

Abstract interpretation of programs as Markov decision processes

David Monniaux

<http://www.di.ens.fr/~monniaux>
École Normale Supérieure
Département d'Informatique
45, rue d'Ulm
75230 Paris cedex 5
France

Abstract. We propose a formal language for the specification of trace properties of probabilistic, nondeterministic transition systems, encompassing the properties expressible in Linear Time Logic. Those formulas are in general undecidable on infinite deterministic transition systems and thus on infinite Markov decision processes. This language has both a semantics in terms of sets of traces, as well as another semantics in terms of measurable functions; we give and prove theorems linking the two semantics. We then apply abstract interpretation-based techniques to give upper bounds on the worst-case probability of the studied property. We propose an enhancement of this technique when the state space is partitioned — for instance along the program points —, allowing the use of faster iteration methods.

1 Introduction

The study of probabilistic programs is of considerable interest for the validation of networking protocols, embedded systems, or simply for compiling optimizations. It is also a difficult matter, due to the undecidability of properties on infinite-state deterministic programs, as well as the difficulties arising from probabilities. In this paper, we provide methods for the analysis of programs represented as infinite-state Markov decision processes.

The analysis of *finite-state* Markov decision processes was originally conducted in the fields of operational research and finance mathematics [23]. More recently, they have been studied from the angle of probabilistic computing systems [1–4, 15, 24]. Effective resolution techniques include linear programming [23, §7.2.7] [7] and newer data structures such as MTBDDs [2]. However, the problem of large- or infinite-state systems has not been so well studied.

In the case of deterministic or nondeterministic systems without a notion of probability, various analysis techniques have been proposed in the last twenty years. Since the problem is undecidable, those techniques are either partially manual (i.e. require the input of *invariants* or similar), either *approximate* (i.e., the analysis takes a pessimistic point of view when it cannot solve the problem

exactly). In this paper, we take the latter approach and build our analysis methods upon the existing framework of *abstract interpretation* [9], a general theory of approximation between semantics.

We have earlier proposed two classes of automatic methods to analyze such system: some *forward* [16,17], some *backward* [19,20]. In this paper, we focus on the backward approach and extend it to a larger class of properties (including those specified by LTL formulas). We also prove that chaotic iterations strategies [8, §2.9] apply to our case, which allows parallel implementations.

In section 2, we give an introduction to probabilistic transition systems, which we extend in section 3 to nondeterministic and probabilistic systems. In section 4, we give a formal language for the specification of trace properties, including those formulated using Büchi or Rabin automata. In section 5, we explain how to analyze those properties backward and in section 6.1 how to apply abstract analyses.

2 Probabilistic transition systems

The natural extension of transition systems to the probabilistic case is *probabilistic transition systems*, also known as *Markov chains* or *discrete-time Markov process*.

2.1 Probabilistic transitions

We assume that the set of states is *finite* or *countable* so as to avoid technicalities. The natural extension of the notion of deterministic state is the notion of probability distribution on the set of states.

Definition 1 *Let Ω be a finite or countable set of states. A function $f : \Omega \rightarrow [0, 1]$ is called a probability distribution if $\sum_{\omega \in \Omega} f(\omega) = 1$. We shall note $D(\Omega)$ the set of probabilistic distributions on Ω .*

Now that we have the probabilistic counterpart of the notion of state, we need to have the counterpart of the notion of transition.

Definition 2 *Let Ω be a finite or countable set of states. Let us consider a function $T : \Omega \times \Omega \rightarrow [0, 1]$ such that for all $\omega_1 \in \Omega$, $\sum_{\omega_2 \in \Omega} T(\omega_1; \omega_2) = 1$. (Ω, T) is called a probabilistic transition system.*

If Ω is finite, the relation can be given by a *probabilistic transition matrix*. Let us assimilate Ω to $\{1, \dots, N\}$. Then, the transition matrix M is defined by $m_{i,j} = T(i, j)$ if $i \rightarrow j$, 0 otherwise.

The intuitive notion of a probabilistic transition is that it maps an *input distribution* to an *output distribution*. It is the probabilistic counterpart of the notion of a successor state.

Definition 3 *Let T be a transition probability between Ω_1 and Ω_2 . Let us define $\vec{T} : D(\Omega_1) \rightarrow D(\Omega_2)$ as follows: $\vec{T}(d)(\omega_2) = \sum_{\omega_1 \in \Omega_1} T(\omega_1, \omega_2)d(\omega_1)$.*

Let us now describe the probabilistic counterpart of the notion of predecessor state. Given a transition probability T between Ω_1 and Ω_2 and a boolean property $\pi : \Omega_2 \rightarrow \{0, 1\}$, the expectation of a state $\omega_1 \in \Omega_1$ to reach π in one step is then $\sum_{\omega_2 \in \Omega_2} T(\omega_1, \omega_2)\pi(\omega_2)$. We have thus defined a function $\Omega_1 \rightarrow [0, 1]$ mapping each state to its corresponding expectation.

A natural extension of this construction is to consider any function $f \in P(\Omega_2) = \Omega_2 \rightarrow [0, 1]$. We call such functions *condition functions*.¹

Definition 4 Let T be a transition probability between Ω_1 and Ω_2 . Let us define $\overleftarrow{T} : P(\Omega_2) \rightarrow P(\Omega_1)$ as follows: $\overleftarrow{T}(f)(\omega_1) = \sum_{\omega_2 \in \Omega_2} T(\omega_1, \omega_2)f(\omega_2)$.

Lemma 1. For all transition probability T , \overleftarrow{T} is ω -continuous.

Those functions are linked by the following *adjunction relation*: if T is a transition probability relative to Ω_1 and Ω_2 , noting $\langle f, \mu \rangle = \sum_{\omega} f(\omega)\mu(\omega)$, then

$$\forall f \in P(\Omega_2) \forall \mu \in D(\Omega_1) \langle f, \overrightarrow{T}.\mu \rangle = \langle \overleftarrow{T}.f, \mu \rangle. \quad (1)$$

2.2 Probability measures on traces

We shall also use probability measures on *sets of traces* arising from probabilistic transition systems. Let us start with a simple example — consider sequences of tosses of a fair coin: the coin has probability 0.5 of giving 0 and 0.5 of giving 1. A trace is then an infinite sequence of zeroes and ones. Let us consider the (regular) set $0^n(0|1)^*$ of sequences starting by at least n zeroes. It is obvious that the probability of falling into that set is 2^{-n} . The probability of the singleton containing the sequence of only zeroes is 0.

We use the theorem of Ionescu Tulcea (Appendix B) to construct the probability measure μ_ω on the set of traces according to the probability distribution μ on the initial states and the transition probability T .

The probability of a property $P : \Omega^{\mathbb{N}} \rightarrow \{0, 1\}$ on the traces is then $\int P d\mu_\omega$.

3 Nondeterministic and probabilistic transition systems

We shall see how to combine the notions of *nondeterministic choice* (sets of possible choices for which we know no probabilistic properties) and *probabilistic choice* (sets of possible choices for which we know probabilistic properties), obtaining *discrete-time Markov decision processes* [23], which has been studied more particularly in the field of operational research and finance mathematics, as well as machine learning.

Let us now consider the case where the system must be able to do both nondeterministic and probabilistic transitions (example in Fig. 1). The system then has the choice between different transition probabilities.

¹ Please note that while those functions look similar to distributions, they are quite different in their meaning and are different mathematical objects when treated in the general, non discrete, case.




Fig. 1. A probabilistic-nondeterministic transition system

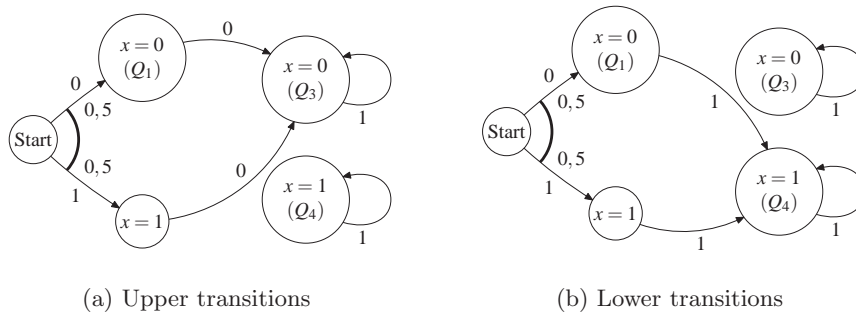


Fig. 2. Two purely probabilistic transition systems that define, when composed together nondeterministically, the same probabilistic-nondeterministic transition system as in Fig. 1.

For instance, on Fig. 1, in state Q_1 , the system has the choice between two partial transition probabilities: the first goes to Q_3 with probability 1, the second goes to Q_4 with probability 1. For an easier intuition, one may think about this choice as if it were made by an adversary willing to induce certain behaviors. The adversary is supposed to follow a strategy or *policy* [23, §2.1.5].

In this paper, we shall assume that the adversary may see the present and past states of the execution. This is realistic if the adversary models a human evildoer as well as the physical environment of an embedded system.

Let us note that other choices can give very different results. For instance, let us consider a program that chooses a Boolean variable x nondeterministically, then chooses a Boolean variable y with uniform probability, then replaces x with the exclusive or of x and y (Fig. 1). Clearly, if the nondeterminism is allowed to “look at the future” and predict the outcome of the random generator, it can always arrange for z to be true. If it can only look at the past, it cannot and z is uniformly distributed.

Let us suppose that our system is defined by a transition probability T from $\Omega \times Y$ to Ω , where Y is the *domain of nondeterministic choices*. For instance, for the system given in Fig. 1, Y is $\{0, 1\}$ (choice between the upper and lower arrows, as seen in Fig. 2). The operation of the adversary for the transition between states n and $n + 1$ is modeled using an unknown transition probability U_n from Ω^n to Y . The whole transition that is executed by the system is then the composition $T_n = T \circ \left[\begin{smallmatrix} Id \\ U_n \end{smallmatrix} \right]$, which is a transition probability between Ω^n and Ω . By this notation, we mean that, using the notation of Def. 2,

$$T_n(x_0, \dots, x_{n-1}; x_n) = T(x_{n-1}, U_n(x_0, \dots, x_{n-1}); x_n). \quad (2)$$

Ionescu Tulcea's theorem (Appendix B) then constructs from the (T_n) a transition probability $G(T, (U_n)_{n \in \mathbb{N}})$ from Ω (the initial state space) to $\Omega^{\mathbb{N}}$. We note

$$S_T(f, (U_n)_{n \in \mathbb{N}}) = t_0 \mapsto \langle \mathbf{t} \mapsto f(t_0, \mathbf{t}), G(T, (U_n)_{n \in \mathbb{N}})(t_0) \rangle \quad (3)$$

(S short for S_T if there is no ambiguity about T) and $R(f)$ the set of all functions $S(T, (U_n)_{n \in \mathbb{N}})$ when $(U_n)_{n \in \mathbb{N}}$ is a sequence of transition probabilities, U_n being a transition probability from Ω^n to Y . The λ is here a functional notation.

Let E_+^T or, if there is no confusion possible, $E_+(f) = \sup R(f)$ be the *worst-case semantics* and $E_-(f) = \inf R(f)$ be the *best-case semantics* (those designations assume that f indicates some kind of failure condition). Intuitively, if f is the characteristic function of a set of “faulty” traces, E_+ expresses a “worst case” analysis, modeling an adversary willing to make the system err and E_- a “best case” analysis, modeling an adversary willing to prevent the system from erring. $E_+(f)$ is often called the *value* of the Markov decision process with respect to the reward function f (even though we use a slightly different framework as the one given in [23]).

Lemma 5 $E_+(1) = 1$ and $E_+(0) = 0$. E_+ is monotone and upper-continuous.

4 The properties to analyze

We consider a property to analyze on the traces. To each initial state we attach its *expectation*, that is, the integral of the property to analyze over the traces starting from this state (or the set of integrals, if considering nondeterminism). The properties to analyze are expressed as measurable functions from the set of (possible) traces of execution; we call these functions *trace valuator*s. We shall actually consider a class of trace valuator defined syntactically by certain formulas.

4.1 Expectation Functions

Let $I = [0, 1]$ or $[0, +\infty]$, depending on the kind of properties to analyze. Let Ω be a finite or countable set of states — we impose this cardinality constraint so as to avoid theoretical complexities. $\mathcal{P}(\Omega)$ is the set of subsets of Ω . Let $\Omega \rightarrow I$ be the set of functions from Ω to I , the *expectation functions* of our system; this set, ordered by \leq point-wise, is a complete lattice.

4.2 Trace Valuators

Let $\Omega^{\mathbb{N}} \rightarrow I$ be the set of measurable functions from $\Omega^{\mathbb{N}}$ to I , ordered point-wise. We shall call such functions “valuators”.

Boolean Trace Valuators We take $I = [0, 1]$ or even $I = \{0, 1\}$.

We shall consider formulas written in the following language:

formula1 ::= *name*
| *constant*
| *name* +_{set} *formula1* where *set* $\subseteq \Omega$
| *constant* +_{set} *formula1*
| lfp(*name* \mapsto *formula1*)
| gfp(*name* \mapsto *formula1*)
| shift(*formula1*)
| let *name* = *formula1* in *formula2*

Let $\text{shift} : \Omega^{\mathbb{N}} \rightarrow \Omega^{\mathbb{N}}$: $(\text{shift}.t)_k = t_{k+1}$.

Let env_t be the set of environments of valuators, mapping each *name* to a valuator, ordered point-wise.

$\llbracket \text{formula} \rrbracket_t : \text{env}_t \rightarrow (\Omega^{\mathbb{N}} \rightarrow I)$ is defined inductively as follows:

$$\llbracket \text{name} \rrbracket_t . \text{env} = \text{env}(\text{name}) \quad (4)$$

$$\llbracket \text{constant} \rrbracket_t . \text{env} = \text{constant} \quad (5)$$

$$\llbracket f_1 +_S f_2 \rrbracket_t . \text{env} = \lambda \mathbf{t}. \chi_S(t_0). (\llbracket f_1 \rrbracket_t . \text{env}) + \chi_{S^C}(t_0). (\llbracket f_2 \rrbracket_t . \text{env}) \quad (6)$$

$$\llbracket \text{lfp}(\text{name} \mapsto f) \rrbracket_t . \text{env} = \text{lfp}(\lambda \phi. \llbracket f \rrbracket_t . \text{env}[\text{name} \mapsto \phi]) \quad (7)$$

$$\llbracket \text{gfp}(\text{name} \mapsto f) \rrbracket_t . \text{env} = \text{gfp}(\lambda \phi. \llbracket f \rrbracket_t . \text{env}[\text{name} \mapsto \phi]) \quad (8)$$

$$\llbracket \text{shift}(f) \rrbracket_t . \text{env} = (\llbracket f \rrbracket_t . \text{env}) \circ \text{shift} \quad (9)$$

$$\llbracket \text{let } \text{name} = f_1 \text{ in } f_2 \rrbracket_t . \text{env} = \llbracket f_2 \rrbracket_t . \text{env}[\text{name} \mapsto \llbracket f_1 \rrbracket_t . \text{env}] \quad (10)$$

χ_S is the characteristic function of S and S^C the complement of S . $t_0 \in \Omega$ is the first state of the sequence of states $\mathbf{t} \in \Omega^{\mathbb{N}}$.

Let us note that we could introduce a negation operator \neg ($\llbracket \neg f \rrbracket_t = 1 - \llbracket f \rrbracket_t$) as syntactic sugar — any formula with negations can be transformed into one without by pushing the negations to the leaves.

Lemma 6 *For all formula f , $\llbracket f \rrbracket_t$ is monotone. For all formula f without lfp or gfp, $\llbracket f \rrbracket_t$ is continuous.*

Some particularly interesting Boolean valuators We shall consider in this section four very important classes of properties, all of which can be shown to be measurable.

- Let A be a set of states. The *reachability* property associated with A defines the set of traces that pass through A at some point. It corresponds to the formula:

$$\text{lfp}(f \mapsto 1 +_A \text{shift} f) \quad (11)$$

- Let A be a set of states. The *liveness* property associated with A defines the set of traces that always remain in A . It corresponds to the formula:

$$\text{gfp}(f \mapsto \text{shift} f +_A 0) \quad (12)$$

- Let A be a (measurable) set of states. The *Büchi acceptance* property associated with A defines the set of traces that pass through A infinitely often; it is written as:

$$\text{gfp}(C \mapsto \text{lfp}(R \mapsto \text{shift}(C) +_A \text{shift}(R))) \quad (13)$$

The *co-Büchi* property is just the negation of this formula.

- Given a sequence $\emptyset \subseteq U_{2k} \subseteq \dots \subseteq U_0 = \Omega$, the *Rabin acceptance property* associated with (U_n) is the set of traces defined by the following temporal property [10, sect. 5]:

$$\mathcal{R} = \bigvee_{i=0}^{k-1} (\Box \Diamond U_{2i} \wedge \neg \Box \Diamond U_{2i+1}) \quad (14)$$

It corresponds to the following formula:

$$\begin{aligned} &\text{gfp}(x_{2k-1} \mapsto \text{lfp}(x_{2k} \mapsto \dots \text{gfp}(x_1 \mapsto \text{lfp}(x_0 \mapsto \\ &((\dots (x_{2k-1} +_{U_{2k-1}} x_{2k-2}) \dots +_{U_1} x_0) +_{U_0} 0) \end{aligned} \quad (15)$$

Summation valuator A related family of trace valutors are the *summing* valutors. The summation valuator associated with a (measurable) function $f : \Omega \mapsto [0, +\infty]$ is the function

$$\llbracket \Sigma A \rrbracket_t : \begin{cases} \Omega^{\mathbb{N}} & \rightarrow [0, +\infty] \\ (x_n)_{n \in \mathbb{N}} & \mapsto \sum_{k=0}^{\infty} f(x_k) \end{cases} \quad (16)$$

Obviously, this function can be formulated as a least fixpoint:

$$\llbracket \Sigma f \rrbracket_t = \text{lfp}(\phi \mapsto f + \phi \circ \text{shift}) \quad (17)$$

This construct has two obvious applications:

- counting the average number of times the program goes through a certain set of states A ; here, f is the characteristic function of A ;
- counting the average time spent in the process; here f maps each state to the amount of time spent in that state (0 for states meaning “termination”).

4.3 Temporal logics

Temporal logics [5, chapter 3] are expressive means of specifying properties of transition systems.

Linear time logic (LTL) and ω -regular conditions A formula F in LTL defines an ω -regular set of traces $\llbracket F \rrbracket_t$, that is, a set of traces recognizable by a (nondeterministic) Büchi automaton B [5, §3.2], or, equivalently, by a deterministic Rabin automaton R [25, §4].

Let us consider a (nondeterministic) probabilistic transition system T , and the according definition of E_+^T . Let us consider the synchronous product $T \times R$, and C the associated Rabin acceptance condition. $E_+^T(\llbracket F \rrbracket_t)$ is then equal to $E_+^{S \times R}(\llbracket C \rrbracket_t)$ [10, §5].

If B is deterministic, we can similarly consider the synchronous product $T \times B$, and C the associated Büchi condition. $E_+^T(\llbracket F \rrbracket_t)$ is then equal to $E_+^{S \times B}(\llbracket C \rrbracket_t)$ [10, §4].²

Branching-time logic: pCTL and pCTL* The branching-time logics CTL and CTL* [5, §3.2] have had much success in the analysis of nondeterministic (albeit non probabilistic) systems. It was therefore quite natural to extend this notion to probabilistic systems. Proposed extensions to the probabilistic case include pCTL [11] and pCTL*. We shall see here briefly how we deal with some pCTL* formulas.

CTL* formulas define sets of states as the starting states of sets of traces defined by LTL path formulas (in which state formulas are CTL* state formulas).

The operation that makes a CTL* state formula out of a LTL path formula is the taking of the initial states: if $\llbracket f \rrbracket_s$ denotes the semantics of f as a state formula and $\llbracket f \rrbracket_p$ its semantics as a path formula, then

$$\llbracket f \rrbracket_s = \{x_0 \in \Omega \mid \exists x_1, \dots \langle x_0, x_1, \dots \rangle \in \llbracket f \rrbracket_p\}. \quad (18)$$

In the case of probabilistic systems, we do not have sets of starting states but expectation functions; such expectation functions are then compared to a threshold value, which gives sets of states. State formulas noted as $f_{\bowtie \alpha}$ are thus obtained, where f is a trace valuator and \bowtie is \leq , $<$, $=$, $>$ or \geq . The semantics of this construct is as follow:

$$\llbracket f_{\bowtie \alpha} \rrbracket = \{x_0 \in \Omega \mid \forall (U_n)_{n \in \mathbb{N}} S(\llbracket f \rrbracket_t, (U_n)_{n \in \mathbb{N}})(x_0) \bowtie \alpha\} \quad (19)$$

In the case of $<$ and \leq , giving an upper bound on those sets is easy provided we have an upper bound $\llbracket f \rrbracket_{e+}^\sharp$ of $E_+(\llbracket f \rrbracket_t)$ (see §5):

$$\forall x_0 \llbracket f \rrbracket_{e+}^\sharp(x_0) \bowtie \alpha \implies x_0 \notin \llbracket f_{\bowtie \alpha} \rrbracket. \quad (20)$$

² Note that this does not hold for nondeterministic Büchi automata, since the automaton is allowed to take its nondeterministic choices with the knowledge of the full sequence of states, not only the past and present states.

5 Backwards Worst Case Analysis

In section 4.2, we gave the syntax and semantics of a logic describing sets of traces, or, more generally, measurable functions over the traces. Given a formula f , one may want to compute its worst-case probability $E_+(\llbracket f \rrbracket_t)$, or at least get an upper bound for it. Unfortunately, the definitions of both $\llbracket f \rrbracket_t$ and E_+ do not yield effective means to do so.

In §5.1 we shall attach to each formula f another semantics $\llbracket f \rrbracket_{e+}$, which we shall show how to abstract in §6. We shall see the abstraction relationship between $E_+(\llbracket f \rrbracket_t)$ and $\llbracket f \rrbracket_{e+}$ in §5.2.

5.1 Backwards Worst Case Semantics on Expectation Functions

A well-known solution to the problem of the optimal value of a Markov decision process is *value iteration* [23, §7.2.4]. This method is mainly of theoretical interest for the analysis of finite state Markov decision processes, since there is little control as to its rate of convergence and much better algorithms are available [23, §7.2.4]. On the other hand, since it is actually a kind of generalization to Markov decision processes of the backwards reachability analysis for nondeterministic systems, we can apply abstract interpretation techniques so as to provide an effective mean to compute upper bounds on the probability of the properties to analyze.

Let env_e be the set of environments of expectation functions (an environment of expectation functions maps each *name* to a expectation function), ordered point-wise.

$\llbracket formula \rrbracket_{e+} : (\Omega \rightarrow I) \rightarrow (\Omega \rightarrow I)$ is defined inductively as follows:

$$\llbracket name \rrbracket_{e+} .env = env(name) \quad (21)$$

$$\llbracket constant \rrbracket_{e+} .env = \lambda x. constant \quad (22)$$

$$\llbracket f_1 +_S f_2 \rrbracket_{e+} .env = \chi_S.(\llbracket f_1 \rrbracket_{e+} .env) + \chi_{S^c}.(\llbracket f_2 \rrbracket_{e+} .env) \quad (23)$$

$$\llbracket \text{lfp}(name \mapsto f) \rrbracket_{e+} .env = \text{lfp}(\lambda \phi. \llbracket f \rrbracket_{e+} .env[name \mapsto \phi]) \quad (24)$$

$$\llbracket \text{gfp}(name \mapsto f) \rrbracket_{e+} .env = \text{gfp}(\lambda \phi. \llbracket f \rrbracket_{e+} .env[name \mapsto \phi]) \quad (25)$$

$$\llbracket \text{shift}(f) \rrbracket_{e+} .env = \sup_{T \in \mathcal{T}} (\overleftarrow{T}(\llbracket f \rrbracket_{e+} .env)) \quad (26)$$

$$\llbracket \text{let } name = f_1 \text{ in } f_2 \rrbracket_{e+} .env = \llbracket f_2 \rrbracket_{e+} .env[name \mapsto \llbracket f_1 \rrbracket_{e+} .env] \quad (27)$$

This semantics is thus some form of μ -calculus, except that “lfp” replaces the μ binder and “gfp” ν ; but since we also use μ to note measures, it would have been confusing to also use it in the syntax of the formulas.

$\llbracket formula \rrbracket_{e+}$ is monotone.

Lemma 7 *Let f be a formula not containing gfp . $\llbracket f \rrbracket_{e+}$ is ω -upper-continuous.*

As for the summation valuator,

$$\llbracket \Sigma f \rrbracket_{e+} = \text{lfp} \left(\phi \mapsto f + \sup_{T \in \mathcal{T}} (\overleftarrow{T} . \phi) \right) \quad (28)$$

5.2 The Abstraction Relation Between the Semantics

Theorem 8 *Let f be a formula not containing gfp . Let env be an environment of valutors. Noting $E_+(\text{env})$ the point-wise application of E_+ to env ,*

$$\llbracket f \rrbracket_{e_+} \cdot (E_+(\text{env})) = E_+(\llbracket f \rrbracket_t \cdot \text{env}) \quad (29)$$

Proof. Proof by induction on the structure of f .

- The cases for “let”, *name* and *constant* are trivial.
- For $f_1 +_S f_2$: Let $t_0 \in X$.

$$\begin{aligned} & \llbracket f_1 +_S f_2 \rrbracket_{e_+} \cdot (E_+(\text{env})).t_0 \\ &= \chi_S(t_0) \cdot (\llbracket f_1 \rrbracket_{e_+} \cdot E_+(\text{env})).t_0 + \chi_{S^c}(t_0) \cdot (\llbracket f_2 \rrbracket_{e_+} \cdot E_+(\text{env})).t_0 \\ &= \chi_S(t_0) \cdot (E_+(\llbracket f_1 \rrbracket_t \cdot \text{env})).t_0 + \chi_{S^c}(t_0) \cdot (E_+(\llbracket f_2 \rrbracket_t \cdot \text{env})).t_0 \quad (\text{induction}) \\ &= E_+(\lambda t. \chi_S(t_0) \cdot (\llbracket f_1 \rrbracket_t \cdot \text{env}) + \chi_{S^c}(t_0) \cdot (\llbracket f_2 \rrbracket_t \cdot \text{env})).t_0 \quad (\text{lemma 2}) \\ &= E_+(\llbracket f_1 +_S f_2 \rrbracket_t \cdot \text{env}).t_0. \end{aligned}$$

- For shift: Let us first fix U_1 . Let us note $T_1 = T \circ [\text{Id}_{U_1}]$ and consider $\overleftarrow{T}_1.E_+(\llbracket f \rrbracket_t)$. From lemma 1, \overleftarrow{T}_1 is a monotonic, ω -continuous, operator; from lemma 3, $R(\llbracket f \rrbracket_t)$ is directed; from lemma 4,

$$\bigsqcup_{f \in R(\llbracket f \rrbracket_t \cdot \text{env})} (\overleftarrow{T}_1 f) = \overleftarrow{T}_1 \left(\bigsqcup_{f \in R(\llbracket f \rrbracket_t \cdot \text{env})} f \right).$$

It follows that (using the λ -notation for functions),

$$\begin{aligned} & \overleftarrow{T}_1.E_+(\llbracket f \rrbracket_t).t_0 = \\ & \sup_{\substack{(U_n)_{n \geq 2} \\ U_n \text{ not depending on } t_0}} \int \lambda \langle t_2, \dots \rangle. (\llbracket f \rrbracket_t \cdot \text{env})(\langle t_1, \dots \rangle) d[G(T, (U_n)_{n \geq 2}).t_1] T_1(t_0, dt_1) \\ &= \sup_{\substack{(U_n)_{n \geq 2} \\ U_n \text{ not depending on } t_0}} \int \lambda \langle t_2, \dots \rangle. (\llbracket f \rrbracket_t \cdot \text{env})(\langle t_1, \dots \rangle) d[G(T, (U_n)_{n \geq 1}).t_0] \end{aligned}$$

Let us apply lemma 11 to that last expression. We obtain

$$\overleftarrow{T}_1.E_+(\llbracket f \rrbracket_t).t_0 = \sup_{(U_n)_{n \geq 2}} \int \lambda \langle t_2, \dots \rangle. (\llbracket f \rrbracket_t \cdot \text{env})(\langle t_1, \dots \rangle) d[G(T, (U_n)_{n \geq 1}).t_0] \quad (30)$$

Let $t_0 \in X$. Let us now consider all U_1 's.

$$\begin{aligned} & E_+(\llbracket \text{shift}(f) \rrbracket_t \cdot \text{env}).t_0 \\ &= \sup_{(U_n)_{n \geq 1}} \int \lambda \langle t_1, \dots \rangle. (\llbracket f \rrbracket_t \cdot \text{env}) \circ \text{shift}(\langle t_0, t_1, \dots \rangle) d[G(T, (U_n)_{n \geq 1}).t_0] \\ &= \sup_{U_1} \sup_{(U_n)_{n \geq 2}} \int \lambda \langle t_1, \dots \rangle. (\llbracket f \rrbracket_t \cdot \text{env})(\langle t_1, \dots \rangle) d[G(T, (U_n)_{n \geq 1}).t_0] \\ &= \left(\sup_{U_1} \left(T \circ [\text{Id}_{U_1}] \right) \cdot E_+(\llbracket f \rrbracket_t) \right) \cdot t_0 \quad (\text{using Equ. 30}) \\ &= \llbracket \text{shift}(f) \rrbracket_{e_+} \cdot E_+(\text{env}) \end{aligned}$$

- $\llbracket \text{lfp}(name \mapsto f) \rrbracket_{e_+} . env = \text{lfp}(\lambda \phi. \llbracket f \rrbracket_{e_+} . env[name \mapsto \phi])$.
From lemma 7, $\lambda \phi. \llbracket f \rrbracket_{e_+} . env[name \mapsto \phi]$ is ω -upper-continuous.
 $\llbracket \text{lfp}(name \mapsto f) \rrbracket_t . env = \text{lfp}(\lambda \phi. \llbracket f \rrbracket_t . env[name \mapsto \phi])$.
From lemma 6, $\lambda \phi. \llbracket f \rrbracket_t . env[name \mapsto \phi]$ is ω -upper-continuous.
From the induction hypothesis,

$$E_+ \circ (\lambda \phi. \llbracket f \rrbracket_t . env[name \mapsto \phi]) = (\lambda \phi. \llbracket f \rrbracket_{e_+} . E_+(env)[name \mapsto \phi]).$$

From lemma 5, E_+ is ω -upper-continuous. The conclusion then follows from lemma 5.

The following theorem guarantees the soundness of the abstract analysis for all formulas.

Theorem 9 *Let f be a formula. Let env be an environment of valuations. Let us suppose that $H \geq E_+(env)$ pointwise. Then*

$$\llbracket f \rrbracket_{e_+} . (H) \geq E_+(\llbracket f \rrbracket_t . env). \quad (31)$$

Proof by induction similar to that of Th. 8. Also similarly we can guarantee the soundness of the analysis of summations:

Theorem 10 *The semantics of the summing operator satisfies:*

$$E_+(\llbracket \Sigma f \rrbracket_t) = \llbracket \Sigma f \rrbracket_{e_+}. \quad (32)$$

6 Abstract Analysis

We shall see here more precisely how to apply abstract interpretation to that backwards semantics.

6.1 General case

We compute safe approximations of $\llbracket f \rrbracket_{e_+}$ by abstract interpretation. We introduce an abstract semantics $\llbracket f \rrbracket_{e_+}^\sharp$ which is an upper approximation of f :

$$\forall env \forall env^\sharp \ env^\sharp \geq env \implies \llbracket f \rrbracket_{e_+}^\sharp . env^\sharp \geq \llbracket f \rrbracket_{e_+} . env. \quad (33)$$

The computations for $\llbracket f \rrbracket_{e_+}^\sharp$ will be done symbolically in an *abstract domain* such as the ones described in [19, 20].

- We shall assume that we have an abstract operation for “shift”. That is, we have a monotone operation pre^\sharp such that

$$\forall f, \forall T \in \mathcal{T}, \text{pre}^\sharp . f^\sharp \geq T^* . f^\sharp. \quad (34)$$

This operation will be supplied by the abstract domain that we use. Then

$$\forall env, \forall env^\sharp, \ env^\sharp \geq env \implies \llbracket \text{shift}(f) \rrbracket_{e_+}^\sharp . env^\sharp \geq \llbracket \text{shift}(f) \rrbracket_{e_+} . env. \quad (35)$$

provided that

$$\forall env, \forall env^\sharp, \ env^\sharp \geq env \implies \llbracket f \rrbracket_{e_+}^\sharp . env^\sharp \geq \llbracket f \rrbracket_{e_+} . env.$$

- We shall approximate least fixpoints using a *widening operator* [9, §4.3]. A widening operator ∇ is a kind of abstraction of the least upper bound that enforces convergence:
 - $f \nabla g \geq \sup(f, g)$ (pointwise);
 - For any ascending sequence $(v_n)_{n \in \mathbb{N}}$, the sequence $(u_n)_{n \in \mathbb{N}}$ defined inductively by $u_{n+1} = u_n \nabla v_n$ is to be ultimately stationary.
 Then the limit L^\sharp of the sequence defined by $u_0 = 0$ and $u_{n+1} = u_n \nabla f^\sharp(u_n)$, where f^\sharp is an upper approximation of f , is an upper approximation to the least fixpoint of f . More precise upper approximations of the least fixpoint of f can then be reached by iterating f^\sharp over L^\sharp using a so-called *narrowing operators* [9, §4.3].
- We shall approximate greatest fixpoints using a limited iteration approach: if f^\sharp is an upper approximation of f , then for any $n \in \mathbb{N}$, $f^{\sharp n}(\top) \geq \text{gfp } f$.

6.2 Partitioning in programs

In the case of programs, the state space is generally $P \times M$, where P is the (finite) set of program points and M the set of possible memory configurations. More generally, P may be a kind of *partitioning* of the program. Non-probabilistic analysis generally operates on abstractions of $\mathcal{P}(P \times M) \simeq P \times M \rightarrow \{0, 1\} \simeq P \rightarrow \mathcal{P}(M)$. Given an abstraction of $\mathcal{P}(M)$ by a lattice L^\sharp , one obtains a pointwise abstraction of $P \rightarrow \mathcal{P}(M)$ by $P \rightarrow L^\sharp$. Elements of $P \rightarrow L^\sharp$ are just vectors of $|P|$ elements of L^\sharp .

This approach can be directly extended to our measurable functions: we shall abstract $P \times M \rightarrow I$ (where $I = [0, 1]$ or $I = [0, +\infty]$) by $P \rightarrow L^\sharp$ if L^\sharp is an abstract domain for $M \rightarrow I$.

The first problem is to get an abstraction of the operation used in the “shift” construct:

$$F : \begin{array}{l} (P \times M \rightarrow I) \rightarrow (P \times M \rightarrow I) \\ h \qquad \qquad \mapsto (l, m) \mapsto \sup_{y \in Y} \sum_{(l', m') \in P \times M} T((l, m), y; (l', m')) . h(l', m') \end{array} \quad (36)$$

Let us take the following form for the program instructions: at program point l , the executed instruction represented by T is the sequence:

1. a nondeterministic choice y is taken in the set Y_l ;
2. a random choice r is taken in set R_l according to distribution \mathcal{R}_p ;
3. the memory state is combined deterministically with the two choices to form the new memory state using a function $F_l : (M \times Y) \times R_l \rightarrow M$;
4. depending on the memory state m , the program takes a deterministic jump to program point $J(l, m)$.

Let us note $\tau_l(l') = \{m \mid J(l, m) = l'\}$ (the set of memory values m that lead to program point l' from program point l ; $\tau_l(l')$ is then essentially the condition for a conditional jump). Then we can rewrite the transition equation as follows

$$(F.h)(l) = \text{choice}_{Y_l}^* \circ \text{random}_{R_l}^* \circ (F_l)_p^* \left(\sum_{l' \in P} \phi_{\tau_l(l')}^* (h(l', \bullet)) \right) \quad (37)$$

using the following building blocks:

$$\text{choice}_{Y_i}^*(h) = m \mapsto \sup_{y \in Y_i} h(m, y) \quad (38)$$

$$\text{random}_{R_i}^*(h) = m \mapsto \int h(m, r) d\mu_{R_i}(r) \quad (39)$$

$$(F_l)_p^*(h) = h \circ F_l \quad (40)$$

$$\phi_A^*(h) = h \cdot \chi_A \quad (41)$$

The reasons for those notations are explained in earlier works on the linear adjoint of Kozen's denotational semantics for probabilistic programs [19].

We shall abstract F as the composition of abstractions for:

- $\text{choice}_{Y_i}^*$, nondeterministic choice;
- $\text{random}_{R_i}^*$, probabilistic choice;
- $F_{l_p}^*$, deterministic run (arithmetic operations and the like);
- ϕ_A^* , test.

Since the function F is ω -upper-continuous, the least fixpoint of F is obtained as the limit of $F^n(0)$ (let us recall that this is the point-wise limit of a sequence of functions from Ω to I). The expression of the iterates using a partition with respect to P is as follows:

$$f_1^{(n+1)} = F_1(f_1^{(n)}, \dots, f_{|P|}^{(n)}) \quad (42)$$

$$\vdots \quad (43)$$

$$f_{|P|}^{(n+1)} = F_{|P|}(f_1^{(n)}, \dots, f_{|P|}^{(n)}) \quad (44)$$

$$(45)$$

In terms of implementation, this means that we update in parallel the $|P|$ elements of the vector representing the iterate. As noted by Cousot [8, §2.9], this parallel update may be replaced by *chaotic iterations* or *asynchronous iterations*. Chaotic iterations allow us to compute the iterations by taking into account the recently updated elements. All these iteration strategies lead to the same limit (the least fixpoint of F).

Let us consider for instance the following strategy:

$$\begin{aligned} f_1^{(n+1)} &= F_1(f_1^{(n)}, \dots, f_{|P|}^{(n)}) \\ f_2^{(n+1)} &= F_2(f_1^{(n+1)}, \dots, f_{|P|}^{(n)}) \\ &\vdots \\ f_{|P|}^{(n+1)} &= F_{|P|}(f_1^{(n+1)}, \dots, f_{|P|}^{(n)}) \end{aligned} \quad (46)$$

This strategy is itself a monotone operator whose least fixpoint is to be determined. It has an obvious abstract counterpart leading to an approximate fixpoint in the usual way (§6.1).

7 Conclusion, Related Works and Discussion

We showed how to apply abstract interpretation techniques to check various temporal properties of (nondeterministic) probabilistic programs, considered as Markov decision processes (small-step semantics).

The most natural point of view on those processes is that the nondeterministic decisions are taken as the program proceeds, taking into account the current state as well as the previous ones. This how Markov decision processes are usually studied [23] and this is the approach we took here.

It can be argued that this model is excessively pessimistic. Indeed, if nondeterminism is used to model the environment of an embedded system, then it is excessive to assume that the behavior of this environment depends on the history of the *internal state* of the system; only the part of this history observable from the environment should be taken into account. This leads to the study of *partially observable Markov decision processes* (POMDP); however, their effective analysis is much more complex than that of fully observable processes [14].

Cleaveland’s work [6] focuses on the model where the nondeterministic choices are taken *after* the probabilistic ones. This simplifies the theory to some extent, since taking the product of the analyzed process with an nondeterministic “observation” process, such as a nondeterministic Büchi automaton, is then easy. We have already proposed a Monte-Carlo method for such semantics [18].

The backwards analysis method we described is a generalization of the value iteration method used in operational research to compute the value of Markov decision processes. Our reachability analysis is related to the study of *positive bounded models* [23, §7.2], where the reward 1 is granted the first time the process runs through the set of states to consider. The liveness analysis is related to the study of *negative models* [23, §7.3], where the reward -1 is granted the first time the process leaves the set of states to consider.

Formal languages similar to the one we consider have been introduced by other authors, such as *quantitative game μ -calculus* [10]. The differences between our approach and this game calculus approach are threefold:

- We give a semantics in terms of traces, then prove its link with a semantics in terms of expectation functions; quantitative μ -calculus only gives the interpretation as expectation functions.
- While we prove a generic link between the semantics as an inequality valid for any formula (or an equation for some class), de Alfaro proves an interpretation for some specific formulas (reachability, liveness, deterministic Büchi and Rabin trace properties). We conjecture that we can extend the equality cases of this link.
- De Alfaro considers random two-player games while we consider random single-player games. We mentioned briefly (§5.2) the differences between Markov decision processes and two-player games. Such problems can model questions such as the choice of an optimal strategy by the program so as to minimize the probability of a problem for all possible environments.

A possible extension of these properties is *discounted models* [23, Ch. 6]. In these, the importance of the future decreases exponentially; for instance, λ -discounted reachability would count passing through A for the first time at step n as λ^n instead of 1 (of course, $0 < \lambda < 1$). The theoretical study of those models is considerably easier than that of non-discounted models, since the fixpoints to study are the fixed points of contraction mappings in Banach spaces. While the extension of the techniques exposed in this paper to discounted models is easy (it suffices to add a multiplication by λ in the semantics of the “shift” operation), the practical interest of such models in the checking of computer programs remains to be seen.

Another possible extension is the computation of averages not only on the space of the program, but also on the time: computing the average value of a certain function as long as the program is running. Since this property is the quotient of two summing properties, there is no obvious method to evaluate it iteratively.

Another possible direction is the study of continuous time probabilistic systems. As usual with continuous-time systems, some kind of reduction to discrete time processes is to be done [12, 13].

References

1. Luca de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, Department of Computer Science, June 1998. CS-TR-98-1601.
2. Christel Baier, Edmund M. Clarke, Vasiliki Hartonas-Garmhausen, and Marta Kwiatkowska. Symbolic model checking for probabilistic processes. In P. Degano, R. Gorrieri, and A. Marchetti-Spaccamela, editors, *Automata, Languages and Programming (ICALP '97)*, volume 1256 of *LNCS*. Springer, 1997.
3. A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *FST TCS 95: Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *Lecture Notes in Computer Science*, pages 499–513. Springer-Verlag, 1995.
4. C. Baier, M. Kwiatkowska, and G. Norman. Computing probability bounds for linear time formulas over concurrent probabilistic systems. *Electronic Notes in Theoretical Computer Science*, 21, 1999.
5. Edmund M. Clarke, Jr, Orna Grumberg, and Doron A. Peled. *Model Checking*. MIT Press, 1999.
6. Rance Cleaveland, Scott A. Smolka, and Amy E. Zwarico. Testing preorders for probabilistic processes. In Werner Kuich, editor, *Automata, Languages and Programming, 19th International Colloquium*, volume 623 of *Lecture Notes in Computer Science*, pages 708–719, Vienna, Austria, 13–17 July 1992. Springer-Verlag.
7. C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. In *Proc. ICALP'90*, volume 443 of *LNCS*, pages 336–349. Springer, 1990.
8. Patrick Cousot. *Méthodes itératives de construction et d'approximation de points fixes d'opérateurs monotones sur un treillis, analyse sémantique de programmes*. Thèse d'état ès sciences mathématiques, Université scientifique et médicale de Grenoble, Grenoble, France, 21 mars 1978.
9. Patrick Cousot and Radhia Cousot. Abstract interpretation and application to logic programs. *J. Logic Prog.*, 2-3(13):103–179, 1992.

10. L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. In *STOC'01, 33rd Annual ACM Symposium on Theory of Computing*. ACM, 2001.
11. Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reability. Technical Report R90-13, Swedish Institute of Computer Science, December 1990.
12. Marta Z. Kwiatkowska, Gethin Norman, Roberto Segala, and Jeremy Sproston. Verifying quantitative properties of continuous probabilistic timed automata. Technical Report CSR-00-6, University of Birmingham, School of Computer Science, March 2000.
13. Marta Z. Kwiatkowska, Gethin Norman, Roberto Segala, and Jeremy Sproston. Verifying quantitative properties of continuous probabilistic timed automata. In C. Palamidessi, editor, *CONCUR 2000 - Concurrency Theory 11th International Conference*, number 1877 in LNCS. Springer, 2000.
14. Michael L. Littman, Anthony R. Cassandra, and Leslie Pack Kaelbling. Efficient dynamic-programming updates in partially observable markov decision processes. Technical Report CS-95-19, Brown University, 1995.
15. A. McIver. Reasoning about efficiency within a probabilistic μ -calculus. In *Proc. of PROBMIV*, pages 45–58, 1998. Technical Report CSR-98-4, University of Birmingham, School of Computer Science.
16. David Monniaux. Abstract interpretation of probabilistic semantics. In *Seventh International Static Analysis Symposium (SAS'00)*, number 1824 in Lecture Notes in Computer Science, pages 322–339. Springer Verlag, 2000. Extended version on the author's web site.
17. David Monniaux. An abstract analysis of the probabilistic termination of programs. In *8th International Static Analysis Symposium (SAS'01)*, number 2126 in Lecture Notes in Computer Science, pages 111–126. Springer Verlag, 2001.
18. David Monniaux. An abstract Monte-Carlo method for the analysis of probabilistic programs (extended abstract). In *28th Symposium on Principles of Programming Languages (POPL '01)*, pages 93–101. Association for Computer Machinery, 2001.
19. David Monniaux. Backwards abstract interpretation of probabilistic programs. In *European Symposium on Programming Languages and Systems (ESOP '01)*, number 2028 in Lecture Notes in Computer Science, pages 367–382. Springer Verlag, 2001.
20. David Monniaux. Abstraction of expectation functions using gaussian distributions. In Lenore D. Zuck, Paul C. Attie, Agostino Cortesi, and Supratik Mukhopadhyay, editors, *Verification, Model Checking, and Abstract Interpretation: VMCAI '03*, number 2575 in Lecture Notes in Computer Science, pages 161–173. Springer Verlag, 2003.
21. Jacques Neveu. *Mathematical Foundations of the Calculus of Probabilities*. Holden-Day, 1965.
22. Jacques Neveu. *Bases mathématiques du calcul des probabilités*. Masson et Cie, Éditeurs, Paris, 1970. Préface de R. Fortet. Deuxième édition, revue et corrigée.
23. Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley series in probability and mathematical statistics. John Wiley & Sons, 1994.
24. Roberto Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Massachusetts Institute of Technology, 1995. Technical report MIT/LCS/TR-676.
25. W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, vol. B*, pages 135–191. Elsevier, 1990.

A Technical lemmas

Lemma 11 For all f , t_0 and U_1 ,

$$\begin{aligned} & \sup_{(U_n)_{n \geq 2}} \int \lambda(t_1, \dots).f(\langle t_1, \dots \rangle) d[G(T, (U_n)_{n \in \mathbb{N}}).t_0] \\ &= \sup_{(U_n)_{n \geq 2}} \int \lambda(t_1, \dots).f(\langle t_1, \dots \rangle) d[G(T, (U_n)_{n \in \mathbb{N}}).t_0] \end{aligned}$$

Lemma 2. For all t_0 in X ,

$$E_+(\lambda_t.\chi_S(t_0).V_1(t) + \chi_{S^c}(t_0).V_2(t)).t_0 = \chi_S(t_0).(E_+(V_1).t_0) + \chi_{S^c}(t_0).(E_+(V_2).t_0).$$

Lemma 3. For any trace valuator f , for any g_1 and g_2 in $R(f)$, for any $A \subseteq \Omega$, the function g_3 defined by $g_3(\mathbf{t}) = g_1(\mathbf{t})$ if $t_0 \in A$, $g_3(\mathbf{t}) = g_2(\mathbf{t})$ otherwise, belongs to $R(f)$.

Lemma 4. Let Y be an ordered set. Let $\phi : (X \rightarrow I) \rightarrow Y$ be a monotonic, ω -upper-continuous function. Let K be a directed subset of $X \rightarrow I$. Then $\phi(\bigsqcup K) = \bigsqcup_{f \in K} \phi(f)$.

Lemma 5. Let T_1 and T_2 be two complete lattices. Let $\alpha : T_1 \rightarrow T_2$ be an ω -upper-continuous operator such that $\alpha(\perp) = \perp$. Let $\psi_1 : T_1 \rightarrow T_1$ and $\psi_2 : T_2 \rightarrow T_2$ be two ω -upper-continuous operators such that $\psi_2 \circ \alpha = \alpha \circ \psi_1$. Then $\alpha(\text{lfp } \psi_1) = \text{lfp } \psi_2$.

B Ionescu Tulcea's Theorem

The intuitive meaning of this theorem [21, 22, proposition V-I-1] is as follows: if $(E_t)_{t \in \mathbb{N}}$ is a sequence of measurable spaces and the $(P_{t+1}^{0, \dots, t})_{t \in \mathbb{N}}$ is a sequence of transition probabilities, respectively from $E_0 \times E_t$ to E_{t+1} , then we can construct a transition probability P from E_0 to $E_1 \times E_2 \times \dots$ such that for each $x_0 \in E_0$, $P(x_0, \cdot)$ is the probability distribution on traces starting from x_0 and following the transition probabilities $(E_t)_{t \in \mathbb{N}}$.

In an even more intuitive fashion: “knowing the starting probability measure, and the transition probabilities to the next states, we can construct the corresponding probability measure on infinite traces”.

Theorem 12 (Ionescu Tulcea) Let $(E_t, \mathcal{F}_t)_{t \in \mathbb{N}}$ be an infinite sequence of measurable spaces and, for any $t \in \mathbb{N}$, let $P_{t+1}^{0, \dots, t}$ be a transition probability relative to the spaces $(\prod_{s=0}^t E_s, \otimes_{s=0}^t \mathcal{F}_s)$ and $(E_{t+1}, \mathcal{F}_{t+1})$. Then there exists for any $x_0 \in E_0$ a unique probability P_{x_0} on

$$(\Omega, \mathcal{A}) = \prod_t (E_t, \mathcal{F}_t)$$

whose value for all measurable Cartesian product $\prod_t F_t$ is given by:

$$\begin{aligned} P_{x_0} \left[\prod_t F_t \right] &= \chi_{A_0}(x_0) \int_{x_1 \in F_1} P_1^0(x_0; dx_1) \int_{x_2 \in F_2} P_2^{0,1}(x_0, x_1; dx_2) \\ &\quad \dots \int_{x_T \in F_T} P_T^{0, \dots, T-1}(x_0, \dots, x_{T-1}; dx_T) \quad (47) \end{aligned}$$

as long as T is sufficiently great such that $F_t = E_t$ if $t > T$ (the second member is then independent of the chosen T). For any positive random variable Y on (Ω, \mathcal{A}) only depending on the coordinates up to T , we have:

$$\int_{\Omega} Y(\omega') P_{x_0}(d\omega') = \int_{F_1} P_1^0(x_0; dx_1) \int_{F_2} P_2^{0,1}(x_0, x_1; dx_2) \cdot \int_{x_T \in F_T} Y(x_0, \dots, x_T) P_T^{0, \dots, T-1}(x_0, \dots, x_{T-1}; dx_T) \quad (48)$$

Furthermore, for any positive random variable Y on (Ω, \mathcal{A}) ,

$$x_0 \mapsto \int Y(\omega') P_{x_0}(d\omega') \quad (49)$$

is a positive random variable on (E_0, \mathcal{F}_0) .