

NEW PERSPECTIVES ON FAIRNESS

Daniele Varacca
Imperial College London, UK

Hagen Völzer
Universität zu Lübeck, Germany

Abstract

We define when a linear-time temporal property is a *fairness property* with respect to a given system. This captures the essence that is shared by most fairness assumptions that are used in the specification and verification of reactive concurrent systems, such as weak fairness, strong fairness, k -fairness, and many others. We give three characterisations for the family of all fairness properties: a language-theoretic, a topological, and a game-theoretic characterisation. It turns out that the fairness properties are the “large” sets from a topological point of view, i.e., they are the *co-meager* sets in the natural topology of runs of a given system. This insight provides a link to probability theory where a set is “large” when it has measure 1. While these two notions of largeness are very similar, they do not coincide in general. However, we show that they coincide for ω -regular properties and bounded Borel measures. That is, an ω -regular temporal property of a finite-state system has measure 1 under a bounded Borel measure if and only if it is a fairness property with respect to that system.

1 Introduction

When we model a concurrent system, we often make use of *nondeterminism*. Nondeterminism abstracts away from different scheduling policies or differences in speed of different parts of the system. Also, if we consider reactive systems, we use nondeterminism to allow different possible interactions with the environment. Furthermore, nondeterminism is used to model freedom of implementation.

A specification for a nondeterministic model of a system must allow several different behaviours. Specifications can thus be seen as sets of behaviours. We then say that a model satisfies the specification if all possible behaviours of the model belong to the specification.

Examples of specifications are *safety* and *liveness*. A safety specification informally requires that “some finite bad thing does not happen”. If a behaviour violates the safety specification, we can recognise this in finite time. Once the “bad thing” has happened, any extension of the behaviour will violate the safety specification. A liveness specification informally requires that “some (possibly infinite) good thing will eventually happen”. No finite behaviour should violate the specification, and therefore at any finite time we still have the possibility to eventually obtain a behaviour that belongs to the liveness specification.

When a model does not satisfy the specification, this can happen for several reasons. The model could be flawed and should be redesigned or the specification could be incorrect. Often, however, some behaviour of the model is not allowed by the specification, but such behaviour is “unlikely” to happen. How do we formalise this notion of unlikelihood?

We introduce nondeterminism to abstract away from some details of the implementation, but in some cases we may be abstracting away too much. For instance, if the nondeterminism is used to abstract away from scheduling policies, we could introduce some behaviour that no concrete scheduling policy would allow. The interaction with the environment can also be considered as a form of scheduling. Also in this case, some patterns of interaction may be allowed by the model, but they might not be happening in practice. This is a first sense in which a behaviour is unlikely.

To deal with this problem, we make use of *fairness assumptions*. Informally, a fairness assumption is an abstract description of a class of schedulers (or environments). A fairness assumption is a set of behaviours that are considered to be “fair”. A model satisfies a specification under a fairness assumption, if all behaviours of the model that violate the specification are “unfair”.

When can a set of behaviours be considered a fairness assumption? Informally, a scheduler is fair with respect to some (finite) behaviour if, whenever the behaviour is sufficiently often possible, then the scheduler guarantees it to happen sufficiently often. But how do we characterise this intuition formally? How do we formalise “possible”, and “sufficiently often”? We will present, by means of examples, different degrees of “possible” and “sufficiently often”. We will then show a formal characterisation of fairness that subsumes all the examples we present.

Another possible formalisation of unlikelihood is by means of probability theory. If the set of behaviours is endowed with a probability measure, we say that a set of behaviours is unlikely, if its probability is 0. In this sense a model satisfies a specification, if the set of the behaviours that violate the specification has probability 0.

We will compare this notion of probabilistic unlikelihood with the above notion of fairness, observing the similarities and the differences.

2 Examples of Fairness

We will show here some simple examples of the use of fairness assumptions.

2.1 Maximality

Consider the following system (Fig. 1), represented as a safe Petri net.




Figure 1: A simple process

As such, the system only says what can and what cannot happen. It does not say that something must happen at all. To say that something must happen, we can use the maximality assumption, which says that the system does not arbitrarily stop the computation. More precisely, a run (i.e., firing sequence) is *maximal* if it is infinite or if its final state does not enable any transition of the system. In the considered system, this means that after every a , there must be a b and that after every b , there must be an a . This leaves only the run $(ab)^\omega$, which is the unique maximal run of the system. Therefore, the system satisfies the property “infinitely often a ” under the maximality assumption.

2.2 Weak fairness

Consider now the following system (Fig. 2) and assume maximality.




Figure 2: Two independent processes

Then, that system does not satisfy “infinitely often a ” because the maximal run $(cd)^\omega$ does not. Although the overall system does not stop in this run, one of its components does.

In order to rule out such a behaviour, we assume *weak fairness* [15]. A run is *weakly fair* with respect to transition t if t is taken infinitely often or t is always eventually disabled. Therefore, the maximal run $(cd)^\omega$ is not weakly fair with respect to a . The system does in fact satisfy “infinitely often a ” under weak fairness with respect to a and b .

Weak fairness with respect to all transitions is strictly stronger than maximality.

2.3 Strong fairness

In the next system below (Fig. 3), weak fairness is not sufficient to establish “infinitely often a ” because the run $(cd)^\omega$ is weakly fair with respect to all transitions of the system. In particular, it is weakly fair with respect to a because a is always eventually disabled.




Figure 3: Two processes sharing a resource

However, we can consider $(cd)^\omega$ unfair with respect to a because a is infinitely often enabled but never taken. This kind of unfairness is captured by the notion of *strong fairness* [15]. A run is *strongly fair* with respect to a transition t if t is taken infinitely often or t is eventually always disabled. Strong fairness with respect to a and weak fairness with respect to b then establish “infinitely often a ” in the system.

Strong fairness is obviously stronger than weak fairness.

2.4 k -Fairness

In the next system below (Fig. 4), strong fairness with respect to all transitions fails to establish “infinitely often e ”, because the run $(abcd)^\omega$ violates it but is

strongly fair. In particular, it is strongly fair with respect to e because e is never enabled.




Figure 4: Two processes sharing an action

Among the fairness notions that establish “infinitely often e ”, there is the notion of *strong k -fairness* [8] for $k \geq 1$. A run is *strongly k -fair* with respect to transition t if t is infinitely often taken or t is eventually never k -enabled, where t is k -enabled in a state s if there is a path of the system of length not more than k that starts in s and ends in a state that enables t . Weak fairness for all transitions and strong 1-fairness for e indeed establish “infinitely often e ”.

Strong $(k + 1)$ -fairness is clearly stronger than strong k -fairness, and strong 0-fairness coincides with strong fairness.

Remark. The “unfairness” arising in the system in Fig. 4 is also known from the variant of the Dining Philosophers in which a philosopher picks up both his forks at the same time to eat. There, a philosopher may starve because his two neighbours “conspire” against him by eating alternately in such a way that his two forks are never available at the same time. Note that transition e in Fig. 4 needs two resources (B and C) at the same time. There are fairness notions that better capture the “unfairness” in this example (cf. [5, 25, 26]). However, we do not introduce them in detail here.

2.5 ∞ -Fairness

Consider now the following infinite-state system (Fig. 5).




Figure 5: A nondeterministic walk on the integer line

Suppose we are interested here in the property “state 0 is visited infinitely often”. This property is not established by strong k -fairness for any k because the diverging run $a_1a_2\dots$ is strongly k -fair with respect to any transition for any $k \geq 0$. However, we can use the stronger notion of *strong ∞ -fairness* [8]. A run is *strongly ∞ -fair* with respect to a transition t , if t is infinitely often taken or t is eventually never ∞ -enabled, where t is ∞ -enabled in a state s if there is a path of the system (of any length) that starts in s and ends in a state that enables t . It is easy to see that ∞ -fairness with respect to a_0 and b_0 establishes the required specification.

2.6 Fairness with respect to words

While strong ∞ -fairness with respect to transitions is very strong, there are still some useful specifications that are not established by it. As an example, consider the following system and the specification “the finite word ba of transitions occurs infinitely often”. The run $(abcd)^\omega$ does not satisfy the specification but it is strongly ∞ -fair with respect to every transition, since every transition is taken infinitely often in this run. In such a case, we can extend the above fairness notions and define them with respect to finite words of transitions rather than with respect to a single transition only. For example, we can see that strong fairness with respect to the word ba establishes the specification considered above.




Figure 6: A recurrent free choice

2.7 Other examples

Another remarkable notion is *equifairness* [9]. Equifairness with respect to a and c in Fig. 6 prescribes that each fair run has infinitely many positions such that the number of previous occurrences of a equals the number of previous occurrences of c .

Fairness notions that were developed for the verification of randomised systems are *extreme fairness* [21] and α -*fairness* [16]. There are many more fairness

notions in the literature, which we cannot all mention here. Overviews can be found in [9, 11, 4, 10, 14].

3 Formal Setting

Most researchers would agree that the above are all examples of fairness assumptions. The intuitive reason is that in all cases, we consider a run to be fair if whenever some transition (or some sequence of transitions) is sufficiently often possible, then it is sufficiently often executed.

This intuitive explanation lacks precision. What is the most general sense of “sufficiently often”? What do we mean by “possible”? Can we consider a notion more general than “transition”? In order to answer these questions, we need first to describe a precise formal setting.

3.1 Systems and runs

Let Σ be a countable set of *states*. Σ^* and Σ^ω denote the set of finite, and infinite sequences over Σ respectively. The set of all sequences $\Sigma^* \cup \Sigma^\omega$ is denoted as Σ^∞ . We use the symbols α, β for denoting finite sequences, and x, y for arbitrary sequences. The length of a sequence x is denoted by $|x|$ ($= \omega$ if x is infinite). Concatenation of sequences is denoted by juxtaposition; \sqsubseteq denotes the usual *prefix order* on sequences. Given a set X of sequences, we denote by $\max(X)$ the set of maximal elements of X under the prefix order. By $x\uparrow = \{y \mid x \sqsubseteq y\}$ and $x\downarrow = \{y \mid y \sqsubseteq x\}$ we denote the set of all *extensions* and *prefixes* of a sequence x respectively. The least upper bound of a sequence $(\alpha_i)_{i=0,1,\dots}$ of finite sequences where $\alpha_i \sqsubseteq \alpha_{i+1}$ is denoted by $\sup_i \alpha_i$. For a sequence $x = s_0, s_1, \dots$ and a position i where $0 \leq i < |x|$ of x , x_i denotes the i -th prefix s_0, \dots, s_i of x .

A *system* M is a tuple $\langle \Sigma, R \subseteq \Sigma \times \Sigma, \Sigma_0 \subseteq \Sigma \rangle$, where R is a *transition relation* between states, and Σ_0 is a set of *initial states*. The system is *finite* if Σ is. A *path* of a system M is a sequence in Σ^∞ that starts in an initial state and every two consecutive states are in the transition relation. The set of all paths of M is denoted by $L(M)$.

3.2 Temporal properties

A *temporal property* (*property* for short) is a set of sequences $E \subseteq \Sigma^\infty$. We say that E is *finitary* if $E \subseteq \Sigma^*$ and E is *infinitary* if $E \subseteq \Sigma^\omega$. Furthermore,

- E is *downward-closed* if $x \in E$ and $y \sqsubseteq x$ implies $y \in E$.
- E is *complete* if $\alpha_i \in E$ for $i \in \mathbb{N}$ with $\alpha_i \sqsubseteq \alpha_{i+1}$ implies $\sup_i \alpha_i \in E$.

We say that some sequence x *satisfies* a property E if $x \in E$, otherwise we say that x *violates* E .

A property S is a *safety property* if for any sequence x violating S , there exists a finite prefix α of x that violates S and each extension of a sequence violating S violates S as well, i.e.:

$$\forall x \notin S : \exists \alpha \sqsubseteq x : \alpha \uparrow \cap S = \emptyset.$$

Equivalently, a property is a safety property precisely when it is *downward-closed* and *complete*. We can think of a safety property S as a tree where nodes are labelled with elements of Σ such that S is the set of all labelled paths starting in the root of the tree. The set $L(M)$ is a safety property for each system M . The set of all sequences Σ^∞ is also a safety property and can be seen as the set of runs of a “universal” system.

Consider a safety property S and a finite sequence $\alpha \in S$. A property E is *live in α* with respect to S , if there exists a sequence $x \in E \cap S$ such that $\alpha \sqsubseteq x$. Intuitively, E is live in a finite run of a system if the system has still a chance to satisfy E in the future¹. A property E is a *liveness property for S* if E is live in every $\alpha \in S \cap \Sigma^*$. In this situation we also say that (S, E) is *machine-closed* [1, 4]. If $S = \Sigma^\infty$, then we simply say that E is a *liveness property*.

A property is *ω -regular* if it is a property accepted by some Büchi automaton, or, equivalently a property definable in Monadic Second Order logic (see e.g. [23]).

Examples. $\Sigma^{\leq k} = \{\alpha \in \Sigma^* \mid |\alpha| \leq k\}$ is a safety property for each $k \in \mathbb{N}$; Σ^* and Σ^ω are examples of liveness properties. While Σ^* is a liveness property with respect to each safety property S , Σ^ω is a liveness property with respect to S only if $\max(S) \subseteq \Sigma^\omega$; $\max(S)$ is always a liveness property with respect to S . Σ^∞ is the only property that is a safety as well as a liveness property.

3.3 Topological notions

A *topology* on a nonempty set Ω is a family $\mathcal{T} \subseteq 2^\Omega$ that is closed under union and finite intersection such that $\Omega, \emptyset \in \mathcal{T}$. The elements of \mathcal{T} are called *open sets*. A family $\mathcal{B} \subseteq \mathcal{T}$ is a *base* for \mathcal{T} if every open set $G \in \mathcal{T}$ is the union of members of \mathcal{B} .

The complement of an open set is called a *closed set*. The *closure* of a set $X \subseteq \Omega$, denoted by \overline{X} , is the smallest closed set that contains X . A set X is closed if and only if $X = \overline{X}$. A set X is *dense* if $\overline{X} = \Omega$. Equivalently, a set X is dense if

¹While there is life, there is hope. –*Cicero*

for every nonempty open set G , $G \cap X$ is nonempty. The family of open sets \mathcal{T} is not closed under countable intersection in general: A G_δ set is a set that is the intersection of countably many open sets.

Given a safety property S , the *Scott topology* on S is the family of sets $G \subseteq S$ such that

$$\forall x \in G : \exists \alpha \sqsubseteq x : \alpha \uparrow \cap S \subseteq G.$$

The family $\{\alpha \uparrow \cap S \mid \alpha \in \Sigma^*\}$ is a basis for the Scott topology. Note that open sets are generated by finitary properties Q by $G = Q \uparrow \cap S = \bigcup_{\alpha \in Q} \alpha \uparrow \cap S$, i.e., there is an exact correspondence between open sets and finitary properties. Open sets can therefore be interpreted as *observations* that can be recognised in finite time.

It is easy to verify that safety properties are exactly the closed sets and that liveness properties are exactly the dense sets of the Scott topology on Σ^∞ . It is a general theorem, that in any topological space, any set is the intersection of a closed and a dense set. Hence every temporal property can be obtained as the intersection of a safety and a liveness property [2].

Given a safety property S , we sometimes concentrate our attention to the set of maximal (finite or infinite) sequences $\max(S)$. Note that, since safety properties are downward closed, the set of maximal sequences $\max(S)$ uniquely identifies the property S , and so we can easily switch between the two points of view.

The *restriction* of the Scott topology on S to $\max(S)$ is the family of sets $(G \cap \max(S))$ where G is an open set of the Scott topology on S . The restriction of the Scott topology to $\max(\Sigma^\infty) = \Sigma^\omega$ is sometimes called the *Cantor topology*.

4 Fairness Properties

We have now all the preliminary tools to formally define fairness. We will present characterisations from three different points of view. Moreover, we will present the properties this notion enjoys.

4.1 First characterisation

In Section 2, we have seen examples of fairness of increasing strength that all fit the informal pattern “if something is sufficiently often possible, then it is sufficiently often taken”. For example, ∞ -fairness with respect to a word w instantiates “is possible” by “is live” and “something” by “the word w ”. Can we find a more general notion of fairness without doing violence to our intuition?

In fact, we can by instantiating the generic term “something” as a finitary property $Q \subseteq \Sigma^*$ where Q is “possible” in a finite run α if Q is live in α and Q is “taken” in α if $\alpha \in Q$. Furthermore, we choose to instantiate “sufficiently often” as “infinitely often”. Hence we say that a finitary property Q is *infinitely often*

satisfied by a sequence x (or that Q is *recurrent* in x) if infinitely many prefixes of x are in Q .

This gives us the following, strong notion of fairness:

Definition 1. Consider a safety property $S \subseteq \Sigma^\infty$. We say that a maximal sequence $x \in \max(S)$ is *fair* in S with respect to a finitary property Q if the following holds:

- if for infinitely many $i \in \mathbb{N}$, the property Q is live in x_i with respect to S , then for infinitely many $j \in \mathbb{N}$, x_j satisfies Q .

The set of fair runs in S w.r.t. Q is denoted as $\text{fair}(S, Q)$.

Note that every finite maximal run is vacuously fair, as Q cannot be infinitely often live in a finite run. Note also that a property Q is infinitely often live in a sequence x if and only if it is *always* live in x , that is if it is live in all prefixes of x .

Examples. If Q is the set of all finite sequences that end with a given transition t , then $\text{fair}(S, Q)$ is exactly strong ∞ -fairness with respect to t as introduced in Section 2.5. This is easily generalised to ∞ -fairness with respect to a word.

Definition 1 presents the strongest form of fairness we consider with respect to some finitary property Q . That notion could also be called *∞ -fairness* with respect to Q . Any weaker form of fairness, such as strong and weak fairness, can be obtained by weakening. We thus define that a property is a fairness property if it *contains* all fair runs with respect to some Q .

Definition 2. We say E is a *fairness property* for S if there exists a finitary property Q such that $\text{fair}(S, Q) \subseteq E$.

The definition easily implies the following observation.

Proposition 3. A property E is a fairness property for S if and only if $E \cap \max(S)$ is.

However it is sometimes convenient to consider non-maximal sequences, as we will discuss in Section 4.3.

Examples. Any property weaker than ∞ -fairness (such as strong k -fairness etc.) is a fairness property according to Definition 2.

Therefore all fairness notions introduced in Section 2 generate fairness properties with respect to a given system. However, could we have chosen a more general definition? We postpone this discussion to Section 4.5. Before, we will provide two further independent characterisations.

4.2 Second characterisation

In the introduction, we argued that unfair runs are unlikely in an intuitive sense. Alternatively we could say that *most* runs are fair. We will later examine a probabilistic interpretation of “most”. But can we formalise the notion of “most runs” without using probabilities? It turns out that we can, using topology.

In a topological space, we say that a set is *nowhere dense* if its closure does not contain any nonempty open set. For an intuition on nowhere dense sets, imagine B to be a set of “dirty” points. If B is a dense set, then it pollutes the whole topological space: wherever you go in the topological space, you will have some dirty point in the neighbourhood. If B is a “somewhere dense” set, then it pollutes part of the space. There are regions where you will be always near a dirty point, but possibly also clean neighbourhoods. Finally, if B is nowhere dense, then every clean point lives in a clean neighbourhood. Intuitively a nowhere dense set is small because the rest of the topological space can stay clear of it.

A set is *meager*, if it is the countable union of nowhere dense sets. Topologically, a countable union of small sets is still small. This was observed by the French mathematician René-Louis Baire, who proved that the unit interval of the real line cannot be obtained as the countable union of nowhere dense sets. This result can be thought of as a generalisation of Cantor’s theorem, which states that the unit interval is not obtained as the countable union of points [19].

The complement of a “small” set is therefore to be thought of as “large”. The complement of a meager set is called *co-meager* (or *residual*). In many topologies, including the Scott topology, co-meager sets can be equivalently characterised as follows:

Proposition 4. *In the Scott topology, a set is co-meager if and only if it contains a dense G_δ set.*

As announced, co-meager sets are precisely the fairness properties:

Theorem 5. *A property E is a fairness property for S if and only if $E \cap S$ is co-meager in the Scott topology of S .*

This point of view formalises the idea that “most” runs are fair. Indeed a property is a fairness property if (topologically) most runs belong to it.

Examples. The set of maximal sequences $\max(S)$ of a safety property S can be obtained as the intersection $\bigcap_{n \in \mathbb{N}} X_n$, where X_n is the set of sequences that are maximal or are longer than n . All such X_n are open in the Scott topology. This shows that maximality is a G_δ set. We already stated that $\max(S)$ is dense, i.e., a liveness property w.r.t. S , hence it is a fairness property.

4.3 Third characterisation

In the 1930ies, a group of Polish mathematicians would meet in a cafe, called the Scottish Café, in the now Ukrainian city of L'viv. During these meetings, they were posing each other problems and seeking the solution together. The minutes of these meetings were kept by the landlord and some of them were published later [18].

One of the problems, posed by Stanisław Mazur, and solved by him together with Stefan Banach involves the following game², since known as the *Banach-Mazur* game.

Let S be a safety property, and E any property. The game $G(S, E)$ is played by the two players called *Alter* and *Ego*. The state of a play is a finite sequence of S . At every move one player extends the current sequence by a finite, possibly empty sequence α_i yielding the sequence $\alpha_0 \dots \alpha_i \in S$. Alter has the first move. The play goes on forever converging to a finite sequence α or infinite sequence x in S . Ego wins if $x \in E$ (resp. $\alpha \uparrow \subseteq E$), otherwise Alter wins.

A *strategy* is a mapping $f : \Sigma^* \rightarrow \Sigma^*$ such that for each $\alpha \in S$, we have $\alpha f(\alpha) \in S$. A strategy f is *winning* for player P , if for each strategy g of the other player, P wins the play that results from P playing according to f and the other player playing according to g .

The question Mazur posed was: how do we characterise the sets for which Ego has a winning strategy? The answer is in the following theorem.

Theorem 6. *Ego has a winning strategy for the game $G(S, E)$ if and only if $E \cap S$ is co-meager in the Scott topology on S .*

Which obviously implies

Theorem 7. *A set E is a fairness property for S if and only if Ego has a winning strategy for the game $G(S, E)$.*

Note that, by Proposition 3, it is not restrictive to consider just target sets E that contain only maximal runs.

The intuition behind this characterisation is that, while fairness restricts the allowed behaviour, it should not restrict it too much. Ego, who wants to produce a fair run, can enforce some (live) choice to be taken infinitely often while she cannot prevent other choices being taken infinitely often (by Alter).

Examples. We can use Theorem 7 to prove that $\text{fair}(S, Q)$ is a fairness property. When Q is not live in α , Ego does nothing. Otherwise, Ego extends to a finite sequence in Q . This is clearly a winning strategy for Ego for the target $\text{fair}(S, Q)$.

²The original definition is slightly different and formulated in a different context: see also [19, 7, 20].

Theorem 7 can also be used to prove that a property is *not* a fairness property. Consider the system M of Section 2.2, and consider the set X of infinite runs of M that have the suffix $(cd)^\omega$. The set X is a liveness but not a fairness property for that system. Ego does not have a winning strategy for the game $G(L(M), X)$, because indeed Alter has a winning strategy: when it is his turn, Alter should just run the left-hand side component of the system, making sure that there are infinitely many a 's and b 's in the resulting sequence.

In the above example, we have shown that Ego does not have a winning strategy by showing that Alter has a winning strategy. A set for which one of the two players has a winning strategy is called *determinate*. The class of determinate properties is quite large. All *Borel sets*³ are determinate [7], of which ω -regular properties constitute, in a sense, a very simple subclass [23]. In order to show the existence of an indeterminate set, one needs the axiom of choice.

4.4 Characteristics of fairness

We have described the same class of properties from three different points of view. We now state some characteristics of this class.

The characteristics we are going to list intuitively confirm our intuition on co-meagerness as “largeness”. To help the intuition we will write “large” for co-meager, and “small” for meager. We will call a set *intermediate* if it is neither large nor small.

1. If a set is large, its complement is not.
2. Any superset of a large set is large.
3. The intersection of countably many large sets is large.
4. Intersection with a large set preserves size, i.e, if A is large and B is small (resp. intermediate, large), then $A \cap B$ is small (resp. intermediate, large).
5. When S is uncountable, every countable set is small, but there are also uncountable sets that are small.
6. Every large set is dense.

Property (6) says that for every fairness property E for S , the pair (S, E) is machine-closed, a property that has been described as the main feature of fairness

³The smallest family of sets that contains the Scott open sets and that is closed under complement and countable union.

by Apt, Francez, and Katz [4] and by Lamport [14]. Property (3) is important for modular specification. Fairness is usually imposed componentwise to the system (with respect to different transitions or processes); (3) assures that the overall fairness assumption, i.e., the intersection of all fairness assumptions for the components is again a fairness assumption.

4.5 Canonicity of the notion

How canonical is our definition of fairness? The fact that it has three independent characterisations makes this notion interesting. But could there be a more liberal definition of fairness?

Roughly, the answer is no if we insist on (3) and (6) in Section 4.4 above. More precisely:

Theorem 8. *Fairness is a maximal class of dense determinate properties that is closed under finite intersection.*

5 Probabilities

We have argued that unfair runs should be unlikely. We have shown a topological view of likelihood. A more common interpretation, however, is by means of probability theory. In this section we present this point of view. We show that probabilistic and topological likelihood are in general different notions, but that under some reasonable conditions, they in fact coincide.

5.1 Definitions

First we recall the standard setting of how probability is adjoined to systems.

A σ -algebra over a nonempty set X is a family \mathcal{A} of subsets of X that contains the empty set and is closed under complementation and countable union. Given a topology, the *Borel* σ -algebra of the topology is the smallest σ -algebra that contains the open sets. A *probability measure* on a σ -algebra \mathcal{A} over X is a function $\mu : \mathcal{A} \rightarrow [0, 1]$ such that $\mu(X) = 1$ and for any sequence of pairwise disjoint sets $(Y_i)_{i \in \mathbb{N}}$, $\mu(\bigcup_{i \in \mathbb{N}} Y_i) = \sum_{i \in \mathbb{N}} \mu(Y_i)$. A *Borel probability measure* of a topology is a probability measure over the Borel σ -algebra of the topology. Given a probability measure μ on \mathcal{A} , and two sets $A, B \in \mathcal{A}$, the *probability of A conditional to B*, is defined as $\mu(A \mid B) = \mu(A \cap B) / \mu(B)$.

Given a safety property S , consider a Borel probability measure μ over the restriction of the Scott topology to $\max(S)$. We say that μ is a *Markov measure* when $\mu(\alpha s s' \uparrow \mid \alpha s \uparrow) = \mu(\beta s s' \uparrow \mid \beta s \uparrow)$ for all $\alpha, \beta \in S \cap \Sigma^*$ and $s, s' \in \Sigma$. We

say that μ is *positive* if $\mu(\alpha\uparrow) > 0$ for each $\alpha \in S$, μ is said to be *bounded* if there exists a $c > 0$ such that $\mu(\alpha s\uparrow \mid \alpha\uparrow) > c$ for each $\alpha s \in S$. A Borel set $X \subseteq \max(S)$ is μ -*large* (or *probabilistically large* when μ is understood from the context) if $\mu(X) = 1$.

Example. Given a finite system M , consider a Markov chain on Σ that assigns positive probabilities to transitions iff they belong to R . This generates a Markov bounded measure on $\max(L(M))$.

5.2 Similarities and differences

Topological and probabilistic largeness are very similar notions. Oxtoby's classic book [19] is devoted to study these similarities. For instance all the properties characterising topological largeness described in Section 4.4 are valid also for probabilistic largeness⁴.

Despite all the common properties, the two notions do not coincide in general: in fact there are sets that are topologically large but not probabilistically large as well as sets where it is the other way around.

As an example, consider the system in Fig. 5 in Section 2.5 together with the Markov measure such that each a_i has probability $p \neq 1/2$ and each b_i has probability $1 - p$, i.e., we are looking at an asymmetric random walk on the integer line. It is well-known that the property $X_1 =$ "state 0 is visited infinitely often" has probability 0. However, it is topologically large as X_1 is established by ∞ -fairness as discussed in Section 2.5. Note that there is also a simple winning strategy for Ego.

We can also reformulate the above example in a finite-state system: Consider the system in Fig. 6 in Section 2.6 together with the Markov measure such that a has probability $p \neq 1/2$ and c has probability $1 - p$. Then, equifairness (cf. Sect. 2.7), i.e., $X_2 =$ "the number of previous a 's equals the number of previous c 's infinitely often" has probability 0 but is topologically large. (Ego's winning strategy consists in evening up the count of the letters.)

5.3 Coincidence theorem

In light of the above examples, it was quite surprising to discover that under not very restrictive hypotheses, the two notions of largeness in fact coincide.

The restrictions we have to impose are the following: we restrict our attention to ω -regular properties on finite systems, and we need to consider only bounded

⁴Property (6) is valid for most probability measures.

measures. Note that all properties that can be described using standard temporal logics such as LTL are ω -regular.

In the first counterexample above, the system is infinite. In the second counterexample, we consider a bounded measure over a finite system, but the property X_2 is not ω -regular.

Theorem 9. *Let M be a finite system, μ a bounded Borel measure on $\max(L(M))$, and X an ω -regular property. Then X is topologically large in $L(M)$ if and only if X is also μ -large.*

The key observation behind the proof is that, for ω -regular properties on finite systems, if Ego has a winning strategy, then she has a memoryless winning strategy [7]. Another important fact is here that, each ω -regular property X is determinate, as already stated in Section 4.3. For the details of the proof see [24].

5.4 Consequences

The above coincidence result has several pleasing consequences. First, it implies that for ω -regular properties, probabilistic scheduling is “fair enough”, i.e., each ω -regular fairness property has probability 1 under such scheduling.

Secondly, the result can be applied to model checking. On the one hand, we can use qualitative probabilistic model checking techniques to decide whether there exists a fairness assumption under which a given system satisfies its linear-time specification. On the other hand, we can use the three characterisations of fairness to further our understanding of probabilistic model checking. We refer the interested reader to our paper [24].

Thirdly, the above result gives a rather nice proof of the folk theorem that “in qualitative probabilistic model checking the actual probability values do not matter”. It has been long well known that a system satisfies an ω -regular specification with probability 1 regardless of what the precise probabilities associated to the local choices are. Theorem 9 is a formalisation of this intuition and allows us to reason about properties having probability 1 without mentioning probabilities at all.

6 Historical Remarks

While safety and liveness have had a formal characterisation for a long time—given by Lamport [13] and Alpern and Schneider [2]—there was no satisfactory characterisation of fairness. Apt, Francez, and Katz [4] gave three criteria that each fairness assumption should meet. Among them, machine-closure⁵ is the most

⁵Called *feasibility* in [4].

prominent. Lamport [14] reviewed their criteria and argues that fairness should be equated with machine-closure (i.e. density). Kwiatkowska [12] proposed to equate fairness with dense G_δ sets.

A couple of papers used the notion that we have described as fairness in different contexts without actually attempting to define fairness: Ben-Eliyahu and Magidor [6] observed that some popular fairness notions describe co-meager sets. Alur and Henzinger [3] propose that machine-closure should be strengthened to what they call *local liveness*, which is the same as fairness defined above. They gave the game-theoretic definition. The Banach-Mazur game has also been considered by Pistore and Vardi [20] as well as Berwanger, Grädel, and Kreutzer [7]. Berwanger et. al. [7] proved the memoryless determinacy result that lead to the coincidence theorem above.

The correspondence of safety and liveness to closed and dense sets given by Alpern and Schneider [2] goes back to G. Plotkin (see [2]) who in turn was motivated by Smyth [22]. Interestingly, Alpern and Schneider [2] write “Plotkin nevertheless is unhappy with our definition of liveness because it is not closed under intersection”. Note that in a sense, fairness with respect to the universal system Σ^∞ is the largest subclass of liveness that is closed under finite intersection as formally stated in Theorem 8. Manna and Pnueli [17] gave an alternative classification of temporal properties that is based on topology.

For more information, we refer the reader to [27, 24].

References

- [1] Martín Abadi and Leslie Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82:253–284, 1991.
- [2] Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21:181–185, October 1985.
- [3] Rajeev Alur and Thomas A. Henzinger. Local liveness for compositional modeling of fair reactive systems. In Pierre Wolper, editor, *CAV*, volume 939 of *Lecture Notes in Computer Science*, pages 166–179. Springer, 1995.
- [4] Krzysztof R. Apt, Nissim Francez, and Shmuel Katz. Appraising fairness in languages for distributed programming. *Distributed Computing*, 2:226–241, 1988.
- [5] Paul C. Attie, Nissim Francez, and Orna Grumberg. Fairness and hyperfairness in multi-party interactions. *Distributed Computing*, 6:245–254, 1993.
- [6] Rachel Ben-Eliyahu and Menachem Magidor. A temporal logic for proving properties of topologically general executions. *Information and Computation*, 124(2):127–144, 1996.
- [7] Dietmar Berwanger, Erich Grädel, and Stephan Kreutzer. Once upon a time in a west - determinacy, definability, and complexity of path games. In Moshe Y. Vardi

- and Andrei Voronkov, editors, *LPAR*, volume 2850 of *Lecture Notes in Computer Science*, pages 229–243. Springer, 2003.
- [8] Eike Best. Fairness and conspiracies. *Information Processing Letters*, 18:215–220, 1984. Erratum ibidem 19:162.
 - [9] Nissim Francez. *Fairness*. Springer, 1986.
 - [10] Yuh-Jzer Joung. On fairness notions in distributed systems, part I: A characterization of implementability. *Information and Computation*, 166:1–34, 2001.
 - [11] Marta Z. Kwiatkowska. Survey of fairness notions. *Information and Software Technology*, 31(7):371–386, 1989.
 - [12] Marta Z. Kwiatkowska. On topological characterization of behavioural properties. In G. Reed, A. Roscoe, and R. Wachter, editors, *Topology and Category Theory in Computer Science*, pages 153–177. Oxford University Press, 1991.
 - [13] Leslie Lamport. Formal foundation for specification and verification. In M.W. Alford, J.P. Ansart, G. Hommel, L. Lamport, B. Liskov, G.P. Mullery, and F.B. Schneider, editors, *Distributed Systems: Methods and Tools for Specification*, volume 190 of *LNCS*. Springer-Verlag, 1985.
 - [14] Leslie Lamport. Fairness and hyperfairness. *Distributed Computing*, 13(4):239–245, 2000.
 - [15] Daniel J. Lehmann, Amir Pnueli, and Jonathan Stavi. Impartiality, justice and fairness: The ethics of concurrent termination. In Shimon Even and Oded Kariv, editors, *ICALP*, volume 115 of *Lecture Notes in Computer Science*, pages 264–277. Springer, 1981.
 - [16] Orna Lichtenstein, Amir Pnueli, and Lenore D. Zuck. The glory of the past. In Rohit Parikh, editor, *Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 196–218. Springer, 1985.
 - [17] Zohar Manna and Amir Pnueli. A hierarchy of temporal properties. In *Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing*, pages 377–408. ACM, 1990.
 - [18] R. Daniel Mauldin. *The Scottish Book: Mathematics from the Scottish Cafe*. Birkhäuser, 1981.
 - [19] John C. Oxtoby. *Measure and Category. A Survey of the Analogies between Topological and Measure Spaces*. Springer-Verlag, 1971.
 - [20] Marco Pistore and Moshe Y. Vardi. The planning spectrum - one, two, three, infinity. In *LICS*, pages 234–243. IEEE Computer Society, 2003.
 - [21] Amir Pnueli. On the extremely fair treatment of probabilistic algorithms. In *STOC*, pages 278–290. ACM, 1983.
 - [22] Michael B. Smyth. Power domains and predicate transformers: A topological view. In Josep Díaz, editor, *ICALP*, volume 154 of *Lecture Notes in Computer Science*, pages 662–675. Springer, 1983.

- [23] Wolfgang Thomas. Automata on infinite objects. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics. Elsevier, 1990.
- [24] Daniele Varacca and Hagen Völzer. Temporal logics and model checking for fairly correct systems. In *LICS*, pages 389–398. IEEE Computer Society, 2006.
- [25] Hagen Völzer. Refinement-robust fairness. In Lubos Brim, Petr Jancar, Mojmír Kretínský, and Antonín Kucera, editors, *CONCUR*, volume 2421 of *Lecture Notes in Computer Science*, pages 547–561. Springer, 2002.
- [26] Hagen Völzer. On conspiracies and hyperfairness in distributed computing. In Pierre Fraigniaud, editor, *DISC*, volume 3724 of *Lecture Notes in Computer Science*, pages 33–47. Springer, 2005.
- [27] Hagen Völzer, Daniele Varacca, and Ekkart Kindler. Defining fairness. In Martín Abadi and Luca de Alfaro, editors, *CONCUR*, volume 3653 of *Lecture Notes in Computer Science*, pages 458–472. Springer-Verlag, 2005.