

MARKOV TWO-COMPONENTS PROCESSES

SAMY ABBES

Université Paris Diderot/Laboratoire PPS, Bâtiment Sophie Germain, avenue de France
75 013 Paris, France
e-mail address: samy.abbes@univ-paris-diderot.fr

ABSTRACT. We propose Markov two-components processes (M2CP) as a probabilistic model of asynchronous systems based on the trace semantics for concurrency. Considering an asynchronous system distributed over two sites, we introduce concepts and tools to manipulate random trajectories in an asynchronous framework: stopping times, an Asynchronous Strong Markov property, recurrent and transient states and irreducible components of asynchronous probabilistic processes. The asynchrony assumption implies that there is no global totally ordered clock ruling the system. Instead, time appears as *partially ordered* and *random*.

We construct and characterize M2CP through a finite family of transition matrices. M2CP have a local independence property that guarantees that local components are independent in the probabilistic sense, conditionally to their synchronization constraints. A synchronization product of two Markov chains is introduced, as a natural example of M2CP.

Dedicated to the memory of Philippe Darondeau (1948–2013)

INTRODUCTION

General settings and requirements. In this paper we introduce a probabilistic framework for a simple asynchronous system distributed over two sites, based on the trace semantics of concurrency. Consider a communicating system consisting of two subsystems, called site 1 and site 2, that need to synchronize with one another from time to time, for example for message exchange. Intended applications are, for instance, simple client-server situations, device-device driver interactions, communication bridge between two asynchronous networks. The synchronization is modeled, for each site, by the fact that the concerned subsystem is entering some *synchronizing state*, corresponding to a synchronization task—there shall be several synchronization states corresponding to different tasks. It is natural to consider that the synchronization states are shared: both subsystems are supposed to

2012 ACM CCS: [Mathematics of computing]: Probability and statistics—Stochastic processes—Markov processes; [Theory of computation]: Models of computation—Concurrency—Parallel computing models.

Key words and phrases: Probabilistic processes, concurrency, Markov chains.

enter together into a shared synchronization state. Beside synchronization states, we assume that each subsystem may evolve between other states that concern the local activity of each subsystem, and seen as *private states*. Hence we consider for each site $i = 1, 2$ some finite set of states S^i , with the intended feature that $Q = S^1 \cap S^2$ is a nonempty set of synchronization states.

Whenever the two subsystems enter some of their private states, the corresponding events are said to be *concurrent*. It is natural to consider that the private activity of a given site should not influence the private activity of the other site. This ought to be reflected by some kind of statistical independence in the probabilistic modeling. Another feature that we are seeking is that the *local* time scales of private activities do not need to be synchronous. Indeed, the local time scale of each subsystem might be driven for instance by the input of a user, by the arrival of network events, or by its internal chipset clock; therefore, it is realistic *not* to assume any correlation between local time scales, but for synchronization. In particular, in a discrete time setting, the synchronization instants counted on the two different local time scales shall not need to be equal, making the two subsystems *asynchronous*.

Sequential probabilistic systems and concurrency. Classically, Markov chains in either discrete or continuous time are a popular model for adding a probabilistic layer to describe the evolution of a transition system. Since the Markov chain formalism is intrinsically sequential, its straightforward application to a concurrent system brings the issue of translating a concurrent system into a sequential one. A solution to this issue, found in the Probabilistic Automata literature for instance [18, 12], is the introduction of a non deterministic scheduler in charge of deciding which subsystem is about to run at each time instant. This defines a Markov Decision Process, a model introduced earlier for control issues in [5]. Other ways of composing probabilistic systems to form a Markov process, with or without non determinism, are usually based on Milner’s CCS [13] or Hoare’s CSP [11], where the synchronization policy for possibly synchronizing processes is either to allow or to force synchronization. In [8] for instance, where both synchronization methods à la CSS and à la CSP are encoded in the model of bundle probabilistic transition systems, renormalization occurs at each step to take into account the selected synchronization paradigm.

Probabilistic trace semantics. Lattice of trajectories. We introduce another way of randomizing our simple concurrent system. We first accept as a basic fact that modeling the evolution of a system as ordered paths of events jeopardizes the concurrency feature of the model. Adopting instead the so-called trace semantics for concurrency (or partial order semantics) [15, 16], lattices replace ordered paths to model trajectories. Unordered events of a trajectory are then intrinsically concurrent. This raises a question on the probabilistic side: which part of Markov chain theory can we rebuild on this new basis?

The aim of this paper is to provide an answer to the question. Our work is thus largely inspired by Markov chain theory; but we try to adapt the theory to the partial order semantics of concurrency, instead of directly turning a concurrent system into a Markov chain (or a variant of it) as in [12, 8].

Let us be precise about what we mean in this paper by a partial order semantics for concurrency, referring to the two sets of local states S^1 and S^2 with synchronization constraint $Q = S^1 \cap S^2$. We will then explain how probability concepts apply in this setting.




Figure 1: *Lattice of subtrajectories of $v = (a \cdot c, e \cdot f \cdot c)$.*

If two sequences of states in $S^1 \cup S^2$ only differ by the interleaving order of private states of different sites, such as $a \cdot e$ and $e \cdot a$ with $a \in S^1 \setminus Q$ and $e \in S^2 \setminus Q$, the trace semantics suggests to simply identify them: $a \cdot e \equiv e \cdot a$. Propagating this identification to sequences of events of arbitrary length, we obtain an equivalence relation on sequences. Sequences that cannot be permuted are those of the form $x \cdot y$ with $x, y \in S^1$ or $x, y \in S^2$, which include those of the form $x \cdot c$ with $c \in Q$ and any $x \in S^1 \cup S^2$. We adopt a simple representation for equivalence classes of sequences by mapping each equivalence class to a *pair* of sequences, where each coordinate is reserved for a given site; synchronization states appear in both coordinates. Hence the equivalence class of $a \cdot e \equiv e \cdot a$ is mapped to (a, e) , the equivalence class of $a \cdot e \cdot c \equiv e \cdot a \cdot c$ is mapped to $(a \cdot c, e \cdot c)$. We define thus a *trajectory* as a pair (s^1, s^2) , where s^i is a sequence of elements in S^i , and such that the sequences of synchronization states extracted from s^1 and from s^2 , and taken in their order of appearance, shall be *equal*.

An *infinite trajectory* is defined as a trajectory $\omega = (\omega^1, \omega^2)$ where both sequences ω^1 and ω^2 are infinite. So for example, if $S^1 = \{a, b, c, d\}$ and $S^2 = \{c, d, e, f\}$, and thus $Q = \{c, d\}$, an infinite trajectory could be $\omega = (\omega^1, \omega^2)$ with ω^1 and ω^2 starting as follows: $\omega^1 = a \cdot c \cdot b \cdot a \cdot b \cdot b \cdot d \cdot (\dots)$ and $\omega^2 = e \cdot f \cdot c \cdot f \cdot d \cdot (\dots)$. The common extracted sequence of synchronization states starts in this example with $c \cdot d$. Note the important feature that each local trajectory ω^i is permitted to have a free evolution between synchronizations: synchronizations occur at instants 2 and 7 for ω^1 , while they occur at instants 3 and 5 for ω^2 ; here, the instants of synchronization are relative to the *local* time scales. The set Ω of infinite trajectories is the natural sample space to put a probability measure on.

There is a natural notion of *subtrajectory*: in the previous example, $v = (a \cdot c, e \cdot f \cdot c)$ is a finite subtrajectory of $\omega = (\omega^1, \omega^2)$. “Being a subtrajectory” defines a binary relation that equips subtrajectories of a given trajectory with a *lattice* structure. For instance, and denoting by ϵ the empty word, the subtrajectories of v are: (ϵ, ϵ) , (a, ϵ) , (ϵ, e) , $(\epsilon, e \cdot f)$, (a, e) , $(a, e \cdot f)$ and $(a \cdot c, e \cdot f \cdot c)$. Their lattice is depicted in Figure 1. Observe that, for a given trajectory, its subtrajectories are naturally identified with two-components “time instants”. In case of v , these time instants are $(0, 0)$, $(1, 0)$, $(0, 1)$, $(0, 2)$, $(1, 1)$, $(1, 2)$ and $(2, 3)$, and they form a sublattice of the lattice $\mathbb{N} \times \mathbb{N}$. However, even if one considers an infinite trajectory ω , the associated lattice of two-components time instants is only a *sublattice* of $\mathbb{N} \times \mathbb{N}$ in general. For instance, if ζ is any infinite trajectory that has v as subtrajectory, then $(2, 2)$ is a time instant that does not correspond to any subtrajectory of ζ , because of the synchronization on state c .

Obviously, considering another trajectory ω' would lead to another lattice of subtrajectories, not necessarily isomorphic to the one associated with ω . We sum up the previous

observations by saying that time is *partially ordered* on the one hand, since time instants form a lattice and not a total order, and *random* on the other hand, since the lattice structure depends on the trajectory considered, that is, on the execution of the system.

Defining M2CP: absence of transition matrix. This has consequences for the way one may construct a probability measure on the space Ω of infinite trajectories. Consider again the finite trajectory encountered above, $v = (a \cdot c, e \cdot f \cdot c)$. The occurrences of a on site 1, and of e on site 2, are *concurrent*. Trying to determine the precise interleaving of a and e is irrelevant for us. This desired feature prevents us from applying the standard recursive construction to assign a probability to trajectory v (that is: the probability that v occurs as a subtrajectory of a sample infinite trajectory ω): starting from the initial state, there is no obvious choice between a and e ; which one should be first plugged in the probability computation?

Therefore the lattice structure of trajectories implies to give up, at least temporarily, the familiar inductive computation of probabilities based on transition matrices. Nevertheless, two important notions can be defined in the asynchronous framework by analogy with Markov chain theory: first, the notion of state reached by (“after”) a finite trajectory (§ 1.1); second, the probabilistic evolution of the system “after” execution of a finite trajectory (Definition 1.6 in § 1.3). We define a *Markov two-components process* (M2CP) as a random system where the probabilistic future after execution of a finite trajectory v only depends on the state reached by v .

Stopping times for M2CP. Recall that a stopping time in Markov chain theory identifies with a random halting procedure that does not need anticipation: an observer can decide whether the stopping time has been reached based on the only knowledge of the process history at each step. Stopping times are a basic tool in Markov chain theory. Important notions such as the first return time to a state, recurrent and transient states are defined by means of stopping times. Stopping times are manipulated with the help of the Strong Markov property, a central result in Markov chain theory. We show that several aspects of the Markovian language carry over to the asynchronous framework. Once an adequate notion of stopping time for asynchronous probabilistic processes has been defined (Definition 3.1 in § 3.1), derived notions such as the first hitting time to a state, and the notions of recurrent and transient states follow by almost literally translating the original ones into the asynchronous language. We show that the Strong Markov property also has an equivalent, called the Asynchronous Strong Markov property, which serves as a basic tool for probabilistic reasoning. Some other notions translate in a more subtle way: the first reaching time of a *set* of states needs some additional care, since the lattice structure of trajectories prevents a straightforward generalization of the analogous notion from Markov chain theory, providing an interesting difference with Markov chain theory. Irreducible processes have an equivalent counterpart in the asynchronous framework, and we detail the decomposition of a M2CP into irreducible components.

The Local Independence Property. Therefore, we have on the one hand these notions obtained as a generalization of analogous notions from Markov chain theory to the asynchronous framework. But on the other hand, we also have other notions specific to the asynchronous framework, and that would not make sense for Markov chains. In particular, the way the two local components behave with respect to one another is a question specific to the asynchronous framework. Since the two local components synchronize with one another, they cannot be fully independent in the probabilistic sense. There is however a weaker notion of independence in probability theory, adapted to our purpose, which is *conditional* independence. We call Local Independence Property (LIP) the property that the two components are independent conditionally to their synchronization constraint. Informally, the LIP says that the local components have the maximal independence that they can have, up to their synchronization constraint. We characterize M2CP with the LIP by a finite family of transition matrices; and we show how to construct a M2CP from an adapted family of such transition matrices. The finite collection of numbers this family of matrices defines is an equivalent, in the asynchronous framework, of the transition matrix for a Markov chain.

Synchronization of systems. The composition of probabilistic systems has always been a challenge, with multiple applications in the theory of network analysis [8, 12, 4]. The main limitation of the theory of probabilistic event structures as it has been developed so far by the author together with A. Benveniste in [3, 2] (another probabilistic model with trace semantics targeting applications to probabilistic 1-safe Petri nets), and by other authors in [19] is the non ability to define a suitable synchronization product. This very limitation has motivated the development of the present framework, by starting with the definition of a synchronization product for two Markov chains. By recursively “forcing” their synchronization, it is shown in this paper how the synchronization of two Markov chains on shared common states naturally leads to a M2CP. Even if one was interested in this construction only (the author is aware of current work on this kind of *a priori* model, simply because it was the only one people could think of), including it inside a more general picture as it is done in this paper is useful to better understand its properties.

Organization of the paper. We describe the model in § 1, defining a general notion of probabilistic two-components process, and then specializing to *Markov* two-components processes. In § 2 we introduce the synchronization product of two Markov chains. This construction provides an example of M2CP, intended to support the intuition for M2CP in general. Next section, § 3, is devoted to Markovian concepts in the asynchronous framework, centered around the Asynchronous Strong Markov property. We introduce recurrence and transience of states and the decomposition of M2CP into irreducible components. The new notions of closed and open processes are studied in this section, as well as the definition of stopping times for asynchronous processes. The Local Independence Property (LIP) is the topic of § 4, and it is shown that the synchronization of Markov chains introduced in § 2 satisfies the LIP. Finally, § 5 is devoted to the construction and characterization of general M2CP with the LIP. A concluding section introduces directions for future work. It discusses limitations imposed by the two-components hypothesis, and possible ways to remove this limiting hypothesis.

1. PROBABILISTIC PROCESSES AND MARKOV PROCESSES ON TWO SITES

1.1. General Framework. A **distributed system** is given by a pair (S^1, S^2) , where S^i for $i = 1, 2$ is a finite set, called the set of **local states** of site i . A **local trajectory** attached to site i is a sequence of local states of this site. For $i = 1, 2$, we denote by Ω^i the set of infinite local trajectories attached to site i .

The two local state sets S^1 and S^2 are intended to have a non empty intersection, otherwise the theory has little interest. We put $Q = S^1 \cap S^2$. Elements of Q are called **common states** or **shared states**. In contrast, states in $S^i \setminus Q$ are said to be **private to site i** , for $i = 1, 2$. From now on, **we will always assume that $S^i \setminus Q \neq \emptyset$** for $i = 1, 2$: each site has at least one private state. This is a convenient technical assumption; removing it would not harm if needed.

Given a sequence $(x_j)_j$ of elements in a set S , either finite or infinite, and given a subset $A \subseteq S$, the **A-sequence induced by $(x_j)_j$** is defined as the sequence of elements of A encountered by the sequence $(x_j)_j$, in their order of appearance. Given two local trajectories $(x_n^1)_{n \geq 0}$ and $(x_n^2)_{n \geq 0}$ on sites 1 and 2 respectively, we will say that they **synchronize** if the two Q -sequences they induce are *equal*. A pair of two synchronizing local trajectories will be called a **global trajectory**, or simply a **trajectory** for brevity. Among them, **finite trajectories** are those whose components are both finite sequences of states.

Trajectories are ordered component by component: if $s = (s^1, s^2)$ and $t = (t^1, t^2)$ are two trajectories, we define $s \leq t$ if $s^1 \leq t^1$ and $s^2 \leq t^2$, where the order on sequences is the usual prefix order. The resulting binary relation on trajectories is a partial order, the maximal elements of which are exactly those whose components are both infinite: this relies on the fact that $S^i \setminus Q \neq \emptyset$ for $i = 1, 2$ (for instance, if $S^1 = \{a, \mathbf{b}\}$ and $S^2 = \{\mathbf{b}\}$ so that $Q = \{\mathbf{b}\}$ and $S^2 \setminus Q = \emptyset$, then $(\mathbf{baaa} \cdots, \mathbf{b})$ is maximal, but the second component is finite). The set of maximal trajectories is denoted by Ω , and we have that $\Omega \subseteq \Omega^1 \times \Omega^2$. For s a *finite* trajectory, the subset of Ω defined by

$$\uparrow s = \{\omega \in \Omega \mid s \leq \omega\} \tag{1.1}$$

is called the **elementary cylinder** of base s —adapting a standard notion from Measure theory to our framework.

Given any trajectory $s = (s^1, s^2)$, the **subtrajectories of s** are those trajectories t such that $t \leq s$. Observe that not any prefix t of s is a subtrajectory; since t could very well not be a trajectory itself.

Given a trajectory (s^1, s^2) , we denote by $(y_j)_j$ the Q -sequence induced by both sequences s^1 and s^2 . It can be finite or infinite, even empty. We refer to $(y_j)_j$ as to the **Q-sequence induced by (s^1, s^2)** .

A **global state** is any pair $\alpha = (x^1, x^2) \in S^1 \times S^2$. We reserve the letters α and β to denote global states. Observe that trajectories are *not* defined as sequences of global states; since the length of the two components may very well differ. Let $\alpha = (x, y)$ be some fixed global state, thought of as the *initial* state of the system. If $s = (s^1, s^2)$ is a finite trajectory, we define

$$\gamma_\alpha(s) = (x^1, x^2) \in S^1 \times S^2 \tag{1.2}$$

as the pair of last states of the two sequences $x \cdot s^1$ and $y \cdot s^2$. We understand $\gamma_\alpha(s)$ as the current global state after the execution of finite trajectory s , starting from α . Note that, with this definition, γ_α is well defined on the empty sequence and $\gamma_\alpha(\emptyset) = \alpha$. By an abuse

of notation, we will omit α and write γ instead of γ_α , the context making clear which initial state α we refer to.

We introduce a notion of length for trajectories. We denote by \mathcal{T} the set

$$\mathcal{T} = (\mathbb{N} \times \mathbb{N}) \cup \{\infty\}.$$

The set \mathcal{T} is partially ordered component by component, with the natural order on each component, and $(m, n) \leq \infty$ for all $(m, n) \in \mathbb{N} \times \mathbb{N}$. If $s = (s^1, s^2)$ is any trajectory, the **length of s** is defined by

$$|s| = \begin{cases} (|s^1|, |s^2|) \in \mathcal{T}, & \text{if } s \text{ is finite,} \\ \infty, & \text{otherwise,} \end{cases}$$

where $|s^1|$ and $|s^2|$ denote the length of sequences. Roughly speaking, lengths can be thought of as time instants; it becomes then clear that time is only partially ordered, and not totally ordered—see random times in § 3.1 for a finer notion.

There is a **concatenation** operation partially defined on trajectories. If $s = (s^1, s^2)$ is a *finite trajectory*, and $t = (t^1, t^2)$ is any trajectory, then the concatenation denoted by $s \cdot t$ and defined by $s \cdot t = (s^1 \cdot t^1, s^2 \cdot t^2)$ is obviously a trajectory. If $t \in \Omega$, then $s \cdot t \in \Omega$ as well. There is an obvious **addition** on lengths, compatible with concatenation of finite trajectories, in the sense that $|s \cdot t| = |s| + |t|$. If we fix s , the concatenation defines a bijection onto the cylinder of base s :

$$\Phi_s : \begin{cases} \Omega \rightarrow \uparrow s \\ \omega \mapsto \Phi_s(\omega) = s \cdot \omega. \end{cases} \quad (1.3)$$

1.2. Trajectory Structure. The fact that we consider only two sites allows to precisely describe the structure of trajectories.

Definition 1.1.

- (1) An **elementary trajectory** is a finite trajectory with a unique synchronization, that occurs at its end. Equivalently, a finite trajectory s is elementary if $\gamma(s) = (x, x)$ for some $x \in Q$, and (x) is the Q -sequence induced by s .
- (2) We say that a trajectory is **synchronization free** if its associated Q -sequence is empty.

We omit the proof of the following proposition, which is elementary, but fundamental for some constructions introduced later in § 2 and in § 5.

Proposition 1.2.

- (1) *Any finite trajectory has a unique decomposition as a concatenation of elementary trajectories, followed by a synchronization free trajectory.*
- (2) *Any maximal trajectory is either, according to its Q -sequence being infinite or finite:*
 - (a) *A countable infinite concatenation of elementary trajectories, and the decomposition as such a concatenation is unique; or*
 - (b) *A finite concatenation of elementary trajectories, followed by a synchronization free trajectory, infinite on both sides. This decomposition is unique.*

Figure 2 depicts the decomposition of global trajectories in cases 2a and 2b. Finally, the following lemma will be useful.




Figure 2: Illustration of the decomposition of a maximal trajectory, according to Proposition 1.2, Cases 2a and 2b. The framed boxes represent the synchronizations, the arrows represent the private paths. In Case 2a, the synchronization pattern keeps repeating on the right.

Lemma 1.3. *For any trajectory v , the set of subtrajectories of v is a well founded and complete lattice. Lower and upper bounds are taken component by component.*

Proof. Let $v = (s^1, s^2)$, and let $\mathcal{I}(s^i)$ denote, for $i = 1, 2$, the set of initial subsequences of s^i . It is well known that $\mathcal{I}(s^i)$ is a total and well-founded order with arbitrary *lubs* (least upper bounds). Therefore the component-wise order on $\mathcal{I}(s^1) \times \mathcal{I}(s^2)$ is a complete lattice, with lower and upper bounds taken component by component.

To prove the lemma, it suffices thus to check that the component-wise upper and lower bounds of subtrajectories of v yield again subtrajectories of v , and this is obvious, hence we are done. \square

1.3. Probabilistic Two-Components Processes. Although time has been abstracted from the framework, the notion of trajectory is still present; this is all we need to introduce a probabilistic layer. We consider the σ -algebra \mathfrak{F} on Ω generated by the countable family of elementary cylinders, defined above in Eq. (1.1). The σ -algebra \mathfrak{F} coincides with the trace on Ω of the product σ -algebra on the infinite product $\Omega^1 \times \Omega^2 = (S^1 \times S^2)^\mathbb{N}$, where of course S^i , as a finite set for $i = 1, 2$, is equipped with the discrete σ -algebra.

Unless stated otherwise, the set Ω will be equipped with the σ -algebra \mathfrak{F} . Assume thus that \mathbf{P} is a probability defined on Ω . By an abuse of notation, if s is a finite trajectory we simply denote by $\mathbf{P}(s)$ the probability of the elementary cylinder of base s , so that $\mathbf{P}(s) = \mathbf{P}(\uparrow s)$. We say that a global state α is **reachable** w.r.t. \mathbf{P} if there exists a finite trajectory s such that $\mathbf{P}(s) > 0$ and $\alpha = \gamma(s)$. A probabilistic two-components process on a distributed system is defined as follows.

Definition 1.4.

- (1) A **probabilistic two-components process**, or **probabilistic process** for brevity, is a family $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ of probability measures on Ω indexed by a set X_0 of global states, and satisfying the following property: for all $\alpha \in X_0$, if β is reachable with respect to \mathbf{P}_α , then $\beta \in X_0$.

- (2) If β is reachable w.r.t. \mathbf{P}_α , we say that β is **reachable from α** .
(3) A **subprocess** of a probabilistic process \mathbb{P} is a subfamily $(\mathbf{P}_\alpha)_{\alpha \in X_1}$, with $X_1 \subset X_0$, that forms a probabilistic process.

The probability \mathbf{P}_α is intended to describe the probabilistic behavior of the system starting from α . However, for technical reasons that will appear later, we consider the evolution of the system *after* α . In other words, we assume that α has already been reached, and we put ourselves just after it. In particular, we do *not* assume that $\mathbf{P}_\alpha(\uparrow \alpha) = 1$, contrary to the usual convention adopted in Markov chain theory.

Definition 1.5. Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a probabilistic two-components process. Let also $*$ be an arbitrary specified value not in $S^1 \cup S^2$.

- (1) For $\omega \in \Omega$, we denote by $Y(\omega) = (Y_n(\omega))_{n \geq 0}$ the Q -sequence induced by ω , followed by the constant value $*$ if the Q -sequence is finite. In all cases, we also put $Y_{-1} = *$. We refer to Y as to the **(random) synchronization sequence**.
(2) We say that $\omega \in \Omega$ **synchronizes infinitely often** if $Y_n(\omega) \neq *$ for all $n \geq 0$.
(3) We say that \mathbb{P} is **closed** if for all $\alpha \in X_0$, $Y_n \neq *$ for all $n \geq 0$ and \mathbf{P}_α -a.s.
(4) We say that \mathbb{P} is **open** if for all $\alpha \in X_0$, $Y_n = *$ for all $n \geq 0$ and \mathbf{P}_α -a.s.

Consider any probability measure \mathbf{P} on Ω , and let s be a finite trajectory. Observe that $\Phi_s : \Omega \rightarrow \uparrow s$ is not only a bijection, it is also bi-measurable. Considering the action of Φ_s^{-1} on measures is thus meaningful. In particular, if $\mathbf{P}(s) > 0$, we define the probability \mathbf{P}_s on Ω as the image of the conditional probability $\mathbf{P}(\cdot | \uparrow s)$. It satisfies, and is characterized by the relations $\mathbf{P}_s(t) = \frac{1}{\mathbf{P}(s)} \mathbf{P}(s \cdot t)$, for t ranging over the set of finite trajectories.

Definition 1.6. If \mathbf{P} is a probability measure on Ω , and if s is a finite trajectory such that $\mathbf{P}(s) > 0$, we define the probability measure \mathbf{P}_s on Ω characterized by:

$$\mathbf{P}_s(t) = \frac{1}{\mathbf{P}(s)} \mathbf{P}(s \cdot t), \quad (1.4)$$

for t ranging over the set of finite trajectories, as the **probabilistic future** of s w.r.t. probability \mathbf{P} .

Markov two-components processes can now be defined as follows, without reference to any explicit notion of time.

Definition 1.7. Given a distributed system, a **Markov two-components process**, abbreviated M2CP, is defined as a probabilistic process $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ over this system, satisfying the following property: for α ranging over X_0 and s ranging over the set of finite trajectories such that $\mathbf{P}_\alpha(s) > 0$, the probabilistic future of trajectory s w.r.t. \mathbf{P}_α only depends on $\gamma(s)$. This is equivalent to saying:

$$\forall \alpha \in X_0 \quad \forall s \quad \mathbf{P}_\alpha(s) > 0 \Rightarrow (\mathbf{P}_\alpha)_s = \mathbf{P}_{\gamma(s)}. \quad (1.5)$$

Equation (1.5) formalizes the intuition that “the probabilistic future only depends on the present state”; we shall refer to it as to the *Markov property*. Some additional comments about Definition 1.7:

1. Markov chains are usually defined by their transition matrix, from which a probability measure on the space of trajectories is derived. Here, on the contrary, the lack of a totally ordered time index leads us to first consider a measure on the space of trajectories with the Markov property already encoded in it. It will be our task to find an equivalent for

the transition matrix, that would characterize the probability measure through a finite number of real parameters with adequate normalization conditions. This is the topic of § 5.

2. Considering the same definition for a probability measure on a space of trajectories with only one component—for instance, taking $S^2 = \{*\}$ a singleton disjoint from S^1 —, would exactly bring us back to the definition of a homogeneous Markov chain on S^1 . The transition matrix $P_{i,j}$ would then be given by $P_{i,j} = \mathbf{P}_{(i,*)}(\uparrow(j, *))$.
3. Contrast this definition with an alternative, naive model consisting of a Markov chain on the state of global states. Note that the Markov property stated in Eq. (1.5) is relative to any “cut” $\gamma(s)$ of the trajectory. However, for a Markov chain, the property would only hold for particular cuts, namely those such that $|s|$ has the form (n, n) for some integer n .

Checking that a probabilistic process \mathbb{P} satisfies the Markov property amounts to verifying the equality:

$$\frac{1}{\mathbf{P}_\alpha(s)} \mathbf{P}_\alpha(s \cdot t) = \mathbf{P}_{\gamma(s)}(t) \quad (1.6)$$

for all finite trajectories s and t such that $\mathbf{P}_\alpha(s) > 0$. The following lemma however shows that, for closed processes, it suffices to verify Eq. (1.6) for elementary trajectories t .

Lemma 1.8. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a closed two-components process, such that:*

$$\forall \alpha \in X_0 \quad (\mathbf{P}_\alpha)_s(t) = \mathbf{P}_{\gamma(s)}(t), \quad (1.7)$$

for every elementary trajectory t and finite trajectory s with $\mathbf{P}_\alpha(s) > 0$. Then \mathbb{P} is a Markov two-components process.

Proof. Let \mathcal{E} denote the set of elementary trajectories (Definition 1.1). We also denote by \mathcal{E}^+ the set of trajectories that are finite concatenations of elementary trajectories, and by \mathcal{V} the set of finite trajectories. We proceed in two steps to show that Eq. (1.7) is valid for $s, t \in \mathcal{V}$.

Step 1: Equation (1.7) is true for $s \in \mathcal{V}$ and $t \in \mathcal{E}^+$. By induction, we show that Eq. (1.7) is true for $s \in \mathcal{V}$ and $t = t_1 \cdot \dots \cdot t_n$ with $t_i \in \mathcal{E}$. The case $n = 1$ is given by the hypothesis of the lemma, assume it is true for all $k < n$. Assume moreover that $\mathbf{P}_\alpha(s \cdot t_1 \cdot \dots \cdot t_k) > 0$ for all $k = 1, \dots, n - 1$. We calculate as follows, using the hypothesis of the lemma and the induction hypothesis:

$$\begin{aligned} (\mathbf{P}_\alpha)_s(t_1 \cdot \dots \cdot t_n) &= \frac{\mathbf{P}_\alpha(s \cdot t_1 \cdot \dots \cdot t_n)}{\mathbf{P}_\alpha(s)} \\ &= (\mathbf{P}_\alpha)_s(t_1 \cdot \dots \cdot t_{n-1}) \cdot (\mathbf{P}_\alpha)_{s \cdot t_1 \cdot \dots \cdot t_{n-1}}(t_n) \\ &= \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_{n-1}) \cdot \mathbf{P}_{\gamma(t_{n-1})}(t_n). \end{aligned} \quad (1.8)$$

We also have, using again the hypothesis of the lemma:

$$\begin{aligned} \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_n) &= \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_{n-1}) \cdot \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_n | t_1 \cdot \dots \cdot t_{n-1}) \\ &= \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_{n-1}) \cdot (\mathbf{P}_{\gamma(s)})_{t_1 \cdot \dots \cdot t_{n-1}}(t_n) \\ &= \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_{n-1}) \cdot \mathbf{P}_{\gamma(t_{n-1})}(t_n). \end{aligned} \quad (1.9)$$

Comparing (1.8) and (1.9), we get $(\mathbf{P}_\alpha)_s(t_1 \cdot \dots \cdot t_n) = \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_n)$, completing the induction in this case.

To be complete, we examine the case where $\mathbf{P}_\alpha(s \cdot t_1 \cdot \dots \cdot t_k) = 0$ for some integer $k \in \{1, \dots, n-1\}$. Then, on the one hand, this implies $\mathbf{P}_\alpha(s \cdot t_1 \cdot \dots \cdot t_n) = 0$ and thus $(\mathbf{P}_\alpha)_s(t_1 \cdot \dots \cdot t_n) = 0$. On the other hand, let i be the smallest integer $1 \leq i < n$ such that $\mathbf{P}_\alpha(s \cdot t_1 \cdot \dots \cdot t_i) = 0$. Then the minimality of i yields $(\mathbf{P}_\alpha)_{s \cdot t_1 \cdot \dots \cdot t_{i-1}}(t_i) = 0$, and by the hypothesis of the lemma this is $\mathbf{P}_{\gamma(t_{i-1})}(t_i) = 0$. Applying again the hypothesis of the lemma: $\mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_i) = \mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_{i-1}) \cdot \mathbf{P}_{\gamma(t_{i-1})}(t_i) = 0$, which implies $\mathbf{P}_{\gamma(s)}(t_1 \cdot \dots \cdot t_n) = 0$. The induction is complete.

Step 2: Equation (1.7) is true for $s, t \in \mathcal{V}$. Let s and t be any finite trajectories. For $\omega \in \uparrow(s \cdot t)$, we put

$$E_\omega = \{v \in \mathcal{E}^+ \mid s \cdot t \leq v \leq \omega\}, \quad \omega_T = \inf E_\omega.$$

On the one hand, $E_\omega \neq \emptyset$ \mathbf{P}_α -a.s. since \mathbb{P} is assumed to be closed. On the other hand, the trajectories of E_ω form a chain, which is well founded by Lemma 1.3; hence $\omega_T = \min E_\omega$ is \mathbf{P}_α -a.s. well defined and $\omega_T \in E_\omega$. It is easy to observe that, for $v = \omega_T$, we have:

$$\{\omega' \in \Omega \mid \omega'_T = v\} = \uparrow v. \quad (1.10)$$

(Later, we will interpret this by saying that $\omega \mapsto \omega_T$ is a stopping time). Since ω_T ranges over finite trajectories, the set of values it can take is countable. Therefore, decomposing with respect to the possible values:

$$\begin{aligned} (\mathbf{P}_\alpha)_s(t) &= \sum_v (\mathbf{P}_\alpha)_s(\uparrow t \cap \{\omega_T = v\}) \\ &= \sum_v (\mathbf{P}_\alpha)_s(\omega_T = v) && \text{since } \uparrow t \subset \{\omega_T = v\} \\ &= \sum_v (\mathbf{P}_\alpha)_s(\uparrow v) && \text{by Eq. (1.10)} \\ &= \sum_v \mathbf{P}_{\gamma(s)}(\uparrow v) && \text{by Step 1 since } v \in \mathcal{E}^+ \\ &= \mathbf{P}_{\gamma(s)}(t) && \text{recomposing.} \end{aligned}$$

The proof is complete. \square

2. SYNCHRONIZATION OF TWO MARKOV CHAINS

In this section we introduce a way of constructing M2CPs. It first shows that our object of study is not empty. It also provides a bridge between M2CPs and usual Markov chains—another, maybe deeper link is developed in § 3.

Consider two Markov chains $(X_n^1)_{n \geq 0}$ and $(X_n^2)_{n \geq 0}$ on S^1 and S^2 respectively. We denote by M_x^i the probability measure on Ω^i associated with the chain X^i starting from state $x \in S^i$, for $i = 1, 2$. We assume for simplicity that both transition matrices have all their coefficients positive. The construction consists of recursively forcing the next synchronization of the chains on a shared state. The formal construction is given in Definition 2.1 below, after an informal explanation. The case where there is only one synchronization state is trivial, in the sense that it reduces to the independent product of the two Markov chains as shown by Proposition 4.6, point 2. A numerical example with two synchronization states is analyzed in § 5.3.

Denoting as above by Q the set $S^1 \cap S^2$ of shared states, let τ^i be the first hitting time of Q for the chain X^i , defined on Ω^i by

$$\tau^i = \inf\{n > 0 \mid X_n^i \in Q\}, \quad \text{noting that } \tau^i < \infty \text{ } M_x^i\text{-a.s.}$$

We consider the subset X_0 of global states given by

$$X_0 = \{(x, z) \in S^1 \times S^2 \mid x \in Q \wedge z \in Q \Rightarrow x = z\}.$$

Introduce also $\Delta = \{(\tau^1 < \infty) \wedge (\tau^2 < \infty) \wedge (X_{\tau^1}^1 = X_{\tau^2}^2)\}$, a measurable subset of $\Omega^1 \times \Omega^2$. Since the transition matrices we consider have all their coefficients positive, we have $M_x^1 \otimes M_z^2(\Delta) > 0$ for any global state (x, z) . We therefore equip the random pair of sequences $\sigma_0 = (X_1^1 X_2^1 \dots X_{\tau^1}^1, X_1^2 X_2^2 \dots X_{\tau^2}^2)$ with the conditional law

$$U_{(x,z)}(\cdot) = M_x^1 \otimes M_z^2(\cdot \mid \Delta).$$

Starting now from the global state $(X_{\tau^1}^1, X_{\tau^2}^2)$, we consider a fresh copy σ_1 of the same random pair of sequences, now equipped with the law $U_{(X_{\tau^1}^1, X_{\tau^2}^2)}$ (observe that, by construction, $X_{\tau^1}^1 = X_{\tau^2}^2$).

We construct inductively in this way a sequence $(\sigma_n)_{n \geq 0}$ of random trajectories, for which the concatenation $\omega = \sigma_0 \cdot \sigma_1 \dots$ is an element of Ω since $|\sigma_k| \geq (1, 1)$ for all $k \geq 0$. Denoting by $\mathbf{P}_{(x,z)}$ the law of ω thus constructed, we obtain a probabilistic two-components process (Definition 1.4), which is a closed process by construction.

Definition 2.1. The **synchronization of the two Markov chains** $(X_n^1)_{n \geq 0}$ and $(X_n^2)_{n \geq 0}$ is the probabilistic process $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$, where:

- (1) α ranges over the set $X_0 = \{(x, z) \in S^1 \times S^2 \mid x \in Q \wedge z \in Q \Rightarrow x = z\}$.
- (2) $\mathbf{P}_{(x,z)}$ is defined as the law of the infinite concatenation $\sigma_1 \cdot \sigma_2 \dots$, where $(\sigma_n)_{n \geq 0}$ is the countable Markov chain on the set \mathcal{E} of elementary trajectories with $U_{(x,z)} = M_x^1 \otimes M_z^2(\cdot \mid \Delta)$ as initial law, and transition kernel K given by:

$$\forall \sigma, \sigma' \in \mathcal{E}, \quad K(\sigma, \sigma') = U_{\gamma(\sigma)}(\sigma').$$

Translating the above definition in the two-components processes language consists of determining the value of $\mathbf{P}_\alpha(\uparrow v)$ for any finite trajectory v . This can be easily done only for v of the following special form:

$$v = \sigma_1 \cdot \dots \cdot \sigma_n, \quad \sigma_i \in \mathcal{E}, \quad \mathbf{P}_\alpha(\uparrow v) = U_\alpha(\sigma_1) \cdot K(\sigma_1, \sigma_2) \cdot \dots \cdot K(\sigma_{n-1}, \sigma_n).$$

Note that this entirely determines the probability \mathbf{P}_α ; since $\mathbf{P}_\alpha(v)$ for any finite trajectory v will be computed as the sum of all $\mathbf{P}_\alpha(w)$, for w of the form $w = \sigma_1 \cdot \dots \cdot \sigma_n$ and $v \leq w$, very much as we did in Step 2 in the proof of Lemma 1.8.

Theorem 2.2. *The synchronization of two Markov chains is a Markov two-components process.*

Proof. Let $\alpha = (x, z)$ be an initial state. Let $t \in \mathcal{E}$ be any elementary trajectory, and let s be any finite trajectory. Denote the coordinates of trajectories on each site by $s = (s^1, s^2)$ and $t = (t^1, t^2)$, and put $\gamma(s) = (x', z')$. Applying Lemma 1.8, we have to show that $(\mathbf{P}_\alpha)_s(t) = \mathbf{P}_{\gamma(s)}(t)$. We proceed in two steps.

- (1) *Step 1: s is synchronization free.* Then $s \cdot t$ is an elementary trajectory, and by construction of σ_0 we have $\uparrow(s \cdot t) = \{\sigma_0 = s \cdot t\}$ and thus $\mathbf{P}_\alpha(s \cdot t) = U_\alpha(s \cdot t)$ by construction. From this we compute:

$$(\mathbf{P}_\alpha)_s(t) = \frac{M_x \otimes M_z(\uparrow(s \cdot t) \cap \Delta)}{M_x \otimes M_z(\uparrow s \cap \Delta)}. \quad (2.1)$$

On the one hand, noting that $\uparrow(s \cdot t) \subset \Delta$, we have $M_x \otimes M_z(\uparrow(s \cdot t) \cap \Delta) = M_x(s^1 \cdot t^1)M_z(s^2 \cdot t^2)$. On the other hand, we have $M_x \otimes M_z(\uparrow s \cap \Delta) = M_x(s^1)M_z(s^2)M_x \otimes M_z(\Delta | \uparrow s)$ and, since M_x and M_z are Markov chains, $M_x \otimes M_z(\Delta | \uparrow s) = M_{x'} \otimes M_{z'}(\Delta)$. Going back to Eq. (2.1) we get:

$$\begin{aligned} (\mathbf{P}_\alpha)_s(t) &= \frac{M_x(s^1 \cdot t^1)}{M_x(s^1)} \times \frac{M_z(s^2 \cdot t^2)}{M_z(s^2)} \times \frac{1}{M_{x'} \otimes M_{z'}(\Delta)} \\ &= M_{x'}(t^1)M_{z'}(t^2) \frac{1}{M_{x'} \otimes M_{z'}(\Delta)} \\ &= M_{x'} \otimes M_{z'}(t | \Delta) = \mathbf{P}_{\gamma(s)}(t). \end{aligned}$$

- (2) *Step 2: s is any finite trajectory.* Let $s = \sigma_0 \cdot \sigma_1 \cdot \dots \cdot \sigma_p \cdot s'$ be the decomposition of s according to Proposition 1.2, case 1, so that $\sigma_0, \dots, \sigma_p$ are elementary trajectories, and s' is a synchronization free trajectory (the case $s' = \emptyset$ is admissible). We compute:

$$\begin{aligned} (\mathbf{P}_\alpha)_s(t) &= \frac{\mathbf{P}_\alpha(\sigma_0 \cdot \dots \cdot \sigma_p \cdot s' \cdot t)}{\mathbf{P}_\alpha(\sigma_0 \cdot \dots \cdot \sigma_p \cdot s')} \\ &= \frac{U_{(x,z)}(\sigma_0)K(\sigma_0, \sigma_1) \dots K(\sigma_{p-1}, \sigma_p)K(\sigma_p, s' \cdot t)}{U_{(x,z)}(\sigma_0)K(\sigma_0, \sigma_1) \dots K(\sigma_{p-1}, \sigma_p)U_{\gamma(\sigma_p)}(\uparrow s')} \\ &= U_{\gamma(\sigma_p)}(s' \cdot t | \uparrow s') \\ &= (\mathbf{P}_{\gamma(\sigma_p)})_{s'}(t) \\ &= \mathbf{P}_{\gamma(s)}(t), \end{aligned}$$

the last equality following from Step 1 together with $\gamma(s') = \gamma(s)$.

Conclusion: Lemma 1.8 applies, and $(\mathbf{P}_\alpha)_{\alpha \in X_0}$ is a M2CP. \square

3. STOPPING TIMES AND THE ASYNCHRONOUS STRONG MARKOV PROPERTY

All the notions and results of this section do not depend on the particular structure of trajectories, and in particular they do not rest on Proposition 1.2. It follows that they have straightforward generalizations to an asynchronous model with an arbitrary number $n \geq 2$ of sites.

3.1. Stopping Times. Stopping times are a fundamental tool in the theory of probabilistic processes in general, and in the theory of Markov chain in particular. Recall that a stopping time associated to a Markov chain is a random integer, maybe infinite and seen as a random time instant, with the following property: an observer aware of the successive values of the chain can decide at each instant whether the stopping time has already been reached or not. Standard examples of stopping times in Markov chain theory are: constant times (trivial since non random); the *first* instant where the chain hits a given state; more generally the

first instant where a chain reaches a given set of states. A standard example of a random time which is *not* a stopping time is the *last* instant where the chain hits a given state.

It is natural to introduce an equivalent notion for two-components processes, and this is the topic of this subsection. We will see that the first instant a process hits a global state defines a stopping time; but, and contrasting with Markov chains, the first instant of reaching a given *set* of global states does not define a stopping time in general, unless one considers special kinds of sets.

Recall that $\mathcal{T} = (\mathbb{N} \times \mathbb{N}) \cup \{\infty\}$ denotes the partially ordered set of “two-components time instants”.

Definition 3.1 (Random times and stopping times). Let $T : \Omega \rightarrow \mathcal{T}$ be an arbitrary mapping. For any $\omega \in \Omega$, we denote by ω_T the prefix of ω of length $T(\omega)$ if $T(\omega) < \infty$, and we put $\omega_T = \omega$ if $T(\omega) = \infty$.

- (1) We say that T is a **random time** if ω_T is a subtrajectory of ω for all $\omega \in \Omega$.
- (2) If T is a random time we say that T is a **stopping time** if furthermore the following property holds:

$$\forall \omega, \omega' \in \Omega \quad \omega' \geq \omega_T \Rightarrow \omega_T = \omega'_T. \quad (3.1)$$

Actually since the space Ω is always implicitly equipped with an initial state α , a more general notion of stopping times would be as for probabilistic processes a *family* of random times $(T_\alpha)_{\alpha \in X_0}$, each one satisfying condition (3.1). But, since we will only be concerned with stopping times independent of α , we prefer limiting ourselves to Definition 3.1 as it is formulated.

Since \mathcal{T} is a countable set, it is naturally equipped with its discrete σ -algebra. It turns out that a stopping time $T : \Omega \rightarrow \mathcal{T}$ is always measurable; and so is the mapping $\omega \in \Omega \mapsto \omega_T$, provided we equip the set of trajectories (either finite or infinite) with the σ -algebra generated by the sets of the form $\{v \mid s \leq v\}$, for s ranging over finite trajectories, and v ranging over trajectories. If the set of trajectories is seen as a DCPO (Directed Complete Partial Order [10]), this is the Borel σ -algebra associated with the Scott topology on the DCPO. Obviously, it induces by restriction the σ -algebra \mathfrak{F} on the subset Ω .

Proposition 3.2. *Let $T : \Omega \rightarrow \mathcal{T}$ be a stopping time. We equip \mathcal{T} with its discrete σ -algebra, and we equip the set of trajectories with its Borel σ -algebra described above.*

- (1) *Then T and ω_T are two measurable mappings.*
- (2) *Let \mathfrak{F}_T denote the σ -algebra generated by ω_T . Then \mathfrak{F}_T is finer than the σ -algebra generated by T , and it is characterized as follows:*

$$\forall A \in \mathfrak{F} \quad A \in \mathfrak{F}_T \iff \forall \omega, \omega' \in \Omega \quad \omega \in A \wedge \omega' \geq \omega_T \Rightarrow \omega' \in A.$$

Proof. If $Y : (\Omega, \mathfrak{F}) \rightarrow (A, \mathfrak{G})$ is a measurable mapping, we denote by $\langle Y \rangle$ the sub- σ -algebra of \mathfrak{F} generated by Y , and given by $\langle Y \rangle = \{Y^{-1}(U) \mid U \in \mathfrak{G}\}$. For $\omega \in \Omega$, let $\zeta(\omega) = \omega_T$. For any finite trajectory v , we put

$$S_v = \{\omega \text{ trajectory} \mid v \leq \omega\}.$$

Since T is a stopping time, $\zeta^{-1}(\{v\})$ is either empty or equal to $\uparrow v$, so $\zeta^{-1}(\{v\})$ is measurable in either cases. Let us denote by \mathcal{V} the set of finite trajectories. Since \mathcal{V} is countable, it follows that $\zeta^{-1}(S_v \cap \mathcal{V}) = \bigcup_{w \in S_v \cap \mathcal{V}} \zeta^{-1}(\{w\})$ is measurable for any $v \in \mathcal{V}$, as well as $\zeta^{-1}(\mathcal{V})$. By definition of $\zeta = \omega_T$, we have that $\zeta(\omega)$ is either finite or maximal. From this,

it follows first that $\zeta^{-1}(\Omega) = \Omega \setminus \zeta^{-1}(\mathcal{V})$ is measurable; and second:

$$\zeta^{-1}(S_v) = \zeta^{-1}(S_v \cap \mathcal{V}) \cup \zeta^{-1}(\uparrow v).$$

But $\zeta^{-1}(\uparrow v) = \zeta^{-1}(\Omega) \cap \uparrow v$, hence $\zeta^{-1}(S_v)$ is the union of two measurable subsets of Ω , and is thus measurable. This shows that ζ is a measurable mapping.

To prove that T is measurable, observe that:

$$\forall (m, n) \in \mathcal{T} \quad \{T = (m, n)\} = \bigcup_{|v|=(m,n)} \zeta^{-1}(v).$$

Since the union is finite, it follows that $\{T = (m, n)\}$ is a $\langle \zeta \rangle$ -measurable subset, from which we deduce that $\{T = \infty\} = \bigcup_{(m,n) \in \mathbb{N} \times \mathbb{N}} \{T \neq (m, n)\}$ is also a $\langle \zeta \rangle$ -measurable subset. Therefore $\langle T \rangle \subset \langle \zeta \rangle$. By the property of stopping times $\omega' \geq \omega_T$ is equivalent to $\omega_T = \omega'_T$, from which follows the characterization of $\mathfrak{F}_T = \langle \zeta \rangle$. \square

Note that any function $f : (\Omega, \mathfrak{F}) \rightarrow (A, \mathfrak{G})$ with value in some measurable space is measurable with respect to \mathfrak{F}_T if and only if it is constant on elementary cylinders of the form $\uparrow v = \{\omega_T = v\}$ with v ranging over the values of ω_T —since it is well known that f is \mathfrak{F}_T -measurable if and only if it can be written as $f(\omega) = g(\omega_T)$ with g some measurable mapping.

3.2. Shift Operators. In Markov chain theory, the “universal” shift operator θ is classically defined on the space of trajectories of a Markov chain by $\theta(x_0 x_1 \dots) = (x_1 x_2 \dots)$. Its iterations θ_n are defined for $n \geq 0$ by $\theta_0 = \text{Id}$ and $\theta_{n+1} = \theta \circ \theta_n$. Allowing the time index n to be random, one defines θ_T , for $T : \Omega \rightarrow \mathbb{N}$ any random variable, by $\theta_T(\omega) = \theta_{T(\omega)}(\omega)$. In our framework, there is no such “universal” shift operator θ . Yet, each stopping time $T : \Omega \rightarrow \mathcal{T}$ induces a shift operator $\theta_T : \Omega \rightarrow \Omega$. Informally $\theta_T(\omega)$ is the queue of trajectory ω that remains “after” the prefix trajectory ω_T .

Definition 3.3. Let $T : \Omega \rightarrow \mathcal{T}$ be a stopping time. The **shift operator** associated with T is the mapping $\theta_T : \Omega \rightarrow \Omega$, which is only partially defined; if $T(\omega) < \infty$, then $\theta_T(\omega)$ is defined as the unique element of Ω such that

$$\omega = \omega_T \cdot \theta_T(\omega),$$

and $\theta_T(\omega)$ is undefined otherwise.

The shift operator allows to define an addition on stopping times, as shown by the following result which mimics an equivalent result widely used in Markov chain theory.

Lemma 3.4. *Let $S, T : \Omega \rightarrow \mathcal{T}$ be two stopping times. Then $U = S + T \circ \theta_S$ is a stopping time (it is understood that $U = \infty$ if $S = \infty$).*

Proof. It is clear that U is a random time. Let $\omega, \omega' \in \Omega$ such that $\omega' \geq \omega_U$, we have to show that $\omega'_U = \omega_U$. If $S(\omega) = \infty$ then $U(\omega) = \infty$ and then $\omega' = \omega$ and $\omega_U = \omega'_U$, trivially.

Hence we assume without loss of generality that $S(\omega) < \infty$, and we put $\zeta = \theta_S(\omega)$ and $\zeta' = \theta_S(\omega')$. We have $\omega_U = \omega_S \cdot \zeta_T$. Hence $\omega' \geq \omega_S$ and thus $\omega'_S = \omega_S$ since S is a stopping time. Therefore: $\omega' = \omega_S \cdot \zeta' \geq \omega_U = \omega_S \cdot \zeta_T$, hence $\zeta' \geq \zeta_T$. Thus $\zeta'_T = \zeta_T$ since T is a stopping time. We have finally $\omega'_U = \omega'_S \cdot \zeta'_T = \omega_S \cdot \zeta_T = \omega_U$, proving that U is a stopping time. \square

Starting from a given stopping time T , we use Lemma 3.4 above to iterate the “addition” of T to itself.

Definition 3.5. Let $T : \Omega \rightarrow \mathcal{T}$ be a stopping time, and let θ_T be the associated shift operator. The sequence $(T^n)_{n \geq 0}$ of mappings $\Omega \rightarrow \mathcal{T}$ defined as follows:

$$T^0 = (0, 0) \quad \forall n \geq 0 \quad T^{n+1} = T^n + T \circ \theta_{T^n}$$

with the convention that $T^{n+1} = \infty$ on $\{T^n = \infty\}$, is a sequence of stopping times, called **iterated stopping times associated with T** .

Remark that $\theta_{T^0} = \text{Id}_\Omega$, and $T^1 = T$.

3.3. Examples of Stopping Times. In this subsection we review some examples of random times, and analyze whether they are stopping times or not. Some of the examples introduced here will be used later in §§ 3.5–3.7.

3.3.1. Constant Times are not Random Times in General. In general, if $(m, n) \in \mathbb{N} \times \mathbb{N}$, then the random variable constant and equal to (m, n) is not a random time. For instance, take $(m, n) = (2, 2)$ and consider as in the Introduction a maximal trajectory ω starting with $(a \cdot \mathbf{c}, e \cdot f \cdot \mathbf{c})$ with \mathbf{c} as synchronization state. Then the prefix of length $(2, 2)$ of ω is $(a \cdot \mathbf{c}, e \cdot f)$, which is not a trajectory. Hence the constant $T = (2, 2)$ is not a random time. This contrasts with Markov chain theory, where constant times are a basic example of stopping times.

However note that any constant time is indeed a random time if the process is open (Definition 1.5). And in this case, it is also a stopping time.

3.3.2. A Random Time which is not a Stopping Time. For ω a maximal trajectory, let v be the first elementary trajectory in the decomposition of ω as in Proposition 1.2, which is defined if ω has at least one synchronization. Then v has the form $v = u \cdot (y, y)$ for some unique finite trajectory u and state $y \in Q$. Put $\omega_T = u$ in this case, and $\omega_T = \omega$ if v is not defined. Time $T(\omega)$ represents the “last instant before first synchronization”. By construction, T is a random time since ω_T is a subtrajectory of ω .

However T is not a stopping time in general. For example, consider $S^1 = \{a, b, \mathbf{c}, \mathbf{d}\}$ and $S^2 = \{\mathbf{c}, \mathbf{d}, e, f\}$, if ω starts with $(a \cdot \mathbf{c}, f \cdot f \cdot e \cdot \mathbf{c})$, then $T(\omega) = (1, 3)$ and $\omega_T = (a, f \cdot f \cdot e)$, corresponding to the last private states a and e before synchronization on (\mathbf{c}, \mathbf{c}) . And if ω' starts with $(a \cdot b \cdot \mathbf{c}, f \cdot f \cdot e \cdot f \cdot \mathbf{c})$, then $T(\omega') = (2, 3) \neq T(\omega)$ although $\omega' \geq \omega_T$. This shows that T is not a stopping time.

3.3.3. The First Return Time of a Global State. Let $\alpha \in X_0$ be a given global state. For $\omega \in \Omega$, consider the following set of finite subtrajectories of ω :

$$N_\alpha(\omega) = \{v \leq \omega \mid v \text{ finite subtrajectory of } \omega \wedge \gamma(v) = \alpha \wedge |v| \geq (1, 1)\}.$$

If nonempty, N_α is a sublattice of the lattice of subtrajectories of ω since, by Lemma 1.3, lower bounds are taken component by component so that $\gamma(v \wedge v') = \alpha$ whenever $v, v' \in N_\alpha(\omega)$ and $|v \wedge v'| \geq (1, 1)$. In particular, if we put $v = \min N_\alpha(\omega)$, which exists whenever $N_\alpha(\omega) \neq \emptyset$, then $\gamma(v) = \alpha$ and $|v| \geq (1, 1)$. We define thus the first return time to α as follows.




Figure 3: A finite trajectory synchronizing on shared states c and d .

Definition 3.6. For any $\alpha \in X_0$, the **first return time to α** is the stopping time $R_\alpha : \Omega \rightarrow \mathcal{T}$ defined by:

$$\forall \omega \in \Omega \quad \omega_{R_\alpha} = \begin{cases} \omega, & \text{if } N_\alpha(\omega) = \emptyset, \text{ and thus } R_\alpha(\omega) = \infty, \\ \min N_\alpha(\omega), & \text{otherwise, and thus } R_\alpha(\omega) = |\min N_\alpha(\omega)|. \end{cases}$$

The **successive return times to α** are the iterated stopping times $(R_\alpha^n)_{n \geq 1}$ associated with R_α as in Definition 3.5.

For any finite subtrajectory v of ω , we have:

$$(\gamma(v) = \alpha \wedge |v| \geq (1, 1)) \implies \omega_{R_\alpha} \leq v,$$

which is consistent with the intuition of what a “first return time” should be. To show that R_α is indeed a stopping time, observe first that ω_{R_α} is clearly a subtrajectory of ω . And second, if $\omega' \in \Omega$ is such that $\omega' \geq \omega_{R_\alpha}$, that implies that $\omega_{R_\alpha} \in N_\alpha(\omega')$, and thus $\omega'_{R_\alpha} \leq \omega_{R_\alpha}$ by minimality of ω'_{R_α} . But then $\omega'_{R_\alpha} \in N_\alpha(\omega)$, and thus $\omega_{R_\alpha} \leq \omega'_{R_\alpha}$ by minimality of ω_{R_α} . Hence $\omega_{R_\alpha} = \omega'_{R_\alpha}$, and this shows that R_α is a stopping time.

As an example, consider $S^1 = \{a, b, c, d\}$ and $S^2 = \{c, d, e, f\}$, and a maximal trajectory ω starting with $(a \cdot b \cdot c \cdot a \cdot d, e \cdot c \cdot e \cdot f \cdot d)$, which is depicted in Figure 3. Consider the global state $\alpha = (a, e)$. Then $R_\alpha(\omega) = (1, 1)$, and $\omega_{R_\alpha} = (a, e)$. Note that, since R_α is indeed a stopping time, we do not need to know the queue of ω to already have information on $R(\omega)$.

Let us determine the value of next return $R_\alpha^2(\omega)$ to α . The shifted trajectory $\theta_{R_\alpha}(\omega)$ starts with $(b \cdot c \cdot a \cdot d, c \cdot e \cdot f \cdot d)$. Therefore $R_\alpha(\theta_{R_\alpha}(\omega)) = (3, 2)$, and $R_\alpha^2(\omega) = (1, 1) + (3, 2) = (4, 3)$. Note that $R^3(\omega)$ is undetermined at this stage.

If ζ is the trajectory $(b \cdot b \cdot \dots, e \cdot e \cdot \dots)$, with only b on the first component and only e on the second component, then $R_\alpha(\zeta) = \infty$ and $\zeta_{R_\alpha} = \zeta$.

3.3.4. Supremum of Stopping Times. If S and T are two stopping times, then the random time $S \vee T$ defined by $\omega_{S \vee T} = \omega_S \vee \omega_T$ is a stopping time. For, if $\omega' \geq \omega_S \vee \omega_T$, then $\omega' \geq \omega_S$ and $\omega' \geq \omega_T$, therefore $\omega'_S = \omega_S$ and $\omega'_T = \omega_T$, hence $\omega'_{S \vee T} = \omega_{S \vee T}$. The same line of proof shows that the supremum of any family of stopping times is a stopping time.

3.3.5. *The Infimum of Stopping Times may not be a Stopping Time.* Contrasting with stopping times from Markov chain theory however, the infimum of two stopping times S and T , defined by $\omega_{S \wedge T} = \omega_S \wedge \omega_T$, may not be a stopping time. Let us consider an example. Let $S^1 = \{a, b, \mathbf{c}\}$ and $S^2 = \{\mathbf{c}, e, f\}$. Let $\alpha = (a, e)$ and $\beta = (b, f)$, and let $S = T_\alpha$ and $T = T_\beta$ be the first return times to α and to β respectively. Consider a trajectory ω starting with $(a \cdot b, f \cdot e)$. Then $\omega_S = (a, f \cdot e)$ and $\omega_T = (a \cdot b, f)$, and thus $\omega_{S \wedge T} = (a, f)$. However, if ω' is the maximal trajectory defined by $\omega' = (a \cdot a \cdots, f \cdot f \cdots)$ we have $\omega' \geq \omega_{S \wedge T}$ on the one hand, and $\omega'_S = \omega'$ and $\omega'_T = \omega'$ on the other hand, so that $\omega'_{S \wedge T} = \omega' \neq \omega_{S \wedge T}$. This show that $S \wedge T$ is not a stopping time.

This example is specific to the asynchronous structure we consider, since it makes use of the *partially* ordered structure of trajectories.

3.3.6. *First Return Time to a Square Set of Global States.* Since the infimum of stopping times is not a stopping time in general, there is an issue for defining the first return time to a set of global states. There is actually no obvious way of defining such a thing in general, as the analysis of the above example reveals. The situation however becomes favorable if one considers a set of states satisfying the following property.

Definition 3.7. We say a subset $A \subset X_0$ of global sets is a **square set** if it has the form $A = X_0 \cap (S'_1 \times S'_2)$ where $S'_1 \subset S_1$ and $S'_2 \subset S_2$.

A first example of a square set is X_0 itself. We will also encounter the square set $(Q \times Q) \cap X_0$. If $\alpha = (x, z)$ and $\beta = (x', z')$, the smallest square set containing α and β is $\{\alpha, \beta, (x, z'), (x', z)\}$.

Assume that A is a square set of global states. Define then, for any $\omega \in \Omega$:

$$N_A(\omega) = \{v \leq \omega \mid v \text{ finite subtrajectory of } \omega \wedge \gamma(v) \in A \wedge |v| \geq (1, 1)\}.$$

Then $N_A(\omega)$ is a sublattice of the lattice of finite subtrajectories of ω whenever it is nonempty. Indeed, since A is a square set. The random time R_A defined by

$$\omega_{R_A} = \min N_A(\omega),$$

and by $R_A = \infty$ as usual when $N_A(\omega)$ is empty, is a stopping time that satisfies $\gamma(\omega_{R_A}) \in A$ whenever or $R_A < \infty$. We define R_A as the *first return time* to the square set A . One furthermore checks that $\omega_{R_A} = \bigwedge_{\alpha \in A} \omega_{R_\alpha}$, providing an example of infimum of stopping times the result of which is indeed a stopping time.

Let us examine the first return times associated with the square sets X_0 and $(Q \times Q) \cap X_0$. In Markov chain theory, R_{X_0} would correspond to the constant time 1. But in the asynchronous framework its action is less simple. Stopping time R_{X_0} can be described as follows: $\omega_{R_{X_0}}$ is the smallest subtrajectory of ω with length $\geq (1, 1)$. In particular, $R_{X_0}(\omega)$ is always *finite*.

We detail the action of R_{X_0} on an example. Consider $S^1 = \{a, b, \mathbf{c}, \mathbf{d}\}$ and $S^2 = \{\mathbf{c}, \mathbf{d}, e, f\}$, and let ω be some maximal trajectory starting with $(a \cdot b \cdot \mathbf{c} \cdot a \cdot \mathbf{d}, e \cdot \mathbf{c} \cdot e \cdot f \cdot \mathbf{d})$, as depicted in Figure 3 above. The exercise consists in finding the values of $R_{X_0}^n(\omega)$ for the first integers n , where $R_{X_0}^n$ denote the iterated stopping times associated with R_{X_0} as in Definition 3.5. Obviously $R_{X_0}(\omega) = (1, 1)$. The shifted trajectory $\theta_{R_{X_0}}(\omega)$ starts with $(b \cdot \mathbf{c} \cdot a \cdot \mathbf{d}, \mathbf{c} \cdot e \cdot f \cdot \mathbf{d})$. The smallest subtrajectory of $\theta_{R_{X_0}}(\omega)$ of length at least $(1, 1)$ is

$(b \cdot \mathbf{c}, \mathbf{c})$, and thus $R_{X_0}(\theta_{R_{X_0}}(\omega)) = (2, 1)$. Hence $R_{X_0}^2(\omega) = (1, 1) + (2, 1) = (3, 2)$. The finite trajectories

$$\omega_{R_{X_0}} = (a, e) \quad \text{and} \quad (\theta_{R_{X_0}}(\omega))_{R_{X_0}} = (b \cdot \mathbf{c}, \mathbf{c})$$

yield the following initial decomposition of ω : $\omega = (a, e) \cdot (b \cdot \mathbf{c}, \mathbf{c}) \cdot \theta_{R_{X_0}^2}(\omega)$. For the next values $n = 3, 4$ we find $R_{X_0}^3(\omega) = (4, 3)$ and $R_{X_0}^4(\omega) = (5, 5)$, corresponding to the initial decomposition $\omega = (a, e) \cdot (b \cdot \mathbf{c}, \mathbf{c}) \cdot (a, e) \cdot (\mathbf{d}, f \cdot \mathbf{d}) \cdots$.

Coming now to the square set $(Q \times Q) \cap X_0$, and denoting by R_Q the first return time associated with it, we may rephrase the definition of infinite synchronization of trajectories (Definition 1.5) as follows: a maximal trajectory ω synchronizes infinitely often if $\omega \in \bigcap_{n \geq 1} \{R_Q^n < \infty\}$. A probabilistic process \mathbb{P} is closed if $R_Q^n < \infty$ for all $n \geq 1$ and \mathbf{P}_α -almost surely, for all $\alpha \in X_0$. It is open if $R_Q = \infty$, \mathbf{P}_α -almost surely and for all $\alpha \in X_0$.

We end this series of examples with the following result which will be useful in the study of recurrence of global states. It makes use of the finitary assumption on the set of global states.

Lemma 3.8. *Let A be a square set. Denoting by $(R_A^n)_{n \geq 1}$ the successive returns to A , i.e., the iterated stopping times associated with the first return time R_A , and by $(R_\alpha^n)_{n \geq 1}$ the successive return times to α for any $\alpha \in A$, we have the following equality of sets:*

$$\bigcap_{n \geq 1} \{R_A^n < \infty\} = \bigcup_{\alpha \in A} \bigcap_{n \geq 1} \{R_\alpha^n < \infty\}.$$

Proof. The \supset inclusion is obvious. For the converse inclusion, let $\omega \in \Omega$ be such that $R_A^n(\omega) < \infty$ for all $n \geq 1$. Since A is a finite set, there exists some state $\alpha \in A$ and a strictly increasing sequence of integers $(n_k)_{k \geq 1}$ such that $\gamma(R_A^{n_k}) = \alpha$ for all k . By induction on k , we show that $R_\alpha^k(\omega) \leq R_A^{n_k}(\omega)$ for all integers $k \geq 1$. The finite trajectory $v = R_A(\omega)$ is a subtrajectory of ω satisfying $\gamma(v) = \alpha$ and $|v| \geq (1, 1)$, and therefore $\omega_{R_\alpha} \leq v$. Since the sequence $(R_A^n(\omega))_{n \geq 1}$ is increasing, as shown by the formula in Definition 3.5 that defines it, we have $\omega_{R_\alpha} \leq v = \omega_{R_A^1} \leq \omega_{R_A^{n_1}}$. Assume for the induction that $R_\alpha^k(\omega) \leq R_A^{n_k}(\omega)$. Then there is some finite trajectory v such that $\omega_{R_A^{n_k}} = \omega_{R_\alpha^k} \cdot v$. Since $n_k < n_{k+1}$, there is also some finite trajectory v' such that $\gamma(v') = \alpha$, $|v'| \geq (1, 1)$ and $\omega_{R_A^{n_{k+1}}} = \omega_{R_A^{n_k}} \cdot v'$. We obtain thus:

$$\omega_{R_A^{n_{k+1}}} = \omega_{R_\alpha^k} \cdot v \cdot v', \quad |v \cdot v'| \geq (1, 1), \quad \gamma(v \cdot v') = \alpha.$$

This implies that $R_\alpha(\theta_{R_\alpha^k}(\omega)) \leq |v \cdot v'|$. By definition, we have $R_\alpha^{k+1} = R_\alpha^k + R_\alpha \circ \theta_{R_\alpha^k}$, whence:

$$|R_\alpha^{k+1}(\omega)| \leq |R_\alpha^k(\omega)| + |v| + |v'| = |R_A^{n_k}(\omega)| + |v'| = |R_A^{n_{k+1}}(\omega)|,$$

completing the induction. This implies in particular that $R_\alpha^k(\omega) < \infty$ for all $k \geq 1$, as expected. \square

3.4. The Asynchronous Strong Markov Property. The Asynchronous Strong Markov Property that we state below has the exact same formulation than the Strong Markov property for Markov chains found in classical references [17, Theorem 3.5 p.23]. The syntactical identity underlines the parallel with Markov chain theory, although the interpretation of symbols must be changed of course: stopping times must be understood in the sense of Definition 3.1, the associated σ -algebra in the sense given in Proposition 3.2, and of course M2CPs replace Markov chains. Nevertheless, once the Asynchronous Strong Markov property has been established, it is possible to transfer verbatim some pieces of Markov chain theory. Examples of such transfers are Lemma 3.10 given just after Theorem 3.9 and the 0-1 law for the infinite return to a given global state, given in point 1 of Proposition-definition 3.11 below.

Theorem 3.9 (Asynchronous Strong Markov property). *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP. For any measurable and non negative function $h : \Omega \rightarrow \mathbb{R}$ and for any stopping time $T : \Omega \rightarrow \mathcal{T}$, we have*

$$\forall \alpha \in X_0 \quad \mathbf{E}_\alpha(h \circ \theta_T | \mathfrak{F}_T) = \mathbf{E}_{\gamma(\omega_T)}(h), \quad \mathbf{P}_\alpha\text{-a.s.}, \quad (3.2)$$

where $\mathbf{E}_\alpha(\cdot | \mathfrak{F}_T)$ denotes the conditional expectation with respect to probability \mathbf{P}_α and σ -algebra \mathfrak{F}_T . By convention, both sides of Eq. (3.2) vanish outside $\{T < \infty\}$.

Note that, as for the Strong Markov Property for Markov chains, both sides of Eq. (3.2) are random variables: the left side, since it is a conditional expectation with respect to σ -algebra \mathfrak{F}_T ; and the right side, since it depends on the random variable $\gamma(\omega_T)$.

Proof. Let Z denote the random variable $Z = \mathbf{E}_{\gamma(\omega_T)}(h)$, which is obviously \mathfrak{F}_T -measurable since $\gamma(\omega_T)$ is. Let ϕ be any non negative, bounded and \mathfrak{F}_T -measurable function. Denote by R the set of finite trajectories taken by ω_T . Then, since R is at most countable:

$$\mathbf{E}_\alpha(\phi \cdot h \circ \theta_T) = \sum_{v \in R} \mathbf{E}_\alpha(\mathbf{1}_{\{\omega_T=v\}} \cdot \phi \cdot h \circ \theta_T). \quad (3.3)$$

Since T is a stopping time, and since $T^{-1}(v) \neq \emptyset$ if $v \in R$, we have $\{\omega_T = v\} = \uparrow v$. Furthermore, ϕ is constant on $\{\omega_T = v\}$, so that if $\phi(v)$ denote this constant, we get:

$$\mathbf{E}_\alpha(\mathbf{1}_{\{\omega_T=v\}} \cdot \phi \cdot h \circ \theta_T) = \phi(v) \mathbf{P}_\alpha(\uparrow v) \frac{\mathbf{E}_\alpha(\mathbf{1}_{\{\uparrow v\}} h \circ \theta_T)}{\mathbf{P}_\alpha(\uparrow v)}$$

Recognizing the conditional expectation defined as the future of v w.t.r. to probability \mathbf{P}_α , we use the Markov property (1.5) of Definition 1.7 to get:

$$\mathbf{E}_\alpha(\mathbf{1}_{\{\omega_T=v\}} \cdot \phi \cdot h \circ \theta_T) = \mathbf{P}_\alpha(\uparrow v) \phi(v) \mathbf{E}_{\gamma(v)}(h).$$

Going back to Eq. (3.3) we obtain:

$$\mathbf{E}_\alpha(\phi \cdot h \circ \theta_T) = \sum_{v \in R} \mathbf{P}_\alpha(\uparrow v) \phi(v) \mathbf{E}_{\gamma(v)}(h) = \mathbf{E}_\alpha(\phi Z).$$

This shows that $Z = \mathbf{E}_\alpha(h \circ \theta_T | \mathfrak{F}_T)$. □

The following result is a typical application of the Strong Markov property in Markov chain theory that applies here too. It intuitively says this: the probability of returning infinitely often to a state β , starting from α , is the product of the probability of hitting β once starting from α , by the probability of returning to β infinitely often, starting from β .

Lemma 3.10. *Let α, β be two global states, and let $(R_\beta^n)_{n \geq 1}$ be the successive return times to β . Let $h = \mathbf{1}_{\{\bigcap_{n \geq 1} \{R_\beta^n < \infty\}\}}$. Then:*

$$\mathbf{E}_\alpha(h) = \mathbf{P}_\alpha(R_\beta < \infty) \cdot \mathbf{E}_\beta(h). \quad (3.4)$$

Proof. Applying the Asynchronous Strong Markov property (Theorem 3.9) with stopping time R_β and function h , we get: $\mathbf{E}_\alpha(h \circ \theta_{R_\beta} | \mathfrak{F}_{R_\beta}) = \mathbf{E}_{\gamma(\omega_{R_\beta})}(h)$. The right side of this equality is simply the constant $\mathbf{E}_\beta(h)$ on $\{R_\beta < \infty\}$. We multiply both sides by $\mathbf{1}_{\{R_\beta < \infty\}}$, which is \mathfrak{F}_{R_β} -measurable by definition of \mathfrak{F}_{R_β} and can therefore be put inside the $\mathbf{E}_\alpha(\cdot | \mathfrak{F}_{R_\beta})$ sign, to obtain:

$$\mathbf{E}_\alpha(\mathbf{1}_{\{R_\beta < \infty\}} h \circ \theta_{R_\beta} | \mathfrak{F}_{R_\beta}) = \mathbf{1}_{\{R_\beta < \infty\}} \mathbf{E}_\beta(h).$$

We observe that $\mathbf{1}_{\{R_\beta < \infty\}} h \circ \theta_{R_\beta} = h$, and therefore $\mathbf{E}_\alpha(h | \mathfrak{F}_{R_\beta}) = \mathbf{1}_{\{R_\beta < \infty\}} \mathbf{E}_\beta(h)$. Taking the \mathbf{E}_α -expectations yields identity (3.4). \square

3.5. Recurrent and Transient Global States. In Markov chain theory, the Strong Markov property is a fundamental tool for studying so-called recurrent states, those states to which the chain returns infinitely often almost surely. There is a strong parallel between Markov chain theory and this part of M2CP theory: recurrence concerns global states, and the infinite return is defined through the successive return times defined in § 3.3. And the Asynchronous Strong Markov property is the fundamental tool in this study.

Denoting as in Definition 3.6 by $(R_\alpha^n)_{n \geq 1}$ the successive returns to $\alpha \in X_0$, we say that a global trajectory $\omega \in \Omega$ **returns infinitely often** to α if $R_\alpha^n(\omega) < \infty$ for all integers $n \geq 1$.

Proposition and definition 3.11. *Let $(\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP.*

(1) *For any $\alpha \in X_0$, the set of trajectories that return infinitely often to α has \mathbf{P}_α -probability either 0 or 1. Following Markov chain terminology, we will say that:*

- α is **recurrent** if $\mathbf{P}_\alpha(\bigcap_{n \geq 1} \{R_\alpha^n < \infty\}) = 1$, which is equivalent to:

$$\mathbf{P}_\alpha(R_\alpha < \infty) = 1.$$

- α is **transient** if $\mathbf{P}_\alpha(\bigcap_{n \geq 0} \{R_\alpha^n < \infty\}) = 0$, which is equivalent to:

$$\mathbf{P}_\alpha(R_\alpha < \infty) < 1.$$

(2) *There is at least one recurrent state in X_0 .*

(3) *If α is a recurrent state, then the successive returning trajectory to α defined by $\rho_n = (\theta_{R_\alpha^{n-1}}(\omega))_{R_\alpha}$ for $n \geq 1$, form a sequence of independent and identically distributed finite trajectories w.r.t. probability \mathbf{P}_α .*

(4) *If α is a recurrent state, and if β is reachable from α , then β is recurrent and α is reachable from β .*

Proof.

- (1) The proof is adapted from [17, Proposition 1.2 p.65]. Recall the usual transformation, for a measurable subset A and some sub- σ -algebra \mathfrak{G} of a probability space $(\Omega, \mathfrak{F}, \mathbf{P})$: $\mathbf{P}(A) = \mathbf{E}(\mathbf{1}_A) = \mathbf{E}(\mathbf{E}(\mathbf{1}_A | \mathfrak{G}))$. Putting $R = R_\alpha$ and $R^n = R_\alpha^n$, we apply this transformation to $(\Omega, \mathfrak{F}, \mathbf{P}_\alpha)$ with $A = \{R^n < \infty\}$ and $\mathfrak{G} = \mathfrak{F}_{R^{n-1}}$:

$$\mathbf{P}_\alpha(R^n < \infty) = \mathbf{E}_\alpha(\mathbf{E}_\alpha(\mathbf{1}_{\{R^n < \infty\}} | \mathfrak{F}_{R^{n-1}})).$$

Since $R^n = R^{n-1} + R \circ \theta_{R^{n-1}}$ we have: $\mathbf{1}_{\{R^n < \infty\}} = \mathbf{1}_{\{R^{n-1} < \infty\}} \cdot \mathbf{1}_{\{R \circ \theta_{R^{n-1}} < \infty\}}$. Since $\mathbf{1}_{\{R^{n-1} < \infty\}}$ is $\mathfrak{F}_{R^{n-1}}$ -measurable, the usual property of conditional expectation yields:

$$\mathbf{P}_\alpha(R^n < \infty) = \mathbf{E}_\alpha(\mathbf{1}_{\{R^{n-1} < \infty\}} \mathbf{E}_\alpha(\mathbf{1}_{\{R \circ \theta_{R^{n-1}} < \infty\}} | \mathfrak{F}_{R^{n-1}})). \quad (3.5)$$

Applying the Asynchronous Strong Markov property (Theorem 3.9) with stopping time R^{n-1} and function $\mathbf{1}_{\{R < \infty\}}$ we have:

$$\mathbf{E}_\alpha(\mathbf{1}_{\{R \circ \theta_{R^{n-1}} < \infty\}} | \mathfrak{F}_{R^{n-1}}) = \mathbf{E}_{\gamma(\omega_{R^{n-1}})}(\mathbf{1}_{\{R < \infty\}}). \quad (3.6)$$

Since $\gamma(\omega_{R^{n-1}}) = \alpha$ on $\{R^{n-1} < \infty\}$, multiplying both sides of (3.6) by $\mathbf{1}_{\{R^{n-1} < \infty\}}$ brings:

$$\mathbf{1}_{\{R^{n-1} < \infty\}} \mathbf{E}_\alpha(\mathbf{1}_{\{R \circ \theta_{R^{n-1}} < \infty\}} | \mathfrak{F}_{R^{n-1}}) = \mathbf{1}_{\{R^{n-1} < \infty\}} \mathbf{E}_\alpha(\mathbf{1}_{\{R < \infty\}}). \quad (3.7)$$

We take the \mathbf{P}_α -expectation of both sides of (3.7) and report the result in (3.5) to obtain:

$$\mathbf{P}_\alpha(R^n < \infty) = \mathbf{P}_\alpha(R < \infty) \cdot \mathbf{P}_\alpha(R^{n-1} < \infty). \quad (3.8)$$

It follows from Borel-Cantelli Lemma that α is recurrent if $\mathbf{P}_\alpha(R < \infty) = 1$, and transient if $\mathbf{P}_\alpha(R < \infty) < 1$.

- (2) Pick any $\alpha \in X_0$, and let $(R^n)_{n \geq 1}$ be the successive return times to the square set X_0 (cf. § 3.3.6). With \mathbf{P}_α -probability 1, $R^n < \infty$ for all $n \geq 1$. It follows from Lemma 3.8 applied with $A = X_0$ that, for some $\beta \in X_0$:

$$\mathbf{P}_\alpha\left(\bigcap_{n \geq 1} \{R_\beta^n < \infty\}\right) > 0. \quad (3.9)$$

It remains to show that (3.9) is still valid with \mathbf{P}_β in place of \mathbf{P}_α . Let h be the non negative function $h = \mathbf{1}_{\{\bigcap_{n \geq 1} \{R_\beta^n < \infty\}\}}$. Then $\mathbf{E}_\alpha(h) = \mathbf{P}_\alpha(R_\beta < \infty) \cdot \mathbf{E}_\beta(h)$ by Lemma 3.10. Since $\mathbf{E}_\alpha(h) > 0$ by Eq. (3.9), it follows that $\mathbf{E}_\beta(h) > 0$, showing that β is recurrent.

- (3) Observe that the ρ_n are related to R_α^n through the identity: $\omega_{R_\alpha^{n+1}} = \omega_{R_\alpha^n} \cdot \rho_n$. Now, let v_1, \dots, v_n be n finite trajectories in the range of R_α . Since the R_α^i are stopping times, we have the equality $\{\rho_1 = v_1, \dots, \rho_n = v_n\} = \{\omega_{R_\alpha^n} = v_1 \cdot \dots \cdot v_n\} = \uparrow(v_1 \cdot \dots \cdot v_n)$. The chain rule yields:

$$\mathbf{P}_\alpha(\rho_1 = v_1, \dots, \rho_n = v_n) = \mathbf{P}_\alpha(\uparrow v_1 \cdot \dots \cdot v_{n-1}) \times \mathbf{P}_\alpha(\uparrow v_1 \cdot \dots \cdot v_n | \uparrow v_1 \cdot \dots \cdot v_{n-1}). \quad (3.10)$$

The Markov property (1.5) combines with $\gamma(v_1 \cdot \dots \cdot v_{n-1}) = \alpha$ to rewrite the conditional probability in Eq. (3.10) as follows:

$$\mathbf{P}_\alpha(\uparrow v_1 \cdot \dots \cdot v_n | \uparrow v_1 \cdot \dots \cdot v_{n-1}) = \mathbf{P}_\alpha(\uparrow v_n).$$

Since v_n is in the range of R_α , and since R_α is a stopping time, $\uparrow v_n = \{R_\alpha = v_n\}$. We replace thus the conditional probability in Eq. (3.10) by $\mathbf{P}_\alpha(R_\alpha = v_n)$, and apply n times the same transformation to finally obtain the identity:

$$\mathbf{P}_\alpha(\rho_1 = v_1, \dots, \rho_n = v_n) = \mathbf{P}_\alpha(R_\alpha = v_1) \cdot \dots \cdot \mathbf{P}_\alpha(R_\alpha = v_n),$$

showing that the ρ_n are i.i.d. random variables, with the law of R_α .

(4) Consider the two measurable and non negative functions:

$$h_\alpha = \mathbf{1}_{\{\bigcap_{n \geq 1} \{R_\alpha^n < \infty\}\}}, \quad h_\beta = \mathbf{1}_{\{\bigcap_{n \geq 1} \{R_\beta^n < \infty\}\}}.$$

Since the successive returns to α are i.i.d. by virtue of point 3 above, each ρ_n one has positive \mathbf{P}_α -probability of hitting β , otherwise the \mathbf{P}_α -probability of ever hitting β would be zero, contradicting the assumption that β is reachable from α . Hence, by Borel-Cantelli Lemma, $\mathbf{E}_\alpha(h_\beta) > 0$. Since $\mathbf{E}_\alpha(h_\beta) = \mathbf{P}_\alpha(R_\beta < \infty) \cdot \mathbf{E}_\beta(h_\beta)$ by Lemma 3.10, this implies that $\mathbf{E}_\beta(h_\beta) > 0$ and thus β is recurrent.

To prove that α is reachable from β , we apply the Asynchronous Strong Markov property (Theorem 3.9) with stopping time \mathfrak{F}_{R_β} and function h_α . We then multiply the resulting identity by $\mathbf{1}_{\{R_\beta < \infty\}}$, and take into account that $\mathbf{1}_{\{R_\beta < \infty\}}$ is \mathfrak{F}_{R_β} -measurable on the one hand, and that $\gamma(\omega_{R_\beta}) = \beta$ on $\{R_\beta < \infty\}$ on the other hand to obtain:

$$\mathbf{E}_\alpha(\mathbf{1}_{\{R_\beta < \infty\}} h_\alpha \circ \theta_{R_\beta} | \mathfrak{F}_{R_\beta}) = \mathbf{1}_{\{R_\beta < \infty\}} \cdot \mathbf{E}_\beta(h_\alpha). \quad (3.11)$$

Observe that $h_\alpha \circ \theta_{R_\beta} = h_\alpha$ on $\{R_\beta < \infty\}$, therefore the following identity is valid everywhere: $\mathbf{1}_{\{R_\beta < \infty\}} h_\alpha \circ \theta_{R_\beta} = \mathbf{1}_{\{R_\beta < \infty\}} h_\alpha$. By assumption, α is recurrent, hence $h_\alpha = 1$ \mathbf{P}_α -almost surely, and finally $\mathbf{1}_{\{R_\beta < \infty\}} h_\alpha \circ \theta_{R_\beta} = \mathbf{1}_{\{R_\beta < \infty\}}$ \mathbf{P}_α -almost surely. Replacing thus $\mathbf{1}_{\{R_\beta < \infty\}} h_\alpha \circ \theta_{R_\beta}$ by $\mathbf{1}_{\{R_\beta < \infty\}}$ in Eq. (3.11), and taking the \mathbf{P}_α -expectations of both sides yields:

$$\mathbf{P}_\alpha(R_\beta < \infty) = \mathbf{P}_\alpha(R_\beta < \infty) \cdot \mathbf{E}_\beta(h_\alpha).$$

But β is assumed to be reachable from α , hence $\mathbf{P}_\alpha(R_\beta < \infty) > 0$, and thus $\mathbf{E}_\beta(h_\alpha) = 1$, implying in particular that α is reachable from β . \square

3.6. Irreducible Components. With the notion of recurrent state at hand, it is now possible to introduce the notions of irreducible process and irreducible components of a M2CP.

Definition 3.12. Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP. We say that \mathbb{P} is **irreducible** if every $\alpha \in X_0$ is reachable from every $\beta \in X_0$.

Proposition 3.13. *If a M2CP is irreducible, then every global state is recurrent.*

Proof. By Proposition 3.11, point 2, there is some recurrent state $\alpha \in X_0$. But then, since any $\beta \in X_0$ is reachable from α , β is recurrent by point 4 of the same proposition. \square

The result in Proposition 3.15 below says that the study of Markov two-components processes essentially reduces to the study of irreducible processes, especially if one is interested in asymptotic properties (so-called limit theorems from probability theory such as the Law of Large Numbers or the Central Limit Theorem). For this we use the notion of subprocess introduced in Definition 1.4, and introduce irreducible components for M2CPs.

Definition 3.14. An **irreducible component** of a probabilistic two-components process $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ is any subset $X_1 \subset X_0$ such that, for all $\alpha \in X_1$:

- (1) any $\beta \in X_1$ is reachable from α ; and
- (2) if $\beta \in X_0$ is reachable from α , then $\beta \in X_1$.

Point 2 in Definition 3.14 ensures that $(\mathbf{P}_\alpha)_{\alpha \in X_1}$ is indeed a probabilistic process (Definition 1.4). Therefore, if X_1 is an irreducible component of M2CP $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$, then the family $(\mathbf{P}_\alpha)_{\alpha \in X_1}$ forms a subprocess of \mathbb{P} , which is obviously an irreducible M2CP. It follows from Proposition 3.13 that any element α of an irreducible component is recurrent. Any two irreducible components are disjoint. Finally, if α is recurrent, then α belongs to a unique irreducible component, namely the set X_1 of those β which are reachable from α (the fact that X_1 is indeed an irreducible component follows from Proposition 3.11). Since recurrent states exist by Proposition 3.11, this implies that any M2CP has at least one irreducible component.

Proposition 3.15. *If $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ is a M2CP, there exists a stopping time $T : \Omega \rightarrow \mathcal{T}$ such that T is almost surely finite and $\gamma(\omega_T)$ belongs to some irreducible component of \mathbb{P} .*

Proof. We fix an initial state $\alpha \in X_0$. Let $(R^n)_{n \geq 1}$ denote the successive return times to the square set X_0 (cf. § 3.3.6). As already observed several times, $\mathbf{P}_\alpha(R^n < \infty) = 1$ for all $n \geq 1$, and therefore, if we put $B = \{\beta \in X_0 \mid \mathbf{P}_\alpha(\bigcap_{n \geq 1} R_\beta^n < \infty)\} > 0\}$, it follows from Lemma 3.8 that:

$$\Omega = \bigcup_{\beta \in B} \bigcap_{n \geq 1} \{R_\beta^n < \infty\}. \quad (3.12)$$

Pick exactly one global state α_i for each irreducible component. Let T_i be the first hitting time of α_i , and put:

$$\forall \omega \in \Omega, \quad \omega_T = \inf_i \{\omega_{T_i}\}.$$

For each $\beta \in B$, let $\tilde{\beta}$ be the unique recurrent state α_i of the same irreducible component. Then $\tilde{\beta}$ is reachable from β , and therefore:

$$\omega \in \bigcap_{n \geq 1} \{R_\beta^n < \infty\} \Rightarrow R_{\tilde{\beta}}(\omega) < \infty \quad \mathbf{P}_\alpha\text{-a.s.} \quad (3.13)$$

From Eqs. (3.12)(3.13) we deduce that $\omega_T < \infty$ \mathbf{P}_α -almost surely. Hence, on the one hand, at least one T_i is finite \mathbf{P}_α -almost surely. On the other hand, only one of them is finite, since the α_i have been chosen in different irreducible components. Therefore $\omega_T = \omega_{T_i}$ for some T_i , and thus $\gamma(\omega_T)$ does belong to some irreducible component, as claimed. \square

3.7. Open and Closed Markov Two-Components Processes. Besides the classical application of the Strong Markov Property to recurrence and transience, it also applies to the notion of open and closed processes which is specific to the two-components framework. Open and closed processes have been defined in Definition 1.5.

Proposition 3.16. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP.*

- (1) *Let $\alpha \in X_0$ be a recurrent state. Then a global trajectory ω synchronizes infinitely often with \mathbf{P}_α -probability 1 if at least some synchronization state is reachable from α , and with \mathbf{P}_α -probability 0 otherwise.*
- (2) *If \mathbb{P} is irreducible, then \mathbb{P} is closed or open.*

Proof.

- (1) Let R_Q be first return time to the square set $X_0 \cap (Q \times Q)$ (see § 3.3.6), and consider the stopping time $U = R_Q + R_\alpha \circ \theta_{R_Q}$, corresponding to reaching α after having reached Q . This is indeed a stopping time by virtue of Lemma 3.4. Let $(U_n)_{n \geq 1}$ be the iterated stopping times associated with U as in Definition 3.5. If $h_n = \mathbf{1}_{\{U_n < \infty\}}$, the same technique involving the Asynchronous Strong Markov property (Theorem 3.9) than in the proof of Proposition 3.11, point 1, shows that: $\mathbf{E}_\alpha(h_n | \mathfrak{F}_{U_{n-1}}) = h_{n-1} \mathbf{E}_\alpha(\mathbf{1}_{\{U < \infty\}})$. Therefore, if $a = \mathbf{P}_\alpha(U < \infty)$ one has $\mathbf{P}_\alpha(U_n < \infty) = a^n$. Since α is recurrent, Proposition 3.11, point 1 implies that R_α is \mathbf{P}_α -almost surely finite, hence $\mathbf{P}_\alpha(U < \infty) = \mathbf{P}_\alpha(R_Q < \infty)$. Since a trajectory ω synchronizes infinitely often if and only if $U_n < \infty$ for all $n \geq 1$, Borel-Cantelli Lemma implies that ω has \mathbf{P}_α -probability 1 of synchronizing infinitely often if $\mathbf{P}_\alpha(R_Q < \infty) = 1$, and 0 otherwise.

It remains to show that $\mathbf{P}_\alpha(R_Q < \infty) = 1$ if and only if some state of the form (x, x) with $x \in Q$ is reachable from α . Since $R_Q = \bigwedge_{x \in Q} R_{(x,x)}$, obviously if no (x, x) is reachable from α then $\mathbf{P}_\alpha(R_Q < \infty) = 0$. Conversely, assume that some (x, x) with $x \in Q$ is reachable from α . Then (x, x) is recurrent, by point 4 of Proposition 3.11, and Lemma 3.10 implies that $R_{(x,x)} < \infty$ \mathbf{P}_α -almost surely. But $R_Q \leq R_{(x,x)}$, hence $\mathbf{P}_\alpha(R_Q < \infty) = 1$, as claimed.

- (2) If \mathbb{P} is irreducible, then by Proposition 3.11, every $\alpha \in X_0$ is recurrent, therefore point 1 above applies to any $\alpha \in X_0$. Assume that the \mathbf{P}_α -probability of synchronizing infinitely often is 0 for some $\alpha \in X_0$, and let $\beta \in X_0$. Consider a finite trajectory v such that $\mathbf{P}_\alpha(v) > 0$ and $\gamma(v) = \beta$; such a v exists since any β is reachable from α . Then \mathbf{P}_α -a.s. every trajectory $\omega \in \uparrow v$ has no synchronization. But the \mathbf{P}_α probability measure on $\uparrow v$ coincides, up to the factor $\mathbf{P}_\alpha(v) \neq 0$, with \mathbf{P}_β on Ω . Hence \mathbf{P}_β -a.s. every $\omega \in \Omega$ has no synchronization, and since this is true for every $\beta \in X_0$, the process \mathbb{P} is open. The same method applies to show that \mathbb{P} is closed if the probability of synchronizing infinitely often is 1 for some $\alpha \in X_0$. This concludes the proof. \square

4. THE LOCAL INDEPENDENCE PROPERTY

Having adapted Markovian concepts from Markov chain theory, we now focus on a topic specific to the asynchronous framework, without equivalent in Markov chain theory: the probabilistic correlation between private behaviors of local components. It is desirable to have a kind of probabilistic independence between private parts of trajectories: otherwise, hidden synchronization constraints would be encoded in the probabilistic structure, while we expect synchronization to occur only on explicit synchronization states. Probabilistic independence of random variables ω^1 and ω^2 however is too much to ask; their synchronization is an obstacle to their mere probabilistic independence. This is easy to understand from an information theoretic viewpoint: the knowledge of ω^1 gives indeed information on ω^2 , since it precisely determines the Q -sequence of ω^2 . The weaker notion of conditional independence proves to be adapted to our purpose. The Local Independence Property that we introduce informally states that the two local components have the maximal probabilistic independence they can have, considering their natural synchronization constraints.

Recall that $Y = (Y_n)_{n \geq 0}$ has been defined in Definition 1.5 as the Q -sequence induced by some trajectory $\omega \in \Omega$, to which we have added $Y_{-1} = *$ and $Y_n = *$ for large n if the Q -sequence is finite, for some fixed specified value $*$. We proceed in a similar way to define

the sequence $(\sigma_n)_{n \geq 0}$ of random elementary trajectories, referring to the decomposition of a trajectory ω as a concatenation of elementary trajectories from Proposition 1.2. If σ_n is defined only until some integer N (that is, in case 2b of Proposition 1.2) we define σ_{N+1} as the synchronization free trajectory such that $\omega = \sigma_1 \cdot \dots \cdot \sigma_N \cdot \sigma_{N+1}$ and $\sigma_n = *$ for $n > N+1$.

Then we observe the following property:

Proposition 4.1. *Let \mathbb{P} be the synchronization product of two Markov chains. Decomposing σ_n as $\sigma_n = (\sigma_n^1, \sigma_n^2)$ we have: for all $\alpha \in X_0$ and for every integer $n \geq 0$, σ_n^1 and σ_n^2 are two random variables independent conditionally on the pair (Y_{n-1}, Y_n) with respect to \mathbf{P}_α .*

Proof. Since \mathbb{P} satisfies the Markov property, the statement is equivalent to the independence of σ_n^1 and σ_n^2 , conditionally on Y_n , and with respect to $\mathbf{P}_{Y_{n-1}}$. But this follows from the construction of the law of $\sigma_n = (\sigma_n^1, \sigma_n^2)$ given in § 2. \square

In order to generalize the above property to processes which may not be closed, and at the cost of a little more abstraction, we introduce the following definition.

Definition 4.2. Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP, let Y be the associated random synchronization sequence. Let ω^1 and ω^2 denote the local components of global trajectories, so that $\omega = (\omega^1, \omega^2)$ for $\omega \in \Omega$. We say that \mathbb{P} has the **local independence property** (abbreviated **LIP**) if ω^1 and ω^2 are independent conditionally¹ to Y with respect to \mathbf{P}_α , for all $\alpha \in X_0$.

The following theorem relates this definition with the previous property stated in Proposition 4.1 for the synchronization of Markov chains.

Theorem 4.3. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP. Then \mathbb{P} satisfies the LIP if and only if the random variables σ_n^1 and σ_n^2 are independent conditionally on the pair (Y_{n-1}, Y_n) , with respect to \mathbf{P}_α for all $n \geq 0$ and for all $\alpha \in X_0$.*

Proof. Let (a) be the property that ω^1 and ω^2 are independent conditionally on Y , and let (b) be the property stated in the theorem.

Proof of (a) \Rightarrow (b). Thanks to the Markov property, it is enough to consider $n = 1$. We denote σ_1^1 and σ_1^2 by σ^1 and σ^2 , and we put: $Z^1 = \mathbf{E}_\alpha(\mathbf{1}_{\{\sigma^1=z^1\}} | Y)$, $Z^2 = \mathbf{E}_\alpha(\mathbf{1}_{\{\sigma^2=z^2\}} | Y)$, and $Z = \mathbf{E}_\alpha(\mathbf{1}_{\{\sigma^1=z^1, \sigma^2=z^2\}} | Y)$. These three random variables are constant on $\{Y_1 = b\}$ and, by (a), satisfy $Z = Z^1 \cdot Z^2$, whence:

$$\begin{aligned} \mathbf{P}_\alpha(\sigma^1 = z^1, \sigma^2 = z^2 | Y_1 = b) &= Z|_{\{Y_1=b\}} \\ &= Z^1|_{\{Y_1=b\}} \cdot Z^2|_{\{Y_1=b\}} \\ &= \mathbf{P}_\alpha(\sigma^1 = z^1 | Y_1 = b) \times \mathbf{P}_\alpha(\sigma^2 = z^2 | Y_1 = b), \end{aligned}$$

as expected.

¹Recall that two random variables X_1 and X_2 are independent w.r.t. a σ -algebra \mathfrak{G} if $\mathbf{E}(\varphi_1 \cdot \varphi_2 | \mathfrak{G}) = \mathbf{E}(\varphi_1 | \mathfrak{G}) \cdot \mathbf{E}(\varphi_2 | \mathfrak{G})$, for all non negative and bounded variables φ_1 and φ_2 , measurable with respect to X_1 and to X_2 respectively. See e.g. [14, Chapter IV]. Here, the independence conditionally to Y means the independence w.r.t. the σ -algebra $\langle Y \rangle$ generated by Y .

Proof of (b) \Rightarrow (a). From (b) used in conjunction with the Markov property and the chain rule, we get for integers $m \geq n$ and with short notations:

$$\mathbf{P}_\alpha(\sigma_1^1, \dots, \sigma_m^1, \sigma_1^2, \dots, \sigma_m^2 | Y_1, \dots, Y_n) = \mathbf{P}_\alpha(\sigma_1^1, \dots, \sigma_m^1 | Y_1, \dots, Y_n) \times \mathbf{P}_\alpha(\sigma_1^2, \dots, \sigma_m^2 | Y_1, \dots, Y_n).$$

The σ -algebra generated by the random trajectories $(\sigma_k^i, k \geq 1)$ for $i = 1, 2$ coincides with the σ -algebra generated by ω^i , since ω^i is obtained as the concatenation of these—the concatenation being finite or infinite. Hence, for any bounded non negative and measurable functions h^1 and h^2 :

$$\mathbf{E}_\alpha(h^1(\omega^1) \cdot h^2(\omega^2) | Y_1, \dots, Y_n) = \mathbf{E}_\alpha(h^1(\omega^1) | Y_1, \dots, Y_n) \cdot \mathbf{E}_\alpha(h^2(\omega^2) | Y_1, \dots, Y_n).$$

The sequence of σ -algebras $\langle Y_1, \dots, Y_n \rangle$ is increasing, and converges to $\langle Y \rangle$. Therefore by the special case [7, Theorem 35.6 p.470] of the Martingale convergence theorem, we get by taking the limit $n \rightarrow \infty$:

$$\mathbf{E}_\alpha(h^1(\omega^1) \cdot h^2(\omega^2) | Y) = \mathbf{E}_\alpha(h^1(\omega^1) | Y) \cdot \mathbf{E}_\alpha(h^2(\omega^2) | Y),$$

completing the proof. \square

Corollary 4.4. *The synchronization product of Markov chains satisfies the LIP.*

Having the specified value $*$ assigned to some Y_k and σ_k described above has the following effect with regard to Theorem 4.3: the statement is trivial if both σ_k, Y_{k-1} and Y_k assume their constant values $*$; but it implies the probabilistic independence of σ_{N+1}^1 and σ_{N+1}^2 with respect to \mathbf{P}_{Y_N} , where N is the last synchronization index. In other words, the local trajectories are independent after their last synchronization.

It is useful to examine a degenerated case of Definition 4.2, where the conditional independence reduces to probabilistic independence.

Lemma 4.5. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP. Let ω^1 and ω^2 denote the local components of global trajectories. Assume that, with respect to \mathbf{P}_α for some state $\alpha \in X_0$, the two components ω^1 and ω^2 are independent. Then ω^1 and ω^2 are the sample paths of two independent Markov chains.*

Proof. Fix $\alpha \in X_0$, and for each $i = 1, 2$ let P^i denote the law of ω^i , characterized by $P^i(\omega^i \geq s^i) = \mathbf{P}_\alpha(\omega^i \geq s^i)$, with s^i ranging over the finite local trajectories on site i . We show that the conditional law $P^i(s^i \cdot \bullet | \uparrow s^i)$ only depends on the last state of s^i , which is enough to obtain that ω^i follows the law of a homogeneous Markov chain. Consider $i = 1$, the case $i = 2$ is identical. Consider s^1 a finite sequence in S^1 such that $P^1(\omega^1 \geq s^1) > 0$. It implies that there exists some sequence in S^2 , say s^2 , such that $\mathbf{P}_\alpha(\uparrow(s^1, s^2)) > 0$. Put $s = (s^1, s^2)$ and let $(x^1, x^2) = \gamma(s)$. For any finite sequence σ in S^1 , we have:

$$\begin{aligned} P^1(\omega^1 \geq s^1 \cdot \sigma | \omega^1 \geq s^1) &= \frac{\mathbf{P}_\alpha(\omega^1 \geq s^1 \cdot \sigma)}{\mathbf{P}_\alpha(\omega^1 \geq s^1)} \\ &= \frac{\mathbf{P}_\alpha(\omega^1 \geq s^1 \cdot \sigma, \omega^2 \geq s^2)}{\mathbf{P}_\alpha(\omega^1 \geq s^1, \omega^2 \geq s^2)} && \text{by independence} \\ &= \mathbf{P}_{(x^1, x^2)}(\omega^1 \geq \sigma). \end{aligned} \tag{4.1}$$

Obviously, the expression $P^1(\omega^1 \geq s^1 \cdot \sigma | \omega^1 \geq s^1)$ does not depend on x^2 , since x^2 is the last state of the arbitrary chosen sequence s^2 . Therefore, the right member of (4.1) does not depend on x^2 neither, hence it only depends on x^1 and σ , which was to be proved. \square

Proposition 4.6. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP with the LIP. Let ω^1 and ω^2 denote the local components of global trajectories.*

- (1) *If $Q = \emptyset$, then ω^1 and ω^2 are two independent Markov chains, with respect to \mathbf{P}_α for any $\alpha \in X_0$.*
- (2) *If Q is a singleton, and if α is a recurrent state, then ω^1 and ω^2 are two independent Markov chains, with respect to \mathbf{P}_α .*

Proof.

- (1) Since $Q = \emptyset$, the synchronization sequence Y is constant, $Y = (*, *, *, \dots)$. The conditional independence in the definition of the LIP reduces to probabilistic independence. The result follows then by Lemma 4.5.
- (2) Let $Q = \{\beta\}$. By Proposition 3.16, point 1, ω synchronizes infinitely often with \mathbf{P}_α -probability either 0 or 1. If it is with probability 0, then $Y = (*, *, *, \dots)$ \mathbf{P}_α -a.s., and the same method than in point 1 above applies. If it is with probability 1, then Y is still constant, now $Y = (*, \beta, \beta, \beta, \dots)$. The same method applies again. \square

Corollary 4.7. *An open M2CP with the LIP identifies with two independent homogeneous Markov chains.*

5. CHARACTERIZATION OF MARKOV TWO-COMPONENTS PROCESSES WITH THE LIP

The topic of this section is to characterize a M2CP with the LIP by means of a finite family of real numbers, very much as the transition matrix of a Markov chain does. It turns out that the law of a M2CP with the LIP is entirely specified by a finite family of transition matrices. We will also investigate, conversely, if such a family of transition matrices always induces a M2CP with the LIP, providing a more general way of constructing M2CPs than the synchronization product of Markov chains. We show through a numerical example at the end of the section that not any M2CP can be obtained as the synchronization product of two Markov chains.

5.1. Technical Preliminaries. We begin with two lemmas.

Lemma 5.1. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a closed M2CP, and let Y denote the associated synchronization sequence. Then for any $\alpha \in X_0$, Y is a homogeneous Markov chain with respect to \mathbf{P}_α .*

Proof. The formulation of Definition 1.7 applies to Y as follows: for any two finite sequences s and u in Q , the conditional probability $\mathbf{P}_\alpha(Y \geq s \cdot u | Y \geq s)$ only depends on u and on the last state of s . This shows that Y is a homogeneous Markov chain. \square

Lemma 5.2. *Let $\mathbb{P} = (\mathbf{P}_\alpha)_{\alpha \in X_0}$ be a M2CP with the LIP, let Y denote the associated synchronization sequence, and let $(\sigma_n)_{n \geq 0}$ denote the sequence of elementary trajectories that decompose global trajectories (see § 4).*

Then for every $n \geq 0$ and for $i = 1, 2$, the sequence of states that appear in σ_n^i is a stopped Markov chain with respect to the conditional probability $\mathbf{P}_\alpha(\cdot | Y_{n-1}, Y_n)$.

Proof. By the Markov property, there is no loss of generality in assuming that $n = 0$. Using the notation $\sigma^i = \sigma_0^i$ for short, we thus have to prove that σ^i is a stopped Markov chain with respect to $\mathbf{P}_\alpha(\cdot | Y_0 = y)$, for any value $y \in Q$. We consider $i = 1$ only, the case $i = 2$ is similar. Let (X_1, \dots, X_τ) be the sequence of states in S^1 , and let \mathbf{Q} denote the conditional probability $\mathbf{Q} = \mathbf{P}_\alpha(\cdot | Y_0 = y)$. Let x_1, \dots, x_n be values in $S^1 \setminus Q$, let $x_{n+1} \in Q \cup \{y\}$, and put $\delta = \mathbf{Q}(X_{n+1} = x_{n+1} | X_1 = x_1, \dots, X_n = x_n)$. We claim that δ only depends on x_n and x_{n+1} . Put $\alpha = (x, z)$. We calculate:

$$\begin{aligned} \delta &= \frac{\mathbf{Q}(X_1 = x_1, \dots, X_{n+1} = x_{n+1})}{\mathbf{Q}(X_1 = x_1, \dots, X_n = x_n)} \\ &= \frac{\mathbf{P}_\alpha(X_1 = x_1, \dots, X_{n+1} = x_{n+1}, X_\tau = y)}{\mathbf{P}_\alpha(X_1 = x_1, \dots, X_n = x_n, X_\tau = y)}. \end{aligned}$$

We can rephrase $\{X_1 = x_1, \dots, X_{n+1} = x_{n+1}\}$ in the two-components framework as $\{\omega^1 \geq (x_1 \cdot \dots \cdot x_{n+1})\} = \uparrow(x_1 \cdot \dots \cdot x_{n+1}, \epsilon)$, observing that $(x_1 \cdot \dots \cdot x_{n+1}, \epsilon)$ is indeed a trajectory. The same applies to $\{X_1 = x_1, \dots, X_n = x_n\} = \uparrow(x_1 \cdot \dots \cdot x_n, \epsilon)$. Therefore the calculation continues as follows:

$$\begin{aligned} \delta &= \frac{\mathbf{P}_\alpha(\omega^1 \geq (x_1 \cdot \dots \cdot x_{n+1}), X_\tau = y)}{\mathbf{P}_\alpha(\uparrow(x_1 \cdot \dots \cdot x_n, \epsilon), X_\tau = y)} \\ &= \frac{\mathbf{P}_\alpha(\omega^1 \geq (x_1 \cdot \dots \cdot x_{n+1}), X_\tau = y \mid \uparrow(x_1 \cdot \dots \cdot x_n, \epsilon))}{\mathbf{P}_\alpha(\uparrow(x_1 \cdot \dots \cdot x_n, \epsilon), X_\tau = y \mid \uparrow(x_1 \cdot \dots \cdot x_n, \epsilon))} \\ &= \frac{\mathbf{P}_{(x_n, z)}(\omega^1 \geq x_{n+1}, X_\tau = y)}{\mathbf{P}_{(x_n, z)}(X_\tau = y)} = \mathbf{P}_{(x_n, z)}(\omega^1 \geq x_{n+1} | X_\tau = y). \end{aligned}$$

On the last expression, it is clear that δ only depends on x_n and x_{n+1} , and not on x_1, \dots, x_n , showing our claim. This is enough to imply that X_1, \dots, X_τ are the terms of a homogeneous Markov chain. \square

5.2. Adapted Family of Transition Matrices. The two above lemmas suggest the following construction for M2CP with the LIP. First consider a Markov chain Y on the set of shared states; then for any two consecutive values y_{n-1} and y_n of Y , consider two independent stopped Markov chains σ_n^1 and σ_n^2 , with σ_n^i taking values in $\{y_n\} \cup (S^i \setminus Q)$, that reaches y_n with probability one and which is stopped at the first hitting time of y_n . This description is formalized in Theorem 5.4 below. It is first convenient to introduce the following definition.

Definition 5.3. An **adapted family of transition matrices** is given by two families $(R_y^i)_{y \in Q_0}$, one for each $i = 1, 2$ and with Q_0 some subset of Q , such that:

- (1) For each $y \in Q_0$ and $i = 1, 2$, R_y^i is a stochastic matrix on $\{y\} \cup (S^i \setminus Q)$;
- (2) With respect to the transition matrix R_y^i , the state y is reachable from any state in $S^i \setminus Q$.

Using this definition, the existence and uniqueness result concerning M2CP with the LIP states as follows. We focus on closed processes only, as suggested by Proposition 3.15, Proposition 3.16 and Corollary 4.7.

Theorem 5.4. *Any closed M2CP \mathbb{P} with the LIP induces the following elements, that entirely characterize \mathbb{P} :*

- (1) *A transition matrix R on the set Q of shared states, defined as the transition matrix of the synchronization sequence Y from Definition 1.5;*
- (2) *An adapted family of transition matrices $(R_y^i)_{y \in Q_0}$, for $i = 1, 2$, where Q_0 is the essential set of values of Y . For each $i = 1, 2$, and for $y \in Q_0$, R_y^i is the transition matrix of the Markov chain σ_n^i with respect to the conditional probability $\mathbf{P}_\alpha(\cdot | Y_{n-1}, Y_n = y)$, which is independent of the integer n and of $\alpha \in X_0$, provided it is defined for these values.*

Conversely, given a set of global states

$$X_0 \subset \{(x, z) \in S^1 \times S^2 \mid (x \in Q) \wedge (z \in Q) \Rightarrow x = z\},$$

such that the set

$$Q_0 = \{y \in Q \mid (y, y) \in X_0\}.$$

is nonempty; and considering:

- (1) *a transition matrix R on the set Q_0 ; and*
- (2) *an adapted family of transition matrices $(R_y^i)_{y \in Q_0}$,*

then there exists a unique M2CP with the LIP, defined on X_0 and inducing R and $(R_y^i)_{y \in Q_0}$. This M2CP is closed.

Proof. The first part of the theorem follows from Lemmas 5.1 and 5.2. For the second part, assume that the considered data are given. The construction of the process \mathbb{P} is essentially the same as the construction of the synchronization product of Markov chains, therefore we omit the routine arguments showing the existence and uniqueness of \mathbb{P} . What we need to show is that the two-components process obtained is indeed a M2CP with the LIP. The LIP is obvious from the construction of \mathbb{P} combined with Theorem 4.3, hence we focus on the Markov property. Since the process is closed by construction, we rely on Lemma 1.8 for this. Hence, let $\alpha \in X_0$, let t be any elementary trajectory and let s be any finite trajectory. The proof then follows the same steps than the proof of Theorem 2.2:

- (1) *Step 1: s synchronization free.* Then $s \cdot t$ is an elementary trajectory. Put $\alpha = (x_0, z_0)$, $\gamma(s) = (x_1, z_1)$ and $\gamma(t) = (y, y)$. We have: $\uparrow(s \cdot t) = \{\sigma_1^1 = s^1 \cdot t^1, \sigma_1^2 = s^2 \cdot t^2\}$. Let \mathbf{Q}_b^i denote the probability associated with the Markov chain starting from b and with transition matrix R_y^i , for $i = 1, 2$ and $b \in Q$. We compute using the independence conditionally on Y_1 :

$$\begin{aligned} (\mathbf{P}_\alpha)_s(\uparrow t) &= (\mathbf{P}_\alpha)_s(\uparrow t \wedge Y_1 = y) \\ &= \frac{\mathbf{P}_\alpha(\uparrow(s \cdot t) | Y_1 = y)}{\mathbf{P}_\alpha(\uparrow s | Y_1 = y)} \\ &= \frac{\mathbf{Q}_{x_0}^1(\uparrow(s^1 \cdot t^1))}{\mathbf{Q}_{x_0}^1(\uparrow s^1)} \cdot \frac{\mathbf{Q}_{z_0}^2(\uparrow(s^2 \cdot t^2))}{\mathbf{Q}_{z_0}^2(\uparrow s^2)} \\ &= \mathbf{Q}_{x_1}^1(\uparrow t^1) \cdot \mathbf{Q}_{z_1}^2(\uparrow t^2). \end{aligned}$$

The last quantity only depends on $(x_1, z_1) = \gamma(s)$ and t . In particular, as expected, we have $(\mathbf{P}_\alpha)_s(\uparrow t) = \mathbf{P}_{\gamma(s)}(\uparrow t)$.

- (2) *Step 2: s is any finite trajectory.* Using Step 1, as in the proof of Theorem 2.2. □

5.3. A Numerical Example. In this subsection, we show on an example how the synchronization product of Markov chains is to be interpreted in terms of an adapted family of transition matrices. We show that not any M2CP can be obtained from the synchronization of two Markov chains.

Let $S^1 = \{a, b, \mathbf{c}, \mathbf{d}\}$ and $S^2 = \{\mathbf{c}, \mathbf{d}, e, f\}$, and let two transition matrices M^1 and M^2 on S^1 and S^2 respectively. Take for instance:

$$M^1 = \begin{matrix} a \\ b \\ \mathbf{c} \\ \mathbf{d} \end{matrix} \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ \frac{1}{2} & \frac{1}{8} & \frac{1}{8} & \frac{1}{4} \\ \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} \\ 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{pmatrix} \quad M^2 = M^1 = \begin{matrix} e \\ f \\ \mathbf{c} \\ \mathbf{d} \end{matrix} \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ \frac{1}{2} & \frac{1}{8} & \frac{1}{8} & \frac{1}{4} \\ \frac{1}{2} & 0 & \frac{1}{4} & \frac{1}{4} \\ 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}.$$

The matrices contain 0 in some places, but that will not harm.

Computation of the adapted family of transition matrices. We need to compute the matrices $R_{\mathbf{c}}^1 = R_{\mathbf{c}}^2$ and $R_{\mathbf{d}}^1 = R_{\mathbf{d}}^2$. Matrix $R_{\mathbf{c}}^1$ is a stochastic matrix on $\{a, b, \mathbf{c}\}$, and drives the subsystem on site 1, conditionally on “next synchronization is \mathbf{c} ”. Referring to the construction detailed in § 2, $R_{\mathbf{c}}^1$ is simply obtained as follows: starting from matrix M^1 , suppress all lines and columns attached to states in Q different from \mathbf{c} , here, this is only state \mathbf{d} . Finally, renormalize each line to obtain a stochastic matrix. The same process is applied to obtain $R_{\mathbf{d}}^1$:

$$R_{\mathbf{c}}^1 = \begin{matrix} a \\ b \\ \mathbf{c} \end{matrix} \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{2}{3} & \frac{1}{6} & \frac{1}{6} \\ \frac{2}{3} & 0 & \frac{1}{3} \end{pmatrix} \quad R_{\mathbf{d}}^1 = \begin{matrix} a \\ b \\ \mathbf{d} \end{matrix} \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{4}{7} & \frac{1}{7} & \frac{2}{7} \\ 0 & \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

This construction implies that the lines obtained from matrices $R_{\mathbf{c}}^1$ and $R_{\mathbf{d}}^1$ by deleting the lines and columns relative to shared states are *proportional*: $(\frac{1}{3} \ \frac{1}{3})$ is proportional to $(\frac{1}{2} \ \frac{1}{2})$, and $(\frac{2}{3} \ \frac{1}{6})$ is proportional to $(\frac{4}{7} \ \frac{1}{7})$. Indeed, the lines of $R_{\mathbf{c}}^1$ and $R_{\mathbf{d}}^1$ are obtained by renormalization after extraction from the same transition matrix M^1 . We deduce from this observation a way to construct a M2CP with the LIP not obtained as a synchronization product of Markov chains. Replace for example the b line of $R_{\mathbf{c}}^1$ by $(0 \ 0 \ 1)$ and leave $R_{\mathbf{d}}^1$ unchanged. This corresponds to some closed M2CP with LIP according to Theorem 5.4, which cannot be a synchronization product of Markov chains.

We have obtained: *not every M2CP with the LIP can be obtained as the synchronization product of two Markov chains.*

Computation of the matrix of the synchronization chain. It remains to compute the transition matrix of the chain $Y = (Y_n)_{n \geq 1}$, which involves the law of $X_{\tau_1}^1$ and $X_{\tau_2}^2$, where τ^i are the first hitting times to Q of chains X^1 and X^2 respectively, which we do here “by hand”. For a general theory, see for instance [9, Ch. XII §§58–59 *Entrance and exit laws*, p.262ff].

Denoting by M_x^1 the law of chain X^1 starting from x , one has: $M_x^1(X_{\tau_1}^1 = \mathbf{c}) = \sum_w M_x^1(w)$, where w ranges over words of the form $w = v \cdot \mathbf{c}$, and v is any word on $\{a, b\}$. Therefore, if $q_k(x)$ denotes, for any integer $k \geq 0$:

$$q_k(x) = \sum_{l_1, \dots, l_k \in \{a, b\}} M_x^1(l_1 \cdot \dots \cdot l_k \cdot \mathbf{c}),$$

one has $M_x^1(X_{\tau_1}^1 = \mathbf{c}) = \sum_{k \geq 0} q_k(x)$. Decomposing over the two possible values of l_1 yields:

$$q_k(x) = M^1(x, a)q_{k-1}(a) + M^1(x, b)q_{k-1}(b).$$

Therefore the vector $(q_k(a) \quad q_k(b))$ satisfies the following recurrence relation:

$$\begin{pmatrix} q_k(a) \\ q_k(b) \end{pmatrix} = N \begin{pmatrix} q_{k-1}(a) \\ q_{k-1}(b) \end{pmatrix}, \quad \text{with } N = \begin{pmatrix} M^1(a, a) & M^1(a, b) \\ M^1(b, a) & M^1(b, b) \end{pmatrix}.$$

We observe that $\begin{pmatrix} q_0(a) \\ q_0(b) \end{pmatrix} = \begin{pmatrix} M^1(a, \mathbf{c}) \\ M^1(b, \mathbf{c}) \end{pmatrix}$ and therefore:

$$\begin{pmatrix} M_a^1(X_{\tau_1}^1 = \mathbf{c}) \\ M_b^1(X_{\tau_1}^1 = \mathbf{c}) \end{pmatrix} = (I - N)^{-1} \begin{pmatrix} M^1(a, \mathbf{c}) \\ M^1(b, \mathbf{c}) \end{pmatrix}.$$

We find in a similar fashion:

$$\begin{pmatrix} M_a^1(X_{\tau_1}^1 = \mathbf{d}) \\ M_b^1(X_{\tau_1}^1 = \mathbf{d}) \end{pmatrix} = (I - N)^{-1} \begin{pmatrix} M^1(a, \mathbf{d}) \\ M^1(b, \mathbf{d}) \end{pmatrix},$$

with same matrix N . Finally we have:

$$\begin{aligned} M_{\mathbf{c}}^1(X_{\tau_1}^1 = \mathbf{c}) &= M^1(\mathbf{c}, \mathbf{c}) + M^1(\mathbf{c}, a)M_a^1(X_{\tau_1}^1 = \mathbf{c}) + M^1(\mathbf{c}, b)M_b^1(X_{\tau_1}^1 = \mathbf{c}) \\ M_{\mathbf{c}}^1(X_{\tau_1}^1 = \mathbf{d}) &= M^1(\mathbf{c}, \mathbf{d}) + M^1(\mathbf{c}, a)M_a^1(X_{\tau_1}^1 = \mathbf{d}) + M^1(\mathbf{c}, b)M_b^1(X_{\tau_1}^1 = \mathbf{d}). \end{aligned} \quad (5.1)$$

And in a similar fashion:

$$\begin{aligned} M_{\mathbf{d}}^1(X_{\tau_1}^1 = \mathbf{c}) &= M^1(\mathbf{d}, \mathbf{c}) + M^1(\mathbf{d}, a)M_a^1(X_{\tau_1}^1 = \mathbf{c}) + M^1(\mathbf{d}, b)M_b^1(X_{\tau_1}^1 = \mathbf{c}) \\ M_{\mathbf{d}}^1(X_{\tau_1}^1 = \mathbf{d}) &= M^1(\mathbf{d}, \mathbf{d}) + M^1(\mathbf{d}, a)M_a^1(X_{\tau_1}^1 = \mathbf{d}) + M^1(\mathbf{d}, b)M_b^1(X_{\tau_1}^1 = \mathbf{d}). \end{aligned} \quad (5.2)$$

Applying these calculations to our numerical example, we find:

$$\begin{aligned} N &= \begin{pmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{8} \end{pmatrix} & (I - N)^{-1} &= \frac{12}{5} \begin{pmatrix} \frac{7}{8} & \frac{1}{3} \\ \frac{1}{2} & \frac{2}{3} \end{pmatrix} \\ \begin{pmatrix} M_a^1(X_{\tau_1}^1 = \mathbf{c}) \\ M_b^1(X_{\tau_1}^1 = \mathbf{c}) \end{pmatrix} &= \frac{12}{5} \begin{pmatrix} \frac{1}{3} \\ \frac{1}{4} \end{pmatrix} & \begin{pmatrix} M_a^1(X_{\tau_1}^1 = \mathbf{d}) \\ M_b^1(X_{\tau_1}^1 = \mathbf{d}) \end{pmatrix} &= \frac{1}{5} \begin{pmatrix} 1 \\ 2 \end{pmatrix}. \end{aligned}$$

We obtain thus, using Eqs. (5.1)(5.2):

$$\begin{aligned} M_{\mathbf{c}}^1(X_{\tau_1}^1 = \mathbf{c}) &= \frac{13}{20} & M_{\mathbf{c}}^1(X_{\tau_1}^1 = \mathbf{d}) &= \frac{7}{20} \\ M_{\mathbf{d}}^1(X_{\tau_1}^1 = \mathbf{c}) &= \frac{11}{20} & M_{\mathbf{d}}^1(X_{\tau_1}^1 = \mathbf{d}) &= \frac{9}{20} \end{aligned}$$

Since we have taken $M^2 = M^1$, we obtain the same laws depending on the initial state \mathbf{c} or \mathbf{d} for $X_{\tau_2}^2$. The 2×2 transition matrix of Y is now obtained by conditioning the free product $(X_{\tau_1}^1, X_{\tau_2}^2)$ on $X_{\tau_1}^1 = X_{\tau_2}^2$, which yields the following transition matrix:

$$\begin{matrix} \mathbf{c} \\ \mathbf{d} \end{matrix} \begin{pmatrix} \frac{169}{218} & \frac{49}{218} \\ \frac{121}{202} & \frac{81}{202} \end{pmatrix}.$$

CONCLUSION

Summary of results. Following the idea that, in a network, the knowledge a node has about time is related to its local clock, and to its local clock only, we have introduced a probabilistic model based on a simple trace model, that allows private changes of states and synchronizations between two sites. We have focused on a Markov model where local components are independent up to the synchronization constraints, which brought us to the formulation of a Markov property without reference to any time index on the one hand, and to the Local Independence Property on the other hand. Triples $(\Omega, \mathfrak{F}, \mathbf{P})$ where (Ω, \mathfrak{F}) is the space of trajectories and \mathbf{P} is a probability measure satisfying both properties have been constructed and entirely characterized by a finite family of transition matrices, extending the familiar transition matrix from discrete time Markov chain theory.

A singular feature of the model is the absence of *constant times*; instead, only random times may be considered, and among them stopping times play a distinguished role. Note that despite the absence of a totally ordered time index, we can conduct probabilistic reasoning about our two-components models at the level of stopping times.

Potential applications. Open research fields involving asynchronous systems are numerous. In some cases, trace models have proved to be more relevant than interleaving models: distributed observation, supervision and diagnosis of concurrent systems, distributed optimization and planning [6] provide examples. In the formal verification community, people have considered interleaving models for composing probabilistic systems (cf. the discussion in the Introduction). Although product of Probabilistic Automata for instance has shown to be efficient for developing proving techniques based on bisimulation relations, it is worth trying other ways for modeling network system where asynchrony plays an important role.

One can therefore expect new advances in the theory of networked systems through the development of a probabilistic layer for trace models. In this respect, asymptotic analysis of probabilistic trace models may have applications in network dimensioning.

Limitations and extensions. Although the model of Markov concurrent process adopted in this paper is limited to *two* components only, it is important to notice that it has a straightforward generalization to an arbitrary number $n \geq 2$ of components. In this generalized framework, the notion of stopping time, the Asynchronous Strong Markov Property and all the results developed in § 3 carry over without additional difficulty. The LIP may also be expressed for $n \geq 2$ components in a similar way than we did for two components only. However, the mere existence of Markov processes with $n \geq 2$ components is not trivial to prove. This relies on the additional combinatorial complexity that appears when at least four components are involved, since then different synchronization events can occur concurrently. Therefore the simple structure of trajectories given by Proposition 1.2 is no longer valid, making in turn the constructions of this paper found in Sections 2 and 5 ineffective.

Nevertheless, the task of proving the existence of Markov processes with the LIP has been tackled in [1], generalizing the synchronization product of Markov chains. However, this construction is not very natural, and its main advantage is to encourage further study in this direction, since at least it ensures that the object of study is not empty.

Regarding a general theory of Markov multi-components processes, one may retain the following elements from the present paper: firstly, stopping times and the Asynchronous

Strong Markov Property have a straightforward extension to $n \geq 2$ components. These are basic tools that remain unchanged. Secondly, the generalized LIP allows to focus on the synchronization process only, since it implies a conditional decorrelation between the synchronization process on the one hand, and the private parts of each component on the other hand. The core of the remaining challenge is thus the construction and characterization of the synchronization process—we have shown above that, for two components, the synchronization process identifies with a homogeneous Markov chain, a drastic simplification compared to the general case of an arbitrary number of components. Recent work by G. Winskel [20] on probabilistic event structures has shown to be promising in this respect.

ACKNOWLEDGMENTS

Many thanks go to Albert Benveniste from IRISA in Rennes (France) for his support, his help and his friendship. I would like also to thank the anonymous referees for their many comments and suggestions, and the Editor Prakash Panangaden, to whom I am profoundly grateful.

REFERENCES

- [1] S. Abbes. Distributed Markov processes. Prépublication PPS, 2011. <http://hal.archives-ouvertes.fr/hal-00631501/en/>.
- [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. *Theoretical Computer Science*, 390:129–170, 2008.
- [4] F. Baccelli, G. Cohen, G.J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity*. Wiley, 1992.
- [5] E. Bellman. A Markovian decision process. *Journal of Mathematics and Mechanics*, 6:679–684, 1957.
- [6] A. Benveniste and É. Fabre. Partial order techniques for distributed discrete event systems: why you can’t avoid using them. *Discrete Event Dynamic Systems*, 17:355–403, 2007.
- [7] P. Billingsley. *Probability and Measure, 3rd edition*. John Wiley, 1995.
- [8] P.R. d’Argenio, H. Hermanns, and J.-P. Katoen. On generative parallel composition. *Electronic Notes in Computer Science*, 22:30–54, 1999.
- [9] C. Dellacherie and P.-A. Meyer. *Probabilities and Potential C*, volume 151 of *Mathematics Studies*. North-Holland, 1988.
- [10] G. Gierz, K.H. Hofmann, K. Keimel, J.D. Lawson, M.W. Mislove, and D.S. Scott. *Continuous Lattices and Domains*, volume 93 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2003.
- [11] C.A.R Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [12] 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. Springer, 2003.
- [13] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [14] J. Neveu. *Mathematical Foundations of the Calculus of Probability*. Holden-Day, 1965.
- [15] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, part 1. *Theoretical Computer Science*, 13:86–108, 1980.
- [16] M. Nielsen, G. Rozenberg, and P.S. Thiagarajan. Transition systems, event structures, and unfoldings. *Information and Computation*, 118(2):191–207, 1995.
- [17] D. Revuz. *Markov Chains*. North Holland, 1975.
- [18] 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. Springer, 1995.
- [19] 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 484–496. Springer, 2004.
- [20] G. Winskel. Distributed probabilistic strategies. In *29th Conference on the Mathematical Foundations of Programming Semantics*, 2013. To appear.