On the Verification of Qualitative Properties of Probabilistic Processes under Fairness Constraints

Christel Baier
Fakultät für Mathematik & Informatik
Universität Mannheim
68131 Mannheim, Germany
baier@pi1.informatik.uni-mannheim.de

Marta Kwiatkowska*
School of Computer Science
University of Birmingham, Edgbaston
Birmingham B15 2TT, UK
M.Z.Kwiatkowska@cs.bham.ac.uk

February 24, 1998

Abstract

We consider sequential and concurrent probabilistic processes and propose a general notion of fairness with respect to probabilistic choice, which allows to express various notions of fairness such as process fairness and event fairness. We show the soundness of proving the validity of qualitative properties of probabilistic processes under fairness constraints in the sense that whenever all fair executions of a probabilistic process fulfill a certain linear time property E then E holds for almost all executions (i.e. E holds with probability 1). It follows that in order to verify probabilistic processes w.r.t. linear time specifications, it suffices to establish that – for some instance of our general notion of fairness – all *fair* executions satisfy the specification. This generalizes the soundness results for extreme and α-fairness established in [25] and [27] respectively. Furthermore, we show that α-fairness of [27] is the only fairness notion which is complete for validity of qualitative linear time properties.

1 Introduction

Probabilistic techniques have proved successful in the specification and verification of systems that exhibit uncertainty, for example, fault-tolerant systems, distributed systems, and communication protocols. The verification of such a system, often presented as a probabilistic automaton (finite transition system extended with probabilities), usually aims to establish *qualitative properties*, i.e. properties that are fulfilled by almost all executions, which amounts to showing that the property is satisfied with probability 1, see e.g. [22, 16, 25, 15, 31, 26, 7, 1, 27]. The analysis of the average-case behaviour, as well as the investigation of the reliability and performance, require methods for proving *quantitative properties*, i.e. properties that express bounds on the probability of system evolutions, see e.g. [6, 14, 5, 2].

In the verification of non-probabilistic concurrent systems, it is well-known that certain liveness properties can only be established when appropriate fairness assumptions are made. In a model for probabilistic processes in which non-deterministic choice is replaced by probabilistic choice, one might argue that fairness assumptions are superfluous since the decision as to which step is taken next is not made arbitrarily, but according to a probability distribution. Nevertheless, probabilistic systems may behave unfairly, and hence – as in the non-probabilistic case – it is possible that certain liveness probabilities are violated in some executions while they hold in all fair executions.

*Supported in part by EPSRC grant GR/K42028.

In this paper we propose a general notion of fairness for probabilistic processes which allows to express e.g. state fairness, process fairness or event fairness (for a survey of fairness notions see e.g. [21, 28, 8, 18]). We consider (strong) fairness *w.r.t. probabilistic choice* in two classes of models for probabilistic systems. The first class consists of models based on (sequential) Markov chains which exhibit probabilistic choice but no non-deterministic choice; they are appropriate for representing the behaviour of sequential randomized algorithms or processes of a probabilistic process algebra with synchronous parallel composition as in e.g. [9, 11, 10, 20, 30]. The second (more general) class includes models based on Markov decision processes that allow for both probabilistic and non-deterministic choice; these are suitable for specifying concurrent randomized algorithms or processes of an asynchronous calculus with probabilistic and non-deterministic choice [16, 31, 26, 13, 32, 12, 29, 4]. By focussing on fairness w.r.t. probabilistic choice we differ from [16, 31], where fairness of a scheduler, i.e. fairness w.r.t. non-deterministic choice, for concurrent and non-deterministic probabilistic processes, is considered.

We suppose that the alternatives of the probabilistic choices are associated with 'labels', with each label denoting e.g. a process name or action. We say an execution sequence is *fair* if whenever a label is enabled infinitely many times then it is taken infinitely many times. Our main result (Theorem 1) is that, for *every* instance of our general notion of fairness, the set of fair execution sequences in finite-state or bounded systems has probability 1. This implies the soundness of proving qualitative properties under fairness constraints in the following sense: *if all fair execution sequences satisfy a certain linear time property E then E holds with probability 1* (cf. Corollary 1 and 4). Hence, in order to establish a (qualitative) linear time property E for a probabilistic process, it suffices to show that all *fair* execution sequences satisfy E for *some* instance of the fairness notion, which can be achieved with well-known, non-probabilistic, methods, e.g. by means of a deductive system or model-checking algorithms [23, 24]. In a model that distinguishes between non-deterministic and probabilistic choice, *extreme fairness* of [25] and *α-fairness* of [27] are instances of our general notion of fairness (Lemma 3.4). Hence, the soundness results established in [25, 27] are special cases of the theorems presented here. We cannot expect a general completeness result (in the sense that if a linear time property E holds with probability 1 for a probabilistic process then E holds on all fair execution sequences) as in [25] extreme fairness is shown to be incomplete. However, we are able to show that – in some sense – α-fairness (shown to be complete in [27]) is *the only* fairness notion that is complete for proving validity of qualitative linear time properties (Lemma 3.5).

2 Fair probabilistic transition systems

In this section we introduce a model for probabilistic processes based on Markov chains in which non-deterministic choice is replaced with probabilistic choice, and define (strong) fairness w.r.t. probabilistic choice.

A *probabilistic transition system* is a tuple $\mathcal{P} = (S, R)$ where S is a countable set of states and $R : S \times S \to [0,1]$ a function with $\sum_{t \in S} R(s,t) = 1$ for all $s \in S$. A *path* in $\mathcal{P}$ is a non-empty finite or infinite sequence $x = s_0 s_1 \ldots$ consisting of states $s_i \in S$ s.t. $R(s_i, s_{i+1}) > 0$. If x is finite then $last(x)$ denotes the last state of x. A *maximal path* is an infinite path in $\mathcal{P}$. If $x = s_0 s_1 \ldots$ then $|x|$ denotes the length of x, i.e. $|x| = k$ if $x = s_0 s_1 \ldots s_k$, and $|x| = \infty$ if x is infinite. $x^{(k)}$ denotes the prefix of x of length k. If $k \leq |x|$ then $x(k)$ denotes the k-th state of x, i.e. $x(k) = s_k$ if $x = s_0 s_1 \ldots$. $Path(s)$ denotes the set of maximal paths starting in s, $Path_{\rm fin}(s)$ the set of finite paths starting in s. The $[0,1]$-valued function R induces a probabilistic space on $Path(s)$ as follows. The probability $P(y)$, for finite paths y in $\mathcal{P}$, is defined by setting $P(y) = 1$ if $y = s$, and $P(y) = R(s_0, s_1) \cdot R(s_1, s_2) \cdot \ldots \cdot R(s_{n-1}, s_n)$ otherwise, in which case $y = s_0 s_1 \dots s_n$. Let $\Sigma(s)$ be the smallest σ-algebra on $Path(s)$ which contains the sets $\{x \in Path : y \text{ is a prefix of } x\}$, $y \in Path_{\text{fin}}(s)$. Let P be the unique measure on $\Sigma(s)$ such that $P\{x \in Path(s) : y \text{ is a prefix of } x\} = P(y)$.

Definition 2.1 *A* fair probabilistic transition system *is a tuple* $\mathcal{F} = (S, R, L, l)$ *where* (S, R) *is a probabilistic transition system,* L *is a non-empty countable set of labels and* $l : S \times S \to 2^L$ *a function with* $l(s,t) = \emptyset$ *if* $R(s,t) = 0$. *We say a label* $\ell \in L$ *is* enabled *in a state* s *iff* $\ell \in l(s,t)$ *for some* $t \in S$. ℓ *is said to be* taken *in the* i*-th step of a maximal path* $x = s_0 s_1 \dots$ *iff* $\ell \in l(s_i, s_{i+1})$. *A maximal path* x *is called* (strongly) fair *w.r.t.* ℓ *iff either* ℓ *is enabled only finitely many times in* x *or* ℓ *is taken infinitely often in* x. x *is called* fair *iff* x *is fair w.r.t. each label* $\ell \in L$.

The set L of labels should be thought of as an abstraction which allows to express various notions of fairness, of which we demonstrate a few for illustration purposes. First consider the models such as the 'generative' model of [11, 10], i.e. probabilistic transition systems where each step is associated with an action label $a \in \text{Act}$ in addition to a probability. To distinguish action names from their occurrences, we refer to the former as 'actions' and the latter as 'events'. We now define *event fairness* by setting L to the set of action labels, and $l(s,t)$ to (the singleton set consisting of) the action which is performed in the step from s to t. To see why we need sets of labels as the range of l, we consider *process fairness* for systems as above endowed with static, concurrently acting processes that synchronise on common actions. We set L to the set of process names, with each process $p \in L$ capable of performing actions from its alphabet set $\alpha(p) \subseteq \text{Act}$. The alphabets of processes need not be disjoint: each action belonging to the intersection $\alpha(p) \cap \alpha(q)$ requires simultaneous participation of p and q, with the remaining actions performed independently. For a step from s to t resulting from the execution of action a we set $l(s,t)$ to the set of all process names that contain a in its alphabet. Another variant of fairness that can be modelled is *interaction fairness*. Assume for simplicity pairwise interaction, then we can take L to be the set of *pairs* (p,q) where p, q are process names, and define $l(s,t)$ as the singleton set $\{(p,q)\}$ if the transition from s to t represents an interaction between p and q, and $\emptyset$ otherwise. A generalisation to multiway synchronisation is simply obtained by taking tuples of process names instead of pairs.

We use $Fair$ to denote the set of fair paths in $\mathcal{F}$ and $Fair_\ell$ the set of paths which are fair w.r.t. ℓ. For $s \in S$ we define $Fair(s) = Fair \cap Path(s)$ and $Fair_\ell(s) = Fair_\ell \cap Path(s)$. To see that $Fair(s)$ is measurable we first express fairness as a linear time formula, and then use a well-known result [31, 27] stating that the set of paths fulfilling a given linear time formula is measurable. Linear time formulas are built from: the truth values tt and $f\!f$, the atomic propositions $enabled(\ell)$ for each label ℓ, the usual Boolean connectives $\wedge$, $\vee$, $\neg$, $\to$, and the temporal operators $\Box$ ('always'), $\Diamond$ ('eventually') and a next-step operator X_ℓ for each label ℓ. The formulas are interpreted over paths of a fair probabilistic transition system. Given such a system $\mathcal{F} = (S, R, L, l)$ we define the satisfaction relation $\models$ as follows. Let $x = s_0 s_1 \dots$ be a maximal path. Then, $x \models enabled(\ell)$ iff ℓ is enabled in s_0 and $x \models X_\ell\, \varphi$ iff $\ell \in l(s_0, s_1)$ and $x' \models \varphi$ where $x' = s_1 s_2 \dots$. The other operators are interpreted in the usual way. As shown e.g. in [27], for a given formula φ the set of maximal paths x starting in a fixed state $s \in S$ such that $x \models \varphi$ is measurable. We define $taken(\ell) = X_\ell\, tt$ and $\varphi_\ell = \Box\Diamond enabled(\ell) \to \Box\Diamond taken(\ell)$. Then, $Fair_\ell(s) = \{x \in Path(s) : x \models \varphi_\ell\}$ and $Fair(s) = \bigcap_{\ell \in L} Fair_\ell(s)$ are measurable.

Our results rely on the boundedness [22, 15, 19] of (possibly infinite-state) probabilistic transition systems. A probabilistic transition system $\mathcal{P} = (S, R)$ is called *bounded* iff there exists a real number c with $0 < c < 1$ such that, for all $s, t \in S$, if $R(s,t) > 0$ then $R(s,t) \geq c$. Clearly, each finite-state probabilistic transition system is bounded.

Theorem 1 *If $\mathcal{F}$ is a bounded fair probabilistic transition system then $P(Fair(s)) = 1$ for all states s.*

Proof: Let $\mathcal{F} = (S, R, L, l)$ and $c > 0$ be a real number such that $R(s,t) > 0$ implies $R(s,t) \geq c$. It suffices to show that $P(Fair_\ell(s)) = 1$ for all $\ell \in L$. (Note that $P(\Gamma_i) = 1$ implies $P(\bigcap_i \Gamma_i) = 1$ which holds in each probabilistic space.) Let ℓ be a fixed label. For convenience, let $x\Gamma$ denote the set of paths $s_0 s_1 \ldots s_k t_1 t_2 \ldots$ such that $s_k t_1 t_2 \ldots \in \Gamma$, where $x = s_0 s_1 \ldots s_k$ is a finite path and Γ a set of infinite paths starting in $last(x) = s_k$. For $s \in S$ let Γ_s be the set of all paths $x \in Path(s)$ where ℓ is enabled infinitely often and which totally ignore ℓ-steps, i.e. Γ_s is the set of maximal paths $x \in Path(s)$ with $x \models \Box\Diamond enabled(\ell) \wedge \Box\neg taken(\ell)$. We show that $P(\Gamma_s) = 0$ and that $Path(s) \setminus Fair_\ell(s)$ can be written as a countable union of sets of the form $x\Gamma_t$ where x is a finite path with $last(x) = t$.

First we show that $P(\Gamma_s) = 0$ for all $s \in S$. Let $T = \{t \in S : t \models enabled(\ell)\}$. For $s \in S$ let Ω_s be the set of finite paths $x \in Path_{\text{fin}}(s)$ such that $|x| \geq 1$, $\ell \notin l(s, x(1))$, $x(i) \notin T$, $i = 1, \ldots, |x| - 1$, and $last(x) \in T$. For $t \in T$, $s \in S$, we define $\Omega_s^t = \{x \in \Omega_s : last(x) = t\}$. Clearly, Ω_s is countable and $\Omega_s = \bigcup_{t \in T} \Omega_s^t$ where $\Omega_s^t \cap \Omega_s^{t'} = \emptyset$ if $t \neq t'$. Thus,

$$(1) \qquad \sum_{x \in \Omega_s} P(x) \; = \; \sum_{t \in T} \sum_{x \in \Omega_s^t} P(x).$$

We have $\Gamma_s = \bigcup_{t \in T} \bigcup_{x \in \Omega_s^t} x\Gamma_t$ for all $s \in S$. As $x\Gamma_t \cap x'\Gamma_{t'} = \emptyset$ if $(x,t) \neq (x',t')$, and as $x\Gamma_t$ is a measurable set with $P(x\Gamma_t) = P(x) \cdot P(\Gamma_t)$ we obtain:

$$(2) \qquad P(\Gamma_s) \; = \; \sum_{t \in T} \sum_{x \in \Omega_s^t} P(x) \cdot P(\Gamma_t) \qquad \text{for all } s \in S.$$

Let $t \in T$. As $t \models enabled(\ell)$ there is some $s_t \in S$ with $\ell \in l(t, s_t)$. Since $R(t, s_t) \geq c$ we obtain:

$$(3) \qquad \sum_{x \in \Omega_t} P(x) \; \leq \; \sum_{s \neq s_t} R(t,s) \; \leq \; 1 \; - \; R(t, s_t) \; \leq \; 1 - c \quad \text{for all } t \in T.$$

We show by induction on k that $P(\Gamma_t) \leq (1-c)^k$ for all $t \in T$. In the basis of induction ($k = 0$) there is nothing to show. In the induction step ($k \Longrightarrow k+1$) we suppose that $P(\Gamma_t) \leq (1-c)^k$ for all $t \in T$. By (1), (2), (3) and the induction hypothesis we get for all $t \in T$:

$$P(\Gamma_t) = \sum_{u \in T} \sum_{x \in \Omega_t^u} P(x) \cdot P(\Gamma_u) \leq (1-c)^k \sum_{u \in T} \sum_{x \in \Omega_t^u} P(x) = (1-c)^k \sum_{x \in \Omega_t} P(x) \leq (1-c)^{k+1}.$$

We conclude $P(\Gamma_t) = 0$ for all $t \in T$. It follows that $P(\Gamma_s) = 0$ for all $s \in S$ (by (2)).

Next we show that $P(Fair_\ell(s)) = 1$ for all $s \in S$. It is clear that $Path(s) \setminus Fair_\ell(s) = \bigcup_{t \in S} \bigcup_{x \in \Lambda_s^t} x\Gamma_t$ where Λ_s^t denotes the set of all finite paths starting in s and ending in t. Note that Λ_s^t is countable and that $x\Gamma_t$ is measurable with $P(x\Gamma_t) = P(x) \cdot P(\Gamma_t) = 0$. Hence,

$$P(Path(s) \setminus Fair_\ell(s)) \; \leq \; \sum_{t \in S} \sum_{x \in \Lambda_s^t} P(x\Gamma_t) \; = \; 0$$

and $P(Fair_\ell(s)) = 1$. ■

If we drop the assumption that $\mathcal{F}$ is bounded then the probability of the fair paths might be less than 1. As a counter-example consider $\mathcal{F} = (S, R, L, l)$ where $S = \{t\} \cup \{s_0, s_1, \ldots\}$, $L = \{\ell\}$ and

$$R(s_i, v) \; = \; \begin{cases} 2^{-r_i} & : \text{ if } v = s_{i+1} \\ 1 - 2^{-r_i} & : \text{ if } v = t \\ 0 & : \text{ otherwise} \end{cases} \qquad l(s_i, v) \; = \; \begin{cases} \emptyset & : \text{ if } v = s_{i+1} \\ \{\ell\} & : \text{ if } v = t \end{cases}$$

and $R(t, s_i) = 0$, $R(t,t) = 1$, $l(t,t) = \emptyset$. Here, $(r_i)_{i \geq 0}$ is a sequence of positive reals where $\sum_{i \geq 0} r_i$ is convergent.

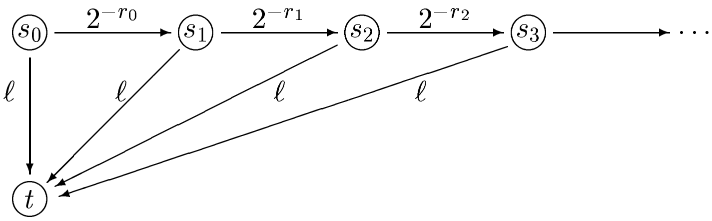

The maximal path $x = s_0 s_1 s_2 \ldots$ is unfair as ℓ is continuously enabled but never taken in x. Any other maximal path $y = s_0 s_1' s_2' \ldots$ starting in s_0 is fair as $s_i' = t$ for almost all i and ℓ is disabled in t. Hence, $P(Fair(s_0)) = 1 - 2^{-r} < 1$ where $r = \sum_{i \geq 0} r_i$.

Theorem 1 yields soundness of proving the validity of linear time formulas under fairness constraints in the following sense. We suppose a logic $\mathcal{L}$ and, for a fixed bounded probabilistic transition system (S, R), a satisfaction relation $\models \subseteq Path \times \mathcal{L}$ such that for each $s \in S$ and each formula φ of $\mathcal{L}$ the set $\{x \in Path(s) : x \models \varphi\}$ is measurable. s is called *φ-valid* iff $P\{x \in Path(s) : x \models \varphi\} = 1$. By Theorem 1 we obtain the following.

Corollary 1 *If $\mathcal{F}$ is a bounded fair probabilistic transition system such that $x \models \varphi$ for all $x \in Fair(s)$ then s is φ-valid.*

Corollary 2 *If $\mathcal{F}$ is a bounded fair probabilistic transition system, s a state of $\mathcal{F}$ and E a measurable subset of $Path(s)$, then $P(E \cap Fair(s)) = P(E)$.*

In particular, whenever φ is a linear time formula then the probability of the set of maximal paths fulfilling φ *equals* the probability of the set of fair maximal paths fulfilling φ. In other words, whether or not a (qualitative or quantitative) linear time property holds for a probabilistic process does not depend on whether fairness w.r.t. probabilistic choice is required.

3 Fairness for concurrent probabilistic processes

It was pointed out in [31] that certain states of a concurrent system whose components work asynchronously are inherently non-deterministic. The non-deterministic choices are beyond the control of the process and are assumed to be resolved by a scheduler or external intervention by the environment, whereas the convention is that the probabilistic choices are made by the system itself (e.g. by tossing a coin). In this section we introduce a model for concurrent processes which distinguishes non-deterministic and probabilistic choices. This model can be used to describe the interleaving behaviour of concurrent randomized systems [16, 31, 26, 13, 32, 12, 29, 4], but it is also suitable for representing underspecification that might be resolved in further refinement steps (cf. [17]). As in [25, 26, 27], and in contrast to [16, 31], our notion imposes fairness w.r.t. the internal decisions of the processes (the probabilistic choices), rather than fairness w.r.t. the (external) decisions of the scheduler (the non-deterministic choices).

Let S be a finite set and let $\mathcal{D}(S)$ denote the set of (probability) distributions on S, i.e. the set of functions $\mu : S \to [0, 1]$ such that $\sum_{t \in S} \mu(t) = 1$. A *concurrent probabilistic process* is a pair $\mathcal{C} = (S, \to)$ where S is a finite set of states and $\to$ a finite and total transition relation, i.e. a finite subset of $S \times \mathcal{D}(S)$ such that for each $s \in S$ there is some $\mu \in \mathcal{D}(S)$ with $(s, \mu) \in \to$. We write $s \to \mu$ rather than $(s, \mu) \in \to$. Let *Steps*(s) be the set of distributions μ with $s \to \mu$. (By the finiteness and the totality of $\to$, *Steps*(s) is non-empty and finite.) *Step* denotes the set $\bigcup_{s \in S}$ *Steps*(s). Intuitively, the transition relation $\to$ represents the non-deterministic alternatives in each state: given a state $s \in S$, a scheduler chooses a transition $s \to \mu$. The process itself resolves the probabilistic choice, i.e. selects some state t with $\mu(t) > 0$. Execution sequences (which we call paths) arise by resolving both the non-deterministic and probabilistic choices. Formally, a *path* in a concurrent probabilistic process $\mathcal{C}$ is a non-empty (finite or infinite) sequence

$$\pi = s_0 \stackrel{\mu_1}{\to} s_1 \stackrel{\mu_2}{\to} s_2 \ldots$$

where s_i are states, $\mu_{i+1} \in Steps(s_i)$ and $\mu_{i+1}(s_i) > 0$ for all i. (The case $\pi = s_0$ is allowed.) If π is finite then the last state of π is denoted by $last(\pi)$. The length $|\pi|$ of a path is defined in the usual way as follows: if $\pi = s_0 \in S$ then $|\pi| = 0$; otherwise, $\pi = s_0 \stackrel{\mu_1}{\to} s_1 \stackrel{\mu_2}{\to} \ldots \stackrel{\mu_n}{\to} s_n$, in which case $|\pi| = n$. For infinite π we put $|\pi| = \infty$. A path π is called a *maximal path* iff it is infinite. Let $\pi = s_0 \stackrel{\mu_1}{\to} s_1 \stackrel{\mu_2}{\to} \ldots$ be a maximal path in $\mathcal{C}$ and $k \geq 0$. We put $step(\pi, k) = \mu_k$, $\pi(k) = s_k$ and $\pi^{(k)} = s_0 \stackrel{\mu_1}{\to} s_1 \stackrel{\mu_2}{\to} \ldots \stackrel{\mu_k}{\to} s_k$ (the k-th prefix of π). $Path$ denotes the set of all maximal paths in $\mathcal{C}$, and $Path_{\text{fin}}$ the set of all finite paths in $\mathcal{C}$. If $E \subseteq Path$ and $s \in S$ then $E(s)$ denotes the set of paths $\pi \in E$ which start in s.

An *adversary* (or *scheduler* or *policy*) of $\mathcal{C}$ is a function A mapping every finite path ω of $\mathcal{C}$ to a distribution $A(\omega)$ on S such that $A(\omega) \in Steps(last(\omega))$. Intuitively, an adversary resolves the non-deterministic choices (but not the probabilistic choices) by choosing, for every finite path ω in $\mathcal{C}$, an outgoing transition from $last(\omega)$. Every adversary induces a 'sequential component' which is also a probabilistic transition system. If A is an adversary of $\mathcal{C}$ then we put $\mathcal{P}^A = (Path_{\text{fin}}, R^A)$ where $R^A(\omega, \omega \stackrel{A(\omega)}{\to} s) = A(\omega)(s)$ and $R^A(\omega, \omega') = 0$ for all finite paths ω' which are not of the form $\omega \stackrel{A(\omega)}{\to} s$. $\mathcal{P}^A$ is a bounded probabilistic transition system. The boundedness follows by the finiteness of S and $\to$, and the fact that $R^A(\omega, \omega') = 0$ or $R^A(\omega, \omega') \geq \min\{\mu(s) : \mu \in Steps, s \in S, \mu(s) > 0\}$. Let $Path^A$ be the set of all paths $\pi \in Path(\mathcal{C})$ with $step(\pi, i) = A(\pi^{(i)})$. If $E \subseteq Path$ then $E^A = E \cap Path^A$.

Definition 3.1 *A* fairness condition *on $\mathcal{C}$ is a pair (L, l) consisting of a non-empty countable set L of labels and a function l that assigns to each pair (ω, ω') of finite paths in $\mathcal{C}$, where ω' is of the form $\omega \stackrel{\mu}{\to} s$ (with $\mu \in Steps(last(\omega))$ and $\mu(s) > 0$), a subset $l(\omega, \omega')$ of L. If π is a maximal path then we say a label $\ell \in L$ is* enabled *in the i-th step of π iff $\ell \in l(\pi^{(i)}, \pi^{(i)} \stackrel{\mu}{\to} s)$ for some state s where $\mu = step(\pi, i)$. We say ℓ is* taken *in the i-th step of π iff $\ell \in l(\pi^{(i)}, \pi^{(i+1)})$. If $\ell \in L$ then π is called (L, l)-*fair w.r.t. *ℓ iff either ℓ is enabled only finitely many times in π or ℓ is taken infinitely many times in π. π is called (L, l)-*fair *iff π is (L, l)-fair w.r.t. each label $\ell \in L$.*

This definition of "enabled" reflects the assumption that the non-deterministic choices are not under the control of the system. If the adversary chooses a step μ such that ℓ can be performed in some of the probabilistic alternatives offered by μ, then ℓ is viewed to be enabled. $Fair_{(L,l,\ell)}$ denotes the set of maximal paths π which are (L, l)-fair w.r.t. ℓ, and $Fair_{(L,l)}$ the set of maximal paths π which are (L, l)-fair. It is easy to see that a maximal path $\pi \in Path^A$ is (L, l)-fair (w.r.t. ℓ) if and only if π is fair (w.r.t. ℓ) as a path in the fair probabilistic transition system $(\mathcal{P}^A, L, l)$, where we identify π with the path $x = \pi^{(0)}\pi^{(1)}\pi^{(2)}\ldots$ in $\mathcal{P}^A$. Theorem 1 yields $P(Fair^A_{(L,l)}(s)) = 1$ for all $s \in S$. Thus, $P(E \cap Fair_{(L,l)}) = P(E)$ for each measurable subset E of $Path^A(s)$.

As in the previous section, we suppose a logic $\mathcal{L}$ and a satisfaction relation $\models \subseteq Path \times \mathcal{L}$ such that, for each $s \in S$, each adversary A and each formula φ of $\mathcal{L}$, the set $\{x \in Path^A(s) : x \models \varphi\}$ is measurable. We then obtain:

Corollary 3 *If (L, l) is a fairness condition and φ a formula of $\mathcal{L}$ then $P\{x \in Fair^A_{(L,l)}(s) : x \models \varphi\} = P\{x \in Path^A(s) : x \models \varphi\}$ for all adversaries A.*

The state s is called *φ-valid* iff $P\{x \in Path^A(s) : x \models \varphi\} = 1$ for all adversaries A. Furthermore, the soundness of proving the validity of linear time formulas under (L,l)-fairness follows.

Corollary 4 *Whenever (L,l) is a fairness condition and φ is a formula of $\mathcal{L}$ such that $\pi \models \varphi$ for all maximal paths $\pi \in Fair_{(L,l)}(s)$ then s is φ-valid.*

In [25] and [27] notions of extreme fairness and α-fairness are introduced for concurrent probabilistic processes. We show that extreme and α-fairness are instances of fairness conditions as defined above. In the approach of [25, 27], each concurrent probabilistic process $(S, \rightarrow)$ is assigned a set of *transitions*, where each transition τ is associated with an enabling predicate $S_\tau \subseteq S$ and a set $\{\tau^1, \ldots, \tau^k\}$ of *modes*. Each mode τ^i is associated with a probability and a set of possible next steps. The definition of extreme fairness [25] employs a collection $\mathcal{Q}$ of *state predicates* (described by first order formulas), whereas α-fairness [27] uses linear time logic with past operators. We now define extreme and α-fairness. Let $\mathcal{C} = (S, \rightarrow)$ be a concurrent probabilistic process which is equipped with a function $\mathcal{T} : Steps \rightarrow 2^S$ such that $\mu \in Steps(s)$ iff $s \in \mathcal{T}(\mu)$. The transitions correspond to the distributions $\mu \in Steps$, while the modes of a transition μ are the states $s \in S$ with $\mu(s) > 0$. (That is, the set of possible next steps of a mode consists of a single state.) For the definition of extreme fairness we suppose a set $\mathcal{Q} \subseteq 2^S$ (where each element $Q \in \mathcal{Q}$ represents a state predicate).

Definition 3.2 *A maximal path π in $\mathcal{C}$ is called* extremely fair *iff, for each $Q \in \mathcal{Q}$, each $\mu \in Steps$ and each mode s of μ, whenever $step(\pi, i) = \mu$ for infinitely many $i \geq 0$ with $\pi(i) \in Q$ then there are infinitely many indices $i \geq 0$ with $\pi(i) \in Q$, $step(\pi, i) = \mu$ and $\pi(i+1) = s$.*

To define α-fairness we suppose $\mathcal{X}$ to be a set consisting of subsets of $Path_{\text{fin}}$. We may assume that the elements χ of $\mathcal{X}$ represent past formulas of some linear time logic; if χ is a past formula then we identify χ with the set of all finite paths ω such that each maximal path π, where ω is a prefix of π, fulfills χ.

Definition 3.3 *Let $\mu \in Steps$ and $\chi \in \mathcal{X}$. A maximal path π in $\mathcal{C}$ is called* α-fair w.r.t. μ and χ *iff whenever there are infinitely many indices i with $\pi^{(i)} \in \chi$ and $step(\pi, i) = \mu$ then for each mode s of μ there are infinitely many indices j with $\pi^{(j)} \in \chi$, $step(\pi, j) = \mu$ and $\pi(j+1) = s$. π is called* α-fair *iff π is α-fair w.r.t. each $\chi \in \mathcal{X}$ and $\mu \in Steps$.*

The next lemma shows that extreme and α-fairness are instances of fairness conditions in our sense:

Lemma 3.4 *Let π be a maximal path in $\mathcal{C}$. Then:*

(a) π is extremely fair if and only if π is (L_{efair}, l_{efair})-fair where L_{efair} is the set of tuples (Q, μ, s) such that $Q \in \mathcal{Q}$, $\mu \in Steps$ and $\mu(s) > 0$ and where l_{efair} is given by:

$$l_{efair}\left(\omega, \omega \xrightarrow{\mu} s\right) \;=\; \{\ (Q, \mu, s) \in L_{efair} \ :\ last(\omega) \in Q\ \}.$$

(b) π is α-fair if and only if π is $(L_{\alpha fair}, l_{\alpha fair})$ where $L_{\alpha fair}$ is the set of tuples (χ, μ, s) with $\chi \in \mathcal{X}$, $\mu \in Steps$ and $\mu(s) > 0$ and where $l_{\alpha fair}$ is given by:

$$l_{\alpha fair}(\omega, \omega \xrightarrow{\mu} s) \;=\; \{\ (\chi, \mu, s) \in L_{\alpha fair}\ :\ \omega \in \chi\ \}.$$

Proof: We only show (a) as (b) can be shown similarly. For simplicity, we write L and l instead of L_{efair} and l_{efair} respectively. Let π be a maximal path in $\mathcal{C}$.

'only if': Let π be extremely fair and let $\ell = (Q, \mu, s) \in L$ such that ℓ is infinitely often enabled in π. Let I be the set of indices $i \geq 0$ such that ℓ is enabled in the i-th state of π. Then,

$\pi(i) \in Q$ and $step(\pi, i) = \mu$ for all $i \in I$. As π is extremely fair there exists an infinite subset J of I such that $\pi(j+1) = s$ for all $j \in J$. Hence, $\ell \in l(\pi^{(j)}, \pi^{(j+1)})$ for all $j \in J$, i.e. ℓ is taken infinitely often in π.

'if': We suppose π to be (L,l)-fair and $step(\pi, i) = \mu$ for infinitely many indices i with $\pi(i) \in Q$. Let s be a mode of μ and let $\ell = (Q, \mu, s)$. Then, ℓ is enabled infinitely often in π. Hence, ℓ is taken infinitely often in π, i.e. there are infinitely many indices j with $\ell \in l(\pi^{(j)}, \pi^{(j+1)})$. For each such index j, $\pi(j) \in Q$, $\mu = step(\pi, j)$ and $\pi(j+1) = s$. Thus, π is extremely fair. ■

From Lemma 3.4(a) we can deduce that our soundness result (Corollary 4) is a generalization of the result of [25] which states the soundness of proving qualitative properties under extreme fairness. In [27] it is shown that, for each state s and each linear time formula φ, s is φ-valid if and only if $\pi \models \varphi$ holds for all α-fair paths $\pi \in Path(s)$. The 'if'-part is an instance of Corollary 4, whereas the 'only-if'-part (the completeness of the α-fairness approach) is not. The reason for this is that a general completeness result cannot be established, as it is shown in [25] that extreme fairness is not a necessary condition for the validity of linear time formulas. In the remainder of this section, we show that α-fairness is the only fairness notion which is complete for verifying qualitative properties expressed by linear time formulas.

We suppose that formulas of the linear time logic $\mathcal{L}$ are built from the truth values tt and $f\!f$, atomic propositions, the usual Boolean connectives, and the temporal operators $\mathcal{U}$ ('until'), $\mathcal{S}$ ('since'), X^{-1} ('previous step') and labelled next-step operators X_μ, $\mu \in Steps$. (The usual next-step operator X can be derived from the labelled next-step operators by putting $X\varphi = \bigvee_\mu X_\mu \varphi$.) $\mathcal{L}_{past}$ denotes the set of *past formulas*, i.e. formulas which are built from atomic propositions, the Boolean connectives and the operators $\mathcal{S}$ and X^{-1}. We fix a concurrent probabilistic process $\mathcal{C} = (S, \rightarrow)$ together with a satisfaction relation $\models \subseteq Path \times I\!N \times \mathcal{L}$ (where $I\!N$ is the set of non-negative integers) with $(\pi, j) \models X_\mu \varphi$ iff $step(\pi, j) = \mu$ and $(\pi, j+1) \models \varphi$. The remaining operators are interpreted in the usual way. The satisfaction relation $\models \subseteq Path \times \mathcal{L}$, as used earlier, is given by $\pi \models \varphi$ iff $(\pi, 0) \models \varphi$. Let $E_\varphi = \{\pi \in Path : \pi \models \varphi\}$. Then, s is *φ-valid* iff $P(E_\varphi^A(s)) = 1$ for all adversaries A. For a past formula ψ and a finite path ω with $|\omega| = j$, we define $\omega \models \psi$ iff $(\pi, j) \models \psi$ for all maximal paths π with $\pi^{(j)} = \omega$ (or equivalently, iff $(\pi, j) \models \psi$ for some maximal path π with $\pi^{(j)} = \omega$). Let χ_ψ be the set of finite ω with $\omega \models \psi$ and $\mathcal{X} = \{\chi_\psi : \psi \in \mathcal{L}_{past}\}$ and let $(L_{\alpha fair}, l_{\alpha fair})$ be as in Lemma 3.4(b). We write $\alpha Fair$ instead of $Fair_{(L_{\alpha fair}, l_{\alpha fair})}$.

Let (L, l) be a fairness condition. (L, l) is called *complete* (for verifying qualitative properties expressed as formulas of $\mathcal{L}$) iff, for all linear time formulas φ of $\mathcal{L}$ and all states $s \in S$, if s is φ-valid then $\pi \models \varphi$ for all $\pi \in Fair_{(L,l)}$. It is easy to see that the completeness result of [27] (where labelled next-step operators are not used) carries over to $\mathcal{L}$, i.e. if s is φ-valid then $\alpha Fair(s) \subseteq E_\varphi(s)$. Thus, $(L_{\alpha fair}, l_{\alpha fair})$ is complete. $\mathcal{L}$ is called *expressive for* (L, l) iff for each $\ell \in L$ there exists a formula φ of $\mathcal{L}$ with $E_\varphi = Fair_{(L,l,\ell)}$. We may assume that for each state $s \in S$ there is an atomic proposition a_s with $(\pi, j) \models a_s$ iff $\pi(j) = s$. Then, $\mathcal{L}$ is expressive for $(L_{\alpha fair}, l_{\alpha fair})$ as, for $\ell = (\chi_\psi, \mu, s)$ and $\varphi_\ell = \Box\Diamond enabled(\ell) \rightarrow \Box\Diamond(\psi \wedge X_\mu a_s)$ where $enabled(\ell) = \psi \wedge X_\mu tt$, we have that $\pi \models \varphi_\ell$ iff π is $(L_{\alpha fair}, l_{\alpha fair})$-fair w.r.t. ℓ.

Lemma 3.5 *If $\mathcal{L}$ is expressive for (L, l) then (L, l) is complete iff $Fair_{(L,l)} = \alpha Fair$.*

Proof: It suffices to show that if (L, l), (L', l') are fairness conditions such that $\mathcal{L}$ is expressive for (L, l) and (L', l') is complete, then $Fair_{(L',l')}(s) \subseteq Fair_{(L,l,\ell)}(s)$ for all $s \in S$ and $\ell \in L$. Since $\mathcal{L}$ is expressive for (L, l) there is a formula φ with $E_\varphi = Fair_{(L,l,\ell)}$. Since $P(E_\varphi^A(s)) = P(Fair_{(L,l,\ell)}^A(s)) = 1$ for all adversaries A we obtain φ-validity of s. Hence, $Fair_{(L',l')}(s) \subseteq Fair_{(L,l,\ell)}(s)$ by completeness of (L', l'). ■

4 Conclusion

We presented a general notion of fairness for probabilistic systems which imposes strong fairness w.r.t. probabilistic choices, and showed soundness of proving qualitative properties, i.e. those satisfied with probability 1, under fairness constraints (Corollary 1 and 4). It should be noted that all our results (Theorem 1 and its corollaries) carry over to weaker fairness notions (e.g. weak fairness) as each superset of the set of fair maximal paths has probability 1.

The import of our results is that in order to demonstrate the validity of qualitative linear time properties φ for probabilistic processes it suffices to show – for some instance of our general fairness notion – that φ holds for all *fair* execution sequences. This allows one, given an instance of our fairness notion, to reduce the verification of qualitative linear time properties of probabilistic processes to the non-probabilistic case: rather than compute the exact probabilities of the set of paths fulfilling φ, it is sufficient to establish by means of well-known *non-probabilistic* methods (e.g. model checking [23, 24]) that φ holds for all fair execution sequences. Moreover, we have shown that the probability of the set of maximal paths fulfilling a certain linear time formula φ does not depend on whether or not fairness w.r.t. probabilistic choice is required (Corollary 2 and 3). In other words, whether a certain (qualitative or quantitative) temporal property holds for a probabilistic process is independent of the presence or absence of fairness assumptions w.r.t. probabilistic choice. One might wonder why such a result is possible, since in the non-probabilistic case it is folklore knowledge that certain liveness properties cannot be established without suitable fairness assumptions. It is worth noting that φ-validity of a state s in a probabilistic transition system is weaker than φ-validity in the corresponding non-probabilistic transition system. (In the non-probabilistic case, a state s of a transition system is said to be φ-valid iff all maximal paths starting in s satisfy φ, whereas in the probabilistic case, φ-validity requires that φ holds for almost all maximal paths starting in s.) This observation no longer applies when fairness w.r.t. non-deterministic choices (fairness of the scheduler) is considered [16, 31, 3]: as in the non-probabilistic case, it is possible that certain liveness properties for a concurrent probabilistic process cannot be established unless fairness of the scheduler is assumed. Hence, from a purely descriptive point of view, fairness w.r.t. probabilistic choice is irrelevant, whereas fairness w.r.t. non-deterministic choice is not.

References

[1] R. Alur, C. Courcoubetis, D. Dill: Model-Checking for Probabilistic Real-Time Systems, Proc. ICALP'91, Lecture Notes in Computer Science 510, pp 115-127, Springer, 1991.

[2] L. de Alfaro: Formal Verification of Performance and Reliability of Real-Time Systems, Technical Report, Stanford University, 1996.

[3] C. Baier, M.Z. Kwiatkowska: Model Checking for a Probabilistic Branching Time Logic with Fairness, to appear in Distributed Computing.

[4] C. Baier, M.Z. Kwiatkowska: Domain Equations for Probabilistic Processes, Proc. EXPRESS'97, Electronic Notes in Theoretical Computer Science, Vol. 7, URL: http://www.elsevier.nl/locate/entcs/volume7.html, 1997.

[5] A. Bianco, L. de Alfaro: Model Checking of Probabilistic and Nondeterministic Systems, Proc. Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science 1026, pp 499-513, Springer, 1995.

[6] L. Christoff, I. Christoff: Reasoning about Safety and Liveness Properties for Probabilistic Processes, Proc. Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science 652, pp 342-355, Springer, 1992.

[7] C. Courcoubetis, M. Yannakakis: Verifying Temporal Properties of Finite-State Probabilistic Programs, Proc. FOCS'88, pp 338-345, ACM Press, 1988.

[8] N. Francez: Fairness, Springer-Verlag, New York, 1988.

[9] A. Giacalone, C. Jou, S. Smolka: Algebraic Reasoning for Probabilistic Concurrent Systems, Proc. IFIP TC2 Working Conference on Programming Concepts and Methods, 1990.

[10] R. van Glabbeek, S. Smolka, B. Steffen: Reactive, Generative, and Stratified Models for Probabilistic Processes, Information and Computation, Vol. 121, No. 1, pp 59-80, 1995.

[11] R. van Glabbeek, S. Smolka, B. Steffen, C. Tofts: Reactive, Generative, and Stratified Models for Probabilistic Processes, Proc. LICS'90, pp 130-141, IEEE CS Press, 1990.

[12] H. Hansson: Time and Probability in Formal Design of Distributed Systems, Real-Time Safety Critical Systems, Elsevier, 1994.

[13] H. Hansson, B. Jonsson: A Calculus for Communicating Systems with Time and Probabilities, Proc. 11th IEEE Symposium on Real Time Systems, 1990.

[14] H. Hansson, B. Jonsson: A Logic for Reasoning about Time and Probability, Formal Aspects of Computing, Vol. 6, pp 512-535, 1994.

[15] S. Hart, M. Sharir: Probabilistic Temporal Logic for Finite and Bounded Models, Proc. 16th ACM Symposium on Theory of Computing, pp 1-13, 1984.

[16] S. Hart, M. Sharir, A. Pnueli: Termination of Probabilistic Concurrent Programs, ACM Transactions on Programming Languages, Vol. 5, pp 356-380, 1983.

[17] B. Jonsson, W. Yi: Compositional Testing Preorders for Probabilistic Processes, Proc. LICS'95, pp 431-443, IEEE CS Press, 1995.

[18] M. Kwiatkowska: Survey of Fairness Notions, Information and Software Technology, Vol. 31, No. 7, pp 371-386, 1989.

[19] K. Larsen, A. Skou: Bisimulation through Probabilistic Testing, Information and Computation, Vol. 94, pp 1-28, 1991.

[20] K. Larsen, A. Skou: Compositional Verification of Probabilistic Processes, Proc. CONCUR'92, Lecture Notes in Computer Science 630, pp 456-471, Springer, 1992.

[21] D. Lehmann, A. Pnueli, J. Stavi: Impartiality, Justice and Fairness: The Ethics of Concurrent Termination, Proc. ICALP'81, Lecture Notes in Computer Science 115, Springer, 1981.

[22] D. Lehmann, S. Shelah: Reasoning with Time and Chance, Information and Control, Vol. 53, pp 165-198, 1982.

[23] O. Lichtenstein, A. Pnueli: Checking that Finite State Concurrent Programs Satisfy Their Linear Specification, Proc. 12th ACM Symp. on Principles of Programming Languages, pp 97-107, 1985.

[24] Z. Manna, A. Pnueli: Temporal Verification of Reactive Systems: Safety, Springer-Verlag, 1995.

[25] A. Pnueli: On the Extremely Fair Treatment of Probabilistic Algorithms, Proc. 15th ACM Symposium on Theory of Computing, 1983.

[26] A. Pnueli, L. Zuck: Verification of Multiprocess Probabilistic Protocols, Distributed Computing, Vol. 1, No. 1, pp 53-72, 1986.

[27] A. Pnueli, L. Zuck: Probabilistic Verification, Information and Computation, Vol. 103, pp 1-29, 1993.

[28] J. Queille, J. Sifakis: Fairness and Related Properties in Transition Systems – a Temporal Logic to deal with Fairness, Acta Informatica, Vol. 19, pp 195-220, 1983.

[29] R. Segala: Modelling and Verification of Randomized Distributed Real-Time Systems, Ph.D.Thesis, Massachusetts Institute of Technology, 1995.

[30] C. Tofts: Processes with Probabilities, Formal Aspects of Computing, Vol. 6, No. 5, 1994.

[31] M. Vardi: Automatic Verification of Probabilistic Concurrent Finite-State Programs, Proc. FOCS'85, pp 327-338, ACM Press, 1985.

[32] W. Yi, K. Larsen: Testing Probabilistic and Nondeterminitic Processes, Proc. Protocol Specification, Testing and Verification XII, Elsevier Science Publishers B.V. (North-Holland), pp 47-61, 1992.