

Quantitative Safety and Liveness

Thomas A. Henzinger, Nicolas Mazzocchi, and N. Ege Saraç

Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria
{tah,nmazzocc,esarac}@ist.ac.at

Abstract. Safety and liveness are elementary concepts of computation, and the foundation of many verification paradigms. The safety-liveness classification of boolean properties characterizes whether a given property can be falsified by observing a finite prefix of an infinite computation trace (always for safety, never for liveness). In quantitative specification and verification, properties assign not truth values, but quantitative values to infinite traces (e.g., a cost, or the distance to a boolean property). We introduce quantitative safety and liveness, and we prove that our definitions induce conservative quantitative generalizations of both (1) the safety-progress hierarchy of boolean properties and (2) the safety-liveness decomposition of boolean properties. In particular, we show that every quantitative property can be written as the pointwise minimum of a quantitative safety property and a quantitative liveness property. Consequently, like boolean properties, also quantitative properties can be min-decomposed into safety and liveness parts, or alternatively, max-decomposed into co-safety and co-liveness parts. Moreover, quantitative properties can be approximated naturally. We prove that every quantitative property that has both safe and co-safe approximations can be monitored arbitrarily precisely by a monitor that uses only a finite number of states.

1 Introduction

Safety and liveness are elementary concepts in the semantics of computation [39]. They can be explained through the thought experiment of a *ghost monitor*—an imaginary device that watches an infinite computation trace at runtime, one observation at a time, and always maintains the set of *possible prediction values* to reflect the satisfaction of a given property. Let Φ be a boolean property, meaning that Φ divides all infinite traces into those that satisfy Φ , and those that violate Φ . After any finite number of observations, **True** is a possible prediction value for Φ if the observations seen so far are consistent with an infinite trace that satisfies Φ , and **False** is a possible prediction value for Φ if the observations seen so far are consistent with an infinite trace that violates Φ . When **True** is no possible prediction value, the ghost monitor can reject the hypothesis that Φ is satisfied. The property Φ is *safe* if and only if the ghost monitor can always reject the hypothesis Φ after a finite number of observations: if the infinite trace that is being monitored violates Φ , then after some finite number of observations, **True** is no possible prediction value for Φ . Orthogonally, the property Φ is *live* if and only if the ghost monitor can never reject the hypothesis Φ after a finite number of

observations: for all infinite traces, after every finite number of observations, `True` remains a possible prediction value for Φ .

The safety-liveness classification of properties is fundamental in verification. In the natural topology on infinite traces—the “Cantor topology”—the safety properties are the closed sets, and the liveness properties are the dense sets [4]. For every property Φ , the location of Φ within the Borel hierarchy that is induced by the Cantor topology—the so-called “safety-progress hierarchy” [17]—indicates the level of difficulty encountered when verifying Φ . On the first level, we find the safety and co-safety properties, the latter being the complements of safety properties, i.e., the properties whose falsehood (rather than truth) can always be rejected after a finite number of observations by the ghost monitor. More sophisticated verification techniques are needed for second-level properties, which are the countable boolean combinations of first-level properties—the so-called “response” and “persistence” properties [17]. Moreover, the orthogonality of safety and liveness leads to the following celebrated fact: *every* property can be written as the intersection of a safety property and a liveness property [4]. This means that every property Φ can be decomposed into two parts: a safety part—which is amenable to simple verification techniques, such as invariants—and a liveness part—which requires heavier verification paradigms, such as ranking functions. Dually, there is always a disjunctive decomposition of Φ into co-safety and co-liveness.

So far, we have retold the well-known story of safety and liveness for *boolean* properties. A boolean property Φ is formalized mathematically as the *set* of infinite computation traces that satisfy Φ , or equivalently, the characteristic *function* that maps each infinite trace to a truth value. Quantitative generalizations of the boolean setting allow us to capture not only correctness properties, but also performance properties [31]. In this paper we reveal the story of safety and liveness for such *quantitative* properties, which are functions from infinite traces to an arbitrary set \mathbb{D} of *values*. In order to compare values, we equip the value domain \mathbb{D} with a partial order $<$, and we require $(\mathbb{D}, <)$ to be a complete lattice. The membership problem [18] for an infinite trace f and a quantitative property Φ asks whether $\Phi(f) \geq v$ for a given threshold value $v \in \mathbb{D}$. Correspondingly, in our thought experiment, the ghost monitor attempts to reject hypotheses of the form $\Phi(f) \geq v$, which cannot be rejected as long as all observations seen so far are consistent with an infinite trace f with $\Phi(f) \geq v$. We will define Φ to be a *quantitative safety* property if and only if every hypothesis of the form $\Phi(f) \geq v$ can always be rejected by the ghost monitor after a finite number of observations, and we will define Φ to be a *quantitative liveness* property if and only if some hypothesis of the form $\Phi(f) \geq v$ can never be rejected by the ghost monitor after any finite number of observations. We note that in the quantitative case, after every finite number of observations, the set of possible prediction values for Φ maintained by the ghost monitor may be finite or infinite, and in the latter case, it may not contain a minimal or maximal element.

Let us give a few examples. Suppose we have four observations: observation `rq` for “request a resource,” observation `gr` for “grant the resource,” observation `tk` for “clock tick,” and observation `oo` for “other.” The boolean property

`Resp` requires that every occurrence of `rq` in an infinite trace is followed eventually by an occurrence of `gr`. The boolean property `NoDoubleReq` requires that no occurrence of `rq` is followed by another `rq` without some `gr` in between. The quantitative property `MinRespTime` maps every infinite trace to the largest number k such that there are at least k occurrences of `tk` between each `rq` and the closest subsequent `gr`. The quantitative property `MaxRespTime` maps every infinite trace to the smallest number k such that there are at most k occurrences of `tk` between each `rq` and the closest subsequent `gr`. The quantitative property `AvgRespTime` maps every infinite trace to the lower limit value \liminf of the infinite sequence $(v_i)_{i \geq 1}$, where v_i is, for the first i occurrences of `tk`, the average number of occurrences of `tk` between `rq` and the closest subsequent `gr`. Note that the values of `AvgRespTime` can be ∞ for some computations, including those for which the value of `Resp` is `True`. This highlights that boolean properties are not embedded in the limit behavior of quantitative properties.

The boolean property `Resp` is live because every finite observation sequence can be extended with an occurrence of `gr`. In fact, `Resp` is a second-level liveness property (namely, a response property), because it can be written as a countable intersection of co-safety properties. The boolean property `NoDoubleReq` is safe because if it is violated, it will be rejected by the ghost monitor after a finite number of observations, namely, as soon as the ghost monitor sees a `rq` followed by another occurrence of `rq` without an intervening `gr`. According to our quantitative generalization of safety, `MinRespTime` is a safety property. The ghost monitor always maintains the minimal number k of occurrences of `tk` between any past `rq` and the closest subsequent `gr` seen so far; the set of possible prediction values for `MinRespTime` is always $\{0, 1, \dots, k\}$. Every hypothesis of the form “the `MinRespTime`-value is at least v ” is rejected by the ghost monitor as soon as $k < v$; if such a hypothesis is violated, this will happen after some finite number of observations. Symmetrically, the quantitative property `MaxRespTime` is co-safe, because every wrong hypothesis of the form “the `MaxRespTime`-value is at most v ” will be rejected by the ghost monitor as soon as the smallest possible prediction value for `MaxRespTime`, which is the maximal number of occurrences of `tk` between any past `rq` and the closest subsequent `gr` seen so far, goes above v . By contrast, the quantitative property `AvgRespTime` is both live and co-live because no hypothesis of the form “the `AvgRespTime`-value is at least v ,” nor of the form “the `AvgRespTime`-value is at most v ,” can ever be rejected by the ghost monitor after a finite number of observations. All nonnegative real numbers and ∞ always remain possible prediction values for `AvgRespTime`. Note that a ghost monitor that attempts to reject hypotheses of the form $\Phi(f) \geq v$ does not need to maintain the entire set of possible prediction values, but only the sup of the set of possible prediction values, and whether or not the sup is contained in the set. Dually, updating \inf (and whether it is contained) suffices to reject hypotheses of the form $\Phi(f) \leq v$.

By defining quantitative safety and liveness via ghost monitors, we not only obtain a conservative and quantitative generalization of the boolean story, but also open up attractive frontiers for quantitative semantics, monitoring, and verification. For example, while the approximation of boolean properties reduces to

adding and removing traces to and from a set, the approximation of quantitative properties offers a rich landscape of possibilities. In fact, we can approximate the notion of safety itself. Given an error bound α , the quantitative property Φ is α -safe if and only if for every value v and every infinite trace f whose value $\Phi(f)$ is less than v , all possible prediction values for Φ are less than $v + \alpha$ after some finite prefix of f . This means that, for an α -safe property Φ , the ghost monitor may not reject wrong hypotheses of the form $\Phi(f) \geq v$ after a finite number of observations, once the violation is below the error bound. We show that every quantitative property that is both α -safe and β -co-safe, for any finite α and β , can be monitored arbitrarily precisely by a monitor that uses only a finite number of states.

We are not the first to define quantitative (or multi-valued) definitions of safety and liveness [41,27]. While the previously proposed quantitative generalizations of safety share strong similarities with our definition (without coinciding completely), our quantitative generalization of liveness is entirely new. The definitions of [27] do not support any safety-liveness decomposition, because their notion of safety is too permissive, and their liveness too restrictive. While the definitions of [41] admit a safety-liveness decomposition, our definition of liveness captures strictly fewer properties. Consequently, our definitions offer a stronger safety-liveness decomposition theorem. Our definitions also fit naturally with the definitions of emptiness, equivalence, and inclusion for quantitative languages [18].

Overview. In Section 2, we introduce quantitative properties. In Section 3, we define quantitative safety as well as safety closure, namely, the property that increases the value of each trace as little as possible to achieve safety. Then, we prove that our definitions preserve classical boolean facts. In particular, we show that a quantitative property Φ is safe if and only if Φ equals its safety closure if and only if Φ is upper semicontinuous. In Section 4, we generalize the safety-progress hierarchy to quantitative properties. We first define limit properties. For $\ell \in \{\text{inf}, \text{sup}, \text{lim inf}, \text{lim sup}\}$, the class of ℓ -properties captures those for which the value of each infinite trace can be derived by applying the limit function ℓ to the infinite sequence of values of finite prefixes. We prove that inf-properties coincide with safety, sup-properties with co-safety, lim inf-properties are suprema of countably many safety properties, and lim sup-properties infima of countably many co-safety properties. The lim inf-properties generalize the boolean persistence properties of [17]; the lim sup-properties generalize their response properties. For example, `AvgRespTime` is a lim inf-property. In Section 5, we introduce quantitative liveness and co-liveness. We prove that our definitions preserve the classical boolean facts, and show that there is a unique property which is both safe and live. As main result, we provide a safety-liveness decomposition that holds for every quantitative property. In Section 6, we define approximate safety and co-safety. We generalize the well-known unfolding approximation of discounted properties for approximate safety and co-safety properties over the extended reals. This allows us to provide a finite-state approximate monitor for these properties. In Section 7, we conclude with future research directions. For complete proofs of all results, we refer the reader to the full version of the paper.

Related Work. The notions of safety and liveness for boolean properties appeared first in [39] and were later formalized in [4], where safety properties were characterized as closed sets of the Cantor topology on infinite traces, and liveness properties as dense sets. As a consequence, the seminal decomposition theorem followed: every boolean property is an intersection of a safety property and a liveness property. A benefit of such a decomposition lies in the difference between the mathematical arguments used in their verification. While safety properties enable simpler methods such as invariants, liveness properties require more complex approaches such as well-foundedness [42,5]. These classes were characterized in terms of Büchi automata in [5] and in terms of linear temporal logic in [46].

The safety-progress classification of boolean properties [17] proposes an orthogonal view: rather than partitioning the set of properties, it provides a hierarchy of properties starting from safety. This yields a more fine-grained view of nonsafety properties which distinguishes whether a “good thing” happens at least once (co-safety or “guarantee”), infinitely many times (response), or eventually always (persistence). This classification follows the Borel hierarchy that is induced by the Cantor topology on infinite traces, and has corresponding projections within properties that are definable by finite automata and by formulas of linear temporal logic.

Runtime verification, or monitoring, is a lightweight, dynamic verification technique [6], where a monitor watches a system during its execution and tries to decide, after each finite sequence of observations, whether the observed finite computation trace or its unknown infinite extension satisfies a desired property. The safety-liveness dichotomy has profound implications for runtime verification as well: safety is easy to monitor [28], while liveness is not. An early definition of boolean monitorability was equivalent to safety with recursively enumerable sets of bad prefixes [35]. The monitoring of infinite-state boolean safety properties was later studied in [26]. A more popular definition of boolean monitorability [44,8] accounts for both truth and falsehood, establishing the set of monitorable properties as a strict superset of finite boolean combinations of safety and co-safety [23]. Boolean monitors that use the set possible prediction values can be found in [7]. The notion of boolean monitorability was investigated through the safety-liveness lens in [43] and through the safety-progress lens in [23].

Quantitative properties (a.k.a. “quantitative languages”) [18] extend their boolean counterparts by moving from the two-valued truth domain to richer domains such as real numbers. Such properties have been extensively studied from a static verification perspective in the past decade, e.g., in the context of model-checking probabilistic properties [38,37], games with quantitative objectives [10,15], specifying quantitative properties [11,1], measuring distances between systems [2,16,22,29], best-effort synthesis and repair [9,20], and quantitative analysis of transition systems [47,14,21,19]. More recently, quantitative properties have been also studied from a runtime verification perspective, e.g., for limit monitoring of statistical indicators of infinite traces [25] and for analyzing resource-precision trade-offs in the design of quantitative monitors [33,30].

To the best of our knowledge, previous definitions of (approximate) safety and liveness in nonboolean domains make implicit assumptions about the spec-

ification language [48,34,24,45]. We identify two notable exceptions. In [27], the authors generalize the framework of [43] to nonboolean value domains. They provide neither a safety-liveness decomposition of quantitative properties, nor a fine-grained classification of nonsafety properties. In [41], the authors present a safety-liveness decomposition and some levels of the safety-progress hierarchy on multi-valued truth domains, which are bounded distributive lattices. Their motivation is to provide algorithms for model-checking properties on multi-valued truth domains. We present the relationships between their definitions and ours in the relevant sections below.

2 Quantitative Properties

Let $\Sigma = \{a, b, \dots\}$ be a finite alphabet of observations. A *trace* is an infinite sequence of observations, denoted by $f, g, h \in \Sigma^\omega$, and a *finite trace* is a finite sequence of observations, denoted by $s, r, t \in \Sigma^*$. Given $s \in \Sigma^*$ and $w \in \Sigma^* \cup \Sigma^\omega$, we denote by $s \prec w$ (resp. $s \preceq w$) that s is a strict (resp. nonstrict) prefix of w . Furthermore, we denote by $|w|$ the length of w and, given $a \in \Sigma$, by $|w|_a$ the number of occurrences of a in w .

A *value domain* \mathbb{D} is a poset. Unless otherwise stated, we assume that \mathbb{D} is a nontrivial (i.e., $\perp \neq \top$) complete lattice and, whenever appropriate, we write $0, 1, -\infty, \infty$ instead of \perp and \top for the least and the greatest elements. We respectively use the terms minimum and maximum for the greatest lower bound and the least upper bound of finitely many elements.

Definition 1 (Property). A quantitative property (or simply property) is a function $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ from the set of all traces to a value domain.

A boolean property $P \subseteq \Sigma^\omega$ is defined as a set of traces. We use the boolean domain $\mathbb{B} = \{0, 1\}$ with $0 < 1$ and, in place of P , its *characteristic property* $\Phi_P : \Sigma^\omega \rightarrow \mathbb{B}$, which is defined by $\Phi_P(f) = 1$ if $f \in P$, and $\Phi_P(f) = 0$ if $f \notin P$.

For all properties Φ_1, Φ_2 on a domain \mathbb{D} and all traces $f \in \Sigma^\omega$, we let $\min(\Phi_1, \Phi_2)(f) = \min(\Phi_1(f), \Phi_2(f))$ and $\max(\Phi_1, \Phi_2)(f) = \max(\Phi_1(f), \Phi_2(f))$. For a domain \mathbb{D} , the *inverse* of \mathbb{D} is the domain $\bar{\mathbb{D}}$ that contains the same elements as \mathbb{D} but with the ordering reversed. For a property Φ , we define its *complement* $\bar{\Phi} : \Sigma^\omega \rightarrow \bar{\mathbb{D}}$ by $\bar{\Phi}(f) = \Phi(f)$ for all $f \in \Sigma^\omega$.

Some properties can be defined as limits of value sequences. A *finitary property* $\pi : \Sigma^* \rightarrow \mathbb{D}$ associates a value with each finite trace. A *value function* $\ell : \mathbb{D}^\omega \rightarrow \mathbb{D}$ condenses an infinite sequence of values to a single value. Given a finitary property π , a value function ℓ , and a trace $f \in \Sigma^\omega$, we write $\ell_{s \prec f} \pi(s)$ instead of $\ell(\pi(s_0)\pi(s_1)\dots)$, where each s_i fulfills $s_i \prec f$ and $|s_i| = i$.

3 Quantitative Safety

Given a property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$, a trace $f \in \Sigma^\omega$, and a value $v \in \mathbb{D}$, the quantitative membership problem [18] asks whether $\Phi(f) \geq v$. We define quantitative safety as follows: the property Φ is safe iff every wrong hypothesis of the form $\Phi(f) \geq v$ has a finite witness $s \prec f$.

Definition 2 (Safety). A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is safe iff for every $f \in \Sigma^\omega$ and value $v \in \mathbb{D}$ with $\Phi(f) \not\leq v$, there is a prefix $s \prec f$ such that $\sup_{g \in \Sigma^\omega} \Phi(sg) \not\leq v$.

Let us illustrate this definition with the *minimal response-time* property.

Example 3. Let $\Sigma = \{\mathbf{rq}, \mathbf{gr}, \mathbf{tk}, \mathbf{oo}\}$ and $\mathbb{D} = \mathbb{N} \cup \{\infty\}$. We define the minimal response-time property Φ_{\min} through an auxiliary finitary property π_{\min} that computes the minimum response time so far. In a finite or infinite trace, an occurrence of \mathbf{rq} is *granted* if it is followed, later, by a \mathbf{gr} , and otherwise it is *pending*. Let $\pi_{\text{last}}(s) = \infty$ if the finite trace s contains a pending \mathbf{rq} , or no \mathbf{rq} , and $\pi_{\text{last}}(s) = |r|_{\mathbf{tk}} - |t|_{\mathbf{tk}}$ otherwise, where $r \prec s$ is the longest prefix of s with a pending \mathbf{rq} , and $t \prec r$ is the longest prefix of r without pending \mathbf{rq} . Intuitively, π_{last} provides the response time for the last request when all requests are granted, and ∞ when there is a pending request or no request. Given $s \in \Sigma^*$, taking the minimum of the values of π_{last} over the prefixes $r \preceq s$ gives us the minimum response time so far. Let $\pi_{\min}(s) = \min_{r \preceq s} \pi_{\text{last}}(r)$ for all $s \in \Sigma^*$, and $\Phi_{\min}(f) = \lim_{s \prec f} \pi_{\min}(s)$ for all $f \in \Sigma^\omega$. The limit always exists because the minimum is monotonically decreasing.

The minimal response-time property is safe. Let $f \in \Sigma^\omega$ and $v \in \mathbb{D}$ such that $\Phi_{\min}(f) < v$. Then, some prefix $s \prec f$ contains a \mathbf{rq} that is granted after $u < v$ ticks, in which case, no matter what happens in the future, the minimal response time is guaranteed to be at most u ; that is, $\sup_{g \in \Sigma^\omega} \Phi_{\min}(sg) \leq u < v$. If you recall from the introduction the ghost monitor that maintains the sup of possible prediction values for the minimal response-time property, that value is always π_{\min} ; that is, $\sup_{g \in \Sigma^\omega} \Phi_{\min}(sg) = \pi_{\min}(s)$ for all $s \in \Sigma^*$. Note that in the case of minimal response time, the sup of possible prediction values is always realizable; that is, for all $s \in \Sigma^*$, there exists an $f \in \Sigma^\omega$ such that $\sup_{g \in \Sigma^\omega} \Phi_{\min}(sg) = \Phi_{\min}(sf)$. \square

Remark 4. Quantitative safety generalizes boolean safety. For every boolean property $P \subseteq \Sigma^\omega$, the following statements are equivalent: (i) P is safe according to the classical definition [4], (ii) its characteristic property Φ_P is safe, and (iii) for every $f \in \Sigma^\omega$ and $v \in \mathbb{B}$ with $\Phi_P(f) < v$, there exists a prefix $s \prec f$ such that for all $g \in \Sigma^\omega$, we have $\Phi_P(sg) < v$.

We now generalize the notion of safety closure and present an operation that makes a property safe by increasing the value of each trace as little as possible.

Definition 5 (Safety closure). The safety closure of a property Φ is the property Φ^* defined by $\Phi^*(f) = \inf_{s \prec f} \sup_{g \in \Sigma^\omega} \Phi(sg)$ for all $f \in \Sigma^\omega$.

We can say the following about the safety closure operation.

Proposition 6. For every property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$, the following statements hold.

1. Φ^* is safe.
2. $\Phi^*(f) \geq \Phi(f)$ for all $f \in \Sigma^\omega$.
3. $\Phi^*(f) = \Phi^{**}(f)$ for all $f \in \Sigma^\omega$.
4. For every safety property $\Psi : \Sigma^\omega \rightarrow \mathbb{D}$, if $\Phi(f) \leq \Psi(f)$ for all $f \in \Sigma^\omega$, then $\Psi(g) \not\leq \Phi^*(g)$ for all $g \in \Sigma^\omega$.

3.1 Alternative Characterizations of Quantitative Safety

Consider a trace and its prefixes of increasing length. For a given property, the ghost monitor from the introduction maintains, for each prefix, the sup of possible prediction values, i.e., the least upper bound of the property values for all possible infinite continuations. The resulting sequence of monotonically decreasing suprema provides an upper bound on the eventual property value. Moreover, for some properties, this sequence always converges to the property value. If this is the case, then the ghost monitor can always dismiss wrong lower-bound hypotheses after finite prefixes, and vice versa. This gives us an alternative definition for the safety of quantitative properties which, inspired by the notion of Scott continuity, was called *continuity* [33]. We now believe that *upper semicontinuity* is a more appropriate term, as becomes clear when we consider the Cantor topology on Σ^ω and the value domain $\mathbb{R} \cup \{-\infty, +\infty\}$.

Definition 7 (Upper semicontinuity [33]). *A property Φ is upper semicontinuous iff $\Phi(f) = \lim_{s \prec f} \sup_{g \in \Sigma^\omega} \Phi(sg)$ for all $f \in \Sigma^\omega$.*

We note that the minimal response-time property is upper semicontinuous.

Example 8. Recall the minimal response-time property Φ_{\min} from Example 3. For every trace $f \in \Sigma^\omega$, the Φ_{\min} value is the limit of the π_{\min} values for the prefixes of f . Therefore, Φ_{\min} is upper semicontinuous. \square

In general, a property is safe iff it maps every trace to the limit of the suprema of possible prediction values. Moreover, we can also characterize safety properties as the properties that are equal to their safety closure.

Theorem 9. *For every property Φ , the following statements are equivalent:*
 1. Φ is safe. 2. Φ is upper semicontinuous. 3. $\Phi(f) = \Phi^*(f)$ for all $f \in \Sigma^\omega$.

3.2 Related Definitions of Quantitative Safety

In [41], the authors consider the model-checking problem for properties on multi-valued truth domains. They introduce the notion of multi-safety through a closure operation that coincides with our safety closure. Formally, a property Φ is *multi-safe* iff $\Phi(f) = \Phi^*(f)$ for every $f \in \Sigma^\omega$. It is easy to see the following.

Proposition 10. *For every property Φ , we have Φ is multi-safe iff Φ is safe.*

Although the two definitions of safety are equivalent, our definition is consistent with the membership problem for quantitative automata and motivated by the monitoring of quantitative properties.

In [27], the authors extend a refinement of the safety-liveness classification for monitoring [43] to richer domains. They introduce the notion of verdict-safety through dismissibility of values not less than or equal to the property value. Formally, a property Φ is *verdict-safe* iff for every $f \in \Sigma^\omega$ and $v \not\leq \Phi(f)$, there exists a prefix $s \prec f$ such that for all $g \in \Sigma^\omega$, we have $\Phi(sg) \neq v$.

We demonstrate that verdict-safety is weaker than safety. Moreover, we provide a condition under which the two definitions coincide. To achieve this, we reason about sets of possible prediction values: for a property Φ and $s \in \Sigma^*$, let $P_{\Phi,s} = \{\Phi(sf) \mid f \in \Sigma^\omega\}$.

Lemma 11. *A property Φ is verdict-safe iff $\Phi(f) = \sup(\lim_{s \prec f} P_{\Phi,s})$ for all $f \in \Sigma^\omega$.*

Notice that Φ is safe iff $\Phi(f) = \lim_{s \prec f}(\sup P_{\Phi,s})$ for all $f \in \Sigma^\omega$. Below we describe a property that is verdict-safe but not safe.

Example 12. Let $\Sigma = \{a, b\}$. Define Φ by $\Phi(f) = 0$ if $f = a^\omega$, and $\Phi(f) = |s|$ otherwise, where $s \prec f$ is the shortest prefix in which b occurs. The property Φ is verdict-safe. First, observe that $\mathbb{D} = \mathbb{N} \cup \{\infty\}$. Let $f \in \Sigma^\omega$ and $v \in \mathbb{D}$ with $v > \Phi(f)$. If $\Phi(f) > 0$, then f contains b , and $\Phi(f) = |s|$ for some $s \prec f$ in which b occurs for the first time. After the prefix s , all $g \in \Sigma^\omega$ yield $\Phi(sg) = |s|$, thus all values above $|s|$ are rejected. If $\Phi(f) = 0$, then $f = a^\omega$. Let $v \in \mathbb{D}$ with $v > 0$, and consider the prefix $a^v \prec f$. Observe that the set of possible prediction values after reading a^v is $\{0, v+1, v+2, \dots\}$, therefore a^v allows the ghost monitor to reject the value v . However, Φ is not safe because, although $\Phi(a^\omega) = 0$, for every $s \prec a^\omega$, we have $\sup_{g \in \Sigma^\omega} \Phi(sg) = \infty$. \square

The separation is due to the fact that, for some finite traces, the sup of possible prediction values cannot be realized by any future. Below, we present a condition that prevents such cases.

Definition 13 (Supremum closedness). *A property Φ is sup-closed iff for every $s \in \Sigma^*$ we have $\sup P_{\Phi,s} \in P_{\Phi,s}$.*

We remark that the minimal response-time property is sup-closed.

Example 14. The safety property minimal response-time Φ_{\min} from Example 3 is sup-closed. This is because, for every $s \in \Sigma^*$, the continuation \mathbf{gr}^ω realizes the value $\sup_{g \in \Sigma^\omega} \Phi(sg)$. \square

Recall from the introduction the ghost monitor that maintains the sup of possible prediction values. For monitoring sup-closed properties this suffices; otherwise the ghost monitor also needs to maintain whether or not the supremum of the possible prediction values is realizable by some future continuation. In general, we have the following for every sup-closed property.

Lemma 15. *For every sup-closed property Φ and for all $f \in \Sigma^\omega$, we have $\lim_{s \prec f}(\sup P_{\Phi,s}) = \sup(\lim_{s \prec f} P_{\Phi,s})$.*

As a consequence of the lemmas above, we get the following.

Theorem 16. *A sup-closed property Φ is safe iff Φ is verdict-safe.*

4 The Quantitative Safety-Progress Hierarchy

Our quantitative extension of safety closure allows us to build a Borel hierarchy, which is a quantitative extension of the boolean safety-progress hierarchy [17]. First, we show that safety properties are closed under pairwise min and max.

Proposition 17. *For every value domain \mathbb{D} , the set of safety properties over \mathbb{D} is closed under min and max.*

The boolean safety-progress classification of properties is a Borel hierarchy built from the Cantor topology of traces. Safety and co-safety properties lie on the first level, respectively corresponding to the closed sets and open sets of the topology. The second level is obtained through countable unions and intersections of properties from the first level: persistence properties are countable unions of closed sets, while response properties are countable intersections of open sets. We generalize this construction to the quantitative setting.

In the boolean case, each property class is defined through an operation that takes a set $S \subseteq \Sigma^*$ of finite traces and produces a set $P \subseteq \Sigma^\omega$ of infinite traces. For example, to obtain a co-safety property from $S \subseteq \Sigma^*$, the corresponding operation yields $S\Sigma^\omega$. Similarly, we formalize each property class by a value function. For this, we define the notion of *limit property*.

Definition 18 (Limit property). *A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is a limit property iff there exists a finitary property $\pi : \Sigma^* \rightarrow \mathbb{D}$ and a value function $\ell : \mathbb{D}^\omega \rightarrow \mathbb{D}$ such that $\Phi(f) = \ell_{s \prec f} \pi(s)$ for all $f \in \Sigma^\omega$. We denote this by $\Phi = (\pi, \ell)$, and write $\Phi(s)$ instead of $\pi(s)$. In particular, if $\Phi = (\pi, \ell)$, where $\ell \in \{\inf, \sup, \liminf, \limsup\}$, then Φ is an ℓ -property.*

To account for the value functions that construct the first two levels of the safety-progress hierarchy, we start our investigation with inf- and sup-properties and later focus on lim inf- and lim sup- properties [18].

4.1 Infimum and Supremum Properties

Let us start with an example by demonstrating that the minimal response-time property is an inf-property.

Example 19. Recall the safety property Φ_{\min} of minimal response time from Example 3. We can equivalently define Φ_{\min} as a limit property by taking the finitary property π_{last} and the value function \inf . As discussed in Example 3, the function π_{last} outputs the response time for the last request when all requests are granted, and ∞ when there is a pending request or no request. Then $\inf_{s \prec f} \pi_{\text{last}}(s) = \Phi_{\min}(f)$ for all $f \in \Sigma^\omega$, and therefore $\Phi_{\min} = (\pi_{\text{last}}, \inf)$. \square

In fact, the safety properties coincide with inf-properties.

Theorem 20. *A property Φ is safe iff Φ is an inf-property.*

Defining the minimal response-time property as a limit property, we observe the following relation between its behavior on finite traces and infinite traces.

Example 21. Consider the property $\Phi_{\min} = (\pi_{\text{last}}, \inf)$ from Example 19. Let $f \in \Sigma^\omega$ and $v \in \mathbb{D}$. Observe that if the minimal response time of f is at least v , then the last response time for each prefix $s \prec f$ is also at least v . Conversely, if the minimal response time of f is below v , then there is a prefix $s \prec f$ for which the last response time is also below v . \square

In light of this observation, we provide another characterization of safety properties, explicitly relating the specified behavior of the limit property on finite and infinite traces.

Theorem 22. *A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is safe iff Φ is a limit property such that for every $f \in \Sigma^\omega$ and value $v \in \mathbb{D}$, we have $\Phi(f) \geq v$ iff $\Phi(s) \geq v$ for all $s \prec f$.*

Recall that a safety property allows rejecting wrong lower-bound hypotheses with a finite witness, by assigning a tight upper bound to each trace. We define co-safety properties symmetrically: a property Φ is co-safe iff every wrong hypothesis of the form $\Phi(f) \leq v$ has a finite witness $s \prec f$.

Definition 23 (Co-safety). *A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is co-safe iff for every $f \in \Sigma^\omega$ and value $v \in \mathbb{D}$ with $\Phi(f) \not\leq v$, there exists a prefix $s \prec f$ such that $\inf_{g \in \Sigma^\omega} \Phi(sg) \not\leq v$.*

We note that our definition generalizes boolean co-safety, and thus a dual of Remark 4 holds also for co-safety. Moreover, we analogously define the notions of co-safety closure and lower semicontinuity.

Definition 24 (Co-safety closure). *The co-safety closure of a property Φ is the property $\Phi_*(f)$ defined by $\Phi_*(f) = \sup_{s \prec f} \inf_{g \in \Sigma^\omega} \Phi(sg)$ for all $f \in \Sigma^\omega$.*

Definition 25 (Lower semicontinuity [33]). *A property Φ is lower semicontinuous iff $\Phi(f) = \lim_{s \prec f} \inf_{g \in \Sigma^\omega} \Phi(sg)$ for all $f \in \Sigma^\omega$.*

Now, we define and investigate the *maximal response-time* property. In particular, we show that it is a sup-property that is co-safe and lower semicontinuous.

Example 26. Let $\Sigma = \{\mathbf{rq}, \mathbf{gr}, \mathbf{tk}, \mathbf{oo}\}$ and $\mathbb{D} = \mathbb{N} \cup \{\infty\}$. We define the maximal response-time property Φ_{\max} through a finitary property that computes the current response time for each finite trace and the value function sup. In particular, for all $s \in \Sigma^*$, let $\pi_{\text{curr}}(s) = |s|_{\mathbf{tk}} - |r|_{\mathbf{tk}}$, where $r \preceq s$ is the longest prefix of s without pending \mathbf{rq} ; then $\Phi_{\max} = (\pi_{\text{curr}}, \text{sup})$. Note the contrast between π_{curr} and π_{last} from Example 3. While π_{curr} takes an optimistic view of the future and assumes the \mathbf{gr} will follow immediately, π_{last} takes a pessimistic view and assumes the \mathbf{gr} will never follow. Let $f \in \Sigma^\omega$ and $v \in \mathbb{D}$. If the maximal response time of f is greater than v , then for some prefix $s \prec f$ the current response time is greater than v also, which means that, no matter what happens in the future, the maximal response time is greater than v after observing s . Therefore, Φ_{\max} is co-safe. By a similar reasoning, the sequence of greatest lower bounds of possible prediction values over the prefixes converges to the property value. In other words, we have $\lim_{s \prec f} \inf_{g \in \Sigma^\omega} \Phi_{\max}(sg) = \Phi_{\max}(f)$ for all $f \in \Sigma^\omega$. Thus Φ_{\max} is also lower semicontinuous, and it equals its co-safety closure. Now, consider the complementary property $\overline{\Phi_{\max}}$, which maps every trace to the same value as Φ_{\max} on a domain where the order is reversed. It is easy to see that $\overline{\Phi_{\max}}$ is safe. Finally, recall the ghost monitor from the introduction, which maintains the infimum of possible prediction values for the maximal response-time property. Since the maximal response-time property is inf-closed, the output of the ghost monitor after every prefix is realizable by some future continuation, and that output is $\pi_{\max}(s) = \max_{r \preceq s} \pi_{\text{curr}}(r)$ for all $s \in \Sigma^*$. \square

Generalizing the observations in the example above, we obtain the following characterizations due to the duality between safety and co-safety.

Theorem 27. *For every property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$, the following are equivalent.*

1. Φ is co-safe.
2. Φ is lower semicontinuous.
3. $\Phi(f) = \Phi_*(f)$ for every $f \in \Sigma^\omega$.
4. Φ is a sup-property.
5. Φ is a limit property such that for every $f \in \Sigma^\omega$ and value $v \in \mathbb{D}$, we have $\Phi(f) \leq v$ iff $\Phi(s) \leq v$ for all $s \prec f$.
6. $\overline{\Phi}$ is safe.

4.2 Limit Inferior and Limit Superior Properties

Let us start with an observation on the minimal response-time property.

Example 28. Recall once again the minimal response-time property Φ_{\min} from Example 3. In the previous subsection, we presented an alternative definition of Φ_{\min} to establish that it is an inf-property. Observe that there is yet another equivalent definition of Φ_{\min} which takes the monotonically decreasing finitary property π_{\min} from Example 3 and pairs it with either the value function \liminf , or with \limsup . Hence Φ_{\min} is both a \liminf - and a \limsup -property. \square

Before moving on to investigating \liminf - and \limsup -properties more closely, we show that the above observation can be generalized.

Theorem 29. *Every ℓ -property Φ , for $\ell \in \{\inf, \sup\}$, is both a \liminf - and a \limsup -property.*

An interesting response-time property beyond safety and co-safety arises when we remove extreme values: instead of minimal response time, consider the property that maps every trace to a value that bounds from below, not all response times, but all of them from a point onward (i.e., all but finitely many). We call this property *tail-minimal response time*.

Example 30. Let $\Sigma = \{\mathbf{rq}, \mathbf{gr}, \mathbf{tk}, \mathbf{oo}\}$ and π_{last} be the finitary property from Example 3 that computes the last response time. We define the tail-minimal response-time property as $\Phi_{\text{tmin}} = (\pi_{\text{last}}, \liminf)$. Intuitively, it maps each trace to the least response time over all but finitely many requests. This property is interesting as a performance measure, because it focuses on the long-term performance by ignoring finitely many outliers. Consider $f \in \Sigma^\omega$ and $v \in \mathbb{D}$. Observe that, if the tail-minimal response time of f is at least v , then there is a prefix $s \prec f$ such that for all longer prefixes $s \preceq r \prec f$, the last response time in r is at least v , and vice versa. \square

Similarly as for inf-properties, we characterize \liminf -properties through a relation between property behaviors on finite and infinite traces.

Theorem 31. *A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is a \liminf -property iff Φ is a limit property such that for every $f \in \Sigma^\omega$ and value $v \in \mathbb{D}$, we have $\Phi(f) \geq v$ iff there exists $s \prec f$ such that for all $s \preceq r \prec f$, we have $\Phi(r) \geq v$.*

Now, we show that the tail-minimal response-time property can be expressed as a countable supremum of inf-properties.

Example 32. Let $i \in \mathbb{N}$ and define $\pi_{i,\text{last}}$ as a finitary property that imitates π_{last} from Example 3, but ignores the first i observations of every finite trace. Formally, for $s \in \Sigma^*$, we define $\pi_{i,\text{last}}(s) = \pi_{\text{last}}(r)$ for $s = s_i r$ where $s_i \preceq s$ with $|s_i| = i$, and $r \in \Sigma^*$. Observe that an equivalent way to define Φ_{tmin} from Example 30 is $\sup_{i \in \mathbb{N}} (\inf_{s \prec f} (\pi_{i,\text{last}}(s)))$ for all $f \in \Sigma^\omega$. Intuitively, for each $i \in \mathbb{N}$, we obtain an inf-property that computes the minimal response time of the suffixes of a given trace. Taking the supremum over these, we obtain the greatest lower bound on all but finitely many response times. \square

We generalize this observation and show that every lim inf-property is a countable supremum of inf-properties.

Theorem 33. *Every lim inf-property is a countable supremum of inf-properties.*

We would also like to have the converse of Theorem 33, i.e., that every countable supremum of inf-properties is a lim inf-property. Currently, we are able to show only the following.

Theorem 34. *For every infinite sequence $(\Phi_i)_{i \in \mathbb{N}}$ of inf-properties, there is a lim inf-property Φ such that $\sup_{i \in \mathbb{N}} \Phi_i(f) \leq \Phi(f)$.*

We conjecture that some lim inf-property that satisfies Theorem 34 is also a lower bound on the countable supremum that occurs in the theorem. This, together with Theorem 34, would imply the converse of Theorem 33. Proving the converse of Theorem 33 would give us, thanks to the following duality, that the lim inf- and lim sup-properties characterize the second level of the Borel hierarchy of the topology induced by the safety closure operator.

Proposition 35. *A property Φ is a lim inf-property iff its complement $\bar{\Phi}$ is a lim sup-property.*

5 Quantitative Liveness

Similarly as for safety, we take the perspective of the quantitative membership problem to define liveness: a property Φ is live iff, whenever a property value is less than \top , there exists a value v for which the wrong hypothesis $\Phi(f) \geq v$ can never be dismissed by any finite witness $s \prec f$.

Definition 36 (Liveness). *A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is live iff for all $f \in \Sigma^\omega$, if $\Phi(f) < \top$, then there exists a value $v \in \mathbb{D}$ such that $\Phi(f) \not\geq v$ and for all prefixes $s \prec f$, we have $\sup_{g \in \Sigma^\omega} \Phi(sg) \geq v$.*

An equivalent definition can be given through the safety closure.

Theorem 37. *A property Φ is live iff $\Phi^*(f) > \Phi(f)$ for every $f \in \Sigma^\omega$ with $\Phi(f) < \top$.*

Our definition generalizes boolean liveness. A boolean property $P \subseteq \Sigma^\omega$ is live according to the classical definition [4] iff its characteristic property Φ_P is live according to our definition. Moreover, the intersection of safety and liveness contains only the single degenerate property that always outputs \top .

Proposition 38. *A property Φ is safe and live iff $\Phi(f) = \top$ for all $f \in \Sigma^\omega$.*

We define co-liveness symmetrically, and note that the duals of the observations above also hold for co-liveness.

Definition 39 (Co-liveness). *A property $\Phi : \Sigma^\omega \rightarrow \mathbb{D}$ is co-live iff for all $f \in \Sigma^\omega$, if $\Phi(f) > \perp$, then there exists a value $v \in \mathbb{D}$ such that $\Phi(f) \not\leq v$ and for all prefixes $s \prec f$, we have $\inf_{g \in \Sigma^\omega} \Phi(sg) \leq v$.*

Next, we present some examples of liveness and co-liveness properties. We start by showing that lim inf- and lim sup-properties can be live and co-live.

Example 40. Let $\Sigma = \{a, b\}$ be an alphabet, and let $P = \square \diamond a$ and $Q = \diamond \square b$ be boolean properties defined in linear temporal logic. Consider their characteristic properties Φ_P and Φ_Q . As we pointed out earlier, our definitions generalize their boolean counterparts, therefore Φ_P and Φ_Q are both live and co-live. Moreover, Φ_P is a lim sup-property: define $\pi_P(s) = 1$ if $s \in \Sigma^* a$, and $\pi_P(s) = 0$ otherwise, and observe that $\Phi_P(f) = \limsup_{s \prec f} \pi_P(s)$ for all $f \in \Sigma^\omega$. Similarly, Φ_Q is a lim inf-property. \square

Now, we show that the maximal response-time property is live, and the minimal response time is co-live.

Example 41. Recall the co-safety property Φ_{\max} of maximal response time from Example 26. Let $f \in \Sigma^\omega$ such that $\Phi_{\max}(f) < \infty$. We can extend every prefix $s \prec f$ with $g = \mathbf{rq} \mathbf{tk}^\omega$, which gives us $\Phi_{\max}(sg) = \infty > \Phi(f)$. Equivalently, for every $f \in \Sigma^\omega$, we have $\Phi_{\max}^*(f) = \infty > \Phi_{\max}(f)$. Hence Φ_{\max} is live and, analogously, the safety property Φ_{\min} from Example 3 is co-live. \square

Finally, we show that the *average response-time* property is live and co-live.

Example 42. Let $\Sigma = \{\mathbf{rq}, \mathbf{gr}, \mathbf{tk}, \mathbf{oo}\}$. For all $s \in \Sigma^*$, let $p(s) = 1$ if there is no pending \mathbf{rq} in s , and $p(s) = 0$ otherwise. Define $\pi_{\text{valid}}(s) = |\{r \preceq s \mid \exists t \in \Sigma^* : r = t \mathbf{rq} \wedge p(t) = 1\}|$ as the number of valid requests in s , and define $\pi_{\text{time}}(s)$ as the number of \mathbf{tk} observations that occur after a valid \mathbf{rq} and before the matching \mathbf{gr} . Then, $\Phi_{\text{avg}} = (\pi_{\text{avg}}, \liminf)$, where $\pi_{\text{avg}}(s) = \frac{\pi_{\text{time}}(s)}{\pi_{\text{valid}}(s)}$ for all $s \in \Sigma^*$ with $\pi_{\text{valid}}(s) > 0$, and $\pi_{\text{avg}}(s) = \infty$ otherwise. For example, $\pi_{\text{avg}}(s) = \frac{3}{2}$ for $s = \mathbf{rq} \mathbf{tk} \mathbf{gr} \mathbf{tk} \mathbf{rq} \mathbf{tk} \mathbf{rq} \mathbf{tk}$. Note that Φ_{avg} is a lim inf-property.

The property Φ_{avg} is defined on the value domain $[0, \infty]$ and is both live and co-live. To see this, let $f \in \Sigma^\omega$ such that $0 < \Phi_{\text{avg}}(f) < \infty$ and, for every prefix $s \prec f$, consider $g = \mathbf{rq} \mathbf{tk}^\omega$ and $h = \mathbf{gr} (\mathbf{rq} \mathbf{gr})^\omega$. Since sg has a pending request followed by infinitely many clock ticks, we have $\Phi_{\text{avg}}(sg) = \infty$. Similarly, since sh eventually has all new requests immediately granted, we get $\Phi_{\text{avg}}(sh) = 0$. \square

5.1 The Quantitative Safety-Liveness Decomposition

A celebrated theorem states that every boolean property can be expressed as an intersection of a safety property and a liveness property [4]. In this section, we prove the analogous result for the quantitative setting.

Example 43. Let $\Sigma = \{\text{rq}, \text{gr}, \text{tk}, \text{oo}\}$. Recall the maximal response-time property Φ_{\max} from Example 26, and the average response-time property Φ_{avg} from Example 42. Let $n > 0$ be an integer and define a new property Φ by $\Phi(f) = \Phi_{\text{avg}}(f)$ if $\Phi_{\max}(f) \leq n$, and $\Phi(f) = 0$ otherwise. For the safety closure of Φ , we have $\Phi^*(f) = n$ if $\Phi_{\max}(f) \leq n$, and $\Phi^*(f) = 0$ otherwise. Now, we further define $\Psi(f) = \Phi_{\text{avg}}(f)$ if $\Phi_{\max}(f) \leq n$, and $\Psi(f) = n$ otherwise. Observe that Ψ is live, because every prefix of a trace whose value is less than n can be extended to a greater value. Finally, note that for all $f \in \Sigma^\omega$, we can express $\Phi(f)$ as the pointwise minimum of $\Phi^*(f)$ and $\Psi(f)$. Intuitively, the safety part Φ^* of this decomposition checks whether the maximal response time stays below the permitted bound, and the liveness part Ψ keeps track of the average response time as long as the bound is satisfied. \square

Following a similar construction, we show that a safety-liveness decomposition exists for every property.

Theorem 44. *For every property Φ , there exists a liveness property Ψ such that $\Phi(f) = \min(\Phi^*(f), \Psi(f))$ for all $f \in \Sigma^\omega$.*

In particular, if the given property is safe or live, the decomposition is trivial.

Remark 45. Let Φ be a property. If Φ is safe (resp. live), then the safety (resp. liveness) part of the decomposition is Φ itself, and the liveness (resp. safety) part is the constant property that maps every trace to \top .

For co-safety and co-liveness, the duals of Theorem 44 and Remark 45 hold. In particular, every property is the pointwise maximum of its co-safety closure and a co-liveness property.

5.2 Related Definitions of Quantitative Liveness

In [41], the authors define a property Φ as *multi-live* iff $\Phi^*(f) > \perp$ for all $f \in \Sigma^\omega$. We show that our definition is more restrictive, resulting in fewer liveness properties while still allowing a safety-liveness decomposition.

Proposition 46. *Every live property is multi-live, and the inclusion is strict.*

We provide a separating example on a totally ordered domain below.

Example 47. Let $\Sigma = \{a, b, c\}$, and consider the following property: $\Phi(f) = 0$ if $f \models \Box a$, and $\Phi(f) = 1$ if $f \models \Diamond c$, and $\Phi(f) = 2$ otherwise (i.e., if $f \models \Diamond b \wedge \Box \neg c$). For all $f \in \Sigma^\omega$ and prefixes $s \prec f$, we have $\Phi(sc^\omega) = 1$. Thus $\Phi^*(f) \neq \perp$, which implies that Φ is multi-live. However, Φ is not live. Indeed, for every $f \in \Sigma^\omega$ such that $f \models \Diamond c$, we have $\Phi(f) = 1 < \top$. Moreover, f admits some prefix s that contains an occurrence of c , thus satisfying $\sup_{g \in \Sigma^\omega} \Phi(sg) = 1$. \square

In [27], the authors define a property Φ as *verdict-live* iff for every $f \in \Sigma^\omega$ and value $v \not\leq \Phi(f)$, every prefix $s \prec f$ satisfies $\Phi(sg) = v$ for some $g \in \Sigma^\omega$. We show that our definition is more liberal.

Proposition 48. *Every verdict-live property is live, and the inclusion is strict.*

We provide a separating example below, concluding that our definition is strictly more general even for totally ordered domains.

Example 49. Let $\Sigma = \{a, b\}$, and consider the following property: $\Phi(f) = 0$ if $f \not\models \diamond b$, and $\Phi(f) = 1$ if $f \models \diamond(b \wedge \bigcirc \diamond b)$, and $\Phi(f) = 2^{-|s|}$ otherwise, where $s \prec f$ is the shortest prefix in which b occurs. Consider an arbitrary $f \in \Sigma^\omega$. If $\Phi(f) = 1$, then the liveness condition is vacuously satisfied. If $\Phi(f) = 0$, then $f = a^\omega$, and every prefix $s \prec f$ can be extended with $g = ba^\omega$ or $h = b^\omega$ to obtain $\Phi(sg) = 2^{-(|s|+1)}$ and $\Phi(sh) = 1$. If $0 < \Phi(f) < 1$, then f satisfies $\diamond b$ but not $\diamond(b \wedge \bigcirc \diamond b)$, and every prefix $s \prec f$ can be extended with b^ω to obtain $\Phi(sb^\omega) = 1$. Hence Φ is live. However, Φ is not verdict-live. To see this, consider the trace $f = a^k ba^\omega$ for some integer $k \geq 1$ and note that $\Phi(f) = 2^{-(k+1)}$. Although all prefixes of f can be extended to reach the value 1, the value domain contains elements between $\Phi(f)$ and 1, namely the values 2^{-m} for $1 \leq m \leq k$. Each of these values can be rejected after reading a finite prefix of f , because for $n \geq m$ it is not possible to extend a^n to reach the value 2^{-m} . \square

6 Approximate Monitoring through Approximate Safety

In this section, we consider properties on extended reals $\mathbb{R}^{\pm\infty} = \mathbb{R} \cup \{-\infty, +\infty\}$. We denote by $\mathbb{R}_{\geq 0}$ the set of nonnegative real numbers.

Definition 50 (Approximate safety and co-safety). *Let $\alpha \in \mathbb{R}_{\geq 0}$. A property Φ is α -safe iff for every $f \in \Sigma^\omega$ and value $v \in \mathbb{R}^{\pm\infty}$ with $\Phi(f) < v$, there exists a prefix $s \prec f$ such that $\sup_{g \in \Sigma^\omega} \Phi(sg) < v + \alpha$. Similarly, Φ is α -co-safe iff for every $f \in \Sigma^\omega$ and $v \in \mathbb{R}^{\pm\infty}$ with $\Phi(f) > v$, there exists $s \prec f$ such that $\inf_{g \in \Sigma^\omega} \Phi(sg) > v - \alpha$. When Φ is α -safe (resp. α -co-safe) for some $\alpha \in \mathbb{R}_{\geq 0}$, we say that Φ is approximately safe (resp. approximately co-safe).*

Approximate safety can be characterized through the following relation with the safety closure.

Proposition 51. *For every error bound $\alpha \in \mathbb{R}_{\geq 0}$, a property Φ is α -safe iff $\Phi^*(f) - \Phi(f) \leq \alpha$ for all $f \in \Sigma^\omega$.*

An analogue of Proposition 51 holds for approximate co-safety and the co-safety closure. Moreover, approximate safety and approximate co-safety are dual notions that are connected by the complement operation, similarly to their precise counterparts (Theorem 27).

6.1 The Intersection of Approximate Safety and Co-safety

Recall the ghost monitor from the introduction. If, after a finite number of observations, all the possible prediction values are close enough, then we can simply freeze the current value and achieve a sufficiently small error. This happens for properties that are both approximately safe and approximately co-safe, generalizing the unfolding approximation of discounted properties [13].

Proposition 52. *For every limit property Φ and all error bounds $\alpha, \beta \in \mathbb{R}_{\geq 0}$, if Φ is α -safe and β -co-safe, then the set $S_\delta = \{s \in \Sigma^* \mid \sup_{r_1 \in \Sigma^*} \Phi(sr_1) - \inf_{r_2 \in \Sigma^*} \Phi(sr_2) \geq \delta\}$ is finite for all reals $\delta > \alpha + \beta$.*

Based on this proposition, we show that, for limit properties that are both approximately safe and approximately co-safe, the influence of the suffix on the property value is eventually negligible.

Theorem 53. *For every limit property Φ such that $\Phi(f) \in \mathbb{R}$ for all $f \in \Sigma^\omega$, and for all error bounds $\alpha, \beta \in \mathbb{R}_{\geq 0}$, if Φ is α -safe and β -co-safe, then for every real $\delta > \alpha + \beta$ and trace $f \in \Sigma^\omega$, there is a prefix $s \prec f$ such that for all continuations $w \in \Sigma^* \cup \Sigma^\omega$, we have $|\Phi(sw) - \Phi(s)| < \delta$.*

We illustrate this theorem with a *discounted safety* property.

Example 54. Let $P \subseteq \Sigma^\omega$ be a boolean safety property. We define the finitary property $\pi_P : \Sigma^* \rightarrow [0, 1]$ as follows: $\pi_P(s) = 1$ if $sf \in P$ for some $f \in \Sigma^\omega$, and $\pi_P(s) = 1 - 2^{-|r|}$ otherwise, where $r \preceq s$ is the shortest prefix with $rf \notin P$ for all $f \in \Sigma^\omega$. The limit property $\Phi = (\pi_P, \inf)$ is called *discounted safety* [3]. Because Φ is an inf-property, it is safe by Theorem 20. Now consider the finitary property π'_P defined by $\pi'_P(s) = 1 - 2^{-|s|}$ if $sf \in P$ for some $f \in \Sigma^\omega$, and $\pi'_P(s) = 1 - 2^{-|r|}$ otherwise, where $r \preceq s$ is the shortest prefix with $rf \notin P$ for all $f \in \Sigma^\omega$. Let $\Phi' = (\pi'_P, \sup)$, and note that $\Phi(f) = \Phi'(f)$ for all $f \in \Sigma^\omega$. Hence Φ is also co-safe, because it is a sup-property.

Let $f \in \Sigma^\omega$ and $\delta > 0$. For every prefix $s \prec f$, the set of possible prediction values is either the range $[1 - 2^{-|s|}, 1]$ or the singleton $\{1 - 2^{-|r|}\}$, where $r \preceq s$ is chosen as above. In the latter case, we have $|\Phi(sw) - \Phi(s)| = 0 < \delta$ for all $w \in \Sigma^* \cup \Sigma^\omega$. In the former case, since the range becomes smaller as the prefix grows, there is a prefix $s' \prec f$ with $2^{-|s'|} < \delta$, which yields $|\Phi(s'w) - \Phi(s')| < \delta$ for all $w \in \Sigma^* \cup \Sigma^\omega$. \square

6.2 Finite-state Approximate Monitoring

Monitors with finite state spaces are particularly desirable, because finite automata enjoy a plethora of desirable closure and decidability properties. Here, we prove that properties that are both approximately safe and approximately co-safe can be monitored approximately by a finite-state monitor. First, we recall the notion of abstract quantitative monitor from [30].

A binary relation \sim over Σ^* is an *equivalence relation* iff it is reflexive, symmetric, and transitive. Such a relation is *right-monotonic* iff $s_1 \sim s_2$ implies $s_1 r \sim s_2 r$ for all $s_1, s_2, r \in \Sigma^*$. For an equivalence relation \sim over Σ^* and a finite trace $s \in \Sigma^*$, we write $[s]_\sim$ for the equivalence class of \sim to which s belongs. When \sim is clear from the context, we write $[s]$ instead. We denote by Σ^*/\sim the quotient of the relation \sim .

Definition 55 (Abstract monitor [30]). *An abstract monitor $\mathcal{M} = (\sim, \gamma)$ is a pair consisting of a right-monotonic equivalence relation \sim on Σ^* and a function $\gamma : (\Sigma^*/\sim) \rightarrow \mathbb{R}^{\pm\infty}$. The monitor \mathcal{M} is finite-state iff the relation*

\sim has finitely many equivalence classes. Let $\delta_{\text{fin}}, \delta_{\text{lim}} \in \mathbb{R}^{\pm\infty}$ be error bounds. We say that \mathcal{M} is a $(\delta_{\text{fin}}, \delta_{\text{lim}})$ -monitor for a given limit property $\Phi = (\pi, \ell)$ iff for all $s \in \Sigma^*$ and $f \in \Sigma^\omega$, we have $|\pi(s) - \gamma([s])| \leq \delta_{\text{fin}}$ and $|\ell_{s \prec f}(\pi(s)) - \ell_{s \prec f}(\gamma([s]))| \leq \delta_{\text{lim}}$.

Building on Theorem 53, we identify a sufficient condition to guarantee the existence of an abstract monitor with finitely many equivalence classes.

Theorem 56. *For every limit property Φ such that $\Phi(f) \in \mathbb{R}$ for all $f \in \Sigma^\omega$, and for all error bounds $\alpha, \beta \in \mathbb{R}_{\geq 0}$, if Φ is α -safe and β -co-safe, then for every real $\delta > \alpha + \beta$, there exists a finite-state (δ, δ) -monitor for Φ .*

Due to Theorem 56, the discounted safety property of Example 54 has a finite-state monitor for every positive error bound. We remark that Theorem 56 is proved by a construction that generalizes the unfolding approach for the approximate determinization of discounted automata [12], which unfolds an automaton until the distance constraint is satisfied.

7 Conclusion

We presented a generalization of safety and liveness that lifts the safety-progress hierarchy to the quantitative setting of [18] while preserving major desirable features of the boolean setting, such as the safety-liveness decomposition.

Monitorability identifies a boundary separating properties that can be verified or falsified from a finite number of observations, from those that cannot. Safety-liveness and co-safety-co-liveness decompositions allow us separate, for an individual property, monitorable parts from nonmonitorable parts. The larger the monitorable parts of the given property, the stronger the decomposition. We provided the strongest known safety-liveness decomposition, which consists of a pointwise minimum between a safe part defined by a quantitative safety closure, and a live part which corrects for the difference. We then defined approximate safety as the relaxation of safety by a parametric error bound. This further increases the monitorability of properties and offers monitorability at a parametric cost. In fact, we showed that every property that is both approximately safe and approximately co-safe can be monitored arbitrarily precisely by a finite-state monitor. A future direction is to extend our decomposition to approximate safety together with a support for quantitative assumptions [32].

The literature contains efficient model-checking procedures that leverage the boolean safety hypothesis [36,40]. We thus expect that also quantitative safety and co-safety, and their approximations, enable efficient verification algorithms for quantitative properties.

Acknowledgments. We thank the anonymous reviewers for their helpful comments. This work was supported in part by the ERC-2020-AdG 101020093.

References

1. de Alfaro, L., Faella, M., Henzinger, T.A., Majumdar, R., Stoelinga, M.: Model checking discounted temporal properties. *Theor. Comput. Sci.* **345**(1), 139–170 (2005). <https://doi.org/10.1016/j.tcs.2005.07.033>
2. de Alfaro, L., Faella, M., Stoelinga, M.: Linear and branching metrics for quantitative transition systems. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12–16, 2004. Proceedings. Lecture Notes in Computer Science*, vol. 3142, pp. 97–109. Springer (2004). https://doi.org/10.1007/978-3-540-27836-8_11
3. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the future in systems theory. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) *Automata, Languages and Programming, 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30 - July 4, 2003. Proceedings. Lecture Notes in Computer Science*, vol. 2719, pp. 1022–1037. Springer (2003). https://doi.org/10.1007/3-540-45061-0_79
4. Alpern, B., Schneider, F.B.: Defining liveness. *Inf. Process. Lett.* **21**(4), 181–185 (1985). [https://doi.org/10.1016/0020-0190\(85\)90056-0](https://doi.org/10.1016/0020-0190(85)90056-0)
5. Alpern, B., Schneider, F.B.: Recognizing safety and liveness. *Distributed Comput.* **2**(3), 117–126 (1987). <https://doi.org/10.1007/BF01782772>
6. Bartocci, E., Falcone, Y., Francalanza, A., Reger, G.: Introduction to runtime verification. In: Bartocci, E., Falcone, Y. (eds.) *Lectures on Runtime Verification - Introductory and Advanced Topics, Lecture Notes in Computer Science*, vol. 10457, pp. 1–33. Springer (2018). https://doi.org/10.1007/978-3-319-75632-5_1
7. Bauer, A., Leucker, M., Schallhart, C.: Comparing LTL semantics for runtime verification. *J. Log. Comput.* **20**(3), 651–674 (2010). <https://doi.org/10.1093/logcom/exn075>
8. Bauer, A., Leucker, M., Schallhart, C.: Runtime verification for LTL and TLTL. *ACM Trans. Softw. Eng. Methodol.* **20**(4), 14:1–14:64 (2011). <https://doi.org/10.1145/2000799.2000800>
9. Bloem, R., Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Better quality in synthesis through quantitative objectives. In: Bouajjani, A., Maler, O. (eds.) *Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings. Lecture Notes in Computer Science*, vol. 5643, pp. 140–156. Springer (2009). https://doi.org/10.1007/978-3-642-02658-4_14
10. Bloem, R., Chatterjee, K., Jobstmann, B.: Graph games and reactive synthesis. In: Clarke, E.M., Henzinger, T.A., Veith, H., Bloem, R. (eds.) *Handbook of Model Checking*, pp. 921–962. Springer (2018). https://doi.org/10.1007/978-3-319-10575-8_27
11. Boker, U., Chatterjee, K., Henzinger, T.A., Kupferman, O.: Temporal specifications with accumulative values. *ACM Trans. Comput. Log.* **15**(4), 27:1–27:25 (2014). <https://doi.org/10.1145/2629686>
12. Boker, U., Henzinger, T.A.: Approximate determinization of quantitative automata. In: D’Souza, D., Kavitha, T., Radhakrishnan, J. (eds.) *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2012, December 15–17, 2012, Hyderabad, India. LIPIcs*, vol. 18, pp. 362–373. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2012). <https://doi.org/10.4230/LIPIcs.FSTTCS.2012.362>
13. Boker, U., Henzinger, T.A.: Exact and approximate determinization of discounted-sum automata. *Log. Methods Comput. Sci.* **10**(1) (2014). [https://doi.org/10.2168/LMCS-10\(1:10\)2014](https://doi.org/10.2168/LMCS-10(1:10)2014)

14. Bouyer, P., Fahrenberg, U., Larsen, K.G., Markey, N.: Quantitative analysis of real-time systems using priced timed automata. *Commun. ACM* **54**(9), 78–87 (2011). <https://doi.org/10.1145/1995376.1995396>
15. Bouyer, P., Markey, N., Randour, M., Larsen, K.G., Laursen, S.: Average-energy games. *Acta Informatica* **55**(2), 91–127 (2018). <https://doi.org/10.1007/s00236-016-0274-1>
16. Cerný, P., Henzinger, T.A., Radhakrishna, A.: Simulation distances. *Theor. Comput. Sci.* **413**(1), 21–35 (2012). <https://doi.org/10.1016/j.tcs.2011.08.002>
17. Chang, E., Manna, Z., Pnueli, A.: The safety-progress classification. In: Bauer, F.L., Brauer, W., Schwichtenberg, H. (eds.) *Logic and Algebra of Specification*. pp. 143–202. Springer Berlin Heidelberg, Berlin, Heidelberg (1993). https://doi.org/10.1007/978-3-642-58041-3_5
18. Chatterjee, K., Doyen, L., Henzinger, T.A.: Quantitative languages. *ACM Trans. Comput. Log.* **11**(4), 23:1–23:38 (2010). <https://doi.org/10.1145/1805950.1805953>
19. Chatterjee, K., Henzinger, T.A., Otop, J.: Nested weighted automata. *ACM Trans. Comput. Log.* **18**(4), 31:1–31:44 (2017). <https://doi.org/10.1145/3152769>
20. D’Antoni, L., Samanta, R., Singh, R.: Qlose: Program repair with quantitative objectives. In: Chaudhuri, S., Farzan, A. (eds.) *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II. Lecture Notes in Computer Science*, vol. 9780, pp. 383–401. Springer (2016). https://doi.org/10.1007/978-3-319-41540-6_21
21. Fahrenberg, U., Legay, A.: Generalized quantitative analysis of metric transition systems. In: Shan, C. (ed.) *Programming Languages and Systems - 11th Asian Symposium, APLAS 2013, Melbourne, VIC, Australia, December 9-11, 2013. Proceedings. Lecture Notes in Computer Science*, vol. 8301, pp. 192–208. Springer (2013). https://doi.org/10.1007/978-3-319-03542-0_14
22. Fahrenberg, U., Legay, A.: The quantitative linear-time-branching-time spectrum. *Theor. Comput. Sci.* **538**, 54–69 (2014). <https://doi.org/10.1016/j.tcs.2013.07.030>
23. Falcone, Y., Fernandez, J., Mounier, L.: What can you verify and enforce at runtime? *Int. J. Softw. Tools Technol. Transf.* **14**(3), 349–382 (2012). <https://doi.org/10.1007/s10009-011-0196-8>
24. Faran, R., Kupferman, O.: Spanning the spectrum from safety to liveness. *Acta Informatica* **55**(8), 703–732 (2018). <https://doi.org/10.1007/s00236-017-0307-4>
25. Ferrère, T., Henzinger, T.A., Kragl, B.: Monitoring event frequencies. In: Fernández, M., Muscholl, A. (eds.) *28th EACSL Annual Conference on Computer Science Logic, CSL 2020, January 13-16, 2020, Barcelona, Spain. LIPIcs*, vol. 152, pp. 20:1–20:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). <https://doi.org/10.4230/LIPIcs.CSL.2020.20>
26. Ferrère, T., Henzinger, T.A., Saraç, N.E.: A theory of register monitors. In: Dawar, A., Grädel, E. (eds.) *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*. pp. 394–403. ACM (2018). <https://doi.org/10.1145/3209108.3209194>
27. Gorostiaga, F., Sánchez, C.: Monitorability of expressive verdicts. In: Deshmukh, J.V., Havelund, K., Perez, I. (eds.) *NASA Formal Methods - 14th International Symposium, NFM 2022, Pasadena, CA, USA, May 24-27, 2022, Proceedings. Lecture Notes in Computer Science*, vol. 13260, pp. 693–712. Springer (2022). https://doi.org/10.1007/978-3-031-06773-0_37
28. Havelund, K., Rosu, G.: Synthesizing monitors for safety properties. In: Ka-toen, J., Stevens, P. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems, 8th International Conference, TACAS 2002, Held as Part of*

- the Joint European Conference on Theory and Practice of Software, ETAPS 2002, Grenoble, France, April 8-12, 2002, Proceedings. Lecture Notes in Computer Science, vol. 2280, pp. 342–356. Springer (2002). https://doi.org/10.1007/3-540-46002-0_24
29. Henzinger, T.A.: Quantitative reactive modeling and verification. *Comput. Sci. Res. Dev.* **28**(4), 331–344 (2013). <https://doi.org/10.1007/s00450-013-0251-7>
 30. Henzinger, T.A., Mazzocchi, N., Saraç, N.E.: Abstract monitors for quantitative specifications. In: Dang, T., Stolz, V. (eds.) *Runtime Verification - 22nd International Conference, RV 2022, Tbilisi, Georgia, September 28-30, 2022, Proceedings*. Lecture Notes in Computer Science, vol. 13498, pp. 200–220. Springer (2022). https://doi.org/10.1007/978-3-031-17196-3_11
 31. Henzinger, T.A., Otop, J.: From model checking to model measuring. In: D’Argenio, P.R., Melgratti, H.C. (eds.) *CONCUR 2013 - Concurrency Theory - 24th International Conference, CONCUR 2013, Buenos Aires, Argentina, August 27-30, 2013, Proceedings*. Lecture Notes in Computer Science, vol. 8052, pp. 273–287. Springer (2013). https://doi.org/10.1007/978-3-642-40184-8_20
 32. Henzinger, T.A., Saraç, N.E.: Monitorability under assumptions. In: Deshmukh, J., Nickovic, D. (eds.) *Runtime Verification - 20th International Conference, RV 2020, Los Angeles, CA, USA, October 6-9, 2020, Proceedings*. Lecture Notes in Computer Science, vol. 12399, pp. 3–18. Springer (2020). https://doi.org/10.1007/978-3-030-60508-7_1
 33. Henzinger, T.A., Saraç, N.E.: Quantitative and approximate monitoring. In: *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*. pp. 1–14. IEEE (2021). <https://doi.org/10.1109/LICS52264.2021.9470547>
 34. Katoen, J., Song, L., Zhang, L.: Probably safe or live. In: Henzinger, T.A., Miller, D. (eds.) *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS ’14, Vienna, Austria, July 14 - 18, 2014*. pp. 55:1–55:10. ACM (2014). <https://doi.org/10.1145/2603088.2603147>
 35. Kim, M., Kannan, S., Lee, I., Sokolsky, O., Viswanathan, M.: Computational analysis of run-time monitoring - fundamentals of java-mac. In: Havelund, K., Rosu, G. (eds.) *Runtime Verification 2002, RV 2002, FLoC Satellite Event, Copenhagen, Denmark, July 26, 2002*. *Electronic Notes in Theoretical Computer Science*, vol. 70, pp. 80–94. Elsevier (2002). [https://doi.org/10.1016/S1571-0661\(04\)80578-4](https://doi.org/10.1016/S1571-0661(04)80578-4)
 36. Kupferman, O., Vardi, M.Y.: Model checking of safety properties. *Formal Methods Syst. Des.* **19**(3), 291–314 (2001). <https://doi.org/10.1023/A:1011254632723>
 37. Kwiatkowska, M., Norman, G., Parker, D.: *Probabilistic Model Checking: Advances and Applications*, pp. 73–121. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-319-57685-5_3
 38. Kwiatkowska, M.Z.: Quantitative verification: models techniques and tools. In: Crnkovic, I., Bertolino, A. (eds.) *Proceedings of the 6th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2007, Dubrovnik, Croatia, September 3-7, 2007*. pp. 449–458. ACM (2007). <https://doi.org/10.1145/1287624.1287688>
 39. Lamport, L.: Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.* **3**(2), 125–143 (1977). <https://doi.org/10.1109/TSE.1977.229904>
 40. Latvala, T.: Efficient model checking of safety properties. In: Ball, T., Rajamani, S.K. (eds.) *Model Checking Software, 10th International SPIN Workshop*. Portland, OR, USA, May 9-10, 2003, Proceedings. Lecture Notes in Computer Science,

- vol. 2648, pp. 74–88. Springer (2003). https://doi.org/10.1007/3-540-44829-2_5
41. Li, Y., Droste, M., Lei, L.: Model checking of linear-time properties in multi-valued systems. *Inf. Sci.* **377**, 51–74 (2017). <https://doi.org/10.1016/j.ins.2016.10.030>
 42. Manna, Z., Pnueli, A.: Adequate proof principles for invariance and liveness properties of concurrent programs. *Sci. Comput. Program.* **4**(3), 257–289 (1984). [https://doi.org/10.1016/0167-6423\(84\)90003-0](https://doi.org/10.1016/0167-6423(84)90003-0)
 43. Peled, D., Havelund, K.: Refining the safety-liveness classification of temporal properties according to monitorability. In: Margaria, T., Graf, S., Larsen, K.G. (eds.) *Models, Mindsets, Meta: The What, the How, and the Why Not? - Essays Dedicated to Bernhard Steffen on the Occasion of His 60th Birthday*. *Lecture Notes in Computer Science*, vol. 11200, pp. 218–234. Springer (2018). https://doi.org/10.1007/978-3-030-22348-9_14
 44. Pnueli, A., Zaks, A.: PSL model checking and run-time verification via testers. In: Misra, J., Nipkow, T., Sekerinski, E. (eds.) *FM 2006: Formal Methods, 14th International Symposium on Formal Methods, Hamilton, Canada, August 21-27, 2006, Proceedings*. *Lecture Notes in Computer Science*, vol. 4085, pp. 573–586. Springer (2006). https://doi.org/10.1007/11813040_38
 45. Qian, J., Shi, F., Cai, Y., Pan, H.: Approximate safety properties in metric transition systems. *IEEE Trans. Reliab.* **71**(1), 221–234 (2022). <https://doi.org/10.1109/TR.2021.3139616>
 46. Sistla, A.P.: Safety, liveness and fairness in temporal logic. *Formal Aspects Comput.* **6**(5), 495–512 (1994). <https://doi.org/10.1007/BF01211865>
 47. Thrane, C.R., Fahrenberg, U., Larsen, K.G.: Quantitative analysis of weighted transition systems. *J. Log. Algebraic Methods Program.* **79**(7), 689–703 (2010). <https://doi.org/10.1016/j.jlap.2010.07.010>
 48. Weiner, S., Hasson, M., Kupferman, O., Pery, E., Shevach, Z.: Weighted safety. In: Hung, D.V., Ogawa, M. (eds.) *Automated Technology for Verification and Analysis - 11th International Symposium, ATVA 2013, Hanoi, Vietnam, October 15-18, 2013, Proceedings*. *Lecture Notes in Computer Science*, vol. 8172, pp. 133–147. Springer (2013). https://doi.org/10.1007/978-3-319-02444-8_11