Labelled Markov Processes: Stronger and Faster Approximations

Vincent Danos
CNRS & Université Paris 7, PPS
2 Place Jussieu 75251 PARIS Cedex 05
Vincent.Danos@pps.jussieu.fr

Josée Desharnais
Département d'informatique et de génie logiciel
Université Laval, Québec, Canada, G1K 7P4
Josee.Desharnais@ift.ulaval.ca

Abstract

This paper proposes a measure-theoretic reconstruction of the approximation schemes developed for Labelled Markov Processes: approximants are seen as quotients with respect to sets of temporal properties expressed in a simple logic. This gives the possibility of customizing approximants with respect to properties of interest and is thus an important step towards using automated techniques intended for finite state systems, e.g. model checking, for continuous state systems.

The measure-theoretic apparatus meshes well with an enriched logic, extended with a greatest fix-point, and gives means to define approximants which retain cyclic properties of their target.

1. Introduction

In [5, 6, 7], a measure-theoretic presentation of Labelled Markov Processes (LMPs) has been given. These can be thought of as discrete-time interactive Markov processes. One might think that handling arbitrary measurable spaces in place of finite state probabilistic automata would complicate the theory a lot, but it does not. The state of the art is that bisimulation and simulation between LMPs are now explained in terms of temporal properties written in simple temporal logics, much in the way CCS bisimulation is explained with Hennessy-Milner logic [9], except that those logics are negationless. Based on these logics, an approximation machinery has been set up that will provide a sequence of finite processes sharing more and more properties with the approximated process.

Yet, there remains two soft spots in the approximation theory. First, while approximants clearly ought to be some sort of finite quotients by temporal properties of what it is they try to approximate, nobody so far was able to lay his hands on a precise way of phrasing just this intuition. This is the main conceptual contribution of this paper: we enrich the theory with the possibility of *customizing* approximants with respect to a specific set of properties. One chooses a (finite or infinite) set $\mathcal{F}$ of interesting properties stated in a modal logic and is then provided with an approximant that is $\mathcal{F}$-equivalent to the original process. For example, one could choose for $\mathcal{F}$ all properties of depth $\leq k$ for some k, and get k-step bisimulation (in general, this would not yield a finite approximant). Some, but not much, reworking of the axiomatics is involved, resulting in a simpler theory and there are minor conceptual clarifications along the way. As is already the case in [7], topological considerations are not needed for approximants and we do not have to meddle with the intricacies of analytic spaces.

Second, another motive for action is that when fed with a finite process the approximation machinery is unable to retrieve the process itself at the limit. Instead, its unfolding is obtained with countably many states. For instance, the pure loop process, with one state and one a-transition to itself, will be approximated by all its finite unfoldings, *i.e.* by chains of a-transitions, which seems spectacularly not what one would like to have intuitively, even if it is acceptable on the technical side (the infinite chain of a-transitions is bisimilar to the loop, after all). The second contribution of this paper is a solution to finding approximants showing more respect. They will be faster also in the sense that every finite-state process will be quotiented in a finite number of steps to itself. The same general construction used to posit standard approximations, will accommodate the stronger approximations as well.

We proceed in the following way. First, we pry out the construction of finite LMP approximants from the framework of LMPs [7] to give it larger generality and flexibility. This involves dropping additivity and working with super-additive Markov kernels in what we call a *pre-LMP*. A brief study of quotients under "reasonable" equivalences (technically, equivalence classes have to be in the underlying σ-algebra and countably generated) in this framework is developed and brought to the conclusion that there is a cone of finite pre-LMPs under any LMP (and even pre-LMPs of course) indexed by finite subsets of temporal properties. The measure-theoretic development, though quite routine when all definitions are set up properly, is, as far as we know, original material. The important implication of this first part is that approximants are really quotients by logical properties after all, as the intuition would have it from the beginning.

With this in hand, we extend the space of temporal properties themselves, to get finite predictors that will also converge more accurately, in the sense that cyclic temporal properties will be obtained at finite stages. Greatest fix-point constructions turn out to fit wonderfully in the general measure-theoretic presentation.

Acknowledgements. The first author wishes to thank Prakash Panangaden for the lecture series on Labelled Markov Processes given during his stay in Paris in May 2002 at the PPS lab, and both authors thank him for numerous instructive discussions about the subject as a whole.

2 Preliminaries

This section is a brief reminder of the main objects of the trade with definitions slightly optimized for the development we have in mind. The paper is self-contained, though the reader might find useful to consult a book on basic probability theory, e.g. [11].

Notations. When S is a set and $A \subseteq S$, we write $\mathbf{1}_A$ for A's indicator function. When A, B are disjoint sets, we sometimes write $A+B$ for the (disjoint) union, and conversely each time we write $A+B$ it is understood that A and B are indeed disjoint. We write $\downarrow A_n$ when A_n is a decreasing sequence of sets, that is $A_n \supseteq A_{n+1}$ and $\cap A_n$ for the limit.

When $\mathfrak{R}$ is an equivalence relation over S, and $s \in S$: the equivalence class of s is written either $[s]_{\mathfrak{R}}$ or simply $[s]$, when $\mathfrak{R}$ is clear from the context; and by $\mathfrak{R}(s)$ we mean $\{t \mid (s,t) \in \mathfrak{R}\}$. If A is a set of equivalence classes, one uses the usual set-theoretic notation for union: $\cup A := \{s \in S \mid [s] \in A\}$. Finally, a set A is said to be $\mathfrak{R}$-closed if whenever $s \in A$ and $(s,t) \in \mathfrak{R}$, then $t \in A$.

We fix once and for all a finite alphabet L of actions.

Labelled Markov processes. A *measurable space* is a pair (S, Σ) where S is any set and $\Sigma \subset 2^S$ is a σ-*algebra* over S, that is a set of subsets of S, containing S and closed by countable intersection and complement. Famous examples are $[0,1]$ and $[0,+\infty]$ equipped with their respective *Borelian* σ-algebras $\mathcal{B}$ and $\mathcal{B}^\infty$ generated by the intervals.

A map q between two measurable spaces (S,Σ) and (S',Σ') is said to be *measurable* if for all $A' \in \Sigma'$, $q^{-1}(A') \in \Sigma$. The smallest σ-algebra such that a given q is measurable is the pointwise inverse image of Σ'.

A *subprobability* on (S,Σ) is a map $p : \Sigma \to [0,1]$, such that for any countable collection (A_n) of pairwise disjoint sets in Σ, $p(\sum_n A_n) = \sum_n p(A_n)$. An actual probability is when in addition $p(S) = 1$. The condition on p is called σ-additivity and can be conveniently broken in two parts:
— *finite additivity*: $p(A+B) = p(A) + p(B)$,
— *co-continuity*: $\forall \downarrow A_n \in \Sigma : p(\cap A_n) = \inf_n p(A_n)$.

Definition 1 *A LMP is a triple $\mathcal{S} = (S, \Sigma, h : L \times S \times \Sigma \to [0,1])$ where (S,Σ) is a measurable space, and for all $a \in L$, $s \in S$, $A \in \Sigma$:*
— $h_a(s, .)$ is a subprobability on (S,Σ),
— $h_a(., A)$ is measurable from (S,Σ) to $([0,1],\mathcal{B})$.

After the traditional terminology in Markov chains, the map h is called the *kernel* of $\mathcal{S}$. Most of the time, we will write $h_a(s,A)$ simply as $h_a(s,A)$. It is a measure of the likelihood that being at s and receiving a the LMP will jump in A.

Temporal properties and simulation. LMPs depart from usual Markov chains in that kernels also depend on an auxiliary set L of actions, but most importantly, because one thinks of them differently. They're construed as interactive processes and therefore one is interested in various notion of bisimulations and simulations as in non-deterministic process algebras [10].

The following "bisimulation logic" $\mathcal{L}_0$ is a central tool for asserting properties of LMPs:

$$\theta := \top,\ \theta \wedge \theta,\ \langle a \rangle_r \theta.$$

The *depth* $|\theta|$ of a formula θ is defined as: $|\top| = 0$, $|\theta_0 \wedge \theta_1| = \max(|\theta|_0, |\theta|_1)$ and $|\langle a \rangle_r \theta| = |\theta| + 1$.

Definition 2 *Given an LMP $\mathcal{S}$, one may inductively define the map $[\![.]\!]_{\mathcal{S}} : \mathcal{L}_0 \to \Sigma$ as:*
— $[\![\top]\!]_{\mathcal{S}} = S$,
— $[\![\theta_0 \wedge \theta_1]\!]_{\mathcal{S}} = [\![\theta_0]\!]_{\mathcal{S}} \cap [\![\theta_1]\!]_{\mathcal{S}}$,
— $[\![\langle a \rangle_r \theta]\!]_{\mathcal{S}} = \{s \in S \mid h_a(s, [\![\theta]\!]_{\mathcal{S}}) \geq r\}$.

We write $s \models \theta$ to mean $s \in [\![\theta]\!]_{\mathcal{S}}$ and $\theta' \leq \theta$ to mean that θ' is a subformula of θ. Monoidal equations: $\theta_0 \wedge (\theta_1 \wedge \theta_2) = (\theta_0 \wedge \theta_1) \wedge \theta_2$, $\theta_0 \wedge \theta_1 = \theta_1 \wedge \theta_0$, $\theta \wedge \top = \theta$ all clearly preserve $[\![.]\!]_{\mathcal{S}}$.

The logic induces a form of *simulation* between states in the sense that a state can be said to simulate another one if it satisfies at least the same formulas as the other does. The concept can be cast in behavioural terms as in the following definition.

Definition 3 ([7]) *Let $\mathcal{S} = (S, \Sigma, h)$ be a LMP. A relation $\mathfrak{R}$ on S is a* simulation *if whenever $s\mathfrak{R}s'$, we have that for all $a \in L$ and every $\mathfrak{R}$-closed measurable set $A \in \Sigma$, $h_a(s, A) \leq h_a(s', A)$. We say s is simulated by s' if $s\mathfrak{R}s'$ for some simulation relation $\mathfrak{R}$.*

This definition can be extended easily to simulation between states of different LMPs.

The notion of simulation meshes properly with the logic in the sense of the following proposition.

Proposition 4 ([7]) *If s simulates s', then for all formulas $\theta \in \mathcal{L}_0$, $s' \models \theta$ implies $s \models \theta$.*

In [7], it is shown that if we add disjunction to $\mathcal{L}_0$, the converse of this result is also true; that is, the simulation induced by the logic is equivalent to Definition 3.

3 Abstract Approximations

In this section we will show how to quotient a LMP by a set of $\mathcal{L}_0$ formulas. A natural candidate for the quotient kernel is to take the infimum of the original kernel over equivalent states. Unfortunately, additivity is lost in so doing. Example 15 shows why. Consequently, we are led to weaken the notion of LMPs.

3.1. Pre-LMPs

The difference between a pre-LMP and an LMP lies in the following definition.

Definition 5 *Given a measurable space (S, Σ), a function $f : \Sigma \to [0, 1]$ is called a* pre-measure *if:*
— $\forall A, B \in \Sigma$ disjoint: $f(A + B) \geq f(A) + f(B)$;
— $\forall \downarrow A_n \in \Sigma : f(\cap A_n) = \inf_n f(A_n)$.

Easy consequences of the first condition are $f(\varnothing) = 0$ and monotonicity: $A \subseteq B \Rightarrow f(A) \leq f(B)$. If one replaces the inequation in the first clause by an equation, the definition is equivalent to f being a sub-probability. This definition is weaker than Choquet capacities [1].

Definition 6 *A pre-LMP is a triple $\mathcal{S} = (S, \Sigma, h : L \times S \times \Sigma \to [0, 1])$ where (S, Σ) is a measurable space, and for all $a \in L$, $s \in S$, $A \in \Sigma$: $h_a(s, .)$ is a pre-measure, and $h_a(., A)$ is measurable.*

The intent of this definition is to use pre-LMPs as predictors for LMPs. It is not necessary that the prediction engine be of the same nature as what it tries to predict. What we are interested in is how easy it is to handle and how well it predicts. Pre-LMPs turn out to be better predictors than LMP as will be illustrated in Proposition 22.

3.2. Temporal properties

Semantics of $\mathcal{L}_0$ still makes sense with pre-LMPs.

Lemma 7 *For all pre-LMP $\mathcal{S}$ and $\theta \in \mathcal{L}_0$: $[\![\theta]\!]_\mathcal{S} \in \Sigma$.*

Proof: Easy induction on $\mathcal{L}_0$.□

To the modal operator of $\mathcal{L}_0$, namely $\langle a \rangle_r$, a family of maps is naturally associated, still written $\langle a \rangle_r : \Sigma \to \Sigma$ and called the *shifts*:

$$\langle a \rangle_r(A) \quad := \quad \{s \in S \mid h_a(s, A) \geq r\}$$

Clearly $\langle a \rangle_r(A) = h_a(., A)^{-1}([r, 1])$, and $h_a(., A)$ being measurable for all $A \in \Sigma$, one has that $\langle a \rangle_r(A) \in \Sigma$. One can also define the strict shifts as $\{a\}_r(A) := \{s \mid h_a(s, A) > r\}$, which are endomaps of Σ as well.

With this new notation: $[\![\langle a \rangle_r \theta]\!]_\mathcal{S} = \langle a \rangle_r([\![\theta]\!]_\mathcal{S})$.

Actually a much stronger statement than the lemma above can be made:

Theorem 8 *Let (S, Σ, h) be a LMP, the σ-algebra generated by $([\![\theta]\!]_\mathcal{S})_{\theta \in \mathcal{L}_0}$ is the smallest sub-σ-algebra of Σ which is stable under the shifts $\langle a \rangle_r$.*

We skip the proof (see [2]) since it is not used in the rest of the paper. Nevertheless, the theorem deserves mention because it gives purely measure-theoretic status to $\mathcal{L}_0$.

3.3 Co-simulation morphisms

The following notion of morphism between pre-LMPs will witness the relation between a process and its approximant. Recall that our goal is to define approximants as quotients of pre-LMPs under equivalence relations. Such quotients are usually related to the original process $\mathcal{S}$ through a measurable map from $\mathcal{S}$ to its quotient; this map will be proven to be a co-simulation morphism.

Definition 9 *Given $\mathcal{S}$, $\mathcal{S}'$ two pre-LMPs, a map $q : S \to S'$ is said to be a* co-simulation *iff it is surjective, measurable and for all $a \in L$, $s \in S$, $A' \in \Sigma'$:*

$$h_a(s, q^{-1}A') \geq h'_a(q(s), A').$$

Caveat: we are changing [4]'s definition of simulation morphisms, reversing the inequation and requiring surjectivity. We can thus use the proof of [4, Proposition 3.6.7] to show that co-simulation morphisms define simulation relations.

Proposition 10 *If $q : \mathcal{S} \to \mathcal{S}'$ is a co-simulation morphism, then every $s \in S$ simulates $q(s)$.*

Proof: The proof of the dual result with simulation morphisms of [4] does not use the additivity property. □

This proposition will allow us to make sure that the approximant is simulated by (or *below*) $\mathcal{S}$.

Proposition 4 can also be extended to pre-LMPs and give us the following.

Corollary 11 *Let $q : \mathcal{S} \to \mathcal{S}'$ be a co-simulation, then for all $\theta \in \mathcal{L}_0$, $s \in S$: $q(s) \in [\![\theta]\!]_{\mathcal{S}'} \Rightarrow s \in [\![\theta]\!]_{\mathcal{S}}$.*

Proof: The statement can be restated as $q^{-1}[\![\theta]\!]_{\mathcal{S}'} \subseteq [\![\theta]\!]_{\mathcal{S}}$. The proof is by induction on $\mathcal{L}_0$:
— for $\top$, one has $[\![\theta]\!]_{\mathcal{S}'} = S'$ and $q^{-1}S' = S = [\![\theta]\!]_{\mathcal{S}}$;
— $q^{-1}[\![\theta \wedge \psi]\!]_{\mathcal{S}'} = q^{-1}([\![\theta]\!]_{\mathcal{S}'} \cap [\![\psi]\!]_{\mathcal{S}'}) = q^{-1}([\![\theta]\!]_{\mathcal{S}'} \cap [\![\psi]\!]_{\mathcal{S}'}) = q^{-1}[\![\theta]\!]_{\mathcal{S}'} \cap q^{-1}[\![\psi]\!]_{\mathcal{S}'} \subseteq [\![\theta]\!]_{\mathcal{S}} \cap [\![\psi]\!]_{\mathcal{S}}$;
— if $q(s) \in [\![\langle a \rangle_r \theta]\!]_{\mathcal{S}'}$, then $s \in [\![\langle a \rangle_r \theta]\!]_{\mathcal{S}}$ because:

$$\begin{aligned} r &\le h'_a(q(s), [\![\theta]\!]_{\mathcal{S}'}) \\ &\le h_a(s, q^{-1}[\![\theta]\!]_{\mathcal{S}'}) \\ &\le h_a(s, [\![\theta]\!]_{\mathcal{S}}). \qquad \square \end{aligned}$$

3.4. The infimum construction

Proposition 13 below, which says that "one can do infs" on measurable equivalence classes, is the most important in a sense, for without it we could not construct any quotient.

Lemma 12 *Let (S, Σ) be a measurable space and $\mathfrak{R}$ be an equivalence relation on S. If there is a finite number of $\mathfrak{R}$ equivalence classes and if they are all in Σ, then for all measurable function $g : S \to [0, +\infty]$, the function $g_{\mathfrak{R}}(s) := \inf_{t \in [s]} g(t)$ is measurable.*

Proof: The inverse image under $g_{\mathfrak{R}}$ of a measurable set is a union of equivalence classes. $\square$

This result can be extended to equivalence relations that are generated by a countable subset $\mathcal{F}$ of measurable sets: that is, two elements are $\mathcal{F}$-equivalent if they belong to exactly the same sets of $\mathcal{F}$.

Proposition 13 *Given (S, Σ) a measurable space, $\mathfrak{R}$ an equivalence relation on S generated by a countable subset of Σ, then for all measurable function $g : S \to [0, +\infty]$, the function $g_{\mathfrak{R}}$ is measurable.*

Proof: Consider $\mathcal{F}_i$ a sequence of finite subsets of $\mathcal{F}$, increasing to $\mathcal{F}$. Each $\mathcal{F}_i$ defines an equivalence $\mathfrak{R}_i$ with finitely many classes, all in Σ. By the preceding lemma, $g_{\mathfrak{R}_i}$ is measurable for all i, and hence $\sup_i g_{\mathfrak{R}_i}$ is also measurable. Since $\mathfrak{R} = \cap_i \mathfrak{R}_i$, the sup is $g_{\mathfrak{R}}$. $\square$

We obtain a particular case of the construction above with $\mathfrak{R} = B \times B + B^c \times B^c$ for some $B \subset S$ and $g = \mathbf{1}_A$ for some A. If $B \subsetneq A \subset S$ (strict inclusion is important) then $g_{\mathfrak{R}} = \mathbf{1}_B$. Therefore, if $B \notin \Sigma$, $g_{\mathfrak{R}}$ is not Σ-measurable and one sees that the measurability assumption in the proposition above is essential.

3.5. Quotients and Simulations

In the rest of the paper we will simply say that $\mathfrak{R}$ is an equivalence on a given pre-LMP $\mathcal{S} = (S, \Sigma, h)$, meaning it is an equivalence on S generated by a countable subset of measurable sets.

Definition 14 *Given an equivalence $\mathfrak{R}$ on $\mathcal{S}$, we define the* quotient *pre-LMP, written $\mathcal{S}_{\mathfrak{R}}$, as the following triple $(S_{\mathfrak{R}}, \Sigma_{\mathfrak{R}}, h_{\mathfrak{R}})$:*
—$S_{\mathfrak{R}}$ is the set of $\mathfrak{R}$ equivalence classes,
—$\Sigma_{\mathfrak{R}}$ is the quotient σ-algebra isomorphic to $\Sigma \cap \Gamma_{\mathfrak{R}}$,
—$h_{\mathfrak{R}}(a, [s], A) = \inf_{t \in [s]} h_a(t, \cup A)$ for $a \in L$, $s \in S$ and $A \in \Sigma_{\mathfrak{R}}$.

When the kernel and the equivalence are compatible, then $h_{\mathfrak{R}}([s], A) = h(t, \cup A)$ for all $t \in [s]$ and the construction downs to an ordinary quotient.

We now present an example showing that this quotient does not always define an LMP.

Example 15 Consider the following LMP, where unweighted transition are of probability 1.

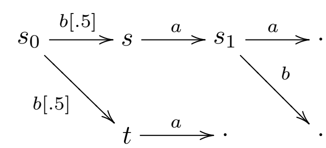

We want to quotient it with respect to the equivalence defined by all formulas of the form $\langle a \rangle_r \top$ and $\langle b \rangle_r \top$ for all $r \in [0, 1]$. The result is as follows:

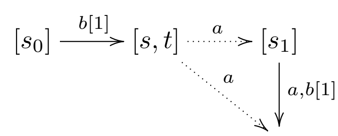

Both dotted transitions are given value 0, for $\inf_{t \in [s]} h_a(t, \cup[s_1]) = \inf_{t \in [s]} h_a(t, s_1) = 0$ and similarly for dead states. However $\inf_{t \in [s]} h_a(t, \cup[s_1, \cdot]) = 1$. Hence h is not a measure and hence the quotient is not an LMP. However, one can see that it is a pre-LMP.

Lemma 16 *$\mathcal{S}_{\mathfrak{R}}$ as defined above is a pre-LMP.*

Proof: We have three things to verify according to Definition 6 above. The first is obvious. For the second condition, the verification that $h_{\mathfrak{R}}$ is a pre-measure breaks down in two subconditions.
Super-additivity. If A, B are disjoint sets in $\Sigma_{\mathfrak{R}}$:

$$\begin{aligned} h(s, \cup(A + B)) &= h(s, \cup A + \cup B) \\ &\ge h(s, \cup A) + h(s, \cup B) \\ &\ge h_{\mathfrak{R}}([s], A) + h_{\mathfrak{R}}([s], B). \end{aligned}$$

Co-continuity. Let $\downarrow A_n$ be a decreasing sequence of sets in $\Sigma_{\mathfrak{R}}$, then $\downarrow \cup A_n$ is also a decreasing sequence in $\Sigma \cap \Gamma_{\mathfrak{R}}$ and:

$$\begin{array}{rcl} h_{\mathfrak{R}}([s], \cap A_n) & := & \inf_{t\mathfrak{R}s} h(t, \cap(\cup A_n)) \\ & = & \inf_{t\mathfrak{R}s} \inf_n h(t, \cup A_n) \\ & = & \inf_n \inf_{t\mathfrak{R}s} h(t, \cup A_n) \\ & =: & \inf_n h_{\mathfrak{R}}([s], A_n) \end{array}$$

so indeed $h_{\mathfrak{R}}(s, .)$ is a pre-measure.

Finally for the third, we verify that for all $A \in \Sigma_{\mathfrak{R}}$ and $r \in \mathbb{Q}$, the set $\{h_{\mathfrak{R}}(., A) \geq r\}$ is in $\Sigma_{\mathfrak{R}}$. Writing q for the canonical projection from S to $S_{\mathfrak{R}}$, we can write our set as:

$$\{[s] \mid h_{\mathfrak{R}}([s], A) \geq r\} = q(\{s \mid \inf_{t\mathfrak{R}s} h(t, q^{-1}A) \geq r\})$$

i.e., as the projection of a set which is clearly $\mathfrak{R}$-closed and, by Proposition 13 applied to $h(., q^{-1}A)$ (which indeed is a measurable function, since q is measurable and therefore $q^{-1}A \in \Sigma$), belongs to Σ. □

Now that we know the quotient $\mathcal{S}_{\mathfrak{R}}$ exists, we need to bring up the properties it might share with $\mathcal{S}$.

Proposition 17 *$\mathcal{S}_{\mathfrak{R}}$ is simulated by $\mathcal{S}$. Specifically, the canonical surjection $q : \mathcal{S} \to \mathcal{S}_{\mathfrak{R}}$ is a co-simulation.*

The proof is obvious.

3.6. Quotients and Logical Properties

Piecing Proposition 17 with Corollary 11, we get that each property $\mathcal{S}_{\mathfrak{R}}$ has, $\mathcal{S}$ also has.

Corollary 18 *Let $\mathfrak{R}$ be an equivalence on $\mathcal{S}$, then for all $\theta \in \mathcal{L}_0$, and $s \in S$: $[s] \in [\![\theta]\!]_{\mathcal{S}_{\mathfrak{R}}} \Rightarrow s \in [\![\theta]\!]_{\mathcal{S}}$.*

We now need a converse to this, that will quantify how good the approximation given by the quotient is, and say how much of the $\mathcal{L}_0$ properties of s in $\mathcal{S}$ are still properties of $[s]$ in $\mathcal{S}_{\mathfrak{R}}$.

Definition 19 *We will say $\mathfrak{R}$* refines *a property θ iff all interpretations of subformulas of θ are $\mathfrak{R}$-closed.*

In other words: for all $\theta' \leq \theta$ and all $(s, t) \in \mathfrak{R}$, if $s \models \theta'$ then $t \models \theta'$.

Proposition 20 *Let $\mathfrak{R}$ be an equivalence on $\mathcal{S}$, then $\forall \theta \in \mathcal{L}_0$ that $\mathfrak{R}$ refines, $s \in S$: $s \in [\![\theta]\!]_{\mathcal{S}} \Rightarrow [s] \in [\![\theta]\!]_{\mathcal{S}_{\mathfrak{R}}}$.*

Proof: The lemma can be rephrased as $q^{-1}[\![\theta]\!]_{\mathcal{S}_{\mathfrak{R}}} \supseteq [\![\theta]\!]_{\mathcal{S}}$. We prove it by induction on θ, the only interesting case being $\theta = \langle a \rangle_r \psi$, and then for all a, s:

$$\begin{array}{rcl} h'_a([s], [\![\psi]\!]_{\mathcal{S}_{\mathfrak{R}}}) & = & \inf_{t \in [s]} h_a(t, q^{-1}[\![\psi]\!]_{\mathcal{S}_{\mathfrak{R}}}) \\ & = & \inf_{t \in [s]} h_a(t, [\![\psi]\!]_{\mathcal{S}}), \end{array}$$

where the second equation is by induction (since $\mathfrak{R}$ refines also ψ, $q^{-1}[\![\psi]\!]_{\mathcal{S}_{\mathfrak{R}}} \supseteq [\![\psi]\!]_{\mathcal{S}}$ and by the corollary above, these two subsets of S are actually equal). It follows that if $s \models \theta$ and $[s] \not\models \theta$, there must be a $t \in [s]$ close enough to the infimum, such that $t \not\models \theta$ either, which means $\mathfrak{R}$ actually does not refine θ. □

Subformulas *have* to be included in the refinement condition and this can be seen on a small example. Say L is reduced to the empty symbol and $r_2 \geq r_1$, here are $\mathcal{S}$ and a quotient $\mathcal{S}_{\mathfrak{R}}$:

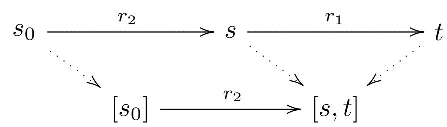

Now set $\theta = \langle\rangle_{r_2} \langle\rangle_{r_1} \top$; one has:

$$\begin{array}{l} [\![\langle\rangle_{r_1} \top]\!]_{\mathcal{S}} = \{s, s_0\},\ [\![\theta]\!]_{\mathcal{S}} = \{s_0\}, \\ [\![\langle\rangle_{r_1} \top]\!]_{\mathcal{S}_{\mathfrak{R}}} = \{[s_0]\},\ [\![\theta]\!]_{\mathcal{S}_{\mathfrak{R}}} = \varnothing, \end{array}$$

so though $[\![\theta]\!]_{\mathcal{S}}$ is $\mathfrak{R}$-closed and $s_0 \in [\![\theta]\!]_{\mathcal{S}}$, yet $q_{\mathfrak{R}}(s_0) = [s_0] \notin [\![\theta]\!]_{\mathcal{S}'_{\mathfrak{R}}}$.

Combining the last two statements in the particular case of logically generated approximations, we get:

Theorem 21 (abstract approximants) *Let $\mathcal{F}$ be a downward closed subset of $\mathcal{L}_0$, and $\mathfrak{F}$ be the associated equivalence on $\mathcal{S}$:*

$$\forall \theta \in \mathcal{F}, \forall s \in S : s \in [\![\theta]\!]_{\mathcal{S}} \Leftrightarrow [s]_{\mathfrak{F}} \in [\![\theta]\!]_{\mathcal{S}_{\mathfrak{F}}}.$$

Proof: $\mathfrak{F}$ is an equivalence on $\mathcal{S}$ because it is generated by a countable subset of Σ, and hence Corollary 18 applies. Now $\mathfrak{F}$ refines $\mathcal{F}$ and is even the coarsest such equivalence, and one may apply Proposition 20. □

A particular case is $\mathcal{F} = \{\top\}$, and then $\mathcal{S}_{\mathfrak{F}} = (\{*\}, \{\varnothing, \{*\}\}, h_a)$ with $h_a(*, \varnothing) = 0$ and $h_a(*, \{*\}) = \inf_{s \in S} h_a(s, S) =: \alpha_a$. So $\mathcal{S}_{\mathfrak{F}}$ is the loop with coefficients $(\alpha_a)_{a \in L}$. Of course very few properties are retained here, namely the combinations of $\langle a \rangle_r$ with $r \leq \alpha_a$. This trivial approximation can be thought of as a quite blunt abstract interpretation of $\mathcal{S}$. The theorem above explains, in essence, how to construct arbitrarily sharper ones.

3.7. Abstract approximants are optimal

A noteworthy observation is that if one wants the quotient map q to generate a simulation, the choice made above for $h_{\mathfrak{R}}$ is optimal, $\Sigma_{\mathfrak{R}}$ is the largest σ-algebra that will make q measurable and all other kernels would be pointwise smaller:

Proposition 22 *Given $\mathcal{S}$, $\mathcal{S}'$ two pre-LMPs and $q : \mathcal{S} \to \mathcal{S}'$ a co-simulation morphism, and defining the*

equivalence relation generated by q on S as $(s,t) \in \mathfrak{R}$ iff $q(s) = q(t)$, one has:
1) Σ' is a sub-σ-algebra of $\Sigma_{\mathfrak{R}}$;
2) for all $s' \in S'$, $A' \in \Sigma'$: $h'(s', A') \leq h_{\mathfrak{R}}(s', A')$.

Yet another way of saying this is: the identity $\iota : \mathcal{S}_{\mathfrak{R}} \to \mathcal{S}'$ is a co-simulation which decomposes q as $\iota \circ q_{\mathfrak{R}}$. The proof of 1) is left to the reader. Point 2) is obvious. Notice that $\mathfrak{R}$ is an equivalence on $\mathcal{S}$ because it is generated by sets of the form $q^{-1}[r, \infty)$ where $r \in [0,1]$ is rational.

For one thing, pre-LMP support what seems the natural construction, as summarized in Theorem 21, whereas with plain LMPs one has to restrict to finite quotients. Moreover, the approximant construction of [7] needed a logic $\mathcal{L}_0$ where the inequality sign was strict in the semantic of the formula $\langle a \rangle_q \phi$. Secondly, pre-LMP also give more accurate finite predictors in the sense of the last proposition.

3.8 The sup-quotient

We have shown how to construct quotients of pre-LMPs using infima of measurable functions. One could be interested in the dual construction using suprema. All the results above can be dualized to their sup counterpart with little modification. Basically, one has to reverse inequality signs and replace co-continuity with continuity. The resulting model could be called conveniently *sub-pre-LMP*, since suprema generate subadditive kernels (and our pre-LMPs shoud then be called super-pre-LMPs since they have super-additive kernels as we know). The quotient of a sub-pre-LMP is above the original process instead of below. Consequently, we have a simulation morphism instead of a co-simulation in the equivalent of Proposition 17. The semantics of $\langle a \rangle_r$ has to be adapted as well, $\langle a \rangle_r(A)$ meaning now the set of states having a probability *strictly* greater than r to jump in A.

4 Stronger Approximants

4.1. Extended logic

We introduce an extended logic $\mathcal{L}_0^*$ to capture cyclic temporal properties. To deal with mutual fixpoint equations, extended formulas are conveniently presented in automaton-style.

Definition 23 (cyclic temporal properties) *An $\mathcal{L}_0^*$ formula is a pair (I, λ), with I a finite indexing set and λ a partial map from $L \times I \times I$ to $[0,1]$.*

We write $\mathrm{dom}(\lambda)$ for the domain of λ; working with total maps, by extending λ to be zero outside $\mathrm{dom}(\lambda)$, turns out to be inconvenient. We will use freely the automaton terminology and talk about I as the state space and λ as the transition map. Notice that there is no condition on the transition function: it may not be a subprobability distribution. One should understand the transitions as if they were non-deterministic ones.

Definition 24 (mapping $\mathcal{L}_0$ to $\mathcal{L}_0^*$) *One defines a map $(.)^*$ from $\mathcal{L}_0$ to $\mathcal{L}_0^*$ as follows:*
— I is the set of θ's (occurrences of) maximal conjunctive sub-formulas,
— $\lambda(a, \theta_0, \theta_1) = r$ iff $\theta_0 = \langle a \rangle_r \theta_1 \wedge \theta'$ for some θ', up to the monoidal equations associated to $\wedge$.

So, for instance:
—$\top^* = (\{\top\}, \varnothing)$,
—$(\langle a \rangle_{.5} \top)^* = (\{\langle a \rangle_{.5}\top, \top\}, \{(a, \langle a \rangle_{.5}, \top, .5)\})$.

i.e. $= \langle a \rangle_{.5} \top \xrightarrow{a[0.5]} \top$

—$(\langle a \rangle_1 \phi \wedge \langle b \rangle_{.5} \top)^* =$

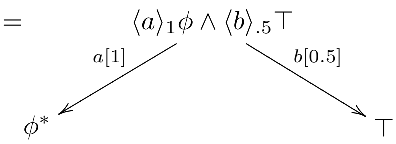

This correspondence is one-one, up to monoidal equations, and θ^* is always a tree.

Now, given $\mathcal{S}$ an LMP, we would like to extend the map $[\![.]\!]_{\mathcal{S}}$ to $\mathcal{L}_0^*$-formulas, or in other words, to make sense of $s \models \theta$ for our new formulas. This will be done using two independent approaches that will prove equivalent. One will be the definition of a suitable fixpoint in the category $\mathbf{C}_{\mathcal{S}}$ defined below, and the other one will be with the help of simulation relations.

Semantics of $\mathcal{L}_0^*$ via fixpoints. Let $\mathbf{C}_{\mathcal{S}}$ be the sub-Cartesian category of **Set** generated by:
— shifts $\langle a \rangle_r : \Sigma \to \Sigma$,
— intersection $\cap : \Sigma \times \Sigma \to \Sigma$.

Products in $\mathbf{C}_{\mathcal{S}}$ are *ordinary* set-theoretic products, not products of measurable spaces, and Σ^ns are ordered with the product ordering: $(A_1, \ldots, A_n) \leq (B_1, \ldots, B_n)$ if for all $1 \leq i \leq n$, $A_i \subseteq B_i$. If one restricts to shifts with rational coefficients, there is only a countable number of arrows left in $\mathbf{C}_{\mathcal{S}}$.

The key to the extension of $[\![.]\!]_{\mathcal{S}}$ is the following:

Lemma 25 *Morphisms of $\mathbf{C}_{\mathcal{S}}$ are all monotonic and co-continuous; endomorphisms of $\mathbf{C}_{\mathcal{S}}$ all have greatest fixpoints.*

Proof: First of all, we observe that shifts are indeed returning in Σ by definition of a pre-LMP. Secondly, all generators are clearly monotonic increasing. Thirdly, if $\downarrow A_n$ is a decreasing sequence in Σ then:

$$\begin{aligned}\langle a \rangle_r(\cap A_n) &= \{s \mid h_a(s, \cap A_n) \geq r\} \\ &= \{s \mid \inf_n h_a(s, A_n) \geq r\} \\ &= \cap \langle a \rangle_r(A_n)\end{aligned}$$

where the second equation uses $h_a(s,.)$ co-continuity on Σ, given by definition of pre-LMP kernels. So shifts are co-continuous, and so are evidently projections, intersections and all cartesian combinations of them.

Lastly, suppose ψ is an endomorphism, since it is monotonic increasing it has a greatest fixpoint in $(2^{\mathcal{S}})^n$, and since ψ is also co-continuous, this fixpoint can be written as $\cap_n \psi^n(S,\ldots,S)$ and hence is in Σ^n. $\square$

Actually, $\mathbf{C}_{\mathcal{S}}$ has a structure of traced Cartesian category (or Cartesian category with fixpoints as discussed for instance in [8]). We will write $\mathsf{Y}\psi$ for ψ's fixpoint.

More generators could be added to the lot while keeping the key lemma above. We could have added unions, countable unions and countable intersections to the generators (and therefore the countable power $\Sigma^{\mathbb{N}}$ as an object of $\mathbf{C}_{\mathcal{S}}$). This might indeed prove useful at some later stage, but for now we don't do this. We could *not* have added maps such as:

$$\psi(A) = \langle a\rangle_r(A) \setminus \langle a\rangle_{r'}(A) = \{s \mid h_a(s,A) \in [r,r')\}$$

which is not monotonic; having only positive operators in the basic logic is crucial here. More subtly, strict shifts which are monotonic cannot be added because they're not co-continuous and we need greatest fixpoints (as made clear below).

So, strict shifts are not continuous and neither are shifts continuous. Here is an example: $([0,1],\mathcal{B},h)$ with $h_a(s,B) = \lambda(B)$, where λ is Lebesgue measure on $\mathcal{B}$, and while $\cup_n[0,1-1/n] = [0,1]$, but for no n can we be Lebesgue sure to hit $[0,1-1/n]$, there is always a $1/n$ chance that we do not, and so $\langle a\rangle_1([0,1-1/n]) = \varnothing$. A similar case can be made against strict shifts being co-continuous with intervals $\downarrow\!\!(0,1/n]$.

Definition 26 *Given $\theta = (I,\lambda) \in \mathcal{L}_0^*$, we define in turn $\{\!|\theta|\!\}_{\mathcal{S}} \in \mathbf{C}_{\mathcal{S}}[\Sigma^I,\Sigma^I]$ and $[\![\theta]\!]_{\mathcal{S}} \in \Sigma^I$ by:*

$$\{\!|\theta|\!\}_{\mathcal{S}}(\tau)(i) := \bigcap_{(a,i,j)\in dom(\lambda)} \langle a\rangle_{\lambda(a,i,j)}(\tau(j)),$$

with $\tau \in \Sigma^I$ a I-indexed tuple in Σ, and:

$$[\![\theta]\!]_{\mathcal{S}} := \mathsf{Y}\{\!|\theta|\!\}_{\mathcal{S}} = \cap_p \downarrow\!\!\{\!|\theta|\!\}^p_{\mathcal{S}}(S,\ldots,S).$$

This somewhat pedantic notation comes handy when one wants to access states by their names, not their indices. We will use concrete tuple notation in examples, but not in proofs. Symbol τ sounds like "tuple" and is supposed to be suggestive of what τ is: a tuple. When dom(λ) is empty (which happens exactly when the corresponding state is dead in θ), we take the convention that the intersection is equal to the full set S.

Each component map $\lambda\tau.\{\!|\theta|\!\}_{\mathcal{S}}(\tau)(i)$ is in $\mathbf{C}_{\mathcal{S}}[\Sigma^I,\Sigma]$ indeed, since it is clearly expressed as a finite intersection of shifts; therefore the lemma above applies, and $[\![\theta]\!]_{\mathcal{S}} = \mathsf{Y}\{\!|\theta|\!\}_{\mathcal{S}}$ is well-defined and lies in Σ^I:

Lemma 27 *For all $\theta = (I,\lambda) \in \mathcal{L}_0^*$, $[\![\theta]\!]_{\mathcal{S}} \in \Sigma^I$.*

Least fixpoints are not interesting here, since one has to use strict shifts to have them in Σ, but $\{a\}_r(\varnothing) = 0$ for all pre-LMPS, so these would always be empty.

Example 28 Here is a simple LMP example $\mathcal{S}$ with state space $S = \{s_0,s_1,s_2\}$ followed by a cyclic formula θ in $\mathcal{L}_0^*$:

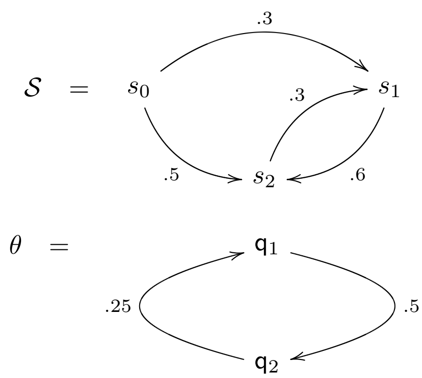

The fixpoint converges in two steps:

$$\{\!|\theta|\!\}_{\mathcal{S}}(A_1,A_2) = (\langle\rangle_{.5}A_2, \langle\rangle_{.25}A_1),$$
$$\mathsf{Y}\{\!|\theta|\!\}_{\mathcal{S}} = \{\!|\theta|\!\}^2_{\mathcal{S}}(S,S) = (\{s_0,s_1\},\{s_0,s_2\}),$$
$$[\![\theta]\!]_{\mathcal{S}}(q_1) = \{s_0,s_1\}.$$

Intuitively $\mathsf{Y}\{\!|\theta|\!\}_{\mathcal{S}}$ is finding the biggest state-sets in $\mathcal{S}$ showing the behaviour described by θ.

Now with Definition 26 and a formula $\theta \in \mathcal{L}_0$, we can build both $[\![\theta]\!]_{\mathcal{S}}$ and $[\![\theta^*]\!]_{\mathcal{S}}$, so obviously we have to say something ! (Reminder: formulas are used as their own indexing sets when coerced in $\mathcal{L}_0^*$.)

Lemma 29 *Definitions 26 and 2 of $[\![\theta]\!]_{\mathcal{S}}$ agree, in the sense that for all $\theta \in \mathcal{L}_0$: $[\![\theta^*]\!]_{\mathcal{S}}(\theta) = [\![\theta]\!]_{\mathcal{S}}$.*

Proof: The proof is an induction on $\mathcal{L}_0$, where we prove in addition that the fixpoint $\mathsf{Y}\{\!|\theta|\!\}_{\mathcal{S}}$ is obtained in $|\theta|$ steps (and therefore "convergence time" for $\mathcal{L}_0$ formulas is independent of $\mathcal{S}$).
— $\theta = \top$: then $I = \{\top\}$, $\lambda = \varnothing$ and $\{\!|\top^*|\!\}_{\mathcal{S}}(\tau)(\top) = S$, $[\![\top^*]\!]_{\mathcal{S}}(\top) = S$ which is the correct answer obtained in $0 = |\top|$ steps;
— $\theta = \theta_0 \wedge \theta_1$: $I = I_0 \cdot I_1$ is the smashed sum of I_0 and I_1, obtained by fusing the initial states θ_0 and θ_1 into θ (since θ_0 and θ_1 are no longer maximal conjunctive) and taking the disjoint union otherwise; the only state where λ changes value is precisely θ itself, and

$\lambda(a,\theta,i) = \lambda_0(a,\theta_0,i) + \lambda_1(a,\theta_1,i)$; so that, by definition:

$$\{\!\!\{\theta^*\}\!\!\}_{\mathcal{S}}(\tau_0 \cdot \tau_1)(\theta) = \{\!\!\{\theta_0^*\}\!\!\}_{\mathcal{S}}(\tau_0)(\theta_0) \cap \{\!\!\{\theta_1^*\}\!\!\}_{\mathcal{S}}(\tau_1)(\theta_1)$$
$$[\![\theta^*]\!]_{\mathcal{S}}(\theta) = [\![\theta_0^*]\!]_{\mathcal{S}}(\theta_0) \cap [\![\theta_1^*]\!]_{\mathcal{S}}(\theta_1)$$

and the answer is obtained in $\max(|\theta_0|, |\theta_1|)$.
— $\theta = \langle a\rangle_r\theta_0$: $I = I_0 + \{\theta\}$; λ takes now one more value, namely $\lambda(a,\theta,\theta_0) = r$ and:

$$\{\!\!\{\theta^*\}\!\!\}_{\mathcal{S}}(\tau)(\theta) = \langle a\rangle_r(\tau(\theta_0)),$$

again the correct answer, and obtained in $|\theta_0|+1$ steps, as expected. □

Semantics of $\mathcal{L}_0^*$ via simulations. The fixpoint definition of $\mathcal{L}_0^*$'s semantics, while being optimal for measure-theoretic considerations, shows a bit clumsy when it comes to the rest. So we bring in a nice and easy rephrasing of $s \in [\![\theta]\!]_{\mathcal{S}}(i)$.

Observe that when we say that a state satisfies a logical property, we expect this state to satisfy *at least* this property, and that it may satisfy other properties as well. Now that our properties are stated in a labelled transition setting, it is tempting to use the corresponding algebraic notion, that is, simulation. Indeed, if we look back to Example 28, we can observe that (the reflexive and transitive closure of) the relation $(q_1,s_0),(q_1,s_1),(q_2,s_0),(q_2,s_2)$ is a simulation relation.

Definition 3 of simulation must be extended to include systems that are not pre-LMPs. Recall that even if we see formulas of $\mathcal{L}_0^*$ as automata, they are not pre-LMPs because of the fact that a-transitions probabilities may sum up to some number > 1. This problem is easily rapped out by considering formulas as non-deterministic systems, and thus every transitions as being distinct sub-probability transitions.

Definition 30 *Given $\theta = (I,\lambda) \in \mathcal{L}_0^*$, $\mathcal{S} = (S,\Sigma,h)$ a pre-LMP, a relation $\mathfrak{S} \subseteq I \times S$ is a* non-deterministic simulation *if for all $a \in L$, $(i,s) \in \mathfrak{S}$ and $j \in I$: $\lambda(a,i,j) \le h_a(s,\mathfrak{S}(j))$.*

It is understood above that $\forall i$, $\mathfrak{S}(i) \in \Sigma$. However, this requirement is not in Definition 3 of simulation between LMPs, for if it was, bisimulation would not be a simulation. This issue is technically complex and not addressed here. The reader must keep in mind that this definition of simulation is safe only if we deal with countable state-space processes or if we manipulate simulated processes related by a co-simulation morphism, as will be argued in Lemma 32 below[1].

[1] We conjecture that requiring that $\mathfrak{S}(i) \in \Sigma$ be an analytic set in $\mathcal{S}$ would solve the problem.

Proposition 31 *Given $\theta = (I,\lambda) \in \mathcal{L}_0^*$, $\mathcal{S} = (S,\Sigma,h)$ a pre-LMP, $\mathfrak{S} \subseteq I \times S$ is a simulation iff for all i:*
— $\mathfrak{S}(i) \in \Sigma$,
— $\mathfrak{S}(i) \subseteq \{\!\!\{\theta\}\!\!\}_{\mathcal{S}}(\lambda j.\mathfrak{S}(j))(i)$.

Proof: The proof is by trivial manipulation of the various definitions.□

Now say s *simulates* i, when there exists a non-deterministic simulation $\mathfrak{S}$, with $(i,s) \in \mathfrak{S}$; and here is the rephrasing: $s \in [\![\theta]\!]_{\mathcal{S}}(i)$, or $s \models \theta(i)$ in shorthand notation, iff s simulates i. We also observe that $[\![\theta]\!]_{\mathcal{S}}$, seen as a relation on $I \times S$, is the coarsest simulation.

4.2. Quotients with $\mathcal{L}_0^*$

We can now prove the analog of Corollary 11: co-simulation morphisms preserve formulas of $\mathcal{L}_0^*$.

Lemma 32 *Let $\mathcal{S}$, $\mathcal{S}'$ be pre-LMPs and $q : \mathcal{S} \to \mathcal{S}'$ be a co-simulation morphism, then for all $\theta = (I,\lambda) \in \mathcal{L}_0^*$, $i \in I$, $s \in S$: $q(s) \in [\![\theta]\!]_{\mathcal{S}'}(i) \Rightarrow s \in [\![\theta]\!]_{\mathcal{S}}(i)$.*

Proof: Composing a non-deterministic simulation on $I \times S'$ with the co-simulation q gives a simulation for $\mathcal{S}$. □

And it remains now to prove the analog of Proposition 20, that is, that quotient states satisfy the same formulas of $\mathcal{F}$ as the states they $\mathcal{F}$-approximate. But before we have to explain what it means now for an equivalence $\mathfrak{R}$ over $\mathcal{S}$ to refine a formula $\theta \in \mathcal{L}_0^*$.

Definition 33 *Let $\mathcal{S}$ be a pre-LMP, $\mathfrak{R}$ be an equivalence over $\mathcal{S}$, and $\theta = (I,\lambda) \in \mathcal{L}_0^*$, then $\mathfrak{R}$ refines θ if for all $i \in I$, $[\![\theta]\!]_{\mathcal{S}}(i)$ is $\mathfrak{R}$-closed.*

By Lemma 29, this second definition coincides with the definition given before for $\mathcal{L}_0$ (to be exact, only maximal conjunctive subformulas have to be $\mathfrak{R}$-closed, so next proposition is marginally better).

With our definition in place we can home in on our proposition:

Proposition 34 *Let $\mathcal{S}$ be a pre-LMP and $\mathfrak{R}$ be an equivalence on $\mathcal{S}$ which refines θ, then for all $i \in I$, $s \in S$: $s \in [\![\theta]\!]_{\mathcal{S}}(i) \Rightarrow [s]_{\mathfrak{R}} \in [\![\theta]\!]_{\mathcal{S}_{\mathfrak{R}}}(i)$.*

Proof: Let $\mathfrak{R}$ be an equivalence relation refining θ, and assume that $s \models \theta(i)$. Then there is an associated simulation relation $\mathfrak{S}$ between θ and $\mathcal{S}$ such that for all $i\mathfrak{S}s$, if $\lambda(a,i,j) = r$ then $h_a(s,\mathfrak{S}(j)) \ge r$. Now let us prove that the corresponding relation between θ and $\mathcal{S}_{\mathfrak{R}}$ is a simulation relation. This relation $\mathfrak{R}^*$ is defined as $i\mathfrak{R}^*[s]$ if $i\mathfrak{R}'t$ for all $t \in [s]$. Since $\mathfrak{R}$ refines θ, and by definition of $\mathfrak{R}^*$ in terms of $\mathfrak{S}$, we have that $\mathfrak{S}(j) = q^{-1}\mathfrak{R}^*(j)$ where q is the quotient function. But now if $\lambda(a,i,j) = r$ then $h_a(t,\mathfrak{S}(j)) \ge r$ for all $t \in [s]$ and hence $h_{\mathfrak{R}}(a, [s], \mathfrak{R}^*(j)) = \inf_{t\in[s]} h_a(t, q^{-1}\mathfrak{R}^*(j)) = \inf_{t\in[s]} h_a(t, \mathfrak{S}(j)) \geq r$. $\square$

We can neatly summarize the results of this section in a statement paralleling Theorem 21:

Theorem 35 (strong approximants) *Let $\mathcal{S}$ be a pre-LMP, $\mathcal{F}$ be a subset of $\mathcal{L}_0^*$, and $\mathfrak{F}$ the associated equivalence on $\mathcal{S}$, then for all $\theta = (I, \lambda) \in \mathcal{F}$, $s \in S$ and $i \in I$:*

$$s \in [\![\theta]\!]_{\mathcal{S}}(i) \Leftrightarrow [s]_{\mathfrak{F}} \in [\![\theta]\!]_{\mathcal{S}_{\mathfrak{F}}}(i).$$

Proof: As in the parallel statement, $\mathfrak{F}$ equivalence classes are in Σ, because $\mathfrak{F}$ has countably many generators, namely the $[\![\theta]\!]_{\mathcal{S}}(i)$, for $\theta \in \mathcal{F}$, $i \in I_\theta$. So again $\mathfrak{F}$ is an equivalence on $\mathcal{S}$, the quotient $\mathcal{S}_{\mathfrak{F}}$ is well-defined by Lemma 16, and the projection is a co-simulation morphism by Proposition 17, so Lemma 32 applies, and this gives the left to right implication. Besides and by definition, $\mathfrak{F}$ is the coarsest equivalence on $\mathcal{S}$ refining all θs in $\mathcal{F}$, so one may apply Proposition 34 and obtain the other implication. $\square$

Note that, even if $\mathcal{S}$ is itself infinite state, the quotient will be finite, as soon as $\mathcal{F}$ is, just as in the $\mathcal{L}_0$ case.

The following result, which now follows easily, is one of the main motivations for using a logic with loops. Finite-state processes are eventually approximated by themselves up to bisimulation (see [5] for bisimulation for LMPs and characterization of bisimulation). We first need to prove that simulation relation between finite LMPs preserve formulas of $\mathcal{L}_0^*$.

Lemma 36 *If two states s and t of a finite LMP are related by a simulation relation, then every formula of $\mathcal{L}_0^*$ that s satisfies is also satisfied by t.*

Proof: The proof lies on simple manipulations of relations and inequalities and on the fact that every set in a finite LMP is measurable.$\square$

Theorem 37 *For every finite-state LMP, there is a finite set of formulas $\mathcal{F}$ of $\mathcal{L}_0^*$ such that the quotient with respect to $\mathcal{F}$ is bisimilar to the process itself.*

Proof: The logic $\mathcal{L}_0^*$ clearly characterizes bisimulation of LMPs[2]. Indeed, it is an extension of $\mathcal{L}_0$ and since simulation preserves satisfaction of formulas of $\mathcal{L}_0^*$ (by the preceding lemma), so does bisimulation. This implies that if two states are not bisimilar, then there is a formula of $\mathcal{L}_0^*$ that will distinguish them. There are finitely many pairs, and taking all formulas that distinguish pairs of non-bisimilar states and closing this set under subformulas yields a finite set of formulas. This set defines a quotient which is bisimilar to the original finite-state process. Indeed, since non-bisimilar states belong to different equivalence classes, we have that every state of the quotient is made of bisimilar states of the original process. These states have the same transition probability to every bismulation-closed set, and hence to every equivalence class. $\square$

[2]Note that for uncountable processes, this result needs an assumption that the state-space is analytic.

5 Conclusion

We have brought in the theory of LMPs two simple ideas: first, LMP approximants should be quotients with respect to the LMP bisimulation logic $\mathcal{L}_0$, yielding *stronger* approximants. Second, the same quotient construction, supposing there is one, should be doable with a logic enriched with greatest fixpoints and produce families of approximants sharing cyclic behaviours with the approximation target, resulting in a *faster* approximation construction, since finite processes are approximated with themselves at some finite stage.

Not only do these two ideas carry through but some of the known constructions and definitions have to be relaxed in so doing, resulting in pretty harmonious mathematics.

We believe that the present work is an important step towards model-checking LMPs. For example, if one knows what are the properties that a given continuous process should satisfy, one would prefer to check for these properties on a finite *faithful* approximant of the process instead of checking each property on the process itself. Our construction achieves this goal since it theoretically ensures exact satisfaction of formulas.

Observe that in Example 15, if one was interested specifically in the initial state s_0 one could live with the approximant: $[s_0] \xrightarrow{b} [s, t]$ because $[s_0]$ is equivalent to s_0, if we consider only formulas of depth 1 – of course there is a loss for other states like s and t which are *not* equivalent to $[s, t]$. This suggests that there may still be a way of quotienting with formulas and obtain an LMP. We want to investigate this possibility, which we think will be a fairly easy task. More interestingly, observe that the quotient we have defined does not depend only on the satisfied formulas, we crucially use probability information from the system itself. This implies that two processes that are $\mathcal{F}$-equivalent may not have the same quotient. We plan to investigate the possibility of using the quotient construction without using the actual values of the transition probabilities in the original process, but only values provided by formulas that are satisfied. An important application of this would be a way to construct a process by using only the formulas that it has to satisfy, that is, the automated design of probabilistic models from specifications. We believe that in this case, there will be no LMP that will satisfy the same property (even for finite quotients), showing that pre-LMPs are essential for the design of probabilistic systems.

On the practical side, the effective construction of these pre-LMPs could be costly in time or inconvenient. One has to choose a set of formulas which will be used to quotient the state space. A pre-LMP is then produced by computation of an inf from every equivalence class to possibly *every* union of states (precisely: unions which are defined with formulas). This last step increases complexity significantly. However, we do have an even faster version of approximants which works exactly as the construction in [7] (and hence produce LMPs) except that it allows loops in the approximants. The loss from the present work is that one does not have the choice of formulas, except for their depth and a desired precision. The same properties are satisfied: every formula satisfied by a state is eventually satisfied by the state's approximant, and finite processes are eventually approximated by themselves.

$\mathcal{S}_{\_}$ as a Functor Let $\mathcal{S}$ be a fixed LMP, then if $\mathfrak{R}_1 \subset \mathfrak{R}_2$, one can define a co-simulation morphism $q_{12} : \mathcal{S}_{\mathfrak{R}_1} \to \mathcal{S}_{\mathfrak{R}_2}$ in the obvious way: $q_{12}([s]_1) = [s]_2$, an assignment which is independent of the choice of $s \in [s]_1$. This correspondence is functorial from the poset of equivalences over $\mathcal{S}$ ordered by inclusion to the category of quotients over $\mathcal{S}$ (with co-simulation as morphisms). We further observe that when $\mathfrak{R}_1 \subset \mathfrak{R}_2$, $\mathfrak{R}_2$ can also be viewed as an equivalence over $\mathfrak{R}_1$ since its classes are $\mathfrak{R}_1$-closed. Therefore it makes sense to define $(\mathcal{S}_{\mathfrak{R}_1})_{\mathfrak{R}_2}$ and expectedly $(\mathcal{S}_{\mathfrak{R}_1})_{\mathfrak{R}_2} = \mathcal{S}_{\mathfrak{R}_2}$. This functor can be precomposed with the contravariant functor from the poset of *finite* subformula-closed subsets of $\mathcal{L}_0$ to the poset of equivalences over $\mathcal{S}$, thus obtaining the *logical* functor $\mathcal{L}$. We leave the question of whether $\mathcal{S}$ is the category-theoretic limit of the functor $\mathcal{L}$ (with target category that of LMPs taken up to bisimulation and co-simulation morphisms) to further investigation. Addressing this problem might demand some of the analytic space machinery to be put back in the picture.

Ongoing research is also trying to apply this theory of approximants to many probabilistic models like continuous time Markov chains and to extend it to richer logic. One potential application are Markov Decision Processes that one finds in the machine learning field. Approximants have been studied in this field, but always with a focus on partitioning the state-space without taking account of the behaviour of processes, that is, of the actual transitions that states can take. As a result, bisimilar or behaviourally close states can be split in the process, whereas our constructions always partition the state-space with respect to satisfaction of formulas. This is the only related work that we are aware of, as far as approximants are concerned.

Concluding with a different matter, we have also explored a more probablistically-minded approach to approximation, based on conditional expectations. Given a probability p on (S, Σ), and a sub-σ-algebra Σ' of Σ, it is possible to define the conditional expectation given Σ' of any integrable function according to p. Applied to finite-state systems, the idea downs to taking the quotient kernel to be an average rather than an infimum. We have showed that this construction works and gives access to more robust approximations (see [3]).

References

[1] G. Choquet. Theory of capacities. *Ann. Inst. Fourier (Grenoble)*, 5:131–295, 1953.

[2] V. Danos and J. Desharnais. Note sur les chaînes de Markov étiquetées. Unpublished (in french). Available from `www.logique.jussieu.fr/www.danos`, 2002.

[3] V. Danos, J. Desharnais, and P. Panangaden. Conditional expectation and the approximation of labeled Markov processes. Available from `www.logique.jussieu.fr/www.danos`, 2003.

[4] J. Desharnais. *Labelled Markov Processes.* PhD thesis, McGill University, November 1999.

[5] J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled Markov processes. *Information and Computation*, 179(2):163–193, Dec 2002. Available from `http://www.ift.ulaval.ca/~jodesharnais`.

[6] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labeled Markov processes. In *Proceedings of CONCUR99*, Lecture Notes in Computer Science. Springer-Verlag, 1999.

[7] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating continuous Markov processes. In *Proceedings of the 15th Annual IEEE Symposium On Logic In Computer Science, Santa Barbara, Californie, USA*, 2000. pp. 95-106.

[8] M. Hasegawa. Recursion from cyclic sharing: traced monoidal categories and models of cyclic lambda calculi. In *3rd International Conference on Typed Lambda Calculi and Applications (TLCA'97)*, volume 1210 of *LNCS*, pages 196–213. Springer, 1997.

[9] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM (JACM)*, 32(1):137–161, 1985.

[10] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes I and II. *Information and Computation*, 100:1–41, 42–78, 1992.

[11] D. Williams. *Probability with Martingales.* CUP, Cambridge, 1991.