approx_{\text{may}}, \approx_{\text{must}}$ are congruences. Note that P, Q in the above proposition are not required to have linear labelings, P might include multiple occurrences of the same label thus limiting the power of the schedulers S_P . This shows the locality of the scheduler's restriction: some choices inside P are hidden from the scheduler but the rest of the context is fully visible.

If we remove the freshness condition then Proposition 5 is no longer true. Let $P = l_1 : a.l_2 : b$, $Q = l_3 : a.l_4 : b$ and $C = l : (l_1 : a.l_2 : c +_p [])$. We have $P \approx_{\text{may}} Q$ but $C[P], C[Q]$ can be separated by the test $O = \bar{a}.b.\omega \mid \bar{a}.\bar{c}.\omega$ (The labeling is omitted for simplicity since tests always have fresh labelings.) It is easy to see that $C[Q]$ can pass the test with probability 1 by selecting the correct branch of O based on the outcome of the probabilistic choice. In $C[P]$ this is not possible because of the labels l_1, l_2 that are common in P, C .

We can now state the result that we announced in Section 1.

Theorem 1. *Let P, Q be CCS_σ processes and C a context with a fresh labeling and without occurrences of bang. Then*

$$\begin{aligned} l : (C[l_1 : \tau.P] +_p C[l_1 : \tau.Q]) &\approx_{\text{may}} C[l : (P +_p Q)] \quad \text{and} \\ l : (C[l_1 : \tau.P] +_p C[l_1 : \tau.Q]) &\approx_{\text{must}} C[l : (P +_p Q)] \end{aligned}$$

The proof is given in the appendix.

There are two crucial points in the above Theorem. The first is that the labels of the context are copied, thus the scheduler cannot distinguish between $C[l_1 : \tau.P]$ and $C[l_1 : \tau.Q]$ based on the labels of the context. The second is that P, Q are protected by a τ action labeled by the same label l_1 . This is to ensure that in the case of a nondeterministic sum ($C = R + []$) the scheduler cannot find out whether the second operand of the choice is P or Q unless it commits to selecting the second operand. For example let $R = l_1 : (l_2 : a +_{0.5} 0)$, $P = l_3 : a$, $Q = 0$. Then $R_1 = (R + a) +_{0.1} (R + 0)$ is not testing equivalent to $R_2 = R + (a +_{0.1} 0)$ since they can be separated by $O = \bar{a}.\omega$ and a scheduler that resolves $R + a$ to a and $R + 0$ to R . However, if we take $R'_1 = (R + l : \tau.a) +_{0.1} (R + l : \tau.0)$ then R'_1 is testing equivalent to R_2 since the scheduler will have to resolve both branches of R'_1 in the same way (even though we still have non-determinism).

The problem with replication is simply the persistence of the processes. It is clear that $!P +_p !Q$ cannot be equivalent in any way to $!(P +_p Q)$ since the first replicates only one of P, Q while the second replicates both. However Theorem 1 together with Proposition 5 imply that

$$C'[l : (C[l_1 : \tau.P] +_p C[l_1 : \tau.Q])] \approx_{\text{may}} C'[C[l : (P +_p Q)]] \quad (6)$$

where C is a context without bang and C' is a context without $+$. The same is also true for \approx_{must} . This means that we can lift the sum towards the root of the context until we reach a bang. Intuitively we cannot move the sum outside the bang since each replicated copy must perform a different probabilistic choice with a possibly different outcome.

Theorem 1 shows that the probabilistic choice is indeed private to the process and invisible to the scheduler. The process can perform it at any time, even in the very beginning of the execution, without making any difference to an outside observer.

6 An application to security

In this section we discuss an application of our framework to anonymity. In particular, we show how to specify the Dining Cryptographers protocol [21] so

$$\begin{aligned}
Master &\triangleq l_1 : \sum_{i=0}^2 p_i (\underbrace{\bar{m}_0 \langle i == 0 \rangle}_{l_2} \mid \underbrace{\bar{m}_1 \langle i == 1 \rangle}_{l_3} \mid \underbrace{\bar{m}_2 \langle i == 2 \rangle}_{l_4}) \\
Crypt_i &\triangleq \underbrace{m_i(\text{pay})}_{l_{5,i}} \cdot \underbrace{c_{i,i}(\text{coin}_1)}_{l_{6,i}} \cdot \underbrace{c_{i,i \oplus 1}(\text{coin}_2)}_{l_{7,i}} \cdot \underbrace{\overline{\text{out}}_i(\text{pay} \otimes \text{coin}_1 \otimes \text{coin}_2)}_{l_{8,i}} \\
Coin_i &\triangleq l_{9,i} : ((\underbrace{\bar{c}_{i,i} \langle 0 \rangle}_{l_{10,i}} \mid \underbrace{\bar{c}_{i \oplus 1,i} \langle 0 \rangle}_{l_{11,i}}) +_{0.5} (\underbrace{\bar{c}_{i,i} \langle 1 \rangle}_{l_{10,i}} \mid \underbrace{\bar{c}_{i \oplus 1,i} \langle 1 \rangle}_{l_{11,i}})) \\
Prot &\triangleq (\nu m)(Master \mid (\nu c)(\prod_{i=0}^2 Crypt_i \mid \prod_{i=0}^2 Coin_i))
\end{aligned}$$

Fig. 5. Encoding of the dining cryptographers with probabilistic master

that it is robust to scheduler-based attacks. We first propose a method to encode *secret value passing*, which will turn out to be useful for the specification.

6.1 Encoding secret value passing

We propose to encode the passing of a secret message as follows:

$$l : c(x).P \triangleq \sum_i l : cv_i.P[v_i/x] \quad (7)$$

$$l : \bar{c}\langle v \rangle.P \triangleq l : \bar{c}\bar{v}.P \quad (8)$$

This is the usual encoding of value passing in CSS except that we use the same label in all the branches of the nondeterministic sum. To ensure that the resulting labeling will be deterministic we should restrict the channels cv_i and make sure that there will be at most one output on c . We will write $(\nu c)P$ for $(\nu cv_1) \dots (\nu cv_n)P$. For example, the labeling of the following process is deterministic:

$$(\nu c)(l_1 : c(x).P \mid l : (l_2 : \bar{c}\langle v_1 \rangle +_p l_2 : \bar{c}\langle v_2 \rangle))$$

This case is a combination of the cases (2) and (4) of Proposition 4. The two outputs on c are on different branches of the probabilistic sum, so during an execution at most one of them will be available. Thus there is no ambiguity in scheduling the sum produced by $c(x)$. The scheduler $l.(l_1, l_2)$ will perform a synchronization on cv_1 or cv_2 , whatever is available after the probabilistic choice. In other words, using the labels we manage to hide the information about which value was transmitted to P .

6.2 Dining cryptographers with probabilistic master

The problem of the Dining Cryptographers is the following: Three cryptographers are dining together. At the end of the dinner, the bill has to be paid by either one of them or by another agent called the master. The master decides

who will pay and then informs each of them separately whether he has to pay or not. The cryptographers would like to find out whether the payer is the master or one of them. However, in the latter case, they also wish to keep the payer anonymous.

The Dining Cryptographers Protocol (DCP) solves the above problem as follows: each cryptographer tosses a fair coin which is visible to himself and his neighbor to the right. Each cryptographer checks the two adjacent coins and, if he is not paying, announces *agree* if they are the same and *disagree* otherwise. However, the paying cryptographer will say the opposite. It can be proved that if the number of *disagrees* is even, then the master is paying; otherwise, one of the cryptographers is paying [21].

An external observer O is supposed to see only the three announcements $\overline{out_i} \langle \dots \rangle$. As discussed in [22], DCP satisfies anonymity if we abstract from their order. If their order is observable, on the contrary, then a scheduler can reveal the identity of the payer to O simply by forcing the payer to make his announcement first. Of course, this is possible only if the scheduler is unrestricted and can choose its strategy depending on the decision of the master (or on the results of the coins).

In our framework we can solve the problem by giving a specification of the DCP in which the choices of the master and of the coins are made invisible to the scheduler. The specification is shown in Figure 5. We use some meta-syntax for brevity: The symbols \oplus and \ominus represent the addition and subtraction modulo 3, while \otimes represents the addition modulo 2 (xor). The notation $i == n$ stands for 1 if $i = n$ and 0 otherwise.

There are many sources of nondeterminism: the order of communication between the master and the cryptographers, the order of reception of the coins, and the order of the announcements. The crucial points of our specification, which make the nondeterministic choices independent from the probabilistic ones, are: (a) all communications internal to the protocol (master-cryptographers and cryptographers-coins) are done by secret value passing, and (b) in each probabilistic choice the different branches have the same labels. For example, all branches of the master contain an output on m_0 , always labeled by l_2 , but with different values each time.

Thanks to the above independence, the specification satisfies strong probabilistic anonymity. There are various equivalent definitions of this property, we follow here the version presented in [22]. Let \mathbf{o} represent an observable (the sequence of announcements), and $p_S(\mathbf{o} \mid \overline{m_i}(1))$ represent the conditional probability, under the scheduler S , that the protocol produces \mathbf{o} given that the master has selected Cryptographer i as the payer.

Proposition 6 (Strong probabilistic anonymity). *The protocol in Figure 5 satisfies the following property: for all schedulers S and for all observables \mathbf{o} :*
 $p_S(\mathbf{o} \mid \overline{m_0}(1)) = p_S(\mathbf{o} \mid \overline{m_1}(1)) = p_S(\mathbf{o} \mid \overline{m_2}(1))$.

Note that different schedulers will produce different traces (we still have nondeterminism) but they will not depend on the choice of the master.

$$\begin{array}{l}
P ::= \dots \mid l:\{P\} \\
CP ::= P \parallel S, T
\end{array}
\quad \text{INDEP} \frac{P \parallel T \xrightarrow{\alpha} \mu}{l:\{P\} \parallel l.S, T \xrightarrow{\alpha} \mu'}$$

where $\mu'(P' \parallel S, T') = \mu(P' \parallel T')$

Fig. 6. Adding an “independent” scheduler to the calculus

Some previous treatment of the DCP, including [22], had solved the problem of the leak of information due to too-powerful schedulers by simply considering as observable sets of announcements instead than sequences. Thus one could think that using a true concurrent semantics, for instance event structures, would solve the problem. We would like to remark that this is false: true concurrency would weaken the scheduler enough in the case of the DCP, but not in general. For instance, it would not help in the anonymity example in the introduction.

6.3 Dining cryptographers with nondeterministic master

We sketch here a method to hide also certain nondeterministic choices from the scheduler, and we show an application to the variant of the Dining Cryptographers with nondeterministic master.

First we need to extend the calculus with the concept of a second *independent* scheduler T that we assume to resolve the nondeterministic choices that we want to make transparent to the main scheduler S . The new syntax and semantics are shown in Figure 6. $l : \{P\}$ represents a process where the scheduling of P is protected from the main scheduler S . The scheduler S can “ask” T to schedule P by selecting the label l . Then T resolves the nondeterminism of P as expressed by the INDEP rule. Note that we need to adjust also the other rules of the semantics to take T into account, but this change is straightforward. We assume that T does not collaborate with S so we do not need to worry about the labels in P .

To model the dining cryptographers with nondeterministic master we replace the *Master* process in Figure 5 by the following one.

$$Master \triangleq l_1 : \left\{ \sum_{i=0}^2 l_{12,i} : \tau \cdot \underbrace{\overline{m}_0 \langle i == 0 \rangle}_{l_2} \mid \underbrace{\overline{m}_1 \langle i == 1 \rangle}_{l_3} \mid \underbrace{\overline{m}_2 \langle i == 2 \rangle}_{l_4} \right\}$$

Essentially we have replaced the probabilistic choice by a *protected* nondeterministic one. Note that the labels of the operands are different but this is not a problem since this choice will be scheduled by T . Note also that after the choice we still have the same labels l_2, l_3, l_4 , however the labeling is still deterministic, similarly to the case 3 of Proposition 4.

In case of a nondeterministic selection of the culprit, and a probabilistic anonymity protocol, the notion of strong probabilistic anonymity has not been established yet, although some possible definitions have been discussed in [22]. Our framework makes it possible to give a natural and precise definition.

Definition 3 (Strong probabilistic anonymity for nondeterministic selection of the culprit). *A protocol with nondeterministic selection of the culprit satisfies strong probabilistic anonymity iff for all observables \mathbf{o} , schedulers S , and independent schedulers T_1, T_2 which select different culprits, we have: $p_{S, T_1}(\mathbf{o}) = p_{S, T_2}(\mathbf{o})$.*

We can prove the above property for our protocol:

Proposition 7. *The DCP with nondeterministic selection of the culprit specified in this section satisfies strong probabilistic anonymity.*

7 Conclusion and Future work

We have proposed a process-calculus approach to the problem of limiting the power of the scheduler so that it does not reveal the outcome of hidden random choices, and we have shown its applications to the specification of information-hiding protocols. We have also discussed a feature, namely the distributivity of certain contexts over random choices, that makes our calculus appealing for verification. Finally, we have considered the probabilistic testing preorders and shown that they are precongruences in our calculus.

Our plans for future work are in two directions: (a) we would like to investigate the possibility of giving a game-theoretic characterization of our notion of scheduler, and (b) we would like to incorporate our ideas in some existing probabilistic model checker, for instance PRISM.

Acknowledgments. We would like to thank Vincent Danos for having pointed out to us an attack to the Dining Cryptographers protocol based on the order of the scheduler, which has inspired this work. We also thank Roberto Segala and Daniele Varacca for their valuable comments on a previous version of this paper.

References

1. Vardi, M.Y.: Automatic verification of probabilistic concurrent finite-state programs. In: Proceedings of the 26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, IEEE Computer Society Press (1985) 327–338
2. Hansson, H., Jonsson, B.: A framework for reasoning about time and reliability. In: Proceedings of the 10th IEEE Symposium on Real-Time Systems, Santa Monica, California, USA, IEEE Computer Society Press (1989) 102–111
3. Yi, W., Larsen, K.G.: Testing probabilistic and nondeterministic processes. In: Proceedings of the 12th IFIP International Symposium on Protocol Specification, Testing and Verification, Florida, USA, North Holland (1992)
4. Segala, R.: Modeling and Verification of Randomized Distributed Real-Time Systems. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (June 1995) Available as Technical Report MIT/LCS/TR-676.

5. Segala, R., Lynch, N.: Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing* **2**(2) (1995) 250–273 An extended abstract appeared in *Proceedings of CONCUR '94*, LNCS 836: 481-496.
6. Hansson, H., Jonsson, B.: A calculus for communicating systems with time and probabilities. In: *Proceedings of the Real-Time Systems Symposium - 1990*, Lake Buena Vista, Florida, USA, IEEE Computer Society Press (1990) 278–287
7. Bandini, E., Segala, R.: Axiomatizations for probabilistic bisimulation. In: *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*. Volume 2076 of *Lecture Notes in Computer Science.*, Springer (2001) 370–381
8. Andova, S.: Probabilistic process algebra. PhD thesis, Technische Universiteit Eindhoven (2002)
9. Mislove, M., Ouaknine, J., Worrell, J.: Axioms for probability and nondeterminism. In Corradini, F., Nestmann, U., eds.: *Proc. of the 10th Int. Wksh. on Expressiveness in Concurrency (EXPRESS '03)*. Volume 96 of *Electronic Notes in Theoretical Computer Science.*, Elsevier (2004) 7–28
10. Palamidessi, C., Herescu, O.M.: A randomized encoding of the π -calculus with mixed choice. *Theoretical Computer Science* **335**(2-3) (2005) 373–404 http://www.lix.polytechnique.fr/~catuscia/papers/prob_enc/report.pdf.
11. Deng, Y., Palamidessi, C., Pang, J.: Compositional reasoning for probabilistic finite-state behaviors. In Middeldorp, A., van Oostrom, V., van Raamsdonk, F., de Vrijer, R.C., eds.: *Processes, Terms and Cycles: Steps on the Road to Infinity*. Volume 3838 of *Lecture Notes in Computer Science*. Springer (2005) 309–337 <http://www.lix.polytechnique.fr/~catuscia/papers/Yuxin/BookJW/par.pdf>.
12. Sokolova, A., Vink, E.d.: Probabilistic automata: system types, parallel composition and comparison. In Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P., Siegle, M., eds.: *Validation of Stochastic Systems: A Guide to Current Research*. Volume 2925 of *Lecture Notes in Computer Science*. Springer (2004) 1–43
13. Jonsson, B., Larsen, K.G., Yi, W.: Probabilistic extensions of process algebras. In Bergstra, J.A., Ponse, A., Smolka, S.A., eds.: *Handbook of Process Algebra*. Elsevier (2001) 685–710
14. Chatzikokolakis, K., Palamidessi, C.: A framework for analyzing probabilistic protocols and its application to the partial secrets exchange. *Theoretical Computer Science* (2005) To appear. A short version of this paper appeared in the *Proceedings of the Symposium on Trustworthy Global Computing (TGC)*, volume 3705 of LNCS, pages 146-162. Springer, <http://www.lix.polytechnique.fr/~catuscia/papers/PartialSecrets/TCSreport.pdf>.
15. Canetti, R., Cheung, L., Kaynar, D., Liskov, M., Lynch, N., Pereira, O., Segala, R.: Task-structured probabilistic i/o automata. In: *Proceedings the 8th International Workshop on Discrete Event Systems (WODES'06)*, Ann Arbor, Michigan (2006)
16. Canetti, R., Cheung, L., Kaynar, D.K., Liskov, M., Lynch, N.A., Pereira, O., Segala, R.: Time-bounded task-PIOAs: A framework for analyzing security protocols. In Dolev, S., ed.: *Proceedings of the 20th International Symposium in Distributed Computing (DISC '06)*. Volume 4167 of *Lecture Notes in Computer Science.*, Springer (2006) 238–253
17. Garcia, F.D., van Rossum, P., Sokolova, A.: Probabilistic anonymity and admissible schedulers (2007) [arXiv:0706.1019v1](https://arxiv.org/abs/0706.1019v1).
18. de Alfaro, L., Henzinger, T.A., Jhala, R.: Compositional methods for probabilistic systems. In Larsen, K.G., Nielsen, M., eds.: *Proceedings of the 12th International Conference on Concurrency Theory (CONCUR 2001)*. Volume 2154 of *Lecture Notes in Computer Science.*, Springer (2001)

19. Nicola, R.D., Hennessy, M.C.B.: Testing equivalences for processes. *Theoretical Computer Science* **34**(1-2) (1984) 83–133
20. Abadi, M., Gordon, A.D.: A calculus for cryptographic protocols: The spi calculus. *Information and Computation* **148**(1) (10 January 1999) 1–70
21. Chaum, D.: The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology* **1** (1988) 65–75
22. Bhargava, M., Palamidessi, C.: Probabilistic anonymity. In Abadi, M., de Alfaro, L., eds.: *Proceedings of CONCUR*. Volume 3653 of *Lecture Notes in Computer Science.*, Springer (2005) 171–185 <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/concur.pdf>.

A Proofs

In this appendix we give the proof of the main technical result of our paper.

Theorem 1 Let P, Q be CCS_σ processes and C a context with a fresh labeling and without occurrences of bang. Then

$$\begin{aligned} l:(C[l_0:\tau.P] +_p C[l_0:\tau.Q]) &\approx_{\mathbf{may}} C[l:(P +_p Q)] \quad \text{and} \\ l:(C[l_0:\tau.P] +_p C[l_0:\tau.Q]) &\approx_{\mathbf{must}} C[l:(P +_p Q)] \end{aligned}$$

Proof.

Since we will always use the label l for all probabilistic sum $+_p$, and l_0 for $\tau.P$ and $\tau.Q$, we will omit these labels to make the proof more readable. We will also denote $(1 - p)$ by \bar{p} .

Let $R_1 = C[\tau.P] +_p C[\tau.Q]$ and $R_2 = C[P +_p Q]$. We will prove that for all tests O and for all schedulers $S_1 \in \text{Syn}((\nu)(R_1 \mid O))$ there exists $S_2 \in \text{Syn}((\nu)(R_2 \mid O))$ such that $p_\omega(R_1, S_1, O) = p_\omega(R_2, S_2, O)$ and vice versa. This implies both $R_1 \approx_{\mathbf{may}} R_2$ and $R_1 \approx_{\mathbf{must}} R_2$.

Without loss of generality we assume that tests do not perform internal actions, but only synchronizations with the tested process. First, it is easy to see that

$$\begin{aligned} p_\omega(P +_p Q, l.S, O) &= p p_\omega(P, S, O) + \bar{p} p_\omega(Q, S, O) & (9) \\ p_\omega(l_1:a.P, (l_1, l_2).S, O) &= p_\omega(P, S, O') & (10) \end{aligned}$$

where $(\nu)(l_1:a.P \mid O) \parallel (l_1, l_2).S \xrightarrow{\tau} \delta((\nu)(P \mid O' \parallel S))$.

In order for the scheduler of R_1 to be non-blocking, it has to be of the form $l.S_1$, since the only possible transition of R_1 is the probabilistic choice labeled by l . By (9) we have

$$p_\omega(C[\tau.P] + C[\tau.Q], l.S_1, O) = p p_\omega(C[\tau.P], S_1, O) + \bar{p} p_\omega(C[\tau.Q], S_1, O)$$

The proof will be by induction on the structure of C . Let O range over tests with fresh labelings, let S_1 range over non-blocking schedulers for both $C[\tau.P]$

and $C[\tau.Q]$ (such that $l.S_1$ is a non-blocking scheduler for R_1) and let S_2 range over non-blocking schedulers for R_2 . The induction hypothesis is:

$$\begin{aligned} \Rightarrow) \quad & \forall O \forall S_1 \exists S_2 : \\ & p p_\omega(C[\tau.P], S_1, O) + \bar{p} p_\omega(C[\tau.Q], S_1, O) = p_\omega(C[P +_p Q], S_2, O) \quad \text{and} \\ \Leftarrow) \quad & \forall O \forall S_2 \exists S_1 : \\ & p p_\omega(C[\tau.P], S_1, O) + \bar{p} p_\omega(C[\tau.Q], S_1, O) = p_\omega(C[P +_p Q], S_2, O) \end{aligned}$$

We have the following cases for C :

- Case $C = []$. Trivial.
- Case $C = l_1:a.C'$

The scheduler S_1 of $C[\tau.P]$ and $C[\tau.Q]$ has to be of the form $S_1 = (l_1, l_2).S'_1$ where l_2 is the label of a \bar{a} prefix in O (if no such prefix exists then the case is trivial).

A scheduler of the form $(l_1, l_2).S$ can schedule any process of the form $l_1:a.X$ (with label l_1) giving the transition:

$$(\nu)(l_1:a.X \mid O) \parallel (l_1, l_2).S \xrightarrow{\tau} \delta((\nu)(X \mid O') \parallel S)$$

and producing always the same O' . The probability p_ω for these processes will be given by equation (10).

Thus for (\Rightarrow) we have

$$\begin{aligned} & p p_\omega(l_1:a.C[\tau.P], (l_1, l_2).S'_1, O) + \bar{p} p_\omega(l_1:a.C[\tau.Q], (l_1, l_2).S'_1, O) \\ &= p p_\omega(C'[\tau.P], S'_1, O') + \bar{p} p_\omega(C'[\tau.Q], S'_1, O') \quad (10) \\ &= p_\omega(C'[P +_p Q], S'_2, O') \quad \text{Ind. Hyp.} \\ &= p_\omega(l_1:a.C'[P +_p Q], (l_1, l_2).S'_2, O) \quad (10) \\ &= p_\omega(R_2, S_2, O) \end{aligned}$$

For (\Leftarrow) we can perform the above derivation in the opposite direction, given that a scheduler for $R_2 = l_1:a.C'[P +_p Q]$ must be of the form $S_2 = (l_1, l_2).S'_2$.

- Case $C = C' \mid R$

Since we only consider contexts with fresh labelings, $R \mid O$ is itself a test, and

$$p_\omega(X \mid R, S, O) = p_\omega(X, S, R \mid O) \quad (11)$$

Thus for (\Rightarrow) we have

$$\begin{aligned} & p p_\omega(C'[\tau.P] \mid R, S_1, O) + \bar{p} p_\omega(C'[\tau.Q] \mid R, S_1, O) \\ &= p p_\omega(C'[\tau.P], S_1, R \mid O) + \bar{p} p_\omega(C'[\tau.Q], S_1, R \mid O) \quad (11) \\ &= p_\omega(C'[P +_p Q], S_2, R \mid O) \quad \text{Ind. Hyp.} \\ &= p_\omega(C'[P +_p Q] \mid R, S_2, O) \quad (11) \\ &= p_\omega(R_2, S_2, O) \end{aligned}$$

For (\Leftarrow) we can perform the above derivation in the opposite direction.

– Case $C = l_1:(C' +_q R)$

Since we consider only contexts with fresh labelings, the labels of C' are disjoint from those of R , thus the scheduler of a process of the form $l_1 : (C'[X] +_q R)$ must be of the form $S = l_1.(\mathbf{if } l_C \mathbf{ then } S_C \mathbf{ else } S_R)$ where $l_C \in tl(C'[X])$, S_C is a scheduler containing labels of $C'[X]$ and S_R is a scheduler containing labels of R . Moreover

$$\begin{aligned} & p_\omega(l_1:(C'[X] +_q R), S, O) \\ &= q p_\omega(C'[X], \mathbf{if } l_C \mathbf{ then } S_C \mathbf{ else } S_R, O) + \\ & \quad \bar{q} p_\omega(R, \mathbf{if } l_C \mathbf{ then } S_C \mathbf{ else } S_R, O) \\ &= q p_\omega(C'[X], S_C, O) + \bar{q} p_\omega(R, S_R, O) \end{aligned} \quad (12)$$

As a consequence, the scheduler S_1 of $C[\tau.P]$ and $C[\tau.Q]$ has to be of the form $S_1 = l_1.(\mathbf{if } l_C \mathbf{ then } S_C \mathbf{ else } S_R)$. Note that $tl(C'[\tau.P]) = tl(C'[\tau.Q])$ so the two processes cannot be separated by a test. S_C will schedule both (possibly separating them later).

For (\Rightarrow) we have

$$\begin{aligned} & p p_\omega(l_1:(C'[\tau.P] +_q R), S_1, O) + \bar{p} p_\omega(l_1:(C'[\tau.Q] +_q R), S_1, O) \\ &= q(p p_\omega(C'[\tau.P], S_C, O) + \bar{p} p_\omega(C'[\tau.Q], S_C, O)) + \\ & \quad \bar{q} p_\omega(R, S_R, O) \quad (12) \\ &= q p_\omega(C'[P +_p Q], S'_C, O) + \\ & \quad \bar{q} p_\omega(R, S_R, O) \quad \text{Ind. Hyp.} \\ &= p_\omega(l_1:(C'[P +_p Q] +_q R), l_1.(\mathbf{if } l'_C \mathbf{ then } S'_C \mathbf{ else } S_R), O) \quad (12) \\ &= p_\omega(R_2, S_2, O) \end{aligned}$$

Where $l'_C \in tl(C'[P +_p Q])$ (and thus $l'_C \notin tl(R)$).

For (\Leftarrow) we can perform the above derivation in the opposite direction, given that a scheduler for $R_2 = l_1:(C'[P +_p Q] +_q R)$ must be of the form $S_2 = l_1.(\mathbf{if } l'_C \mathbf{ then } S'_C \mathbf{ else } S_R)$.

– Case $C = C' + R$

Consider the process $C'[l_0:\tau.P] + R$. The scheduler S_1 of this process has to choose between $C'[l_0:\tau.P]$ and R .

There are two cases to have a transition using the SUM1, SUM2 rules.

i) Either $S_1 = S_R$ and

$$\text{SUM2} \frac{(\nu)(R \mid O) \parallel S_R \xrightarrow{\alpha} \mu}{(\nu)(C'[l_0:\tau.P] + R \mid O) \parallel S_R \xrightarrow{\alpha} \mu}$$

In this case

$$p_\omega(C'[l_0:\tau.P] + R, S_R, O) = p_\omega(R, S_R, O) \quad (13)$$

ii) Or $S_1 = S_C$ and

$$\text{SUM1} \frac{(\nu)(C'[l_0:\tau.P] \mid O) \parallel S_C \xrightarrow{\alpha} \mu}{(\nu)(C'[l_0:\tau.P] + R \mid O) \parallel S_C \xrightarrow{\alpha} \mu}$$

In this case

$$p_\omega(C'[l_0:\tau.P] + R, S_C, O) = p_\omega(C'[l_0:\tau.P], S_C, O) \quad (14)$$

Now consider the process $C'[l_0:\tau.Q] + R$. Since P and Q are behind the $l_0:\tau$ action, we have $tl(C'[l_0:\tau.Q]) = tl(C'[l_0:\tau.P])$. Thus S_R and S_C will select R and $C'[l_0:\tau.Q]$ respectively and the equations (13) and (14) will hold.

In the case (i) ($S = S_R$) we have:

$$\begin{aligned} & p p_\omega(C'[\tau.P] + R, S_R, O) + \bar{p} p_\omega(C'[\tau.Q] + R), S_R, O) \\ &= p p_\omega(R, S_R, O) + \bar{p} p_\omega(R, S_R, O) \quad (13) \\ &= p_\omega(R, S_R, O) \\ &= p_\omega(C'[P +_p Q] + R, S_R, O) \\ &= p_\omega(R_2, S_2, O) \end{aligned}$$

In the case (ii) ($S = S_C$) we have:

$$\begin{aligned} & p p_\omega(C'[\tau.P] + R, S_C, O) + \bar{p} p_\omega(C'[\tau.Q] + R), S_C, O) \\ &= p p_\omega(C'[\tau.P], S_C, O) + \bar{p} p_\omega(C'[\tau.Q], S_C, O) \quad (14) \\ &= p_\omega(C'[P +_p Q], S'_C, O) \quad \text{Ind. Hyp.} \\ &= p_\omega(C'[P +_p Q] + R, S'_C, O) \\ &= p_\omega(R_2, S_2, O) \end{aligned}$$

For (\Leftarrow) we can perform the above derivation in the opposite direction.

– Case $C = (\nu a)C'$

The process $(\nu)((\nu a)C'[X] \mid O)$ has the same transitions as $(\nu)(C'[X] \mid (\nu a)O)$. The result follows by the induction hypothesis.

□