



HAL
open science

Distributed Markov Processes

Samy Abbes

► **To cite this version:**

| Samy Abbes. Distributed Markov Processes. 2011. hal-00631501

HAL Id: hal-00631501

<https://hal.science/hal-00631501>

Submitted on 12 Oct 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed Markov Processes

Samy Abbes*

Draft. October 12, 2011

Abstract

This paper introduces the model of *Distributed Markov Processes* (DMP), a probabilistic model of a system with distributed state over $n \geq 2$ sites. The definition of DMP being given, the notion of stopping time in the distributed context is introduced, and a Strong Markov Property is derived.

DMP are then characterized by their *characteristic coefficients*. These play a role similar to the coefficients of the transition matrix of discrete Markov chains, excepted that normalization conditions sufficient to define a DMP are not given here. The characteristic coefficients of a DMP are shown to satisfy the *concurrency equations*.

The main result of the paper is the proof of the existence of DMP on n sites, $n \geq 2$. The proof makes use of the tools introduced, especially the notion of stopping times. The case $n = 2$ has been extensively studied in a previous note.

Contents

1	The Algebraic Model	2
2	The Probabilistic Model	6
3	Stopping Times and the Strong Markov Property for DMPs	8
3.1	Stopping Times	8
3.2	Shift Operator and Iteration of Stopping Times	9
3.3	The (Strong) Markov Property	10
4	Construction of a Class of Distributed Markov Processes	10
4.1	Synchronization of Trajectories	10
4.2	Synchronization Product: Informal Description	11
4.3	Construction of the Synchronization Product	12
4.4	Properties of the Synchronization Product	14
4.5	Inductive Construction of Synchronization Products	17

*PPS/Université Paris Diderot – Paris 7. <http://www.pps.jussieu.fr/~abbes/en>

1 The Algebraic Model

Sites. Local and Global States. Transitions. We consider $n \geq 2$ sites, indexed $1, \dots, n$. Each site i has a set of *local states* denoted X^i . The X^i 's are *not* supposed to be pairwise disjoint. We denote by x^i a local state that belongs to X^i . We assume that “ \emptyset ” is a special symbol that does not belong to any X^i .

We define a *global state* of the system as any n -tuple $(x^1, \dots, x^n) \in X^1 \times \dots \times X^n$. It will be convenient to adopt a special notation for the set of global states, so we put

$$\mathcal{X} = X^1 \times \dots \times X^n.$$

To each local state $x \in \bigcup_{i=1}^n X^i$, we associate a n -tuple $t = (t^1, \dots, t^n)$ as follows:

$$\text{for } j = 1, \dots, n, \quad t^j = \begin{cases} x, & \text{if } x \in X^j, \\ \emptyset, & \text{the empty word, otherwise.} \end{cases} \quad (1)$$

So for instance, if x belongs to S^1 and S^2 only, we have

$$t = (x, x, \emptyset, \dots, \emptyset).$$

DEFINITION 1.1 (transitions and ressources) We call *transition* any n -tuple $t = (t^1, \dots, t^n)$ defined as in Eq. (1), and we denote by E the set of transitions. The *resources* of transition t are defined as those indices $i \in \{1, \dots, n\}$ such that $t^i \neq \emptyset$. The set of resources of t is denoted by $\rho(t)$, so that

$$\rho(t) = \{i : 1 \leq i \leq n, t^i \neq \emptyset\} = \{i : 1 \leq i \leq n, t^i \in X^i\}. \quad (2)$$

DEFINITION 1.2 (independence relation) We define the *independence* relation on transitions by declaring two transitions *independent* if their resources are disjoint. The independence relation is denoted with symbol “ \parallel ”, so we have:

$$\forall u, v \in E \quad u \parallel v \iff \rho(u) \cap \rho(v) = \emptyset.$$

Sequences of Transitions and Finite Trajectories. Let $S = E^*$ be the set of finite sequences of transitions. Then S is equipped with two structures: the prefix ordering, denoted \leq , and the concatenation on sequences, denoted by the dot \cdot . The latter identifies S with the free semigroup generated by E . The two structures are related as follows:

$$\forall u, v \in S \quad u \leq v \iff \exists z \in S \quad v = u \cdot z. \quad (3)$$

Furthermore, they satisfy the following compatibility relation:

$$\forall u, v, w \in S \quad u \leq v \Rightarrow w \cdot u \leq w \cdot v. \quad (4)$$

DEFINITION 1.3 (Adjacency relation) Say that two sequences $u, v \in S$ are immediately adjacent if one can deduce one from the other by inverting two neighbor transitions t, t' such that $t \parallel t'$. The symmetric and transitive closure of the immediate adjacency relation thus defined is called the *adjacency* relation, denoted by \sim .

Let \mathcal{S} denote the quotient set $\mathcal{S} = S / \sim$. Then the partial order and the semigroup structures on S induce similar structures on the quotient \mathcal{S} , which satisfy both properties (3) and (4), in \mathcal{S} . We denote by ϵ the neutral element in \mathcal{S} , which is also the least element of \mathcal{S} , and is the image in \mathcal{S} of the empty word.

DEFINITION 1.4 Elements of \mathcal{S} are called *finite trajectories*.

Projection of Trajectories on Local Sites. For each site $i = 1, \dots, n$, let $\theta^i : E \rightarrow (X^i)^*$ denote the i^{th} natural projection. An inspection of Eq. (2) defining resources of transitions shows that θ^i satisfies the following:

$$\forall i = 1, \dots, n \quad \begin{cases} \theta^i(t) = \emptyset & \iff i \notin \rho(t), \\ \theta^i(t) \in X^i & \iff i \in \rho(t). \end{cases} \quad (5)$$

LEMMA 1.1 For each $i = 1, \dots, n$, the i^{th} projection $\theta^i : E \rightarrow (X^i)^*$ satisfies

$$t \parallel t' \Rightarrow \theta^i(t) \cdot \theta^i(t') = \theta^i(t') \cdot \theta^i(t). \quad (6)$$

It follows from (6) that the semigroup extension $\theta^i : S \rightarrow (X^i)^*$ induces an application that we still denote by

$$\theta^i : \mathcal{S} \rightarrow (X^i)^*.$$

Intuitively, if u is a finite trajectory, then $\theta^i(u)$ represents the local action of u on site i .

PROPOSITION 1.1 1.—The application θ^i is a morphism of semigroups.
2.—For $u, v \in \mathcal{S}$, we have

$$(\theta^1(u), \dots, \theta^n(u)) = (\theta^1(v), \dots, \theta^n(v)) \Rightarrow u = v. \quad (7)$$

It is to be noted that the mapping

$$\theta^1 \times \dots \times \theta^n : \mathcal{S} \rightarrow (X^1)^* \times \dots \times (X^n)^*,$$

which is injective by virtue of Proposition 1.1 above, is not surjective in general.

Action of Finite Trajectories on Global States. If A is any set, and if $s \in A^*$ is a nonempty word on A , denote by $\gamma(s)$ the *last* letter of s . The semigroup A^* comes equipped with a natural right action on A , defined as follows:

$$\forall a \in A, \quad \forall s \in A^* \quad a \cdot s = \begin{cases} a, & \text{if } s \text{ is empty,} \\ \gamma(s), & \text{if } s \text{ non empty.} \end{cases} \quad (8)$$

This is indeed a *right* action, that is, we have

$$\forall a \in A \quad \forall s, s' \in A^* \quad a \cdot (s \cdot s') = (a \cdot s) \cdot s'.$$

Next we extend the above action to a right semigroup action of \mathcal{S} on the set X of global states, as follows. Consider first the extension componentwise of the above action to an action of S on X , and observe that the following holds:

$$\forall x \in X \quad \forall t, t' \in E \quad t \parallel t' \Rightarrow x \cdot (t \cdot t') = x \cdot (t' \cdot t).$$

It follows that the right action of S on X induces a right action of the quotient semigroup \mathcal{S} on X . By construction, the action commutes with the projections $\pi^i : X \rightarrow X^i$ in the following way: for $x \in X$ and $u \in \mathcal{S}$ and putting $y = x \cdot u$ and $y^i = x^i \cdot \theta^i(u)$, then:

$$\begin{array}{ccc} x & \xrightarrow{u} & y \\ \pi^i \downarrow & & \downarrow \pi^i \\ x^i & \xrightarrow{\theta^i(u)} & y^i \end{array}$$

Histories as Limiting Trajectories. Poset \mathcal{S} lacks *lubs* (least upper bounds). But it has a natural completion which is constructed as follows. Consider two *increasing* sequences $u = (u_k)_{k \geq 0}$ and $v = (v_k)_{k \geq 0}$ in \mathcal{S} . Define $u \preceq v$ whenever

$$\forall k \geq 0 \quad \exists p \geq 0 \quad u_k \leq v_p. \quad (9)$$

This is the so-called Egli-Miller order on sequences. Consider also the equivalence relation on increasing sequences in \mathcal{S} defined by

$$u \equiv v \iff (u \preceq v) \wedge (v \preceq u). \quad (10)$$

A limit element is represented by some increasing sequence u of elements in \mathcal{S} ; any two increasing sequences u and v are identified provided they satisfy $u \equiv v$. Denote by $(\overline{\mathcal{S}}, \leq)$ the quotient poset, the partial order of which is inherited from the \preceq order on sequences defined in (9).

There is an obvious embedding of posets $\mathcal{S} \rightarrow \overline{\mathcal{S}}$. Furthermore, $\overline{\mathcal{S}}$ has the property that any increasing sequence in $\overline{\mathcal{S}}$ has a *lub* in \mathcal{S} . And for any element $x \in \overline{\mathcal{S}}$ there is a maximal element $\omega \in \overline{\mathcal{S}}$ such that $x \leq \omega$.

DEFINITION 1.5 We say that a limiting element $\omega \in \overline{\mathcal{S}}$ is a *history* if ω is a maximal element of $\overline{\mathcal{S}}$. The set of histories is denoted by Ω .

Local Projections of Histories. For each $i = 1, \dots, n$, the application $\theta^i : S \rightarrow (X^i)^*$ is increasing since it is a semigroup morphism according to Prop. 1.1. Denote by $\overline{(X^i)^*}$ the set whose elements are either finite or infinite sequences of elements of X^i . For $u = (u_k)_{k \geq 0}$ and $v = (v_k)_{k \geq 0}$ two increasing sequences of elements of S satisfying $u \equiv v$ as in (10), the sequences $(\theta^i(u_k))_{k \geq 0}$ and $(\theta^i(v_k))_{k \geq 0}$ are both increasing in $(X^i)^*$ since θ^i is increasing, and satisfy furthermore

$$(\theta^i(u_k))_{k \geq 0} \equiv (\theta^i(v_k))_{k \geq 0},$$

where the \equiv equivalence is now taken as the equivalence relation relative to the Egli-Milner order on increasing sequences in $(X^i)^*$. They define thus a unique limiting sequence of elements in X^i , finite or infinite, that is to say, a unique element in $\overline{(X^i)^*}$. We will denote this element by $\theta^i(u)$, for u the corresponding class in $\overline{\mathcal{S}}$. Since the data

$$(\theta^1(u), \dots, \theta^n(u)) \in \overline{(X^1)^*} \times \dots \times \overline{(X^n)^*} \quad (11)$$

uniquely determines $u \in \overline{\mathcal{S}}$, we will most often represent an element $u \in \overline{\mathcal{S}}$ by the corresponding tuple (11). In particular the set Ω of maximal elements in $\overline{\mathcal{S}}$ identifies with a *subset* of $\overline{(X^1)^*} \times \dots \times \overline{(X^n)^*}$.

Concatenation of Trajectories with Histories. The concatenation

$$\mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S},$$

seen as a left action of \mathcal{S} on itself, extends to a left action of \mathcal{S} on $\overline{\mathcal{S}}$. Indeed, if $u = (u_k)_{k \geq 0}$ and $v = (v_k)_{k \geq 0}$ are two increasing sequences in \mathcal{S} such that $u \equiv v$ (in the sense of (10)), and if x is some fixed element in \mathcal{S} , then the sequences $x \cdot u$ and $x \cdot v$ defined by

$$x \cdot u = (x \cdot u_k)_{k \geq 0}, \quad x \cdot v = (x \cdot v_k)_{k \geq 0},$$

are increasing and satisfy $x \cdot u \equiv x \cdot v$. It is readily seen that $\cdot : \mathcal{S} \times \overline{\mathcal{S}} \rightarrow \overline{\mathcal{S}}$ is indeed a left action of \mathcal{S} on $\overline{\mathcal{S}}$, that is to say, that $x \cdot (y \cdot u) = (x \cdot y) \cdot u$ for $x, y \in \mathcal{S}$ and $u \in \overline{\mathcal{S}}$; and that it restricts to a left action of \mathcal{S} on Ω .

Furthermore, if $x \in \mathcal{S}$ is finite and if $u \in \overline{\mathcal{S}}$ is such that $x \leq u$ there is a unique element $v \in \overline{\mathcal{S}}$ such that $u = x \cdot v$. This element v is

- finite if and only if u is finite;
- maximal if and only if u is maximal.

Shadows. For any finite trajectory $u \in \mathcal{S}$, we define the *shadow* of u as the subset of Ω denoted by $\uparrow u$ and defined by

$$\uparrow u = \{\omega \in \Omega : u \leq \omega\}.$$

By analogy with the theory of infinite product spaces in measure theory, $\uparrow u$ can be seen as the *elementary cylinder* of base u .

For $u \in \mathcal{S}$ a fixed finite trajectory, the mapping

$$\Phi_u : \Omega \rightarrow \uparrow u, \quad \omega \mapsto u \cdot \omega, \quad (12)$$

where the concatenation $u \cdot \omega$ is defined as above, is a bijection.

2 The Probabilistic Model

General Notions. As a subset of $\overline{(X^1)^*} \times \cdots \times \overline{(X^n)^*}$, the set Ω is naturally equipped with a σ -algebra, that we denote by \mathfrak{F} , inherited from the Borel σ -algebra on $\overline{(X^1)^*} \times \cdots \times \overline{(X^n)^*}$.

The σ -algebra \mathfrak{F} is generated by the collection of shadows $\uparrow u$, for u ranging over \mathcal{S} . It is a result of classical measure theory that if two probabilities P and P' on (Ω, \mathfrak{F}) satisfy $P(\uparrow u) = P'(\uparrow u)$ for all $u \in \mathcal{S}$, then $P = P'$.

Assume P is some probability on (Ω, \mathfrak{F}) , and let $u \in \mathcal{S}$ be some finite trajectory. It is straightforward to check that the bijection $\Phi_u : \Omega \rightarrow \uparrow u$ defined in (12) is not only a bijection, but is also bi-measurable. Therefore, if $P(\uparrow u) > 0$, we shall consider the probability P_u on (Ω, \mathfrak{F}) , image under $\Phi_u^{-1} : \uparrow u \rightarrow \Omega$ of the conditional probability $P(\cdot | \uparrow u)$. The probability measure P_u is characterized by its values on shadows $\uparrow w$, for w ranging over \mathcal{S} , and given by

$$\forall w \in \mathcal{S}, \quad P_u(\uparrow w) = \frac{1}{P(\uparrow u)} P(\uparrow (u \cdot w)). \quad (13)$$

Probabilistic Processes. Instead of merely considering a probability measure on Ω , we will consider a *finite family* $(P_x)_{x \in X_0}$ of probability measures, indexed by some subset X_0 of X , and all defined on Ω . Intuitively, P_x describes the behavior of the system starting from the global state x . We will denote by \mathbb{P} such a family of probabilities, and call it a *probabilistic process*.

Reachable and Locally Reachable States.

DEFINITION 2.1 (reachable states) Let $\mathbb{P} = (P_x)_{x \in X}$ be a probabilistic process defined on n sites. A global state $y \in \mathcal{X}$ is *reachable from some global state* $x \in \mathcal{X}$ if

$$\exists u \in \mathcal{S}, \quad y = x \cdot u, \quad P_x(\uparrow u) > 0.$$

A weaker notion, often more useful, is that of *locally* reachable state.

DEFINITION 2.2 (locally reachable state) Let $\mathbb{P} = (P_x)_{x \in X}$ be a probabilistic process defined on n sites. Let i be a given site, and let $s \in X^i$ be a local state of site i . We say that s is *locally reachable* from some global state $x \in \mathcal{X}$ if there exists a global state $y \in \mathcal{X}$ such that:

1. y is reachable from x , in the sense of Def. 2.1; and
2. $s = \theta^i(y)$.

REMARK 2.1 If $n = 1$, the notion of local state coincides with that of local state; and reachable and locally reachable states are two identical notions.

Distributed Markov Processes.

DEFINITION 2.3 (DMP) A probabilistic process $\mathbb{P} = (P_x)_{x \in X_0}$ is said to be a *distributed Markov process (DMP)* if for all $x \in X_0$, and for all $u \in \mathcal{S}$ such that $\mathbb{P}_x(\uparrow u) > 0$, the two following assertions are true:

1. $x \cdot u \in X_0$, and
2. $(P_x)_u = P_{x \cdot u}$, where $(P_x)_u$ is the probability measure on (Ω, \mathfrak{F}) defined as in (13), with respect to probability P_x .

The following proposition gives a criterion to insure that a probabilistic process is indeed a DMP.

PROPOSITION 2.1 A probabilistic process $\mathbb{P} = (P_x)_{x \in X}$ is a DMP if and only if the value of the following quantities:

$$\frac{1}{P_x(\uparrow u)} P_x(\uparrow (u \cdot t)), \quad \text{for } (u, x, t) \in \mathcal{S} \times \mathcal{X} \times E, \quad (14)$$

only depend on the transition t , and on the resulting state $x \cdot u$.

Characteristic Numbers of a Distributed Markov Process. Consider a DMP \mathbb{P} —provided that it exists, to be proved later. Then define the following family of real numbers:

$$(P_x(\uparrow t)), \quad \text{for } x \in X \text{ and } t \in E. \quad (15)$$

We call it the family of *characteristic numbers* associated with the DMP.

PROPOSITION 2.2 If \mathbb{P} and \mathbb{Q} are two DMPs that induce the same families of characteristic numbers (15), then $\mathbb{P} = \mathbb{Q}$.

Concurrency Equations for Characteristic Numbers. Not any family of numbers $(p_x(t), x \in X, t \in E)$ identifies with the characteristic numbers of some DMP. Sadly, given a family $(p_x(t), x \in X, t \in E)$ of non-negative numbers, it is difficult to find a convenient necessary and sufficient condition for the existence of a DMP $\mathbb{P} = (P_x)_{x \in X}$ such that

$$\forall x \in X, \quad \forall t \in E, \quad p_x(t) = P_x(\uparrow t).$$

However, a first condition is given in the following proposition.

PROPOSITION 2.3 Let $(p_x(t), x \in X, t \in E)$ be the family of characteristic numbers of some DMP. Then the following equations are satisfied:

$$\forall x \in X_0, \quad \forall a, b \in E, \quad a // b \Rightarrow p_{x \cdot a}(b) = p_{x \cdot b}(a). \quad (16)$$

The equations (16), that reflect the partial order structure of trajectories, we call the *concurrency equations*. What is missing are normalization equations that would reflect the structure of the synchronizations.

3 Stopping Times and the Strong Markov Property for DMPs

3.1 Stopping Times

Let $\mathcal{N} = (\mathbb{N} \cup \{\infty\})^n$. This set will describe the “instants” of a system distributed over n sites. The analogy with the usual time must be used carefully, since in particular not all instants make sense; more precisely, the instants which make sense for a given history are random, that is to say, depend on the given history.

Recall that $\bar{\mathcal{S}}$ denotes the set of finite or infinite trajectories. If $u \in \bar{\mathcal{S}}$ is any trajectory, we define the *length* of u as follows:

$$|u| = (|\theta^1(u)|, \dots, |\theta^n(u)|) \in \mathcal{N}, \quad (17)$$

where $|\cdot|$ denotes the usual length of words, finite or infinite. We say that $u \in \bar{\mathcal{S}}$ has *finite length* whenever u is a finite trajectory, denoted by $|u| < \infty$.

NOTATION 3.1 Let $\omega \in \Omega$ and $t \in \mathcal{N}$ with $t = (t^1, \dots, t^n)$. We denote by ω_t the n -tuple

$$\omega_t = (\omega_t^1, \dots, \omega_t^n) \in \overline{(X^1)^*} \times \dots \times \overline{(X^n)^*}$$

such that:

$$\forall i = 1, \dots, n \quad \omega_t^i \leq \theta^i(\omega), \quad |\omega_t^i| = t^i.$$

It must be noted that, in general, ω_t is *not* a prefix of ω .

DEFINITION 3.1 (Random times and stopping times) Let $T : \Omega \rightarrow \mathcal{N}$ be an arbitrary mapping. We denote by the short notation ω_T the n -tuple $\omega_{T(\omega)} \in (\overline{X^1})^* \times \cdots \times (\overline{X^n})^*$ defined for $\omega \in \Omega$.

1. We say that T is a *random time* if $\omega_T \leq \omega$ 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:

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

PROPOSITION 3.1 Let $T : \Omega \rightarrow \mathcal{N}$ be a stopping time.

1. Then T and ω_T are two measurable mappings, \mathcal{N} being equipped with the discrete σ -algebra.
2. The σ -algebras generated by T and by ω_T coincide. This σ -algebra \mathfrak{F}_T is also characterized as follows:

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

Proof. 1. We show that $\zeta = \omega_t$ is measurable. It is enough to show that for every finite trajectory $u \in \mathcal{S}$, we have that $\zeta^{-1}(\{u\})$ is a measurable subset of Ω . But, by the property of stopping times, $\zeta^{-1}(\{u\})$ is either empty or $\uparrow u$, so it is measurable in both cases. It follows that T is measurable as well, since the length mapping $|\cdot| : \overline{\mathcal{S}} \rightarrow \mathcal{N}$ is clearly measurable. \square

3.2 Shift Operator and Iteration of Stopping Times

DEFINITION 3.2 (Shift operator) Let $T : \Omega \rightarrow \mathcal{N}$ 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.

DEFINITION 3.3 (iterated stopping times) Let $T : \Omega \rightarrow \mathcal{N}$ be a stopping time, and let θ_T be the associated shift operator. We define a sequence $(R_k)_{k \geq 0}$ of mappings $\Omega \rightarrow \mathcal{N}$, called *iterated stopping times associated with T* , as follows:

$$R_0 = 0$$

$$\forall k \geq 0 \quad R_{k+1} = \begin{cases} R_k + T \circ \theta_{R_k}, & \text{whenever } R_k \text{ and } T \circ \theta_{R_k} \text{ are finite,} \\ \infty, & \text{elsewhere.} \end{cases}$$

REMARK 3.1 $\theta_0 = \text{Id}_\Omega$, and $R_1 = T$.

PROPOSITION 3.2 Let $T : \Omega \rightarrow \mathcal{N}$ be a stopping time. Then the sequence $(R_k)_{k \geq 0}$ of iterated stopping times is a sequence of stopping times.

LEMMA 3.1 Let $T : \Omega \rightarrow \mathcal{N}$ be a stopping time, and let $(R_k)_{k \geq 0}$ be the associated sequence of iterated stopping times. Let $U : \Omega \rightarrow \mathcal{N}$ be another stopping time. Then the mapping $V : \Omega \rightarrow \mathcal{N}$ defined as follows:

$$\forall \omega \in \Omega, \quad \omega_V = \begin{cases} \bigwedge \{ \omega_{T_k} : \omega_{T_k} \geq \omega_U, k \geq 0 \}, & \text{if } \exists k \geq 0, \quad \omega_U \leq \omega_{T_k}, \\ \infty, & \text{otherwise.} \end{cases}$$

is a stopping time.

3.3 The (Strong) Markov Property

Since constant times do not make sense in our context, a weak Markov property is excluded. Fortunately, stopping times have a meaning for us, and therefore a strong Markov property can be expected, as shown by the following theorem.

THEOREM 3.1 Let $T : \Omega \rightarrow \mathcal{N}$ be a stopping time and let $\mathbb{P} = (P_x)_{x \in X_0}$ be a DMP. Then the following equality holds for all $x \in X_0$ and for every non-negative and measurable function $h : \Omega \rightarrow \mathbb{R}$:

$$E_x(h \circ \theta_T | \mathfrak{F}_T) = E_{\gamma(\omega_T)}(h), \quad P_x\text{-almost surely,} \quad (18)$$

where both members are random variables only defined on $\{T < \infty\}$.

4 Construction of a Class of Distributed Markov Processes

4.1 Synchronization of Trajectories

For u a trajectory on $n \geq 2$ sites, we denote by $\theta^{[1, n-1]}(u)$ the projection of u on the $n - 1$ first sites. In other words, $w = \theta^{[1, n-1]}(u)$ is the unique trajectory on $n - 1$ sites such that

$$\forall i = 1, \dots, n - 1 \quad \theta^i(w) = \theta^i(u).$$

DEFINITION 4.1 Let w be a finite trajectory on $n - 1$ sites with local state spaces X^1, \dots, X^{n-1} , and z be a finite trajectory on a n^{th} site with local state space X^n .

1. We say that w and z are *compatible* if there exists a trajectory u on n sites such that

$$\theta^{[1, n-1]}(u) = w, \quad \text{and} \quad \theta^n(u) = z.$$

2. We say that w and z are *sub-compatible* if there exists a sub-trajectory $w' \leq w$ such that w' and z are compatible. In this case, we put

$$v = \inf\{w' \leq w : w' \text{ and } z \text{ are compatible}\}$$

and we define the *synchronization of w and z* as the unique trajectory on n sites, denoted by $w \parallel z$ such that

$$\theta^{[1,n-1]}(w \parallel z) = v, \quad \text{and} \quad \theta^n(w \parallel z) = z.$$

DEFINITION 4.2 Consider a process defined on $n - 1$ sites with state spaces X^1, \dots, X^{n-1} , and let X^n be another set. Put

$$S = X^{1,n} \cup \dots \cup X^{n-1,n},$$

to be understood as the set of shared states of site n with some other site. Let $\Omega^{[1,n-1]}$ be the space of histories defined on the $n - 1$ first sites, and let Ω^n be the set of trajectories on the n^{th} site. Finally, put

$$\Lambda = \Omega^{[1,n-1]} \times \Omega^n.$$

For $z \in \Omega^n$, we consider $T(z) \in \mathbb{N} \cup \{\infty\}$, defined as the first hitting time of S for z . We keep our usual notation z_T to denote the sub-trajectory of z until time $T(z)$. We define the *first synchronization time* as $\sigma : \Lambda \rightarrow \mathcal{N}$ such that:

$$\text{for } \lambda = (w, z) \in \Lambda, \quad \lambda_\sigma = \begin{cases} \infty, & \text{if } w \text{ and } z_T \text{ are not sub-compatible,} \\ w \parallel z, & \text{if } w \text{ and } z_T \text{ are sub-compatible.} \end{cases}$$

PROPOSITION 4.1 The first synchronization time satisfies the following property:

$$\forall \lambda, \lambda' \in \Omega^{[1,n-1]} \times \Omega^n \quad \lambda_\sigma \leq \lambda'_\sigma \Rightarrow \lambda_\sigma = \lambda'_\sigma. \quad (19)$$

Of course, similarity with Def. 3.1 is to be noted. Def. 3.1 does not directly apply here since $\Lambda = \Omega^{[n-1]} \times \Omega^n$ is not our usual space of trajectories on n sites, precisely because we are in the process of building trajectories on n sites.

4.2 Synchronization Product: Informal Description

Let \mathbb{Q} and \mathbb{M} be two DMP, \mathbb{Q} over $n - 1$ sites and \mathbb{M} over a n^{th} site with set of states X^n . Observe that \mathbb{M} is just a homogeneous Markov chain on X^n .

We will denote by $\mathcal{X}^{[1,n-1]}$ the set $X^1 \times \dots \times X^{n-1}$ of global states of process \mathbb{Q} . We will use the notation \mathcal{X} to denote

$$\mathcal{X} = \mathcal{X}^{[1,n-1]} \times X^n.$$

The notation \mathcal{X} is purposely identical with the one introduced to denote the set of global states on n sites.

We assume that the following assumptions are in force:

1. For any $s, s' \in X^n$, s' is reachable from s with respect to \mathbb{M} ;
2. For any $x \in \mathcal{X}^{[1, n-1]}$, for any site $i \leq n-1$ and for any local state $s \in X^i$, s is a locally reachable state from x ;
3. $X^{i, n} \neq \emptyset$ for all $i = 1, \dots, n-1$.

Informally, the construction we perform intends to force the synchronization of the two processes described by \mathbb{Q} and \mathbb{M} , and proceeds as follows. Consider an element $\lambda \in \Lambda = \Omega^{[n-1]} \times \Omega^n$. Consider the first synchronization time σ described in § 4.1, provided it is finite, yielding λ_σ . To insure λ_σ is well defined, introduce the event

$$\Delta = \{\lambda \in \Lambda : |\lambda_\sigma| < \infty\}, \quad (20)$$

and equip λ_σ with the conditional law

$$K_{(x, y)}(\lambda_\sigma) = Q_x \otimes M_y(\lambda_\sigma | \Delta), \quad (21)$$

where x and y are respectively the initial global states of histories ω and ζ , and where Q_x and M_y are the probability measures that define the processes \mathbb{Q} and \mathbb{M} respectively. By construction, the random variable λ_σ identifies with a finite trajectory on n sites, that brings the system with n sites to a new global state (x', y') ; then repeat this construction, starting from (x', y') this time, and then concatenate the resulting new random trajectory with λ_σ previously obtained.

The concatenation of infinitely many of these random trajectories eventually leads to a random history $\omega \in \Omega$ and the law of this random history is precisely the probability measure on Ω that we are seeking.

4.3 Construction of the Synchronization Product

We now proceed to the rigorous construction of the synchronization product, informally described above. The first step consists in describing the transition matrix and initial measure of the auxiliary Markov chain that has been described in the previous subsection.

Transition Matrix of the Auxiliary Markov Chain. Consider the countable set \mathcal{C} of values of $\lambda_\sigma : \Lambda \rightarrow \mathcal{S}$, equipped with the family of probabilities $K_{(x, y)}$ given by Eq. (21); and then the product $\mathcal{C}' = \mathcal{C} \times \mathcal{X}$. Elements of \mathcal{C}' will be noted as $(w, z, x, y,)$ with w a finite trajectory on $n-1$ sites, z a finite path on X^n , and $(x, y) \in \mathcal{X}^{[1, n-1]} \times X^n$.

Then we introduce the transition matrix L on \mathcal{C}' given by:

$$L(w, z, x, y \rightarrow w', z', x', y') = \mathbf{1}_{\{\gamma(w||z)=(x', y')\}} \times K_{(x', y')}(\lambda_\sigma = (w', z')), \quad (22)$$

where $\mathbf{1}_{\{\cdot\}}$ stands for the characteristic function.

Initial Measures for the Auxiliary Markov Chain. We consider a family $(\mu_{(x,y)})_{(x,y)}$ of initial measures on \mathcal{C}' , for (x, y) ranging over the set of global states $X^1 \times \cdots \times X^n$. For $(x_0, y_0) \in \mathcal{X}$ being given, $\mu_{(x_0, y_0)}$ is defined as follows:

$$\forall (w, z, x, y) \in \mathcal{C}', \quad \mu_{(x_0, y_0)}(w, z, x, y) = \mathbf{1}_{\{x=x_0, y=y_0\}} K_{(x_0, y_0)}(\lambda_\sigma = w \parallel z). \quad (23)$$

Definition of the Synchronization Product.

LEMMA 4.1 If assumptions 1 and 3 above are in force, and if

$$(w_n, z_n, x_n, y_n)_{n \geq 0}$$

is a sample path of the Markov chain with transition matrix L given above, then the concatenation

$$(w_1 \parallel z_1) \cdot (w_2 \parallel z_2) \cdot \dots \quad (24)$$

is almost surely a history, i.e., a maximal trajectory in the space Ω of trajectories over n sites.

Proof. Assume that $(w_n, z_n)_{n \geq 0}$ is a sample such that the concatenation (24) is not maximal. Put

$$\begin{aligned} \forall n \geq 1 \quad r_n &= (w_1 \parallel z_1) \cdot (w_2 \parallel z_2) \cdot \dots \cdot (w_n \parallel z_n), \\ \text{and} \quad r &= (w_1 \parallel z_1) \cdot (w_2 \parallel z_2) \cdot \dots \\ &= \sup_{n \geq 1} r_n. \end{aligned}$$

Since r is assumed to be non maximal, there is a site i such that the $\theta^i(r)$ is of finite length. Since $\theta^i(r) = \sup_{n \geq 1} \theta^i(r_n)$, and since $\theta^i(r)$ has length i at least, it follows that $i < n$. Hence, infinitely often, site n has not synchronized with site i although it has non zero probability according to Assumption 3 (stated at the beginning of § 4.2). By the Borel-Cantelli Lemma, all such r together form a subset of zero probability, what was to be proved. \square

Let Ξ be the natural sample space for Markov chains on \mathcal{C}' . For $(x, y) \in \mathcal{X}$, let $\pi_{(x,y)}$ be the probability measure on Ξ associated with initial measure $\mu_{(x,y)}$ defined in Eq. (23) and with transition matrix L defined in Eq. (22). According to Lemma 4.1, the concatenation of elements of a sample path $\xi \in \Xi$ defines an element $\omega \in \Omega$. The concatenation defines thus a mapping $\Phi : \Xi \rightarrow \Omega$, which is easily checked to be measurable.

DEFINITION 4.3 (synchronization product) Let \mathbb{Q} and \mathbb{M} be two DMP respectively defined on $n - 1$ site for \mathbb{Q} and on a n^{th} site for \mathbb{M} . We assume the two assumptions 1 and 2 are in force. Then the *synchronization product* of \mathbb{Q} and \mathbb{M} is the probabilistic process on n sites $\mathbb{P} = (P_{(x,y)})_{(x,y)}$, with (x, y) ranging over the set \mathcal{X} of global states over n sites, and with $P_{(x,y)}$ the probability on (Ω, \mathfrak{F}) given as the image measure

$$P_{(x,y)} = \Phi\pi_{(x,y)},$$

where $\pi_{(x,y)}$ is the probability on the sample space Ξ defined as before, with initial measure $\mu_{(x,y)}$ and transition matrix L .

4.4 Properties of the Synchronization Product

Let $\omega \in \Omega$ be a history on n sites. Then ω naturally decomposes as $\omega = (w, z)$ with w a history on the $n - 1$ first sites and z a sample path on the n^{th} site, whence a natural injection $\Omega \rightarrow \Lambda$. This injection allows to transpose on Ω whatever mappings previously defined on Λ , such as $\sigma : \Lambda \rightarrow \mathcal{N}$.

LEMMA 4.2 The mapping $S : \Omega \rightarrow \mathcal{N}$ obtained by composition of the natural injection $\Omega \rightarrow \Lambda$ with $\sigma : \Lambda \rightarrow \mathcal{N}$ defined in Def. 4.2, is a stopping time in the sense of Def. 3.1.

COROLLARY 4.1 Let u be a given finite trajectory on n sites. Consider the sequence $(S_k)_{k \geq 0}$ of iterated stopping times associated with S previously defined in Lemma 4.2. The mapping $V : \Omega \rightarrow \mathcal{N}$ defined as follows:

$$\forall \omega \in \Omega, \quad \omega_V = \begin{cases} \bigwedge \{ \omega_{S_k} : \omega_{S_k} \geq u, k \geq 0 \}, & \text{if } \exists k \geq 0, \quad \omega_{S_k} \geq u, \\ \infty, & \text{otherwise,} \end{cases}$$

is a stopping time.

Proof. Consider the stopping time $U : \Omega \rightarrow \mathcal{N}$ defined by:

$$\forall \omega \in \Omega, \quad \omega_U = \begin{cases} u, & \text{if } u \leq \omega, \\ \infty, & \text{otherwise.} \end{cases} \quad (25)$$

By Lemma 4.2, $S : \Omega \rightarrow \mathcal{N}$ is a stopping time. Applying Lemma 3.1 to stopping time S and its iterated sequence $(S_k)_{k \geq 0}$ on the one hand, and to stopping time U defined by Eq. (25) on the other hand, yields the result of the corollary. \square

PROPOSITION 4.2 The synchronization product \mathbb{P} described in Def. 4.3 is a DMP over n sites.

Proof. We denote by x a generic global state of the $n - 1$ first sites, that is, an element of $\mathcal{X}^{[1, n-1]}$, and by y an element of X^n . What needs to be proved is that the quantity

$$\mathbb{P}_{(x,y)}(\uparrow(t \cdot t') | \uparrow t) \quad (26)$$

only depends on $\gamma(t)$ and t' , for t, t' ranging over the finite trajectories on n sites and for (x, y) ranging over $\mathcal{X} = \mathcal{X}^{[1, n-1]} \times X^n$. For this, we proceed by steps. In the remaining of the proof, and accordingly with the notations of (26), (x, y) denotes the initial global state of the system distributed over n sites, and t denotes some finite trajectory over n sites. We also put $(x', y') = \gamma(t)$.

Step 1. Without loss of generality, one may assume that t' has the form $t' = \omega_S$, where S is the stopping time defined in Lemma 4.2. Indeed, assume that the result is true for any t' of this form, then by the chain rule it is clearly true also for t' any finite concatenation of such trajectories. Hence the result is true for t' of the form $t' = \omega_{S_k}$ for any $k \geq 0$, with $(S_k)_{k \geq 0}$ the sequence of iterated stopping times associated with S . Let now t, t' be any finite trajectories over n sites.

By Corollary 4.1, we have that $V : \Omega \rightarrow \mathcal{N}$ defined by

$$\forall \omega \in \Omega, \quad \omega_V = \begin{cases} \bigwedge \{ \omega_{S_k} : \omega_{S_k} \geq t', k \geq 0 \}, & \text{if } \exists k \geq 0, \quad \omega_{S_k} \geq t \cdot t', \\ \infty, & \text{otherwise,} \end{cases}$$

is a stopping time. We decompose over the possible values v of V to get:

$$\begin{aligned} P_{(x,y)}(\uparrow(t \cdot t') | \uparrow t) &= \sum_v P_{(x,y)}(\uparrow(t \cdot t') \cap \{t \cdot \omega_V = t \cdot v\} | \uparrow t) \\ &= \sum_v P_{(x,y)}(\uparrow(t \cdot v) | \uparrow t) \quad \text{since } V \text{ is a stopping time.} \end{aligned}$$

Now by assumption of Step 1 each term of the sum only depends on $\gamma(t)$ and v . And by construction, the range of v only depends on t' . It follows that the sum itself only depends on $\gamma(t)$ and t' .

Step 2. Without loss of generality, one may assume that t decomposes as $t = w \cdot z$ with w over the $n - 1$ first sites, z over X^n , and $w \parallel z$. In other words, one may assume that $\theta^n(t) \cap X^i = \emptyset$ for $i = 1, \dots, n - 1$. For, if t is any finite trajectory, let $z = \theta^n(t)$. Let k be the number of synchronizations of z with the $n - 1$ first sites. Then if $(w_1 \parallel z_1), \dots, (w_k \parallel z_k)$ denote the k successive synchronization times of t , then t can be written as follows:

$$t = (w_1 \parallel z_1) \cdot \dots \cdot (w_k \parallel z_k) \cdot r,$$

with r of the following form:

$$r = u \cdot s,$$

with u over the $n - 1$ first sites, s over X^n and $u \parallel s$. Finally, putting

$$t_0 = (w_1 \parallel z_1) \cdot \dots \cdot (w_k \parallel z_k),$$

we have:

$$\begin{aligned} P_{(x,y)}(\uparrow (t \cdot t') \mid \uparrow t) &= \frac{P_{(x,y)}(\uparrow (t_0 \cdot r \cdot t'))}{P_{(x,y)}(\uparrow (t_0 \cdot r))} \\ &= \frac{P_{(x,y)}(\uparrow (t_0 \cdot r \cdot t') \mid \uparrow t_0)}{P_{(x,y)}(\uparrow (t_0 \cdot r) \mid \uparrow t_0)}. \end{aligned} \quad (27)$$

Since t_0 is a concatenation of finite trajectories of the form $w_i \parallel z_i$, and by the construction of the synchronization product, it follows by the chain rule that

$$P_{(x,y)}(\uparrow (t_0 \cdot r \cdot t') \mid \uparrow t_0) = P_{\gamma(t_0)}(\uparrow (r \cdot t')) \quad (28)$$

$$P_{(x,y)}(\uparrow (t_0 \cdot r) \mid \uparrow t_0) = P_{\gamma(t_0)}(\uparrow r). \quad (29)$$

Reinserting the right members of Eqs. (28) and (29) in (27), we get:

$$P_{(x,y)}(\uparrow (t \cdot t') \mid \uparrow t) = P_{\gamma(t_0)}(\uparrow (r \cdot t') \mid \uparrow r). \quad (30)$$

Now, by the assumption made in this step, and since r has the requested form, we find that the right member of Eq. (30) only depends on $\gamma(r) = \gamma(t)$ and t' , what was to be shown.

Step 3. The basic case. According to the reduction established in Steps 1–2, let t be a finite trajectory over n sites satisfying $\theta^n(t) \cap X^i = \emptyset$ for $i = 1, \dots, n - 1$, and let $t' = w \parallel z$. Then, according to the construction of the synchronization product \mathbb{P} , we have:

$$\begin{aligned} P_{(x,y)}(\uparrow (t \cdot t') \mid \uparrow t) &= \frac{P_{(x,y)}(\uparrow (t \cdot t'))}{P_{(x,y)}(\uparrow t)} \\ &= \frac{Q_x \otimes M_y(\uparrow (t \cdot t') \cap \Delta)}{Q_x \otimes M_y(\uparrow t \cap \Delta)} \\ &= \frac{Q_x \otimes M_y(\uparrow (t \cdot t'))}{Q_x \otimes M_y(\uparrow t)} \cdot \frac{1}{Q_x \otimes M_y(\Delta \mid \uparrow t)} \\ &= \frac{Q_x \otimes M_y(\uparrow (t \cdot t') \mid \uparrow t)}{Q_x \otimes M_y(\Delta \mid \uparrow t)}. \end{aligned}$$

On the last line, since both Q_x and M_y are Markovian, it is clear that the expression only depends on $\gamma(t)$ and t' . \square

4.5 Inductive Construction of Synchronization Products

LEMMA 4.3 Let \mathbb{Q} and \mathbb{M} be two DMP on $n - 1$ and 1 site respectively, with global states $\mathcal{X}^{[1,n-1]} = X^1 \times \dots \times X^{n-1}$ and X^n respectively. We assume that the transition matrix M associated with Markov chain \mathbb{M} has all its coefficients positive, and that Assumptions 1–3 from § 4.2 are in force.

Then the synchronization product of \mathbb{Q} and \mathbb{M} , with global states in $\mathcal{X} = \mathcal{X}^{[n-1]} \times X^n$ satisfies the following property: for every global state $x \in \mathcal{X}$ and for any site $i = 1, \dots, n$, any local state $s \in X^i$ is locally reachable from x (see Def. 2.2).

Proof. We keep the notations of § 4.3–§ 4.4. Let $(x, y) \in \mathcal{X}^{[1,n-1]} \times X^n$ be any initial global state on n sites, and let s be some local state of some site i . We show the existence of some finite trajectory $u \in \overline{\mathcal{S}}$ on n sites such that $P_{(x,y)} > 0$ and $\theta^i \circ \gamma(u) = s$.

1. Assume that $i = n$.

- (a) If $s \notin X^j$ for all $j = 1, \dots, n - 1$. Put $u = (\epsilon, \dots, \epsilon, s)$. Consider any value (w, z) of λ_σ defined in § 4.1. Then $u \cdot (w, z)$ is also a value for λ_σ . Therefore:

$$P_{(x,y)}(\uparrow u) \geq P_{(x,y)}(\uparrow u \cdot (w, z)) = K_{(x,y)}(\lambda_\sigma = u \cdot (w, z)) > 0.$$

- (b) Is $s \in X^j$ for some $j = 1, \dots, n - 1$. By Assumption 2, s is locally reachable from x . Hence there is a finite trajectory $v \in \overline{\mathcal{S}}^{[1,n-1]}$ on $n - 1$ sites such that $\theta^j \circ \gamma(v) = s$, and $Q_x(\uparrow v) > 0$. Next extract from v all states that belong to X^n , and sequentialize them into a path on X^n with s as its last element (such a path exists). Then, by construction, v and z are compatible and by putting $u = v \parallel z$ we have $\theta^{[1,n-1]}(u) = v$ and $\theta^n(u) = z$. Furthermore, if $k = |z|$, and if $(S_k)_{k \geq 0}$ denotes as before the sequence of iterated stopping times associated with stopping time S , then any $\omega \in \Omega$ satisfying $\omega \geq u$ satisfies $S_k(\omega) = u$. By construction of the synchronization product, and since z has positive M_y -probability since by assumption the transition matrix only has positive coefficients, it follows that $P_{(x,y)}(S_k = u) > 0$. Since $\theta^n(u) = z$, we also have that $\gamma \circ \theta^n(u) = s$ since s is the last element of z .

2. Assume that $i < n$. The construction of the element u is similar to the case 1b above.

□

THEOREM 4.1 For any $n \geq 1$ and for finite space states X^1, \dots, X^n of arbitrary finite size, there exists DMP inductively obtained by synchronization product with global states $X^1 \times \dots \times X^n$.

Proof. For each $i = 1, \dots, n$, consider any transition matrix with positive coefficients on X^i . Let \mathbb{M}_i denote the probabilistic process associated on X^i associated with the given transition matrix.

Rename the elements of X^i so that

$$\forall 1 \leq i, j \leq n \quad X^i \cap X^j \neq \emptyset. \quad (31)$$

Then proceed inductively to the construction of the DMP \mathbb{P}_i on $\mathcal{X}^{[1,i]} = X^1 \times \dots \times X^i$, synchronization product of \mathbb{P}_{i-1} with \mathbb{M}_i . By Lemma 4.3, \mathbb{P}_i has the property that any local states are locally reachable. Hence Assumption 2 is fulfilled. Assumption 1 is also fulfilled since the transition matrices all have positive coefficients. Finally, Assumption 3 is also fulfilled thanks to Eq. (31). Hence the construction of \mathbb{P}_{i+1} is allowed, completing the proof. \square