


PUBLICATION
INTERNE
N° 1866



TRULY CONCURRENT SYNCHRONIZATION OF
MARKOV CHAINS

SAMY ABBES , ALBERT BENVENISTE

Truly Concurrent Synchronization of Markov Chains

Samy Abbes^{*}, Albert Benveniste^{**}

Systèmes communicants
Projet DistribCom

Publication interne n° 1866 — June 2007 — 20 pages

Abstract: We study the synchronization of two discrete Markov chains that share common states. Markov chains define transition systems, and we consider the synchronization product of these from the partial orders viewpoint. We propose a randomization of the new, concurrent system. The random concurrent system we construct has two key properties: first it is Markovian, and second it has a local independence property, meaning that local processes are, in the probabilistic sense, conditionally independent. We compare this construction with the product of Probabilistic Automata.

Key-words: True concurrency, probabilistic models, Markov chains, product

(Résumé : [tsvp](#))

^{*} University of Cambridge Computer Laboratory

^{**} IRISA/INRIA, Campus de Beaulieu, 35042 Rennes, France; <http://www.irisa.fr/distribcom/benveniste>

Produit concurrent de chaînes de Markov

Résumé : On étudie la synchronisation de deux chaînes de Markov discrètes à états partagés. Les chaînes de Markov définissent des systèmes de transition, et on en considère le produit de synchronisation d'un point de vue d'ordres partiels. On propose une probabilisation du système concurrent obtenu. Le système probabiliste concurrent proposé a deux propriétés clef : premièrement, il est markovien, et deuxièmement il possède une propriété d'indépendance locale. Les processus locaux sont conditionnellement indépendants en probabilité. On compare cette construction au produit d'automates probabilistes.

Mots clés : concurrence, ordres partiels, modèles probabilistes, chaînes de Markov, produit

1 Introduction

One of the major aim of concurrency theory is to propose models for the synchronization of concurrent systems. The construction of such models highly depends on the semantics in use for the analysis of systems. Two kinds of semantics have attracted interest: the interleaving semantics and the partial order semantics—the latter is also called the true-concurrency semantics. They differ by the way processes are considered: the interleaving semantics considers processes as sequences of events, while the true-concurrency semantics only retains the partial orders spanned by concurrent events. To our knowledge, no solution is available today for the true-concurrency synchronization of probabilistic systems. This lack leaves the interleaving semantics as the only possibility for modeling the synchronization of probabilistic processes, thus resulting in several drawbacks that a true-concurrency semantics would adequately solve. Our paper investigates probabilistic synchronization under true-concurrency semantics.

Nondeterminism plays a special role in concurrency theory. In a complex system composed of several local components, nondeterminism can occur due to different reasons. First, it can occur at the local level due to the local choices performed by each component. Second, in the interleaving semantics, nondeterminism also occurs at the global level to interleave the local moves of the different components. In this paper, we consider the composition of fully probabilistic components, i.e., components in which the (local) choices are solved in a fully probabilistic way. That is, choices performed by each component are entirely probabilistic and no other source of local nondeterminism exists. Since fully probabilistic systems are autonomous, it makes sense to require that their composition shall also be autonomous, i.e., fully probabilistic. Unfortunately, performing system synchronization according to the interleaving semantics still leaves present the second source of nondeterminism.

To explain this, let us consider a framework in which the two types of nondeterminism can be combined. Markov Decision Processes (MDP, originating from [4]) are of this kind. Generalizations of MDP are found in probabilistic process algebra [5] and in Probabilistic Automata (PA, [8, 6]) that also provide a blend of probabilistic and nondeterministic choices. Let us focus on PA for this discussion. Consider two *fully probabilistic* PA, meaning that every choice performed by each PA is probabilistic. Consider then their synchronization on some common actions. The main point is that synchronization involves concurrency of local, non-synchronizable actions, which introduces in turn new non-determinism. The latter is not motivated by modeling purposes, but only implied by the underlying modeling framework, namely the interleavings semantics. Hence, starting from autonomous entites, we end up with a new system which is not autonomous anymore: local choices are still solved internally, but the interleavings of these needs the interference of some “external agent” modeled by non-determinism. This additional non-determinism, in this context, appears as an artifact, and rather as a drawback of the interleaving semantics more than as a desirable feature of it. The analysis we detail below shows that the alternative choice of true-concurrency probabilistic semantics recently developed in [1, 2, 3, 9] might be appropriate to avoid this drawback. We now make this discussion technically precise.

Probabilistic Automata. The definition we give is a simplified version of the PA of [8], enough for our purpose. A PA is defined as a quadruple $\mathcal{A} = (S, i, A, r)$, where:

1. S is a finite set of *states*, and $i \in S$ is the *initial state*.
2. A is a finite set of *actions*.
3. r is a transition relation $r \subseteq S \times (A \times \text{Prob}(S))$, where $\text{Prob}(S)$ denotes the set of probability distributions over S .

At each state $s \in S$, an action a is said to be *enabled* if there exists $m \in \text{Prob}(S)$ such that $(s, a, m) \in r$. In general, several actions are enabled at each state, and each action can lead to different probability distributions over S . There is therefore plenty of room for non-determinism in this model. Intuitively, when the automaton is in state s , a pair $(a, p) \in A \times \text{Prob}(S)$ is first non-deterministically chosen such that $(s, a, p) \in r$; then the new state is picked randomly by using the probability distribution p . A *partial run* of the automaton consists then of alternated sequences $i \xrightarrow{a_1} s_1 \dots \xrightarrow{a_n} s_n$ of actions and states thus chosen.

We say the automaton is *fully probabilistic* if non-determinism is degenerated, *i.e.*, if for any state s , there is a unique action a and a unique probability distribution p on S such that $(s, a, p) \in r$ (this strengthens the definition of [8]).

Markov chains as PA. Remark how a Markov chain with definite initial state is easily transposed into a fully probabilistic automaton. In fact, let S be the set of states of a Markov chain, i its initial state, and P the transition matrix of the chain. Then the corresponding PA is defined as (S, i, A, r) , where A and r are as follows:

1. $A = S$;
2. r is the subset of $S \times S \times \text{Prob}(S)$ consisting of triples of the form $(s, s, P(s, \cdot))$, for s ranging over S .

Therefore, each state $s \in S$ has a unique enabled action, namely itself, which leads to the unique probability distribution $P(s, \cdot)$ coming from the transition probability of the original chain.

Product of PA. We recall now the definition of the product of two PA, which we call *PA-product* in the sequel. The PA-product $\mathcal{A}_1 \times \mathcal{A}_2$ of two PA $\mathcal{A}_j = (S_j, i_j, A_j, r_j)$ for $j = 1, 2$ is the PA $\mathcal{A} = (S, i, A, r)$ defined as follows:

1. $S = S_1 \times S_2$, $i = (i_1, i_2)$.
2. $A = A_1 \cup A_2$, a union that may not be disjoint.
3. Transition r satisfies the following: $((s_1, s_2), a, p) \in r$ if and only if $p = p_1 \otimes p_2$, such that:

- (a) If $a \in A_1$ then $(s_1, a, p_1) \in r_1$, else $p_1 = \delta_{s_1}$.
- (b) If $a \in A_2$ then $(s_2, a, p_2) \in r_2$, else $p_2 = \delta_{s_2}$.

Here we use the notation δ_s to denote the probability distribution with zero values on every singleton but $\{s\}$, and ‘ \otimes ’ denotes the usual product of measures.

An illustrative example. Now that definitions of PA and their product have been recalled, let us have a close look at an example, namely the synchronization of two Markov chains seen as fully probabilistic automata. Consider the two Markov chains depicted in Figure 1, where labels indicate state-actions (recall that states and actions are in bijection for Markov chains). We denote by \mathcal{A}_1 and \mathcal{A}_2 respectively the corresponding fully probabilistic automata, and rely on the notations used above. We then consider the PA-product $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$. It has 6 actions a, b, c, d, e and f , and the set of states $S = S_1 \times S_2$ has 16 elements. The next paragraph discusses some of its features.

A first problem: blocking states. As a first issue and prior to discussing nondeterminism, observe the existence of blocking states in the PA-product $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$, i.e., *states of \mathcal{A} from which no action is enabled*. Consider indeed any state of the form $s = (s_1, s_2)$, such that $s_1 \neq s_2$ and s_1 and s_2 , seen as actions, are both shared. For instance we may pick $s = (c, d)$. Hence the PA-product $\mathcal{A}_1 \times \mathcal{A}_2$ possesses blocking states, even if the two PA we started with had no blocking states. A Borel-Cantelli argument shows then that, with probability 1, the PA $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ will eventually reach such a blocking state. In fact, \mathcal{A} possesses infinite runs, but they together form a set of probability zero. This is embarrassing since it forbids us considering statistical properties for the resulting probabilistic product. For example, we cannot compare how frequently runs of the PA-product traverse two given states $s = (s_1, s_2)$ and $s' = (s'_1, s'_2)$ in the long range. We cannot get anything useful from a law of large numbers. The resulting product cannot exhibit any nontrivial stationary behaviour (as Markov chains do). And so on.

This unpleasant feature can be overcome by ruling out blocking states, from the product probability $p = p_1 \otimes p_2$, which governs the choice of the next state in step 3 of the definition of PA-product. To this end, replace $p = p_1 \otimes p_2$ by $\hat{p}(s) = 0$ if s is blocking, and otherwise by $\hat{p}(s) = p(s)/C_p$ where $C_p = \sum_{s'} p(s')$ and s' ranges over the set of non-blocking states of the PA-product. Note that $\hat{p} \in \text{Prob}(S)$. This construction assumes that $C_p \neq 0$, i.e., if there exist non blocking states that are reachable from the current state. If $C_p = 0$, then one must, using backward induction, rule out the current state, and so on. We shall use renormalization techniques of this kind in the sequel.

Non-determinism and the need for partial orders. Let us come back to our central issue of nondeterminism arising in the product. Observe that the PA-product of the above two Markov chains is not anymore fully probabilistic, i.e., *some states have several actions non-deterministically enabled*—in particular, this PA-product is no longer a Markov chain. Indeed, consider the product state $s = (a, f)$. From the product state s , the automaton can either use action a or action f , leading respectively to probability distributions $P_1(a, \cdot) \otimes \delta_f$




Figure 1: Two Markov chains interpreted as fully probabilistic automata. Initial states are indicated by \Rightarrow . Actions are in bijection with states. Each state enables only its corresponding action; the associated probability distribution simulates a token moving from the current state to the next state through one of the depicted arcs.

or $\delta_a \otimes P_2(f, \cdot)$, which amounts to moving either the “left” or the “right” component by making use of the probability distribution associated to the action that has been picked and leaving the other component as it is.

Non-determinism lies thus in the choice of deciding which component of the automaton is to be “activated”. If the other component is to be “activated” on next move, we can understand the effect of non-determinism in the product as choosing the *order* of actions. In other words, non-determinism resolves the interleaving of local actions.

The drawback we observe in the composition of PA after discussing the above example is that *the class of fully probabilistic systems is not closed under PA-product*. Providing an additional specification in the form of a *scheduler* [6, 8] appears as an artifact, technically justified by the interleaving semantics, but artificial from the modeling viewpoint.

Since the interleaving semantics is the very cause of this issue, it appears natural to abandon this semantics for another one, more adapted to the concurrency features involved with synchronization. The *true-concurrency* semantics, based on partial orders, is the natural semantics that precisely abstracts the interleaving issues. Accordingly, the way we synchronize the two systems is depicted in Figure 2. We place the two components side by side, and link together the share states b and c on each side. Then we consider the *two-tokens transition system* thus obtained, where a global state consists of a pair $(s_1, s_2) \in S_1 \times S_2$.

The runs of the system are described as follows. Informally, we ask that the two components synchronize on shared states. So the rule is that any token can move along any arc, up to the constraint that if one token arrives in a shared state, then the other token must arrive in the same shared state—by doing so we avoid the above discussed issue with blocking states in the product. Otherwise, tokens are free to move locally in an asynchronous fashion one with respect to another. A (partial) run of the system is consists thus of a *pair of (partial) runs* in each local component, provided the two local partial runs are able to marry. So we may consider the following partial run, starting from state (a, f) :

$$\rho = \{(a, f) \rightarrow (b, f) \rightarrow (b, f) \rightarrow (b, e) \rightarrow (d, d) \rightarrow (b, d) \rightarrow (c, c)\}.$$

In this partial run, (d, d) is the first synchronization state, and (c, c) is the second. Sequences of symbols are naturally ordered by the prefix relation; there is therefore a product partial




Figure 2: The two Markov chains of Figure 1 now synchronize on shared states c and d . If one system enters a shared state, it must be matched by the other system.

order on partial runs, defined component by component. Hence the above run is more adequately represented as follows, where asynchrony is better captured:

$$\rho = \left\{ \begin{array}{ccccccc} a & \longrightarrow & b & \longrightarrow & b & \longrightarrow & d & \longrightarrow & b & \longrightarrow & c \\ & & & & & & \parallel & & & & \parallel \\ f & \longrightarrow & e & \longrightarrow & & & d & \longrightarrow & & & c \end{array} \right\}$$

Note the two synchronizing states (d, d) and (c, c) .

Contrasting with the case of a single Markov chain, we observe that sub-runs of a given partial run are *not totally ordered*, but really *partially ordered*. For instance, the partial run ρ contains the two sub-runs $\rho_1 = (a, f) \rightarrow (b, f)$ and $\rho_2 = (a, f) \rightarrow (a, e)$, but neither $\rho_1 \leq \rho_2$ nor $\rho_2 \leq \rho_1$ holds. This expresses the concurrency features of the model.

Our proposal: two-components processes and their randomization. Since we are interested in the partial orders semantics of the model, we shall randomize the runs of the system seen as partial orders, and not the interleaving of sequences. We call such a process a *two-component process*, and we say it is randomized when a probability measure has been defined on the space of *maximal runs*, where maximal refers to the prefix relation.

If the construction can be achieved, this give rise to a probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, from which usual tools from probability theory can be applied. Since we start from Markov chains, it is natural to consider a synchronization that would keep the system globally Markovian. We rely for this on the notion of Markovian true-concurrent systems developed by one of the authors in [1]. Informally, it captures the broad idea that Markovian systems have no memory beyond the present: “the probabilistic future only depends on the current state of the system, and not on the whole history”.

Another key property, natural to require, is a “local independence” property. The space $(\Omega, \mathfrak{F}, \mathbb{P})$ induces two local probabilistic processes, with values in the local state spaces. These two processes are *not* independent, excepted in some very particular conditions, since they synchronize with each other. However, we can require that they are *maximally independent*, up to the synchronization constraint. In the language of probability, the local independence property states that the two local processes are independent, conditionally to

the synchronization. In particular, if there is no synchronization at all, the two processes are just independent in probability.

The aim of the paper is to construct, starting from two Markov chains, a synchronization probability space $(\Omega, \mathfrak{F}, \mathbb{P})$ as above, satisfying both the true-concurrent Markov property and the local independence property. The construction we propose proceeds inductively as follows. First, observe that a partial run, or *trace* as we call them, can be decomposed as a succession of joins in the shared states. This gives rise to parallel *excursions* in non-shared states, separated by *incursions* in shared states. We randomize the first pair of parallel excursions by letting the Markov chains evolve freely, conditionally on the constraint to synchronize at the end of the excursion, i.e., to arrive in the same shared state—this rules out blocking states in the resulting product. We then proceed inductively by randomizing the next pair of parallel excursions, and so on. We show that these randomizations, localized in time, sum up together to produce an appropriate probability state $(\Omega, \mathfrak{F}, \mathbb{P})$ as required. We then show that both the Markov property and the local independence property hold for the so constructed probability space.

Outline of the paper. In § 2, we recall background material on probability and Markov chains and on true-concurrent randomization of concurrent systems. We formalize our problem in § 3. Then § 4 contains our main results with the construction of a synchronization product and its two key properties. Finally, we examine future research directions and related work in § 5.

2 Background on Probabilistic Processes

After some generalities on probability spaces (§ 2.1), we embark on Markov chains (§§ 2.2, 2.3) and on true-concurrent probabilistic processes (§§ 2.4, 2.5).

2.1 Probability Spaces. Conditional Probabilities. Let Ω be a set. A σ -field on Ω is a collection \mathfrak{F} of subsets of Ω , that contains the empty set, and which is closed under countable intersection and under complement in Ω . Elements of \mathfrak{F} are called the \mathfrak{F} -measurable subsets of Ω , or shortly the measurable subsets of Ω if \mathfrak{F} is understood. The pair (Ω, \mathfrak{F}) is called a *measurable space*.

A *probability measure* on a measurable space (Ω, \mathfrak{F}) is a non negative real valued set function $\mathbb{P} : \mathfrak{F} \rightarrow \mathbb{R}$, such that $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(\Omega) = 1$, and $\mathbb{P}(\bigcup_{n \geq 0} E_n) = \sum_{n \geq 0} \mathbb{P}(E_n)$ for any countable collection $(E_n)_{n \geq 0}$ of measurable subsets that are pairwise disjoint. The triple $(\Omega, \mathfrak{F}, \mathbb{P})$ is called a *probability space*.

An application $X : \Omega \rightarrow \Omega'$, where (Ω, \mathfrak{F}) and (Ω', \mathfrak{F}') are two measurable spaces, is called a *random variable* (following the terminology from probability theory) if $X^{-1}(U) \in \mathfrak{F}$ for every $U \in \mathfrak{F}'$. Furthermore, if \mathbb{P} makes $(\Omega, \mathfrak{F}, \mathbb{P})$ a probability space, the probability *law* or *distribution* of random variable $X : \Omega \rightarrow \Omega'$ is the probability measure \mathbb{P}_X on (Ω', \mathfrak{F}') defined by $\mathbb{P}_X(U) = \mathbb{P}(X \in U)$, for every $U \in \mathfrak{F}'$.

Let $(\Omega_i, \mathfrak{F}_i, \mathbb{P}_i)$, $i = 1, 2$, be two probability spaces. The product set $\Omega = \Omega_1 \times \Omega_2$ is equipped with the product σ -field \mathfrak{F} generated by (i.e., the smallest σ -field that contains) the sets of the form $U_1 \times U_2$, with $U_i \in \mathfrak{F}_i$ for $i = 1, 2$. The measurable space (Ω, \mathfrak{F}) is then equipped with the *product probability* denoted by $\mathbb{P} = \mathbb{P}_1 \otimes \mathbb{P}_2$, which is the unique probability such that $\mathbb{P}(U_1 \times U_2) = \mathbb{P}_1(U_1) \times \mathbb{P}_2(U_2)$ for $U_i \in \mathfrak{F}_i$, $i = 1, 2$. If $(\Omega, \mathfrak{F}, \mathbb{P})$ is a probability space and $X : \Omega \rightarrow E$, $Y : \Omega \rightarrow F$ are two random variables, they are called *independent* if $\mathbb{P}_{(X,Y)} = \mathbb{P}_X \otimes \mathbb{P}_Y$, where (X, Y) denotes the join random variable with values in $E \times F$.

Let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a probability space, and let E be a measurable subset of Ω . The subset E is equipped with the σ -field \mathfrak{F}_E defined by $\mathfrak{F}_E = \{U \cap E : U \in \mathfrak{F}\}$. If $\mathbb{P}(E) > 0$, we can equip the measurable space (E, \mathfrak{F}_E) with the *conditional probability* $\mathbb{P}(\cdot | E)$ defined by

$$\forall U \in \mathfrak{F}_E, \quad \mathbb{P}(U|E) = \frac{1}{\mathbb{P}(E)} \mathbb{P}(U). \tag{1}$$

By convention, any finite or countable set E is always equipped with its *discrete σ -field*, simply defined as the powerset of E . In this case, singletons are measurable, and any probability \mathbb{P} on E is uniquely determined by the function $x \in E \mapsto \mathbb{P}(\{x\})$. In particular, the *Dirac measure* on an element $x \in E$, denoted by δ_x , is defined by $\delta_x(\{x'\}) = 1$ if $x = x'$, and 0 otherwise, for x' ranging over E .

Finally, we recall that a *stochastic matrix* of size $n \geq 1$ is a square matrix of size n , each row of which is a probability distribution over a set of size n .

2.2 Markov Chains. Let S be a finite set of *states*, and let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a probability space. A *probabilistic process on S* is defined as a sequence $(X_n)_{n \geq 0}$ of random variables, defined on Ω and with value in S . Let Ω^S denote the set of single-sided infinite sequences with values in S ; such sequences are generically denoted by $\omega = (x_n)_{n \geq 0}$. It is understood that Ω^S is equipped with its natural Borel σ -field \mathfrak{F}^S , generated by the elementary cylinders $S_{s_0, \dots, s_n} = \{\omega = (x_n)_{n \geq 0} \in \Omega^S : x_0 = s_0, \dots, x_n = s_n\}$, for $n \geq 0$ and $s_0, \dots, s_n \in S$; Ω^S is thus considered as a measurable space. The measurable space Ω^S is universal for the definition of probabilistic processes on S , so that we can always assume that $\Omega = \Omega^S$. The random variables $(X_n)_{n \geq 0}$ are then given as the coordinate projections $X_n : \Omega \rightarrow S$, for $n \geq 0$, so that $X_n(\omega) = x_n$ for $\omega = (x_n)_{n \geq 0}$.

We will be quite flexible with our notations, following the usual traditions from probability theory. For instance, $\{X_0 = i\}$ is a short notation for $\{\omega \in \Omega : X_0(\omega) = i\}$, where $i \in S$. Also, $\mathbb{P}(X_{n+1}|X_0, \dots, X_n)$ must be understood as the function $(x_0, \dots, x_{n+1}) \in S \times \dots \times S \mapsto \mathbb{P}(X_{n+1} = x_{n+1}|X_0 = x_0, \dots, X_n = x_n)$. This allows a more comfortable support for the probabilistic intuition.

We say that the process is a (non homogeneous) *Markov chain* if, for any integer $n \geq 0$, $\mathbb{P}(X_{n+1}|X_0, \dots, X_n) = \mathbb{P}(X_{n+1}|X_n)$ holds true. In this case, the process is entirely characterized by 1) its initial distribution, i.e., the probability distribution μ of X_0 over S defined by $\mu(i) = \mathbb{P}(X_0 = i)$ for i ranging over S , and by 2) the sequence of stochastic matrices $(P_n)_{n \geq 0}$ defined by $P_n(i, j) = \mathbb{P}(X_{n+1} = j|X_n = i)$, for $(i, j) \in S \times S$. The

matrices P_n are called the *transition matrices* of the chain. Conversely, given any probability distribution μ over S and any sequence of stochastic matrices $(P_n)_{n \geq 0}$ of size $|S|$, there is a unique (non homogeneous) Markov chain \mathbb{P} with initial distribution μ and with $(P_n)_{n \geq 0}$ as sequence of transition matrices.

Given a Markov chain \mathbb{P} over S , we say that the chain is *homogeneous* if all transition matrices are identical to a given stochastic matrix P . Matrix P is then called *the* transition matrix of the chain. Given a fixed transition matrix P of size $|S|$, we denote by \mathbb{P}_μ the unique homogeneous Markov chain with initial distribution μ over S and with transition matrix P . If x is a state of the chain, we denote by \mathbb{P}_x the homogeneous Markov chain associated with initial Dirac distribution δ_x concentrated on x . A homogeneous Markov chain is usually represented as a graph with S as set of vertexes, and with an arc from state i to state j whenever $P(i, j) > 0$ (i.e., if there is a positive probability to go from state i to state j). The arc is then labeled with the value $P(i, j)$.

2.3 Hitting and Return Times. We will use the notion of *hitting* and *return time* of a Markov chain, defined as follows—hitting and return times are particular cases of stopping times, but we do not need the latter in their whole generality. Let B be a subset of S . The hitting time to B is defined as the integer $\sigma(\omega)$, depending on the random outcome $\omega \in \Omega^S$, given by

$$\sigma(\omega) = \inf\{n \geq 0 : X_n(\omega) \in B\},$$

where, by convention, $\sigma(\omega) = +\infty$ if the set $\{n \geq 0 : X_n(\omega) \in B\}$ is empty. Notation $X_\sigma(\omega)$ denotes the element $X_{\sigma(\omega)}(\omega) \in B$, if $\sigma(\omega) < \infty$.

The *return time* to B is only defined for those $\omega = (x_n)_{n \geq 0}$ such that $x_0 \in B$. It is defined as the random integer $\tau(\omega)$ given by

$$\tau(\omega) = \inf\{n \geq 1 : X_n(\omega) \in B\},$$

with the same convention as above, and X_τ is defined by $X_\tau(\omega) = X_{\tau(\omega)}(\omega)$ if $\tau(\omega) < \infty$.

2.4 Two-Components Processes. Let A and C be two finite sets, and let $B = A \cap C$. Let $a = (a_0, a_1, \dots)$ be a sequence, finite or infinite, with values in A . Let n_1 be the finite or infinite number of visits of a in B and let

$$\begin{cases} (a'_1, \dots, a'_{n_1}), & \text{if } n_1 < \infty, \\ (a'_1, a'_2, \dots), & \text{if } n_1 = \infty, \end{cases}$$

be the states of B visited by a , in this order. Similarly, let $c = (c_0, c_1, \dots)$ be a sequence with values in C . Let n_2 be the finite or infinite number of visits of c in B , and let (c'_1, \dots, c'_{n_2}) or (c'_1, c'_2, \dots) be the states of B visited by c , in this order. We say that the sequences a and c *synchronize* if:

$$n_1 = n_2, \quad \text{and} \quad \begin{cases} \forall k \geq 1, & k \leq n_1 \Rightarrow a'_k = c'_k, & \text{if } n_1 < \infty, \\ \forall k \geq 1, & a'_k = c'_k, & \text{if } n_1 = \infty. \end{cases}$$

Call *trace* a pair (a, c) of sequences that synchronize. We say that a trace $t = (a, c)$ is *finite* whenever both a and c are finite sequences. If $t = (a, c)$ is a finite trace, the *marking reached by* x is defined as the pair $(i, j) \in A \times C$ such that i (respectively, j) is the last element of sequence a (respectively, of sequence b). We denote it by $\gamma(t) = (i, j)$.

Finally, we say that of pair $(i, j) \in A \times C$ is a *synchronizing pair of states* if (i, j) satisfies: $(i \in B \iff j \in B \iff i = j)$. Be careful that the marking of the system shall not be necessarily a synchronizing pair of states.

2.5 Markovian Two-Components Processes. The prefix ordering \leq on sequences induces a partial order on traces defined as follows: $(u, v) \leq (a, c)$ if $u \leq a$ and $v \leq c$. It is well known that this partial ordering makes the set of traces a prime ω -algebraic domain [7]. In particular, it follows that the set of maximal elements of this domain is not empty; we denote it by Ω . Generic elements of Ω are denoted by ω and decompose as $\omega = (\omega^A, \omega^C)$. Set Ω is equipped with the Borel σ -field generated by the topology induced on Ω by the Scott topology on the domain. Equivalently, if we denote by $\uparrow t$ the subset $\uparrow t = \{\omega \in \Omega : t \leq \omega\}$ for any trace $t = (a, c)$, the Borel σ -field is generated by the countable collection of subsets $\uparrow t$, for t ranging over the set of finite traces.

We say that the two-component process defined by the pair (A, C) is *randomized* if a probability measure \mathbb{P} has been defined on Ω .

Observe that if $t_1 = (a_1, c_1)$ and $t_2 = (a_2, c_2)$ are two traces, with t_1 finite, then the concatenated components $a_1 \cdot a_2$ and $c_1 \cdot c_2$ are two sequences that synchronize. We define thus the *concatenation* of t_1 and t_2 by

$$t_1 \cdot t_2 = (a_1 \cdot a_2, c_1 \cdot c_2).$$

Let t be a finite trace such that $\mathbb{P}(\uparrow t) > 0$. Define the *probabilistic future of* t as the conditional probability measure \mathbb{P}^t on Ω defined as follows, for U a measurable subset of Ω :

$$\mathbb{P}^t(U) = \mathbb{P}(t \cdot U | \uparrow t) = \frac{1}{\mathbb{P}(\uparrow t)} \mathbb{P}(t \cdot U),$$

where $t \cdot U = \{t \cdot y, y \in U\}$. Say the two-component process \mathbb{P} is *Markovian* [1] whenever the following holds, for any two finite traces t and t' :

$$\gamma(t) = \gamma(t') \Rightarrow \mathbb{P}^t = \mathbb{P}^{t'}.$$

This definition meets the intuitive notion of a Markovian system, *i.e.*, a system for which the future only depends on the present state ($=\gamma(t)$), and not on the whole past history.

3 Statement of the Problem

The data of our problem is the following. Let A and C be two finite sets, with intersection $B = A \cap C$. Let $(X_n^A)_{n \geq 0}$ and $(X_n^C)_{n \geq 0}$ be two homogeneous Markov chains, with state spaces A and C , and transition matrices P and Q , respectively. We also consider the two

graphs with sets of vertexes A and C respectively, associated with matrices P and Q as in the end of § 2.2, and the space Ω of maximal traces relative to synchronization of A and C with respect to $B = A \cap C$. The problem is to construct a two-component probabilistic process \mathbb{P} on Ω , with the following properties:

1. The two-component probabilistic process is Markovian.
2. For any finite trace $t = (a, c)$ such that $\mathbb{P}(\uparrow t) > 0$, the component sequences a and c are paths in the graphs A and C , respectively.
3. When the two components are *not* synchronized, they must be independent in probability—to be detailed in § 4.
4. The two-components process has to be constructed “from” the two matrices P and Q .

Next section will be devoted to the construction of a two-components process meeting these requirements.

Before presenting our construction, we wish to discuss the following natural approach for constructing a two-components probabilistic process. Let Ω be the space of maximal traces, and let (Ω^A, \mathbb{P}^A) and (Ω^B, \mathbb{P}^B) be the canonical probability spaces associated with Markov chains $(X_n^A)_{n \geq 0}$ and $(X_n^B)_{n \geq 0}$ respectively. To simplify our discussion, assume that for any maximal trace $\omega = (\omega^A, \omega^C)$, sequences ω^A and ω^C are infinite. Then, there is a natural mapping $\Phi : \Omega \rightarrow \Omega^A \times \Omega^B$, which is injective. Hence Ω can be seen as a (measurable) subset of $\Omega^A \times \Omega^B$. By referring to Eq. (1), in order to randomize Ω , it is thus tempting to consider the conditional probability $\mathbb{P}(\cdot) = (\mathbb{P}^A \otimes \mathbb{P}^B)(\cdot | \Omega)$ on Ω , where $\mathbb{P}^A \otimes \mathbb{P}^B$ denotes the product probability on the product space $\Omega^A \times \Omega^B$. However, except in some particular cases—one of them being that B has 0 or 1 element only, see below in § 4.6—, Ω has $\mathbb{P}^A \otimes \mathbb{P}^B$ probability zero, and thus the conditional probability is undefined. Therefore this direct approach does not work properly. Nevertheless, it serves as a guideline for our actual approach.

4 Construction of a Synchronization Product

This section contains our main results. They are based on a decomposition of two-components processes that we explain first.

4.1 Decomposition of Traces. We keep the notations introduced above. For each maximal trace $\omega \in \Omega$, we denote by ω^A and ω^C respectively the left and right components of ω , so that we have: $\omega = (\omega^A, \omega^C)$, with $\omega^A \in \Omega^A$ and $\omega^C \in \Omega^C$. The decomposition of an element ω we propose below is depicted in Fig. 3.

Fix for a while an element $\omega \in \Omega$. Denote by $s = (s^A, s^C)$ the largest prefix trace of ω (i.e., with $s \leq \omega$) such that no element in sequences s^A or s^C belongs to B . Then s enjoys the following property: $s = (s^A, s^C)$ is a strict prefix of ω if and only if *at least one* of s^A or s^C is finite, if and only if *both* of s^A and s^C are finite. If s is indeed a strict prefix of ω ,




Figure 3: Decomposition of a trace through “incursions” in B and “parallel excursions” in $A \setminus B$ and $C \setminus B$. Excursions $s_n = (s_n^A, s_n^C)$ may be empty; incursions always have exactly one element, namely $y_n \in B$.

then there exists a unique element $y \in B$ such that $(s^A \cdot y, s^C \cdot y)$ is a prefix of ω . Element y is the first synchronizing state of the pair (ω^A, ω^C) . Then, let r be the tail of ω , i.e., the unique trace such that $\omega = s \cdot (y, y) \cdot r$. Since the definitions of $s = (s^A, s^C)$, of y and of r depend on ω , they define partial mappings

$$S : \omega \mapsto s, \quad \text{with } S = (S^A, S^C), \quad Y : \omega \mapsto y, \quad \text{and} \quad R : \omega \mapsto r.$$

As Fig. 3 suggests, we will then proceed to define inductively sequences of partial mappings, built upon the partial mappings just defined. For this, define first:

$$S_0 = S, \quad Y_0 = Y, \quad \text{and} \quad R_0 = \text{Id},$$

where Id denotes the identity map $\text{Id} : \Omega \rightarrow \Omega$. Then define inductively, for all $n > 0$:

$$\begin{aligned} R_n &= R \circ R_{n-1}, & Y_n &= Y \circ R_{n-1}, \\ S_n^A &= S^A \circ R_{n-1}, & S_n^C &= S^C \circ R_{n-1}, & S_n &= (S_n^A, S_n^C) \end{aligned} \tag{2}$$

where \circ denotes the composition of functions. Finally, for a generic element $\omega \in \Omega$, we denote by s_n the trace $s_n = S_n(\omega)$, and so on for s_n^A, s_n^C, y_n and r_n . An induction argument shows that the following identity always holds, if all elements are well-defined:

$$\forall n \geq 0, \quad \omega = s_0 \cdot (y_0, y_0) \cdot s_1 \cdot (y_1, y_1) \cdot \dots \cdot s_n \cdot (y_n, y_n) \cdot r_{n+1}.$$

4.2 Construction of the Synchronization Product. We are now ready for the randomization of two-components processes. Consider the following assumption:

Assumption 1 Let τ^A and τ^C denote respectively the return times of $(X_n^A)_{n \geq 0}$ and of $(X_n^C)_{n \geq 0}$ in B , and let σ^A and σ^C denote respectively the hitting times of $(X_n^A)_{n \geq 0}$ and $(X_n^C)_{n \geq 0}$ in B (see § 2.3). For every pair $(x, y) \in A \times C$ of synchronizing states:

1. If $x = y$ (and then $x \in B$), then the return times τ^A and τ^C are defined with probability 1 with respect to \mathbb{P}_x^A and \mathbb{P}_x^C ; we ask that

$$\mathbb{P}_x^A(\tau^A < \infty) = 1, \quad \mathbb{P}_x^C(\tau^C < \infty) = 1, \quad \mathbb{P}_x^A \otimes \mathbb{P}_x^C(X_{\tau^A}^A = X_{\tau^C}^C) > 0.$$

2. If $x \neq y$, we ask that

$$\mathbb{P}_x^A(\sigma^A < \infty) = 1, \quad \mathbb{P}_y^C(\sigma^C < \infty) = 1, \quad \mathbb{P}_x^A \otimes \mathbb{P}_y^C(X_{\sigma^A}^A = X_{\sigma^C}^C) > 0.$$

Comment on Assumption 1. — Intuitively, what we ask is that, if $(X_n^A)_{n \geq 0}$ and $(X_n^C)_{n \geq 0}$ are free to evolve independently, they have at least one chance to meet at some time at B . This is not too much to ask. If this does not hold, then the following case must also be analyzed in our construction below: with positive probability, the two components remain in the private parts of their respective state spaces for the rest of their life. Assumption 1 allows to ignore this case. We consider it for simplicity.

Theorem 4.1 *Assume that Assumption 1 is in force. Let μ be a probability distribution over the set of synchronization pairs of states. Then, there exists a unique probability measure \mathbb{P} on Ω such that:*

1. *The probability law of the pair (X_0^A, X_0^C) over $A \times C$ is μ , i.e.:*

$$\forall (i, j) \in A \times C, \quad \mathbb{P}(X_0^A = i, X_0^C = j) = \mu(i, j).$$

2. *With probability 1, S_n, Y_n and R_n are well defined for all $n \geq 0$;*

3. *The probability distribution of the random trace $S_0 \cdot (Y_0, Y_0)$ is given by the conditional probability distribution*

$$\mathbb{P}_{X_0^A}^A \otimes \mathbb{P}_{X_0^C}^C((X_0^A, \dots, X_{\sigma^A}^A), (X_0^C, \dots, X_{\sigma^C}^C) | X_{\sigma^A}^A = X_{\sigma^C}^C),$$

where σ^A and σ^C denote the hitting times to B as in Assumption 1.

4. *For every $n \geq 0$, the probability distribution of the random trace $S_n \cdot (Y_n, Y_n)$, conditionally on the random trace*

$$S_0 \cdot (Y_0, Y_0) \cdot \dots \cdot S_{n-1} \cdot (Y_{n-1}, Y_{n-1}),$$

coincides with the conditional probability distribution

$$\mathbb{P}_{Y_{n-1}}^A \otimes \mathbb{P}_{Y_{n-1}}^C((X_1^A, \dots, X_{\tau^A}^A, X_1^C, \dots, X_{\tau^C}^C) | X_{\tau^A}^A = X_{\tau^C}^C),$$

where τ^A and τ^C denote the return times to B as in Assumption 1.

Proof. — For the sake of simplicity, we will assume throughout the proof that the initial measure μ is concentrated on those synchronization pairs of states of the form (b, b) , with $b \in B$. In this case $S_0 = (\epsilon, \epsilon)$, where ϵ denotes the empty sequence, $\sigma^A = 0$ with $\mathbb{P}_{X_0^A}^A$ -probability 1, and $\sigma^C = 0$ with $\mathbb{P}_{X_0^C}^C$ -probability 1, and thus $X_{\sigma^A}^A = X_{\sigma^C}^C = b$ with $\mathbb{P}_{X_0^A}^A \otimes \mathbb{P}_{X_0^C}^C$ -probability 1. Therefore condition 3 above reduces to a trivially true condition, so that

we simply ignore it in the remaining of the proof—treating the general case would involve similar techniques as the ones used below.

1. *Existence.* Let E be the set of finite traces $t = (a, c)$ such that all elements in a (respectively, in c) belong to $A \setminus B$ (respectively, to $C \setminus B$), except for the last ones that belong to B —and thus, coincide. Let also x be an element of B . Observe that, thanks to Assumption 1, the set $\{X_{\tau^A}^A = X_{\tau^C}^C\}$ has $\mathbb{P}_x^A \otimes \mathbb{P}_x^C$ probability non zero. Therefore, we legitimately define a probability ν_x on E by setting, for every $t \in E$:

$$\nu_x(t) = \mathbb{P}_x^A \otimes \mathbb{P}_x^C((X_1^A, \dots, X_{\tau^A}^A, X_1^C, \dots, X_{\tau^C}^C) = t | X_{\tau^A}^A = X_{\tau^C}^C).$$

Consider then the Markov chain $(\xi_n)_{n \geq 0}$ with state space E , initial distribution $\sum_{x \in B} \nu_x \mu(x)$, and transition matrix $Q(t, t')$ given by:

$$\forall t, t' \in E, \quad Q(t, t') = \nu_z(t'), \quad \text{where } (z, z) = \gamma(t).$$

Let (Ξ, \mathbb{Q}) denote the canonical probability space associated with Markov chain $(\xi_n)_{n \geq 0}$, so that $\Xi = E^{\mathbb{N}}$. Consider also the canonical mapping $\Phi : \Xi \rightarrow \Omega$ defined by the following infinite concatenation:

$$\Phi((\xi_n)_{n \geq 0}) = \xi_0 \cdot \xi_1 \cdot \xi_2 \cdots,$$

and the image probability \mathbb{P} on Ω induced by Φ (i.e., $\mathbb{P}(F) = \mathbb{Q}(\Phi^{-1}(F))$ for any measurable subset F of Ω). Routine checks show that \mathbb{P} thus defined satisfies the conditions of the theorem.

2. *Uniqueness.* Since all variables S_n are supposed to be defined with probability 1, any $\omega \in \Omega$ can be written as the following infinite concatenation, with probability 1:

$$\omega = s_0 \cdot (y_1, y_1) \cdot s_1 \cdot (y_2, y_2) \cdot s_2 \cdots$$

Therefore \mathbb{P} is entirely determined by the successive conditional laws described in point 4 in the statement of Th. 4.1. Uniqueness of \mathbb{P} follows. |

Definition 4.2 *We define the probability \mathbb{P} given by Theorem 4.1 as the synchronization product of chains $(X_n^A)_{n \geq 0}$ and $(X_n^C)_{n \geq 0}$, starting from initial measure μ .*

4.4 Two Key Properties of the Synchronization Product. The probabilistic synchronization product enjoys two key properties, as stated by Th. 4.4 below. Firstly, it is Markovian, in the sense of true-concurrent probabilistic processes, as explained in § 2.5. Secondly, local concurrent components are “locally independent”. So far this is an informal statement. We formalize it next.

Definition 4.3 *Let \mathbb{P} be a probabilistic two-components process. We use the notations Y_n, S_n^A, S_n^C introduced in Eq. (2) to decompose the two-components process. We say that the two-components process has the local independence property if its local components are*

probabilistically independent, conditionally on their start and end; i.e., if the following holds, for every integer $n \geq 0$:

$$\mathbb{P}(S_n^A, S_n^C \mid Y_{n-1}, Y_n) = \mathbb{P}(S_n^A \mid Y_{n-1}, Y_n) \times \mathbb{P}(S_n^C \mid Y_{n-1}, Y_n), \quad (3)$$

where by convention Y_{-1} is given an arbitrary value null, and thus can be simply dropped from the above equation if $n = 0$.

Theorem 4.4 *For any initial probability distribution μ on B , the synchronization product defines a Markovian two-components probabilistic process that enjoys the local independence property.*

Proof. — Let \mathbb{P} denote the probability driving the two-component process. To avoid the particular case of the beginning, we assume as in the proof of Th. 4.1 that the initial measure μ is concentrated on synchronization pairs of states of the form (b, b) with $b \in B$. Then, by decomposing μ as a weighted sum of Dirac measures, we assume furthermore without loss of generality that μ is itself a Dirac measure $\mu = \delta_{(b,b)}$. We then have $S_0 \cdot (Y_0, Y_0) = (b, b)$ with probability 1.

1. *The probabilistic process (Ω, \mathbb{P}) is Markovian.* We use the notation \mathbb{P}^t to denote the probabilistic future relative to a finite trace t , as defined in § 2.5; we have to show that \mathbb{P}^t depends only on $\gamma(t)$.

By construction of \mathbb{P} , it is clear that \mathbb{P}^t only depends on $\gamma(t)$ if $\gamma(t)$ has the form $\gamma(t) = (z, z)$ for some element $z \in B$. Therefore, we may assume without loss of generality that t has the form $t = (b, b) \cdot (a, c)$ where a and c are sequences of elements in $A \setminus B$ and in $C \setminus B$ respectively. Denote also by (x, y) the pair of synchronizing states given by the marking $\gamma(t) = (x, y)$. What remains to be shown is that, for any finite trace t' , the quantity $\mathbb{P}^t(\uparrow t')$ only depends on $\gamma(t)$ and t' . Using again the Markov property for traces that end in B , we only consider without loss of generality traces

$$t' = (a', c')$$

such that a' and c' have no element in B , except maybe their last element. Accordingly, we investigate two cases for $t' = (a', c')$.

In the first case, we assume that the last elements of a' and c' indeed belong to B . In this case we have:

$$\{S_0 \cdot (Y_0, Y_0) \cdot S_1 \cdot (Y_1, Y_1) = t \cdot t'\} = \uparrow t \cdot t'.$$

Therefore, using the notation $\Delta = \{X_{\tau^A}^A = X_{\tau^C}^C\}$, and using statement 4 of Th. 4.1 we have:

$$\mathbb{P}^t(\uparrow t') = \frac{1}{\mathbb{P}(\uparrow t)} \mathbb{P}(\uparrow (t \cdot t')) = \frac{\mathbb{P}^A \otimes \mathbb{P}^C(\uparrow (t \cdot t') \mid \Delta)}{\mathbb{P}^A \otimes \mathbb{P}^C(\uparrow t \mid \Delta)},$$

where \mathbb{P}^A and \mathbb{P}^C are short notations for \mathbb{P}_b^A and \mathbb{P}_b^C respectively. Observing that $\Delta \cap \uparrow(t \cdot t') = \uparrow(t \cdot t')$, we get:

$$\begin{aligned} \mathbb{P}^t(\uparrow t') &= \frac{\mathbb{P}^A \otimes \mathbb{P}^C(\uparrow(t \cdot t') \cap \Delta)}{\mathbb{P}^A \otimes \mathbb{P}^C(\uparrow t \cap \Delta)} = \frac{\mathbb{P}^A \otimes \mathbb{P}^C(\uparrow(t \cdot t'))}{\mathbb{P}^A \otimes \mathbb{P}^C(\uparrow t)} \times \frac{1}{\mathbb{P}^A \otimes \mathbb{P}^C(\Delta | \uparrow t)} \\ &= \mathbb{P}_x^A \otimes \mathbb{P}_y^C(a', c') \times \frac{1}{\mathbb{P}_x^A \otimes \mathbb{P}_y^C(\Delta)} = \mathbb{P}_x^A \otimes \mathbb{P}_y^C((a', c') | \Delta). \end{aligned}$$

Now it is clear on the last expression that $\mathbb{P}^t(\uparrow t')$ only depends on $t' = (a', c')$ and on $\gamma(t) = (x, y)$.

For the second case, we assume that t' has the form $t' = (a', c')$, where no element neither in a' nor in c' belongs to B . Since S_1 and Y_1 are defined with \mathbb{P} -probability 1 according to point 2 of Th. 4.1, we use the so-called ‘‘total probabilities law’’ to get the following expression for $\mathbb{P}^t(\uparrow t')$:

$$\mathbb{P}^t(\uparrow t') = \sum_{u \geq t'} \frac{1}{\mathbb{P}(\uparrow t)} \mathbb{P}^t(S_0 \cdot (Y_0, Y_0) \cdot S_1 \cdot (Y_1, Y_1) = t \cdot u),$$

where u ranges over the set of finite traces that satisfy the conditions of our first case. Therefore:

$$\mathbb{P}^t(\uparrow t') = \sum_{u \geq t'} \mathbb{P}^t(\uparrow u).$$

Using the first case, we see that this expression that $\mathbb{P}^t(\uparrow t')$ only depends on $\gamma(t)$ and t' . The Markov property is thus established.

2. Local independence property. We have to prove that Eq. (3) from Def. 4.3 is valid for any $n \geq 0$. The case $n = 0$ is a slightly different case from the case $n \geq 1$; but since it uses the same techniques as below, we simply omit it and focus on the case $n \geq 1$. Thanks to the Markovian property just shown, we assume without loss of generality that $n = 1$. We use the short notations \mathbb{P}^A and \mathbb{P}^C for \mathbb{P}_b^A and \mathbb{P}_b^C respectively, where we recall that $\mu = \delta_{(b,b)}$. Consider two generic sequences a and c of elements in A and C respectively, and a generic element $y \in B$. Using the notations $U^A = \{S_1^A = a\}$, $U^C = \{S_1^C = c\}$ and $U = U^A \cap U^C$, showing Eq. (3) reduces to showing the following, for any generic a, c and y :

$$\mathbb{P}(U|Y_1 = y) = \mathbb{P}(U^A|Y_1 = y) \times \mathbb{P}(U^C|Y_1 = y), \tag{4}$$

since $Y_0 = b$ with \mathbb{P} -probability 1. Consider as above the notation $\Delta = \{X_{\tau^A}^A = X_{\tau^C}^C\}$. Then we have, by point 4 of Th. 4.1 and by the definition of conditional probabilities:

$$\mathbb{P}(U|Y_1 = y) = \frac{\mathbb{P}^A \otimes \mathbb{P}^C(U, Y_1 = y | \Delta)}{\mathbb{P}^A \otimes \mathbb{P}^C(Y_1 = y | \Delta)} = \frac{\mathbb{P}^A \otimes \mathbb{P}^C(U, Y_1 = y, \Delta)}{\mathbb{P}^A \otimes \mathbb{P}^C(Y_1 = y, \Delta)}.$$

Observe that we have $\{Y_1 = y\} \cap \Delta = \{X_{\tau^A}^A = y, X_{\tau^C}^C = y\}$, and therefore:

$$\begin{aligned} \mathbb{P}(U|Y_1 = y) &= \frac{\mathbb{P}^A(U^A, X_{\tau^A}^A = y) \times \mathbb{P}^C(U^C, X_{\tau^C}^C = y)}{\mathbb{P}^A(X_{\tau^A}^A = y) \times \mathbb{P}^C(X_{\tau^C}^C = y)} \\ &= \mathbb{P}^A(U^A | X_{\tau^A}^A = y) \times \mathbb{P}^C(U^C | X_{\tau^C}^C = y). \end{aligned} \tag{5}$$

From Eq. (5), it follows in particular the following, by summing over c to get Eq. (6) and by summing over a to get Eq. (7):

$$\mathbb{P}(U^A | Y_1 = y) = \mathbb{P}^A(U^A | X_{\tau^A}^A = y) \quad (6)$$

$$\mathbb{P}(U^C | Y_1 = y) = \mathbb{P}^C(U^C | X_{\tau^C}^C = y). \quad (7)$$

Putting the left members of Eqs. (6) and (7) back in Eq. (5), we get Eq. (4), which was to be shown. \blacksquare

Comment on the local independence property. — Because of the Markovian property, conditioning w.r.t. Y_{n-1} in Eq. (3) is just like considering the probabilistic future, or, to consider that the process has just started. Much more critical is the conditioning w.r.t. Y_n . If we fail to apply it, Eq. (3) does not hold true anymore since processes S_n^A and S_n^C are *not* independent in general. To informally explain this, observe that, from the last two states visited by S_n^A and S_n^C , each component has to jump to the *same* state in B . Note in passing that this reasoning does not work when there is only *one* common state, i.e., $|B| = 1$; we provide more details on this case in the next subsection.

4.6 A Particular Case: Synchronization on One Common State. Assume that the synchronization set B has a unique element. Nothing prevents the construction of Th. 4.1 from working. But the result might be surprising at a first glance.

Proposition 4.5 *Assume that $B = \{b\}$ is a singleton, and let \mathbb{P} be a synchronization product of Markov chains with respect to synchronization space state B , with initial measure μ . Let μ^A and μ^C be the initial distributions induced by \mathbb{P} on A and on C respectively, i.e.:*

$$\forall x \in A, \quad \mu^A(x) = \mathbb{P}(X_0^A = x), \quad \text{and} \quad \forall y \in C, \quad \mu^C(y) = \mathbb{P}(X_0^C = y).$$

Then μ coincides with the product probability $\mu = \mu^A \otimes \mu^C$, and \mathbb{P} coincides with the product probability $\mathbb{P}_{\mu^A}^A \otimes \mathbb{P}_{\mu^C}^C$.

Proof. — We write \mathbb{P}^A and \mathbb{P}^C as short notations for \mathbb{P}_b^A and \mathbb{P}_b^C , respectively. Since B has a unique element, we obviously have $Y_n = b$ for all $n \geq 1$. Moreover, we have: $\mathbb{P}^A(X_{\tau^A}^A = b) = 1$ and $\mathbb{P}^C(X_{\tau^C}^C = b) = 1$, and thus $\mathbb{P}^A \otimes \mathbb{P}^C(X_{\tau^A}^A = X_{\tau^C}^C) = 1$. It follows that, for any $n \geq 0$, the conditional probability law below trivially reduces to a product probability, as follows:

$$\begin{aligned} \mathbb{P}_{Y_n}^A \otimes \mathbb{P}_{Y_n}^C & \left((X_1^A, \dots, X_{\tau^A}^A, X_1^C, \dots, X_{\tau^C}^C) \mid X_{\tau^A}^A = X_{\tau^C}^C \right) \\ & = \mathbb{P}_b^A \otimes \mathbb{P}_b^C \left((X_1^A, \dots, X_{\tau^A}^A, X_1^C, \dots, X_{\tau^C}^C) \right). \end{aligned}$$

By induction, and thanks to point 4 in Th. 4.1, it follows that, for any $n \geq 1$:

$$\mathbb{P}((Y_0, Y_0) \cdot S_0 \cdot \dots \cdot S_{n-1} \cdot (Y_n, Y_n)) = \mathbb{P}^A \otimes \mathbb{P}^C((Y_0, Y_0) \cdot S_0 \cdot \dots \cdot S_{n-1} \cdot (Y_n, Y_n)),$$

and thus $\mathbb{P} = \mathbb{P}^A \otimes \mathbb{P}^C$. \blacksquare

Comment. — Informally, Proposition 4.5 says that synchronization on one single state has *no* effect at all on the probability distribution of the individual components. This does not mean that synchronization has no effect on the two-component process: it of course eliminates lots of pairs of infinite sequences as valid traces. But each individual sequence from one component can always find another one for marrying, from the other component.

Also, observe that although μ^A and μ^C can always be defined, the identity $\mu = \mu^A \otimes \mu^C$ does not hold in general if $|B| \neq 1$.

Another way to say the same thing, in relation to our comment at the end of § 3, is to see Ω as a subset of $\Omega^A \times \Omega^C$, and to observe that $\mathbb{P}^A \otimes \mathbb{P}^C(\Omega) = 1$, and thus:

$$\mathbb{P}(\cdot) = \mathbb{P}^A \otimes \mathbb{P}^C(\cdot | \Omega) = \mathbb{P}^A \otimes \mathbb{P}^C(\cdot).$$

5 Conclusion and Perspectives

In this paper we have proposed a method to randomize two-components processes, from a partial order point of view—*i.e.*, by randomizing traces rather than sequences. We have shown how to construct a two-component process by synchronizing two Markov chains. The so obtained two-component process possesses two key properties: 1/ it is Markovian, and 2/ it has a local independence property.

Conversely, we have studied two-components probabilistic processes having the above two properties and have asked ourselves whether such processes can always be considered as a synchronization of Markov chains. We conjecture that the answer to the latter question is affirmative. We also doubt that there is uniqueness, but expect that there is still a canonical reverse operation to the synchronization product of Markov chains—the categorical properties of which would then be worth detailed.

Several other research directions can also be considered. A generalization to more than two Markov chains is certainly interesting. This would lead to a way of randomizing 1-safe Petri nets that cannot be randomized otherwise—all true-concurrent randomizations that have been proposed so far suffer from limitations, confusion-free in [9], locally finite in [3, 2]. Of course, a natural problem is also to compare such a randomization, if it can indeed be achieved, with the randomization proposed by the authors just cited. The main issue is the property of local independence in case of more than two local processes, which is quite challenging to formulate for the synchronization of an arbitrary number of local processes.

We have been interested in Markovian processes. However, for sequential systems, Markov chains are only one class among several others that deserve interest. In particular, stationary processes (processes that have an identical behavior, in probability, when time is shifted) is a very important class of stochastic processes, that has attracted interest since the work of Poincaré in the beginning of the 20th century. Finding an adequate notion of two-components stationary process is a very promising research direction. On a larger scale, it is maybe time to develop a theory of multi-component stochastic processes. A major goal would be to define the notions of stationary processes, of ergodic processes, and to obtain a related ergodic theorem. In this respect, the notion of stopping operator developed

in [1] seems to be an adequate tool to replace the usual *stopping times*, a fundamental tool in the study of usual stochastic processes.

References

- [1] S. Abbes. The (true) concurrent Markov property and some applications to Markov nets. In G. Ciardo and P. Darondeau, editors, *ICATPN 2005*, volume 3536 of *LNCS*, pages 70–89, 2005.
- [2] S. Abbes and A. Benveniste. Probabilistic true-concurrency models: branching cells and distributed probabilities for event structures. *Information and Computation*, 204(2):231–274, 2006.
- [3] S. Abbes and A. Benveniste. Probabilistic true-concurrency models: Markov nets and a law of large numbers. To appear in *Theoretical Computer Science*, 2007.
- [4] E. Bellman. A Markovian decision process. *Journal of Mathematics and Mechanics*, 6:679–684, 1957.
- [5] P.R. d’Argenio and J.-P. Katoen. A theory of stochastic systems. Part II: Process algebra. *Information and Computation*, 203(1):39–74, 2005.
- [6] N. Lynch, R. Segala, and F. Vaandrager. Compositionality for probabilistic automata. In R. Amadio and D. Lugiez, editors, *CONCUR 2003*, volume 2761 of *LNCS*, pages 208–221, 2003.
- [7] M. Nielsen, Plotkin. G., and G. Winskel. Petri nets, event structures and domains, part 1. *Theoretical Computer Science*, 13:85–108, 1981.
- [8] R. Segala. A compositional trace-based semantics for Probabilistic Automata. In I. Lee and S.A. Smolka, editors, *CONCUR 1995*, volume 962 of *LNCS*, pages 234–248, 1995.
- [9] D. Varacca, H. Völzer, and G. Winskel. Probabilistic event structures and domains. In P. Gardner and N. Yoshida, editors, *CONCUR 2004*, volume 3170 of *LNCS*, pages 481–496, 2004.