A truly concurrent synchronization product of Markov chains

Samy Abbes

A Truly Concurrent Synchronization Product of Labelled Markov Chains

Samy Abbes*

May 2005

Abstract

In this paper we introduce a product operation on labeled Markov chains. Whereas this kind of product is most usually achieved under an interleaving semantics, for instance in the framework of probabilistic automata, our construction stays within the true-concurrent semantics. Hence the product of two labeled Markov chains we define is a so-called probabilistic Petri net, i.e. a safe Petri net where Mazurkiewicz traces are randomized, not interleavings. We show that this construction is not trivial as far as the number of synchronization transitions is greater or equal than 2. Our main result is that the product of Markov chains remains Markovian, in the sense of probabilistic true-concurrent systems.

Key words: Probabilistic Petri nets, synchronization product

Introduction

Probabilistic systems are used to solve the nondeterminism of systems by means of probabilistic choices. Probabilistic extensions of models from Computer science fall thus into the large mathematical class of *dynamical systems*. When the systems in question are actually concurrent systems, such a direct probabilistic extension is still possible. For instance Stochastic Petri Nets (SPNs) are very often defined as Continuous Time Markov Chains [4], and products of probabilistic automata are randomized as Discrete Time Markov Chains [5]. This embedding into dynamical systems involves the so-called interleavings semantics of systems: events are chronologically ordered, and their occurrences are randomized accordingly. However, the cost of this direct transposition is that of a huge state-space, which brings in particular issues for the verification of systems.

Motivated by both practical and theoretical issues, and in particuliar to tackle this state-explosion problem, Concurrency theory has introduced another semantics for concurrent systems, called the *true-concurrent*, or *(Mazurkiewicz)*

*From work at ISR (Institute for Systems Research), University of Maryland at College Park (MD), USA

traces semantics. Here, all events are not necessarily chronologically ordered, so that an execution of the system is constituted by a partial order of events. Probabilistic extensions shall thus directly randomize the partial orders; i.e., traces are randomized, not firing sequences. Probabilistic models studied until now from this viewpoint are safe Petri nets and event structures [8, 2]. Whereas an execution of a Markov chain is a succession of choices that are inductively randomized, we have shown that an execution of a safe Petri net can be regarded as a partial order of choices [2]. Executions of the net are randomized accordingly, with concurrent choices being independent in the probabilistic sense. Using the unfolding theory of safe Petri nets [6], choices are made inside so-called branching cells, which are sub-occurrence nets of the unfolding of a safe Petri net. For a Markov chain, branching cells simply correspond to the possible successors of the current state. For Petri nets, branching cells are finitely many up to isomorphism of labelled occurrence nets, and play the role of *local state* of the net.

Unfortunately, in this construction, the finiteness of local choices is not guaranteed. Although this appears for a large class of Petri nets—in particular, for those whose unfoldings have the so-called *local finiteness* property—, we may indeed be faced with infinite branching cells. Moreover, the local finiteness is not stable under composition of nets. This is unfortunate, since a composition theory of probabilistic concurrent systems would be appreciated in several engineering fields. This paper proposes some advances towards such a theory, by introducing a probabilistic construction for the product of two labelled Markov chains. The product of the transition systems that underlie two Markov chains remain within the category of Petri nets, but may have a non locally finite unfolding. Hence, the probabilistic product that we propose is already a real progress within the theory of true-concurrent probabilistic systems. Ideally, it would be a preliminary work for a general theory of synchronization product of an arbitrary number of probabilistic Petri nets. In turn, this would be the crucial synchronization operation in a language of probabilistic Petri nets "à la CCS".

We are given two ergodic Markov chains $X^1 = (X^1_n)_n$ and $X^2 = (X^2_n)_n$, seen as particles in a finite state space, moving from state to state through arcs. We wish to constraint the chains to synchronize on some common arcs. Let $\mathbb{P}^1$ and $\mathbb{P}^2$ be the probabilities on canonical spaces associated to the Markov chains X^1 and X^2. As a first trial, we may say: a probability is constructed by conditioning the direct product probability $\mathbb{P}^1 \otimes \mathbb{P}^2$ on the property that executions of X^1 and X^2 indeed synchronize. Unfortunately, as soon as X^1 and X^2 are constrained to synchronize on two or more common arcs, this property has in most cases probability zero. Hence, this direct construction cannot work. Instead, we propose a *recursive conditioning* that makes X^1 and X^2 synchronize, based on the study of the return times of X^1 and X^2 to their synchronization transitions. Besides this non-trivial construction, the main contribution of the paper is to show that the resulting system is still Markovian, in the sense of probabilistic true-concurrent systems [1]. Finally, the case where X^1 and X^2 synchronize only on 1 shared arc brings a large and surprising simplification:

although constrained to synchronize, X^1 and X^2 remain independent in the probabilistic sense.

Outline. §1 presents the problem. §2 analyzes the case of synchronization on a unique arc, and §3 is devoted to the general case. Proofs of results are postponed in Appendix A.2.

1 Statement of the Problem

We assume that the reader is familiar with Petri nets and with the unfolding theory of safe Petri nets [6].

Foreword on Probability. We assume some familiarity of the reader with probability and Markov chains theory (for these topics, we refer to [3, 7]). We will underlie the (natural) σ-algebras involved in the following probabilistic constructions. We recall that, if $\mathbb{P}$ is a probability on a space Ω, a property is said to hold $\mathbb{P}$-a.s. ($\mathbb{P}$-almost surely) if it holds with probability 1. If $f : \Omega \to E$ is a (measurable) mapping, the *probability law* of f is the probability $f\mathbb{P}$ on E defined, for D a measurable subset of E, by $f\mathbb{P}(D) = \mathbb{P}\big(f^{-1}(D)\big)$.

Data of the Problem. We consider two disjoint finite sets S^1 and S^2, and two Markov chains $X^1 = (X^1_n)_{n\geq 0}$ and $X^2 = (X^2_n)_{n\geq 0}$, defined on S^1 and S^2 respectively, with transition matrices P^1 and P^2.

For each $i = 1, 2$, we consider the graph G^i whose vertices are indexed by S^i, and with an arc from x to y if $P^i_{x,y} > 0$. For a an arc of G^i, $\partial^i_+(a)$ and $\partial^i_-(a)$ respectively denote the ending and starting vertices of a. a is entirely determined by the values $\partial^i_+(a)$, $\partial^i_-(a)$, so that we write $a = \big(\partial^i_-(a)\partial^i_+(a)\big)$.

Denoting by $\mathbb{N}$ the set of natural integers, let Ξ^i, $i = 1, 2$, be the product space $(S^i)^{\mathbb{N}}$, equipped with its natural σ-algebra. The canonical sample space that we consider for the Markov chain X^i is the subset $\Omega^i \subseteq \Xi^i$ consisting of those sequences $\omega^i = (x^i_n)_{n\geq 0}$ such that $(x^i_n x^i_{n+1})$ is an arc of G^i for all $n \geq 0$. Ω^i can be regarded as the set of infinite paths in G^i. Two equivalent representations for elements of Ω^i are infinite sequences of states or infinite sequences of arcs that draw actual paths in G^i. We will make use of one representation or the other, preferring the more convenient according to the context. We will use the notation $(X^i_n)_{n\geq 0}$ for the sequence of states, and $(Y^i_n)_{n>0}$ to refer to the sequence of arcs, with $Y^i_n = (X^i_{n-1}X^i_n)$ for all $n > 0$. For $x \in S^i$, we denote by $\mathbb{P}^i_x$ the probability measure on Ω^i associated with the Markov chain X^i starting from the initial state x. We assume that both chains X^i, $i = 1, 2$, are ergodic (i.e., aperiodic and with only one recurrence class). It is well known that the $(Y^i_n)_{n>0}$ are then also two ergodic Markov chains.

Synchronization Product. We will perform the synchronization product of G^1 and G^2 along common transitions. For this, we first label the arcs of G^1 and G^2 with labels t^1_j and t^2_k respectively, $j, k = 1, 2, \ldots$, in such a way that for each $i = 1, 2$, different arcs of G^i have different labels. Hence we simply identify arcs with their labels. The main point is that we authorize $t^1_j = t^2_k$ for some indices j, k. We denote by A the set of common labels, i.e. labels t such that $t = t^1_j = t^2_k$ for some pair (j, k)—and then such a pair (j, k) is unique. Elements of A are called *public transitions*; while, for each $i = 1, 2$, arcs of G^i not in A are called *transitions private to* G^i.

We may refer to the labeled graphs G^i as to *transition systems*. Their *synchronization product* is the safe Petri net defined as follows. The set of places P is the union $P = S^1 \cup S^2$, while the set of transitions T is the set of labels. For each $i = 1, 2$, if t is a transition private to G^i, with $t = (xy)$, we draw an arc from place x to transition t, and from transition t to place y. If t is a public transition, let j, k be such that $t = t^1_j = t^2_k$. Then we draw arcs from $\partial^1_-(t^1_j)$ to t and from $\partial^2_-(t^2_k)$ to t, from t to $\partial^1_+(t^1_j)$ and from t to $\partial^1_+(t^2_k)$. This way, we have constructed a Petri net $N = (P, T, F)$, where F is the flow relation between places and transitions. If x^1_0, x^2_0 are initial states in G^1 and G^2, the marking m_0 consisting of two tokens, one in x^1_0 and one in x^2_0, is defined as the initial marking of N. The marked net $\mathcal{N} = (N, m_0)$ thus defined is a particular case of synchronization of Petri nets, as defined in [9]. In particular, $\mathcal{N}$ is safe. We list below some results about synchronization products, particularized for our case.

Let s be a firing sequence of $\mathcal{N}$. Define $\pi^1(s)$ as the sequence of transitions obtained from s by deleting the transitions not in G^1. Then $\pi^1(s)$ defines a path in G^1, starting from the initial state x^1_0. $\pi^2(s)$ is the path in G^2 starting from x^2_0 defined symmetrically.

Let $\mathcal{U}_{m_0}$ be the unfolding of the safe Petri net $\mathcal{N}$. $\mathcal{U}_{m_0}$ is an occurrence net [6]. We denote by $\mathcal{V}_{m_0}$ the poset of configurations of $\mathcal{U}_{m_0}$—configurations are ordered by set inclusion. Let v be a configuration of $\mathcal{U}_{m_0}$, and let s be a linearization of v. Then $\pi^1(s)$ and $\pi^2(s)$ only depend on v, and not on the linearization. Therefore π^1 and π^2 induce two mappings, still denoted π^1 and π^2, defined on $\mathcal{V}_{m_0}$.

A configuration $\omega \in \mathcal{V}_{m_0}$ is said to be *maximal* if $\forall v \in \mathcal{V}_{m_0}, v \supseteq \omega \Rightarrow v = \omega$. Every configuration u is a sub-configuration of a maximal configuration ω. We denote by Ω_{m_0} the set of maximal configurations of $\mathcal{U}_{m_0}$. To keep the model simple, we will assume that the labelled graphs satisfy the following requirement:

Definition 1.1. *The product* $\mathcal{N} = (N, m_0)$ *of* G^1 *and* G^2 *with initial states* x^1_0, x^2_0, *is said to be* max-synchronous *if, for every* $\omega \in \Omega_{m_0}$, $\pi^1(\omega)$ *and* $\pi^2(\omega)$ *are both infinite sequences.*

Hence the product is max-synchronous if the synchronization does not introduce any blocking (see Appendix A.1 for computable conditions that guarantee the product to be max-synchronous). We now state informally our problem.

Statement of the Problem. We say that a *probabilistic net* [2] is a pair $(\mathcal{N}, \mathbb{P})$, where $\mathcal{N} = (N, m_0)$ is a safe Petri net with initial marking m_0, and $\mathbb{P}$ is a probability on the space Ω_{m_0}. Considering the labeled graphs G^1 and G^2 as above, *can we construct a probabilistic net associated with the synchronisation product of G^1 and G^2, defined "by means" of the Markov chains X^1 and X^2?* Furthermore, *is the probabilistic net constructed Markovian in the sense of truly concurrent probabilistic Markovian systems?* ([1], see also §3).

Our problem cannot be directly solved by the composition of probabilistic automata. Indeed, in the construction of [5], different interleavings of a same trace are given different probabilities; hence, the randomization is not truly concurrent. The method for randomizing Petri nets that we proposed in [2] relies on the local finiteness assumption of the unfolding, which is in general not satisfied by product nets (see Ex. 2 below). Hence the product approach that we develop here construct probabilities for cases not covered by the local finiteness hypotheses. We give an example of this application below in Ex. 2.

Natural Embedding of Ω into $\Omega^1 \times \Omega^2$. We consider the marked net $\mathcal{N} = (N, m_0)$ constructed as above from G^1 and G^2, and we suppose that the product is max-synchronous. Let s^1, s^2 be two sequences of transitions of G^1 and G^2. We say that s^1 and s^2 *synchronize* if there is a configuration v of $\mathcal{N}$ such that $\pi^1(v) = s^1$ and $\pi^2(v) = s^2$. Our model is kept simple thanks to the following elementary observation (see § A.2).

Lemma 1.1. *If s^1 and s^2 are two transition sequences of G^1 and G^2 that synchronize, then the configuration v such that $\pi^1(v) = s^1$ and $\pi^2(v) = s^2$ is unique.*

We denote for short $\Omega = \Omega_{m_0}$. Since the product is supposed to be max-synchronous, the restrictions of π^1 and π^2 to Ω define mappings $\Omega \to \Omega^1$, $\Omega \to \Omega^2$, and thus a product mapping $\Omega \to \Omega^1 \times \Omega^2$. Now, by the above lemma, this mapping is injective. Hence Ω identifies with its image in $\Omega^1 \times \Omega^2$. We will thus consider in the following Ω as a subset of $\Omega^1 \times \Omega^2$:

$$\boxed{\Omega \subseteq \Omega^1 \times \Omega^2} \tag{1}$$

Denote for short $\mathbb{P}^i = \mathbb{P}^i_{x^i_0}$, $i = 1, 2$. Of course, $\Omega^1 \times \Omega^2$ is naturally equipped with the probability obtained by taking the direct product of probabilities $\mathbb{P}^1 \otimes \mathbb{P}^2$. From (1), one shall answer the problem of constructing a probability on Ω by simply considering the conditional probability $\mathbb{P}^1 \otimes \mathbb{P}^2(\,\cdot\,|\Omega)$. This requires $\mathbb{P}^1 \otimes \mathbb{P}^2(\Omega) > 0$. Unfortunately, we shall see below that in most cases, as soon as the Markov chains X^1 and X^2 are constrained to synchronized on a set of public transitions of cardinal ≥ 2, $\mathbb{P}^1 \otimes \mathbb{P}^2(\Omega) = 0$. Hence, a brute conditioning like this is hopeless. This works however if the chains are constrained to synchronize on at most one public transition; we detail this next.

2 Synchronization on Zero or One Transition

We still consider the above framework. Recall that A denotes the set of public transitions. We fix an initial marking $m_0 = (x^1_0, x^2_0)$, and we consider the associated space $\Omega = \Omega_{m_0}$ and associated probabilities $\mathbb{P}^i = \mathbb{P}^i_{x_0^i}$ on Ω^i for $i = 1, 2$ as above. In case $\mathrm{Card}(A) \leq 1$, we can state (see the proof in A.2):

Theorem 2.1. *If* $\mathrm{Card}(A) \leq 1$*, we have* $\mathbb{P}^1 \otimes \mathbb{P}^2(\Omega) = 1$*. Hence* Ω *is naturally equipped with the direct product probability* $\mathbb{P}^1 \otimes \mathbb{P}^2$*.*

Comment—The result of Th. 2.1 may look surprising. If $\mathrm{Card}(A) = 1$, although the two chains X^1 and X^2 are actually constrained to synchronize on a transition, they remain free in the probabilistic sense, as shown by the product form of the probability on Ω. The fact that the resulting probability is Markovian in the sense of probabilistic nets [1] is a particular case of the more general result stated below, see Th. 3.1.

Example— Theorem 2.1 can already be used to randomize Petri nets whose unfoldings do not have the local finiteness property. We recall the definition of local finiteness [2]: let $(\mathcal{E}, \preceq, \#)$ be the underlying event structure of the unfolding of a net, with $\preceq$ and $\#$ respectively the causality and conflict relations on $\mathcal{E}$. Denote, for $e \in \mathcal{E}$, by $\downarrow(e)$ the set of events $y \preceq e$. Define the immediate conflict $\#_\mu$ relation on $\mathcal{E}$ by: $e \#_\mu f \iff \big(\downarrow(e) \times \downarrow(f)\big) \cap \# = \{(e, f)\}$. Say that a subset $B \subseteq \mathcal{E}$ is a stopping prefix of $\mathcal{E}$ if B is downard closed ($e \in B \Rightarrow \downarrow(e) \subseteq B$) and $\#_\mu$-closed ($\forall e, f \in \mathcal{E}$, $e \in B$ and $e \#_\mu f \Rightarrow f \in B$). We have introduced an inductive decomposition of maximal configurations through *minimal nonempty* stopping prefixes, called *branching cells*. The decomposition works for every safe Petri net; under some mild conditions, it is associated with the construction of a probability on Ω [2], by means of elementary probabilities defined on branching cells. However, this construction is fully effective only if branching cells are finite. This holds in particular if the unfolding is locally finite, i.e. if every event of $\mathcal{E}$ belongs to a *finite* stopping prefix of $\mathcal{E}$. It turns out that, whereas any transition system has a locally finite unfolding, it may not be the case for products of transition systems.

Consider for instance the Petri net depicted in Fig. 1, left-top, obtained by synchronization of the two transition systems depicted at left-bottom, and let $\mathcal{E}$ be the underlying event structure of its unfolding. A prefix of the infinite event structure $H \subseteq \mathcal{E}$ consisting of events $\{a_n, b_m, c_{n,m},\ n, m \geq 1\}$ is depicted at right. H is the smallest stopping prefix of $\mathcal{E}$ that contains a_1, hence $\mathcal{E}$ is not locally finite. H can be regarded as the elementary event structure that composes the whole unfolding. Indeed, $\mathcal{E}$ is obtained by adding after each event $c_{n,m}$, $n, m \geq 1$, a fresh copy of H, and recursively. Accordingly to this decomposition, a probability on Ω can be constructed for instance by a sequence of i.i.d. (independent identically distributed) random variables with values in Ω_H, each Ω_H being a copy of the space of maximal configurations of H. Since the set Ω_H is infinite, this is not a fully effective construction: how is specified the probability law on Ω_H? The product method furnished by Th. 2.1 says more quickly that we can simply consider, in the probabilistic sense, Ω as the product $\Omega^1 \times \Omega^2$.

Conversally, the product method of Th. 2.1 indeed induces an i.i.d. sequence on the iterated copies of Ω_H. Let us explicitely describe the resulting probability law p on Ω_H. Set $u_\infty = (a_1, b_1, a_2, b_2, \dots)$. Then Ω_H is described as:

$$\Omega_H = \{u_{n,m},\ n, m \geq 1\} \cup \{u_\infty\}, \quad \text{with } u_{n,m} = (a_1, \dots, a_{n-1}, b_1, \dots, b_{m-1}, c_{n.m}).$$

Let α and β denote the probabilities of a and of b respectively in the two initial Markov chains, starting from the initial states depicted with tokens in Fig. 1, left-bottom. Then a simple computation shows that p is the product of two geometric laws, $p(u_{n,m}) = (1-\alpha)\alpha^{n-1}(1-\beta)\beta^{m-1}$ for $n, m \geq 1$, and $p(u_\infty) = 0$.

Remark that, without the probabilistic interpretation, the identification $\Omega = \Omega^1 \times \Omega^2$ does not hold, since for example $\omega^1 = (ccc\dots)$ and $\omega^2 = (bbb\dots)$ are two elements of Ω^1 and of Ω^2 that cannot synchronize. To summarize, we shall say: $\mathbb{P}^1$-almost all elements of Ω^1 synchronize with $\mathbb{P}^2$-almost all elements of Ω^2, but the "almost" part of the sentence cannot be removed .

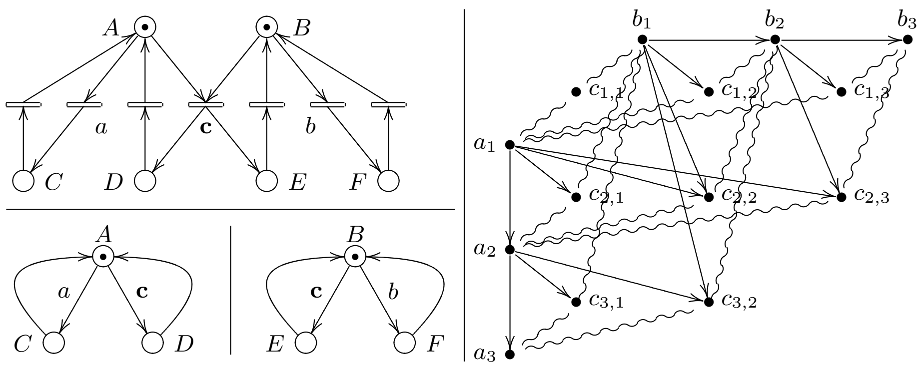

Figure 1: ***A synchronization product with non-locally finite unfolding.*** *Left-top, a safe Petri net obtained by synchronization of the two labelled graphs with shared arc* **c** *(left-bottom). The unfolding of the net contains as a sub-event structure the event structure H depicted at right, to be continued infinitely in both directions (we underlie the transitions that follow states C, D, E and F).*

3 Synchronization on Two or More Transitions

We still consider the above framework. This time, no particular assumption is done on $\text{Card}(A)$, where A is the set of public transitions. $\mathbb{P}^1$ and $\mathbb{P}^2$ are short notations for $\mathbb{P}^1_{x^1_0}$ and $\mathbb{P}^2_{x^2_0}$, defined for some initial states x^1_0 and x^2_0. We introduce the following definition. We note $\overline{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$. For each $i = 1, 2$, and $q > 0$, we denote by R^i_q the q^{th} *return time of Y^i to A*, which is the integer random variable $R^i_q : \Omega^i \to \overline{\mathbb{N}}$ defined inductively, for $\omega^i = (Y^i_n)_{n\geq 0}$ an infinite sequence of arcs, by:

$$R^i_0 = 0, \quad q > 0, \quad R^i_q(\omega^i) = \inf\bigl\{n > R^i_{q-1}(\omega^i) \,:\, Y^i_n \in A\bigr\}, \tag{2}$$

with the convention $\inf(\emptyset) = \infty$. Consider the random variables for $i = 1, 2$ and $q > 0$: $Z^i_q = Y^i_{R^i_q}$, defined by $Z^i_q(\omega^i) = Y^i_{R^i_q(\omega^i)}(\omega^i)$ if $R^i_q(\omega^i) < \infty$, while $Z^i_q(\omega^i)$ is undefined if $R^i_q(\omega^i) = \infty$. Hence, Z^i_q is the transition used by Y^i at its q^{th} return time in A. By construction, if Z^i_q is defined, $Z^i_q \in A$. Since both chains are recurrent with only one recurrence class, $R^i_q < \infty$ with $\mathbb{P}^i$-probability 1 for all $i = 1, 2$ and $q > 0$. Therefore, for each $i = 1, 2$, Z^i_q are random variables defined $\mathbb{P}^i$-almost surely, $Z^i_q : \Omega^i \to A$, for all $q > 0$.

We denote by μ^i_q, for $i = 1, 2$ and $q > 0$, the law of Z^i_q in A, so that each μ^i_q is a probability law on A. Although the chains Y^i are ergodic, the probability laws μ^q_i can be trivial (i.e., concentrated on a unique transition). For example, in the graphs depicted in Fig. 1, left-bottom, decompose the shared arc $\mathbf{c}$ into two shared arcs $\mathbf{c}'$ and $\mathbf{c}''$, with an intermediate state between $\mathbf{c}'$ and $\mathbf{c}''$. Then the measures μ^i_q will be alternately concentrated on $\mathbf{c}'$ and $\mathbf{c}''$, depending on the parity of q. Despite this kind of particular example, the laws μ^i_q are "in general" not trivial. The following result has thus a wide range of applicability.

Proposition 3.1 (§ A.2). *If there exists one pair (i, q), $i = 1, 2$, $q > 0$, such that μ^i_q is non trivial, then $\mathbb{P}^1 \otimes \mathbb{P}^2(\Omega) = 0$.*

Comment—Hence, to construct a probability on Ω, conditioning $\mathbb{P}^1 \otimes \mathbb{P}^2$ on Ω is hopeless in general. If $\mathrm{Card}(A) = 1$, the conditioning works and is even trivial since $\mathbb{P}^1 \otimes \mathbb{P}^2(\Omega) = 1$ thanks to Th. 2.1. But the proposition says that this brute conditioning needs to be refined in general.

Synchronizable Labeled Markov Chains. We first study a property satisfied by a large class of pairs of labeled Markov chains. For each pair of states $(x^1, x^2) \in S^1 \times S^2$, let $\mathbb{P}_{x^1,x^2}$ be the probability measure on $\Omega^1 \times \Omega^2$ defined by $\mathbb{P}_{x^1,x^2} = \mathbb{P}^1_{x^1} \otimes \mathbb{P}^2_{x^2}$. Since the chains Y^1, Y^2 are ergodic, the first return times in A, R^1_1 and R^2_1, are $\mathbb{P}_{x^1,x^2}$-a.s. finite. The following mapping Φ is thus $\mathbb{P}_{x^1,x^2}$-a.s. defined:

$$\forall (\omega^1, \omega^2) \in \Omega^1 \times \Omega^2, \quad \Phi(\omega^1, \omega^2) = \big((Y^1_1, \ldots, Y^1_{R^1_1}), (Y^2_1, \ldots, Y^2_{R^2_1})\big) . \tag{3}$$

Recall that we have introduced the notation $Z^i_1 = Y^i_{R^i_1}$, for $i = 1, 2$. Since we are interested in synchronizing the chains on public transitions, we have to distinguish between two main cases:

1. $\mathbb{P}_{x^1,x^2}(Z^1_1 = Z^2_1) > 0$. The pair (x^1, x^2) is said to be *synchronizable*;
2. $\mathbb{P}_{x^1,x^2}(Z^1_1 = Z^2_1) = 0$. The pair (x^1, x^2) is said to be *non-synchronizable*.

Remark that its easily decidable whether a pair (x^1, x^2) is synchronizable, see Appendix A.1. Hence the following definition is fully effective:

Definition 3.1. *The two labelled Markov chains X^1 and X^2 are said to be* synchronizable *if all pairs $(x^1, x^2) \in S^1 \times S^2$ are synchronizable.*

The reader may convince himself by some examples that the system may be not synchronizable even if the product is max-synchronous (Def. 1.1).

Construction of a Probability. To construct a probability on Ω, we first construct an intermediate Markov chain. We assume that X^1 and X^2 are synchronizable, and that their product is max-synchronous. For each pair of states $(x^1, x^2) \in S^1 \times S^2$, we consider the conditional probability $\mathbb{P}_{x^1,x^2}\big(\,\cdot\,\big|Z^1_1 = Z^2_1\big)$ on $\Omega^1 \times \Omega^2$. This conditional probability is well defined since we assume the chains to be synchronizable. Let K denote the set of pairs of sequences described by $\Phi(\omega^1, \omega^2)$ when (ω^1, ω^2) ranges over $\{Z^1_1 = Z^2_1\}$. Regarding Φ as a random variable, we define μ_{x^1,x^2} as the law of Φ under $\mathbb{P}_{x^1,x^2}\big(\,\cdot\,\big|Z^1_1 = Z^2_1\big)$. Hence μ_{x^1,x^2} is a probability measure on K.

We consider the Markov chain with K as space state, and with transition matrix Q defined as follows. Denote by (v^1, v^2) a generic element of K; denote also by $Z(v^1)$ the last transition of v^1, which coincides by construction with the last transition $Z(v^2)$ of v^2. Then define the stochastic matrix Q, indexed by $K \times K$, by:

$$Q_{(v^1,v^2),(w^1,w^2)} = \mu_{x^1,x^2}(w^1, w^2), \quad \text{with: } \begin{cases} x^1 = \partial^1_+\big(Z(v^1)\big) \\ x^2 = \partial^2_+\big(Z(v^2)\big) \end{cases} \tag{4}$$

Remark that the state space K is in general infinite, but still countable. For $(x^1_0, x^2_0) \in S^1 \times S^2$ a pair of initial states, we define the initial measure ν_0 on K by $\nu_0 = \mu_{x^1_0,x^2_0}$. Then we define the infinite product space $\Lambda = K^{\mathbb{N}}$, equipped with the probability measure $\mathbb{P}'$ that defines a Markov chain on K with starting measure ν_0 and with transition matrix Q. We denote by ξ the generic elements of Λ. A generic element ξ can be written $\xi = (V^1_n, V^2_n)_{n\geq 0}$, with $(V^1_n, V^2_n) \in K$ for all $n \geq 0$. $(\Lambda, \mathbb{P}')$ enjoys the two following properties:

1. For $\mathbb{P}'$-a.s. every $\xi \in \Lambda$, V^1_0 and V^2_0 are two paths in G^1 and G^2 respectively, starting from x^1_0 and x^2_0 respectively.

2. For $\mathbb{P}'$-a.s. every $\xi \in \Lambda$, if we form the concatenations $S^1 = (V^1_0, V^1_1, \ldots)$ and $S^2 = (V^2_0, V^2_1, \ldots)$, then S^1 and S^2 are two transition sequences of G^1 and of G^2 respectively, that can synchronize.

These two properties show that there is a mapping $\Psi : \Lambda \to \Omega_{m_0}$, where m_0 is the initial marking of $\mathcal{N}$ obtained from (x^1_0, x^2_0). Regarding Ψ as a random variable, the law of Ψ under $\mathbb{P}'$ in Ω_{m_0} defines a probability measure on Ω_{m_0}, say $\mathbb{Q}$. To summarize the construction of $\mathbb{Q}$, we shall say that $\mathbb{Q}$ is obtained by *recursively* conditioning on the property that the next public transitions of X^1 and of X^2 coincide.

Definition 3.2. *We say that the probability* $\mathbb{Q}$ *on* Ω_{m_0} *obtained as above is the* synchronization product *of probabilities* $\mathbb{P}^i_{x^i_0}$, $i = 1, 2$.

Remark—If $\mathrm{Card}(A) = 1$, we re-obtain the product probability of Th. 2.1.

An indirect construction was made necessary, due to the negative result of Prop. 3.1. The synchronization product of probabilities that we have thus defined enjoys a fundamental property, that of being *homogeneous*. This implies that the product system, equipped with the synchronization probability, satisfies the Strong Markov property stated for true-concurrent systems [1]. The remaining of the paper is devoted to this topic.

Let $\mathbb{Q}$ be the synchronization product of probabilities $\mathbb{P}^1$ and $\mathbb{P}^2$ (underlying the initial states). Informally, the homogeneity of $\mathbb{Q}$ says that the product system is memory-less: the probabilistic evolution of the system after a finite history that ends to a given marking m, only depends on m, and not on the entire history. Formally, let v be a finite configuration of the unfolding $\mathcal{U}_{m_0}$ of the product net, and let m be the marking of the net reached by v. Let $\mathcal{S}(v)$ denote the set of $\omega \in \Omega_{m_0}$ such that $\omega \supseteq v$. Then $\mathcal{S}(v)$ identifies with Ω_m, the set of maximal configurations of the unfolding of the net (N, m). Hence the conditional probability $\mathbb{Q}\big(\,\cdot\,|\mathcal{S}(v)\big)$ defines a probability on Ω_m. If v' is an other finite configuration of $\mathcal{U}_{m_0}$ ending to the same marking m, $\mathbb{Q}\big(\,\cdot\,|\mathcal{S}(v')\big)$ defines, *a priori*, an other probability on the same space Ω_m. We say that the probability $\mathbb{Q}$ is *homogeneous* if, for any pair (v, v') of such finite configurations, we have $\mathbb{Q}\big(\,\cdot\,|\mathcal{S}(v)\big) = \mathbb{Q}\big(\,\cdot\,|\mathcal{S}(v')\big)$. In this case, the net $\mathcal{N}$ is said to be *Markovian*.

Theorem 3.1. *The synchronization probability of two Markov chains is homogeneous, and thus the product probabilistic net is Markovian.*

Proof. Let v be a configuration of $\mathcal{U}_{m_0}$, and let $v^1 = \pi^1(v)$ and $v^2 = \pi^2(v)$ be the decompositions of v in G^1 and in G^2. Let m be the marking reached by v. m consists of two states, say $x^i \in S^i$, $i = 1, 2$. Let $\mathbb{Q}_v$ be the conditional probability $\mathbb{Q}_v = \mathbb{Q}\big(\,\cdot\,|\mathcal{S}(v)\big)$. As a first case, if the last transition of v^1, and thus of v^2 too, is a public transition, then it is clear from the construction of $\mathbb{Q}$ that $\mathbb{Q}_v$ only depends on m.

Now consider the general case. A measure theoretic argument shows that $\mathbb{Q}_v$ is uniquely determined by the collection of numbers $\mathbb{Q}_v\big(\mathcal{S}(v, w)\big)$, for w ranging over the finite configurations of $\mathcal{U}_m$, where (v, w) denotes the concatenation of configurations v and w. We underlie the symbol "$\mathcal{S}$", and simply write $\mathbb{Q}_v(v, w)$ for $\mathbb{Q}_v\big(\mathcal{S}(v, w)\big)$. Thanks to the above first case, we may assume without loss of generality that both v and w contain no public transitions. Underlying the initial states of X^1 and X^2, we denote the associated probabilities by $\mathbb{P}^1$ and $\mathbb{P}^2$. Becuase of the construction of $\mathbb{Q}$ by condiotnal probability, there is a constant C such that, with w^1 and w^2 the decompositions of w:

$$\mathbb{Q}(v, w) = \frac{1}{C}\mathbb{P}^1 \otimes \mathbb{P}^2(v, w) = \frac{1}{C}\mathbb{P}^1(v^1)\mathbb{P}^1(w^1 \,|\, v^1)\mathbb{P}^2(v^2)\mathbb{P}^2(w^2 \,|\, v^2)$$
$$\mathbb{Q}(v) = \frac{1}{C}\mathbb{P}^1 \otimes \mathbb{P}^2(v) = \frac{1}{C}\mathbb{P}^1(v^1)\mathbb{P}^2(v^2)\,.$$

Taking the ratios, $\mathbb{Q}^v(w) = \mathbb{P}^1(w^1 \,|\, v^1)\mathbb{P}^2(w^2 \,|\, v^2)$. Since $\mathbb{P}^1$ and $\mathbb{P}^2$ are probabilities for homogeneous Markov chains, $\mathbb{P}^i(w^i \,|\, v^i)$ only depends on x^i and w^i, for $i = 1, 2$, and thus $\mathbb{Q}_v(w)$ only depend on $m = (x^1, x^2)$ and on w. Hence $\mathbb{Q}_v$ only depends on m, and this completes the proof. $\square$

A Appendix

A.1 Decidability of Conditions Introduced on Products

We have introduced two conditions on the systems: being max-synchronous (Def. 1.1) and being synchronizable (Def. 3.1). We show that both are decidable.

For $i = 1, 2$, x a state of S^i, and v a path in G^i starting from x, let $l^i(x, v)$ be the first transition of A encountered by v, if such a transition exists. We define $L^i(x)$ as the set of transitions $l^i(x, v)$, with v ranging over the set of paths starting from x. Clearly, $L^i(x)$ is computable in less than $N \times q$ steps, where N is the cardinal of S^1 and q is the number of arcs of the graph G^1. The knowledge of the sets $L^i(x)$ allows to determine the sought properties of G^1 and G^2:

1. A pair $(x^1, x^2) \in S^1 \times S^2$ is synchronizable if and only if $L^1(x) \cap L^2(x) \neq \emptyset$;
2. The product of G^1 and G^2 is not max-synchronous if and only if there is a pair of states (x^1, x^2), reachable by the synchronisation product, such that:
   - the set $R(x^2)$ of arcs a of G^2 with $\partial^2_-(a) = x^2$ is included in A, and $L^1(x^1) \cap R(x^2) = \emptyset$, or
   - symmetrically by exchanging the roles of indices 1 and 2.

Remark that, regarded as reachable marking of a safe Petri net, the set $(x, y) \in S^1 \times S^2$ of pairs of states reachable by the product is computable. Hence, both conditions that we introduced are computable.

A.2 Omitted Proofs

Sketch of proof of Lemma 1.1. It is enough to show it for finite configurations. Then this is seen by induction on the cardinal of configurations, using the fact that the labeling is injective on each component. □

Proof of Theorem 2.1. If $\mathrm{Card}(A) = 0$, we clearly have the equality $\Omega = \Omega^1 \times \Omega^2$, and there is nothing to prove. Hence we assume that $\mathrm{Card}(A) = 1$, and we let t be the unique public transition of the system. Put $\mathbb{P} = \mathbb{P}^1 \otimes \mathbb{P}^2$, and let $H \subseteq \Omega^1 \times \Omega^2$ be the set of pairs $(\omega^1, \omega^2) \in \Omega^1 \times \Omega^2$ such that, for $\omega^i = (Y^i_n)_{n>0}$, $i = 1, 2$, there are infinitely many n such that $Y^1_n = t$ and infinitely many p such that $Y^2_p = t$. Both chains Y^i, $i = 1, 2$, are recurrent with only one recurrent class, hence $\mathbb{P}(H) = 1$. But it is clear that every pair (ω^1, ω^2) in H is a pair of synchronizing sequences. To see it explicitly, denote for $i = 1, 2$ by R^i_q the q^{th} return time of Y^i to t, defined inductively by:

$$R^i_0 = 0, \quad \forall q > 0, \quad R^i_q = \inf\{n > R^i_{q-1} : Y^i_n = t\}. \tag{5}$$

Remark that $R^i_q < \infty$ on H for all $i = 1, 2$ and $q > 0$. Hence for $(\omega^1, \omega^2) \in H$, we construct a synchronization sequence s, firing sequence of $\mathcal{N}$ representative of a trace σ such that $\pi^1(\sigma) = \omega^1$, $\pi^2(\sigma) = \omega^2$, as follows:

$$\begin{aligned}\omega^1 &= \big(Y^1_1, \ldots\ldots, Y^1_{R^1_1-1}, \boxed{t}, Y^1_{R^1_1+1}, \ldots, Y^1_{R^1_2-1}, \boxed{t}, Y^1_{R^1_2+1}, \ldots\big)\\ \omega^2 &= \big(Y^2_1, \ldots, Y^2_{R^2_1-1}, \boxed{t}, Y^2_{R^2_1+1}, \ldots\ldots\ldots, Y^2_{R^2_2-1}, \boxed{t}, Y^2_{R^2_2+1}, \ldots\big)\end{aligned}$$

$$\begin{aligned}s = \big(Y^1_1, \ldots, Y^1_{R^1_1-1}, Y^2_1, \ldots, Y^2_{R^2_1-1}, \boxed{t},&\\ Y^1_{R^1_1+1}, \ldots, Y^1_{R^1_2-1}, Y^2_{R^2_1+1}, \ldots, Y^2_{R^2_1-1}, \boxed{t}, \ldots\big)&\end{aligned}$$

This implies that $H \subseteq \Omega$, and thus, since $\mathbb{P}(H) = 1$, $\mathbb{P}(\Omega) = 1$ as announced. □

Proof of Proposition 3.1. We set $\mathbb{P} = \mathbb{P}^1 \otimes \mathbb{P}^2$. $\mathbb{P}$ is the probability associated with the product Markov chain $V = (V_n)_{n>0}$, defined by $V_n = Y^1_n, Y^2_n$ for $n > 0$. Since both Y^1 and Y^2 are ergodic, so is V. Hence the return times to the initial state of V, defined by $S_0 = 0$, $S_{q+1} = \inf\{n > S_q : X^1_n = x^1_0,\ X^2_n = x^2_0\}$ for $q > 0$, are $\mathbb{P}$-a.s. finite for all $q > 0$.

For all $q > 0$, the random variables are Z^i_q are $\mathbb{P}$-a.s. finite. Therefore, with $\mathbb{P}$-probability 1, a pair $(\omega^1, \omega^2) \in \Omega^1 \times \Omega^2$ may synchronize if and only if $Z^1_q = Z^2_q$ for all $q > 0$. In symbols:

$$\Omega = \bigcap_{q>0} \{Z^1_q = Z^2_q\}, \quad \mathbb{P}\text{-a.s.} \tag{6}$$

Now let (i, r) be a pair such that μ^i_r is non trivial, say with $i = 1$. For $i = 1, 2$, denote by $\theta^i : \Omega^i \to \Omega^i$ the shift operator defined by $\theta^i(Y^i_1, Y^i_2, \ldots) = (Y^i_2, Y^i_3, \ldots)$, and by θ^i_q the q^{th} iterate of θ^i, $q \geq 0$. Since S_n and Z^i_r are stopping times, it is well known that $Q_n = S_n + Z^i_r \circ \theta^i_{S_n}$ are stopping times for all $n > 0$. From (6), we have:

$$\Omega \subseteq \bigcap_{n>0} \left\{Z^1_{Q_n} = Z^2_{Q_n}\right\}. \tag{7}$$

Since the Q_n are stopping times, standard Markov chains techniques show that the random variables $(Z^1_{Q_n})_{n>0}$ and $(Z^2_{Q_n})_{n>0}$ form two sequence of i.i.d. (independent identically distributed) random variables with laws μ^1_r and μ^2_r respectively. Hence we get from (7), by independence:

$$\mathbb{P}(\Omega) \leq \mathbb{P}\big(Z^1_r = Z^2_r\big) \times \mathbb{P}\Big(\bigcap_{n>1} \left\{Z^1_{Q_n} = Z^2_{Q_n}\right\}\Big). \tag{8}$$

To evaluate the first factor at right in (8), we put $\alpha = \max\{\mu^1_r(a),\ a \in A\}$, and $\alpha < 1$ by hypothesis. $\mathbb{P}\big(Z^1_r = Z^2_r\big) = \sum_{a\in A} \mu^2_r(a)\mu^1_r(a)$ can be seen as the expectation, under probability μ^2_r, of the function $a \in A \mapsto \mu^1_r(a)$. Since this function is $\leq \alpha$ on A, the expectation is $\leq \alpha$, and thus $\mathbb{P}\big(Z^1_r = Z^2_r\big) \leq \alpha$. By induction, we get from (8):

$$\begin{aligned}\forall j > 0, \quad \mathbb{P}(\Omega) &\leq \alpha^j \times \mathbb{P}\Big(\bigcap_{n>j} \left\{Z^1_{Q_n} = Z^2_{Q_n}\right\}\Big)\\ &\leq \alpha^j \longrightarrow_{j\to\infty} 0.\end{aligned}$$

Thus $\mathbb{P}(\Omega) = 0$, and this completes the proof. □

References

[1] S. Abbes. A (true) concurrent Markov property and some applications to Markov nets. In *Proc. of ATPN, 26th conference on Th. and App. of Petri nets*, 2005. To appear.

[2] S. Abbes and A. Benveniste. Branching cells as local states for event structrues and nets: probabilistic applications. In *FOSSACS 05*, volume 3441 of *LNCS*, pages 95–109, 2005. Extended version available as Research Report INRIA RR-5347.

[3] L. Breiman. *Probability*. SIAM, 1968.

[4] P.J. Haas. *Stochastic Petri nets*. Sp. V., 2002.

[5] N. Lynch, R. Segala, and F. Vandrager. Compositionality for probabilistic automata. In *CONCUR' 03*, volume 2761 of *LNCS*, pages 208–221, 2003.

[6] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, part 1. *T.C.S.*, 13:86–108, 1980.

[7] D. Revuz. *Markov chains*. North Holland, 1975.

[8] D. Varacca, H. Völzer, and G. Winskel. Probabilistic event structures and domains. In *CONCUR 04*, volume 3170 of *LNCS*, pages 481–496, 2004.

[9] G. Winskel. Petri nets, algebras, morphisms and compositionality. *Information and Computation*, 72(3):197–238, 1987.