

On Safety Properties and Their Monitoring

Grigore Roşu

Department of Computer Science,
University of Illinois at Urbana-Champaign, USA.
`grosu@cs.uiuc.edu`

Abstract. Various definitions of safety properties as sets of execution traces have been introduced in the literature, some over finite traces, others over infinite traces, yet others over both finite and infinite traces. By employing cardinality arguments, this paper first shows that these notions of safety are ultimately equivalent, by showing each of them to have the cardinal of the continuum. It is then shown that all safety properties can be characterized as “always past” properties, and then that the problem of monitoring a safety property can be arbitrarily hard. Finally, two decidable specification formalisms for safety properties are discussed, namely extended regular expressions and past time LTL. It is shown that monitoring the former requires non-elementary space. An optimal monitor synthesis algorithm is given for the latter; the generated monitors run in space linear with the number of temporal operators and in time linear with the size of the formula.

1 Introduction

Informally, a *safety property* is a behavioral property on execution traces which, once violated, cannot be satisfied anymore. For example, a property “always $x > 0$ ” is violated when $x \leq 0$ is observed for the first time; this safety property remains violated even though eventually $x > 0$. That means that one can identify each safety property with a set of “bad” finite execution traces, with the intuition that once one of those is reached the safety property is violated.

There are several apparently different ways to formalize safety. Perhaps the most immediate one is to complement the “bad traces” above and thus to define a safety property as a prefix-closed property over finite traces (containing the “good traces”) – by “property” in this paper we mean a set of finite or infinite traces. Inspired by Lamport [19], Alpern and Schneider [4] define safety properties over infinite traces as ones with the property that if an infinite trace is unacceptable then there must be some finite prefix of it which is already unacceptable, in the sense that there is no acceptable infinite completion of it. Are these two definitions of safety equivalent? We show rather indirectly that they are indeed equivalent, by showing that the corresponding sets of safety properties have the cardinal c of the continuum (i.e., the cardinal of \mathbb{R} , the set of real numbers), so there must be some bijective mapping between the two. Unfortunately, the existence of such a bijection is as little informative as the

existence of a bijection between \mathbb{R} and $\mathbb{R} - \mathbb{Q}$. To capture the relationship between finite-trace and infinite-trace safety properties in a more meaningful way, we introduce a subset of finite-trace safety properties that we call *persistent*, and then construct an explicit bijection between that subset and the infinite-trace safety properties. Interestingly, over finite traces there are as many safety properties as unrestricted properties (finite-traces are enumerable and $\mathcal{P}(\mathbb{N})$ is in bijection with \mathbb{R}), while over infinite traces there are c safety properties versus 2^c unrestricted properties (infinite traces are in bijection with \mathbb{R}).

It is also common to define safety properties as properties over both finite and infinite traces, the intuition for the finite traces being that of unfinished computations. For example, Lamport [20] extends the notion of infinite-trace safety properties to properties over both finite and infinite traces, while Schneider et al. [27, 10] give an alternative definition of safety over finite and infinite traces, called “execution monitoring”. One immediate technical advantage of allowing both finite and infinite traces is that one can define prefix closed properties. We indirectly show that prefix closeness is not a sufficient condition to define safety properties when infinite traces are also allowed, by showing that there are 2^c prefix closed properties versus, as expected, “only” c safety properties.

Another common way to specify safety properties is as “always past” properties, that is, as properties containing only words whose finite prefixes satisfy a given property. If P is a property on finite prefixes, then we write $\Box P$ for the “always past P ” corresponding safety property containing those words whose prefixes are in P . We show that specifying safety properties as “always past” properties is fully justified by showing that, for each of the three types of traces (finite, infinite, and both), the “always past” properties are precisely the safety properties as defined above. It is common to specify P using some logical formalism, for example past time LTL [21]; for example, one can specify the property “ a before b ” in past LTL as $b \rightarrow \diamond a$, and “always a before b ” as $\Box(b \rightarrow \diamond a)$.

The problem of monitoring safety properties is also investigated in this paper. Since there are as many safety properties as real numbers, it is not unexpected that some of them can be very hard to monitor. We show that monitoring safety properties can be arbitrarily hard by showing that it reduces to checking membership of natural numbers to sets of natural numbers. In particular, we can associate to any degree in the arithmetic hierarchy as well as to any complexity class in the decidable universe, a safety property whose monitoring is as hard as that degree or complexity class.

In practice not all ($c = |\mathbb{R}|$) safety properties are meaningful, but only those ($\aleph_0 = |\mathbb{N}|$) which are specifiable using formal specification languages or logics of interest. We also investigate the problem of monitoring safety properties expressed using two common formalisms, namely regular expressions extended with complement, also called extended regular expressions (ERE), and LTL. It is known that both formalisms allow polynomial finite-trace membership checking algorithms [13, 22] if one has random access to the trace, but that both require exponential space if the trace can only be analyzed online [23, 16]. It is also known that LTL can indeed be monitored in exponential space [9] and so

is claimed¹ for EREs in [23]. We show that the claim in [23] is, unfortunately, wrong, by showing that ERE monitoring requires non-elementary space. To do so, we propose for any $n \in \mathbb{N}$ a safety property P_n whose monitoring requires space non-elementary in n , as well as an ERE of size $O(n^3)$. Since the known monitoring algorithms for LTL in its full generality are asymptotically optimal, what is left to do is to consider important fragments of LTL. We focus on the “always past” fragment and give a monitor synthesis algorithm that takes formulae φ and generate monitors for them that need $O(k)$ total space and $O(|\varphi|)$ time to process each event, where k is the number of past operators in φ . This improves over the best known algorithm that needs space $O(|\varphi|)$ (and same time).

2 Preliminaries and Notations

We let \mathbb{N} denote the set of natural numbers including 0 but excluding the infinity symbol ∞ and let \mathbb{N}_∞ denote the set $\mathbb{N} \cup \{\infty\}$. We also let \mathbb{Q} denote the set of rational numbers and \mathbb{R} the set of real numbers; as for natural numbers, the “ ∞ ” subscript can also be added to \mathbb{Q} and \mathbb{R} for the corresponding extensions of these sets. \mathbb{Q}^+ and \mathbb{R}^+ denote the sets of strictly positive (0 not included) rational and real numbers, respectively.

We fix a set Σ of elements called *events* or *states*. We call words in Σ^* *finite traces* and those in Σ^ω *infinite traces*. If $u \in \Sigma^* \cup \Sigma^\omega$ then u_i is the i -th state or event that appears in u . We call *finite-trace properties* sets $P \subseteq \Sigma^*$ of finite traces, *infinite-trace properties* sets $P \subseteq \Sigma^\omega$ of infinite traces, and just *properties* sets $P \subseteq \Sigma^* \cup \Sigma^\omega$ of finite or infinite traces. If the finite or infinite aspect of traces is understood from context, then we may call any of the types or properties above just *properties*. We may write $P(w)$ for a property P and a (finite or infinite) trace w whenever $w \in P$. Traces and properties are more commonly called *words* and *languages* in the literature; we prefer to call them traces and properties to better reflect the intuition that our target application is monitoring and system observance, not formal languages. We take, however, the liberty to also call them words and languages.

Let $\text{prefixes}: \Sigma^* \rightarrow \mathcal{P}(\Sigma^*)$ be the prefix function returning for any finite trace all its prefixes, and let $\text{prefixes}: \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$ be its corresponding closure operator that takes sets of finite traces and closes them under prefixes. Note that $\text{prefixes}: \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$ is indeed a closure operator, that is, it is extensive ($P \subseteq \text{prefixes}(P)$), monotone ($P \subseteq P'$ implies $\text{prefixes}(P) \subseteq \text{prefixes}(P')$), and idempotent ($\text{prefixes}(\text{prefixes}(P)) = \text{prefixes}(P)$). We overload the function prefixes to one that also returns the finite prefixes of infinite traces. Let $\text{prefixes}: \Sigma^\omega \rightarrow \mathcal{P}(\Sigma^*)$ be the function returning for any infinite trace u all its finite prefixes $\text{prefixes}(u)$, and let $\text{prefixes}: \mathcal{P}(\Sigma^\omega) \rightarrow \mathcal{P}(\Sigma^*)$ be its corresponding extension to sets of infinite traces.

In some cases states can be simply identified with their names, or labels, and specifications of properties on traces may just refer to those labels. For

¹ We thank Prasanna Thati for finding a subtle error in the proof in [23].

example, the regular expression $(s_1 \cdot s_2)^*$ specifies all those finite traces starting with state s_1 and in which states s_1 and s_2 alternate. In other cases, one can think of states as sets of atomic predicates, that is, predicates that hold in those states: if s is a state and a is an atomic predicate, then we say that $a(s)$ is true iff a “holds” in s ; thus, if all it matters with respect to states is which predicates hold and which do not hold in each state, then states can be faithfully identified with sets of predicates. We prefer to stay loose with respect to what “holds” means, because, depending on the context, it can mean anything. In conventional software situations, atomic predicates can be: boolean expressions over variables of the program, their satisfaction being decided by evaluating them in the current state of the program; or whether a function is being called or returned from; or whether a particular variable is being written to; or whether a particular lock is being held by a particular thread; and so on. In the presence of atomic predicates, specifications of properties on traces typically only refer to the atomic predicates. For example, the property “always a before b ”, that is, those traces containing no state in which b holds that is not preceded by some state in which a holds (for example, a can stand for “authentication” and b for “resource access”), can be expressed in LTL as the formula $\Box(b \rightarrow \diamond a)$.

Let us recall some basic notions and notations from formal languages, temporarily using the consecrated terminology of “words” and “languages” instead of traces and properties. For an alphabet Σ , let \mathcal{L}_Σ denote the set of languages over Σ , i.e., the powerset $\mathcal{P}(\Sigma^*)$. By abuse of language and notation, let \emptyset denote the empty language $\{\}$ and ϵ the language containing only the empty word, $\{\epsilon\}$. If $L_1, L_2 \in \mathcal{L}_\Sigma$ then $L_1 \cdot L_2$ is the language $\{\alpha_1\alpha_2 \mid \alpha_1 \in L_1 \text{ and } \alpha_2 \in L_2\}$. note that $L \cdot \emptyset = \emptyset \cdot L = \emptyset$ and $L \cdot \epsilon = \epsilon \cdot L = L$. If $L \in \mathcal{L}_\Sigma$ then L^* is $\{\alpha_1\alpha_2 \cdots \alpha_n \mid n \geq 0 \text{ and } \alpha_1, \alpha_2, \dots, \alpha_n \in L\}$ and $\neg L$ is $\Sigma^* - L$.

We next recall some notions related to cardinality. If A is any set, we let $|A|$ denote the *cardinal* of A , which expresses the size of A . When A is finite, $|A|$ is precisely the number of elements of A and we call it a *finite cardinal*. Infinite sets can have different cardinals, called *transfinite* or even *infinite*. For example, natural numbers \mathbb{N} have the cardinal \aleph_0 (pronounced “aleph zero”) and real numbers \mathbb{R} have the cardinal c , also called the *cardinal of the continuum*. Two sets A and B are said to have the same cardinal, written $|A| = |B|$, iff there is some bijective mapping between the two. We write $|A| \leq |B|$ iff there is some injective mapping from A to B .

The famous *Cantor-Bernstein-Schroeder theorem* states that if $|A| \leq |B|$ and $|B| \leq |A|$ then $|A| = |B|$. In other words, to show that there is some bijection between sets A and B , it suffices to find an injection from A to B and an injection from B to A . The two injections need not be bijections. For example, the inclusion of the interval $(0, 1)$ in \mathbb{R}^+ is obviously an injection, so $|(0, 1)| \leq |\mathbb{R}^+|$. On the other hand, the function $x \rightsquigarrow x/(2x + 1)$ from \mathbb{R}^+ to $(0, 1)$ (in fact its codomain is the interval $(0, 1/2)$) is also injective, so $|\mathbb{R}^+| \leq |(0, 1)|$. Neither of the two injective functions is bijective, yet by the Cantor-Bernstein-Schroeder theorem there is some bijection between $(0, 1)$ and \mathbb{R}^+ , that is, $|(0, 1)| = |\mathbb{R}^+|$. We will use this theorem to relate the various types of safety properties; for example,

we will show that there is an injective function from safety properties over finite traces to safety properties over infinite traces and another injective function in the opposite direction. Unfortunately, the Cantor-Bernstein-Schroeder theorem is existential: it only says that some bijection exists between the two sets, but it does not give us an explicit bijection. Since the visualization of a concrete bijection between different sets of safety properties can be very meaningful, we will avoid using the Cantor-Bernstein-Schroeder theorem when we can find an explicit bijection between two sets of safety properties.

If A is a set of cardinal α , then 2^α is the cardinal of $\mathcal{P}(A)$, the power set of A (the set of subsets of A). It is known that $2^{\aleph_0} = c$, that is, there are as many sets of natural numbers as real numbers. The famous, still unanswered *continuum hypothesis*, states that there is no set whose size is strictly between \aleph_0 and c ; more generally, it states that, for any transfinite cardinal α , there is no proper cardinal between α and 2^α . If A and B are infinite sets, then $|A| + |B|$ and $|A| \cdot |B|$ are the cardinals of the sets $A \cup B$ and $A \times B$, respectively. An important property of transfinite cardinals is that of *absorption* – the larger cardinal absorbs the smaller one: if α and β are transfinite cardinals such that $\alpha \leq \beta$, then $\alpha + \beta = \alpha \cdot \beta = \beta$; in particular, $c \cdot 2^c = 2^c$. Besides sets of natural numbers, there are several other important sets that have cardinal c : streams (i.e., infinite sequences) of booleans, streams of reals, non-empty closed or open intervals of reals, as well as the sets of all open or closed sets of reals, respectively.

For our purposes, if Σ is an enumerable set of states, then Σ^* is also enumerable, so it has cardinal \aleph_0 . Also, if Σ is no denser than the set of real numbers, in particular if it is finite, then Σ^ω has the cardinal c , because it is equivalent to streams of states. We can then immediately infer that the set of finite-trace properties over Σ has cardinal $2^{\aleph_0} = c$, while the set of infinite-trace properties has cardinal 2^c .

3 Safety Properties

Intuitively, a safety property of a system is one stating that the system cannot “go wrong”, or, as Lamport [19] put it, that the “bad thing” never happens. In other words, in order for a system to violate a safety property, it should eventually “go wrong” or the “bad thing” should eventually happen. There is a very strong relationship between safety properties and runtime monitoring: if a safety property is violated by a running system, then the violation should happen *during* the execution of the system, in a finite amount of time, so a monitor for that property observing the running system should be able to detect the violation; an additional point in the favor of monitoring is that, if a system violates a safety property at some moment during its execution, then there is no way for the system to continue its execution to eventually satisfy the property, so a monitor needs not wait for a better future once it detects a bad present/past.

State properties or assertions that need only the current state of the running system to check whether they are violated or not, such as “no division by 0”, or “ x positive”, or no deadlock, are common safety properties; once violated, one

can stop the computation or take corrective measures. However, there are also interesting safety properties that involve more than one state of the system, such as “if one uses resource x then one must have authenticated at some moment in the past”, or “any start of a process must be followed by a stop within 10 units of time”, or “take command from user only if the user has logged in at some moment in the past and has not logged out since then”, etc. Needless to say that the atomic events, or states, which form execution traces on which safety properties are defined, can be quite abstract: not all the details of a system execution are relevant for the particular safety property of interest. In the context of monitoring, these relevant events or states can be extracted by means of appropriate instrumentation of the system. For example, runtime monitoring systems such as Tracematches [3] and MOP [7] use aspect-oriented technology to “hook” relevant observation points and appropriate event filters in a system.

It is customary to define safety properties as properties over *infinite traces*, to capture the intuition that they are defined for systems that can potentially run forever, such as reactive systems. A point in favor of infinite traces is that finite traces can be regarded as special cases of infinite traces, namely ones that “stutter” indefinitely in their last state (see, for example, Abadi and Lamport [1, 2]). Infinite traces are particularly desirable when one specifies safety properties using formalisms that have infinite-trace semantics, such as linear temporal logics or corresponding automata.

While “infinity” is a convenient abstraction that is relatively broadly-accepted nowadays in mathematics and in theoretical foundations of computer science, there is no evidence so far that a system can have an infinite-trace behavior (we have not seen any). A disclaimer is in place here: we do *not* advocate finite-traces as a foundations for safety properties; all we try to do is to argue that, just because they can be seen as a special case of infinite traces, finite traces are not entirely uninteresting. For example, an airplane engine has a finite life-span and an obvious safety property for an airplane engine is precisely the one stating that its behaviors should end with a unique “unsafe” state (most likely issued by FAA). Also, a safety property associated to a one-time-access key issued to a client can be “activate, then use at most once, then close”. Using regular patterns over the alphabet of relevant events $\Sigma = \{activate, use, close\}$, this safety property can be expressed as “ $activate \cdot (\epsilon + use) \cdot close$ ”; any trace that is not a prefix of the language of this regular expression violates the property, including any other activation or use of the key after it was closed. While these finite-trace safety properties can easily be expressed as infinite-trace safety properties, we believe that that would be more artificial than simply accepting that in practice we deal with many finite-trace safety properties.

In this section we discuss various approaches to formalize safety properties and show that they are ultimately directly or indirectly equivalent. We categorize them into finite-trace safety properties, infinite-trace safety properties, and finite- and infinite-trace safety properties:

1. Section 3.1 defines safety properties over finite traces as prefix closed properties. A subset of finite-trace safety properties, that we call *persistent*, contain

only traces that “have a future” within the property, that is, finite traces that can be continued into other finite traces that are also in the safety property. Persistent safety properties appear to be the right finite-trace variant that corresponds faithfully to the more conventional infinite-trace safety properties. Even though persistent safety properties form a proper subset of finite-trace safety properties and each finite-trace safety property has a largest persistent safety property included in it, we show that there is in fact a bijection between safety properties and persistent safety properties by showing them both to have the cardinal of the continuum c .

2. In Section 3.2, we consider two standard infinite-trace definitions of a safety property, one based on the intuition that violating behaviors must manifest so after a finite number of events and the other based on the intuition of a safety property as a closed set in an appropriate topology over infinite-traces. We show them both equivalent to persistent safety properties over finite traces, by constructing an explicit bijection (as opposed to using cardinality arguments and infer the existence of a bijection); consequently, infinite-trace safety properties also have the cardinal of the continuum c . Since closed sets of real numbers are in a bijective correspondence with the real numbers, we indirectly rediscover Alpern and Schneider’s result [4] stating that infinite-trace safety properties correspond to closed sets in infinite-trace topology.
3. Section 3.3 considers safety properties defined over both finite and infinite traces. We discuss two definitions of such safety properties encountered in the literature, and, using cardinality arguments, we show their equivalence with safety properties over only finite traces. In particular, safety properties over finite and infinite traces also have the cardinality of the continuum c . We also show that prefix-closeness is not a sufficient condition to characterize (not even bijectively) such safety properties, by showing that there are significantly more (2^c) prefix-closed properties over finite and infinite traces than safety properties.

Therefore, each of the classes of safety properties is in bijection with the real numbers. Since there are so many safety properties, we can also insightfully conclude that there is *no* enumerable mechanism to define all the safety properties, because $\aleph_0 \lesssim c$. Therefore, particular logical or syntactic recursive formalisms can only define *some* of the safety properties, but not all of them.

3.1 Safety Properties over Finite Traces

One of the most common intuitions for a safety property is as a prefix-closed set of finite traces. This captures best the intuition that once something bad happened, there is no way to recover: if $w \notin P$ then there is no u such that $P(wu)$, which is equivalent to saying that if $P(wu)$ then $P(w)$, which is equivalent to saying that P is prefix closed. From a monitoring perspective, a prefix closed property can be regarded as one containing all the good (complete or partial) behaviors of the observed system: once a state is encountered that does not form a good behavior together with the previously observed states, then a violation can be reported.

Definition 1. Let \mathbf{Safety}^* be the set of finite-trace prefix-closed properties, that is, the set $\{P \in \mathcal{P}(\Sigma^*) \mid P = \text{prefixes}(P)\}$. In other words, \mathbf{Safety}^* is the set of fixed points of the prefix operator $\text{prefixes} : \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$.

The star superscript in the notation \mathbf{Safety}^* reflects that its traces are finite; in the next section we will define a set \mathbf{Safety}^ω of infinite-trace safety properties. Since $\text{prefixes}(P) \in \mathbf{Safety}^*$ for any $P \in \mathcal{P}(\Sigma^*)$, we can assume from here on that $\text{prefixes} : \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$ is actually a function $\mathcal{P}(\Sigma^*) \rightarrow \mathbf{Safety}^*$.

Example 1. Consider the one-time-access key safety property discussed above, saying that a client can “activate, then use at most once, and then close” the key. If $\Sigma = \{\text{activate}, \text{use}, \text{close}\}$, then this safety property can be expressed as the finite set of finite words $\{\epsilon, \text{activate}, \text{activate close}, \text{activate use}, \text{activate use close}\}$; no other behavior is allowed. Now suppose that the safety policy is extended to allow multiple uses of the key once activated, but still no further events once it is closed. The extended safety property has now infinitely many finite-traces:

$$\{\epsilon\} \cup \{\text{activate}\} \cdot \{\text{use}^n \mid n \in \mathbb{N}\} \cdot \{\epsilon, \text{close}\}.$$

Note that this property is indeed prefix-closed. A monitor in charge of online checking this safety property would report a violation if the first event is not *activate*, or if it encounters any second *activate* event, or if it encounters any event after a *close* event is observed, including another *close* event.

It is interesting to note that this finite-trace safety property encompasses both finite and infinite aspects. For example, it does not preclude behaviors in which one sees an *activate* event and then an arbitrary number of *use* events; *use* events can persist indefinitely after an *activate* event without violating the property. On the other hand, once a *close* event is encountered, no other event can be further seen. We will shortly see that the safety property above properly includes the *persistent* safety property $\{\epsilon\} \cup \{\text{activate use}^n \mid n \in \mathbb{N}\}$, which corresponds to the infinite-trace safety property $\{\text{activate use}^\omega\}$.

While prefix closeness seems to be the right requirement for a safety property, one can argue that it is not sufficient. For example, in the context of reactive systems that supposedly run forever, one may think of a safety property as one containing safe finite traces, that is, ones for which the reactive system can always find a way to continue its execution safely. The definition of safety properties above includes, among other safety properties, the empty set of traces as well as all prefix-closed *finite* sets of finite traces; any reactive system will eventually violate such safety properties, so one can say that the definition of safety property above is too generous.

We next define *persistent safety properties* as ones that always allow a future; intuitively, an observed reactive system that is in a safe state can always (if persistent enough) find a way to continue its execution to a next safe state.

Definition 2. Let $\mathbf{PersistentSafety}^*$ be the set of finite-trace persistent safety properties, that is, safety properties $P \in \mathbf{Safety}^*$ such that if $P(w)$ for some $w \in \Sigma^*$ then there is some $a \in \Sigma$ such that $P(wa)$.

If a persistent safety property is non-empty, then note that it must contain an infinity number of words. The persistency aspect of a finite-trace safety property can be regarded, in some sense, as a liveness argument. Indeed, assuming that it is a “good thing” for a trace to be indefinitely continued, then a persistent safety property is one in which the “good thing” always eventually happens. If one takes the liberty to regard “stuck” computations as unfair, then the persistency aspect above can also be regarded as a fairness argument.

Another way to think of persistent safety properties is as a means to refer to infinite behaviors by means of finite traces. This view is, in some sense, dual to the more common approach to regard finite behaviors as infinite behaviors that stutter infinitely in a “last” state (see, for example, Abadi and Lamport [1, 2] for a formalization of such last-state infinite stuttering).

Note that if Σ is a degenerate set of events containing only one element, that is, if $|\Sigma| = 1$, then $|\text{Safety}^*| = \aleph_0$ and $|\text{PersistentSafety}^*| = 1$; indeed, if $\Sigma = \{a\}$ then Safety^* contains precisely the properties $a^{\leq n} = \{a^i \mid 0 \leq i \leq n\}$ for each $n \in \mathbb{N}$ and $\text{PersistentSafety}^*$ contains precisely one property, namely $\{a^n \mid n \in \mathbb{N}\}$. The case when there is only one event or state in Σ is neither interesting nor practical. Therefore, from here on in this paper we take the liberty to assume that $|\Sigma| \geq 2$. Since in practice Σ contains states or events generate by a computer, for simplicity in stating some of the subsequent results, we also take the liberty to assume that $|\Sigma| \leq \aleph_0$; therefore, Σ can be any finite or recursively enumerable set, including \mathbb{N} , \mathbb{N}_∞ , \mathbb{Q} , etc., but cannot be \mathbb{R} or any set “larger” than \mathbb{R} . With these assumptions, it follows that $|\Sigma^*| = \aleph_0$ (finite words are recursively enumerable) and $|\Sigma^\omega| = c$ (infinite streams have the cardinality of the continuum).

Proposition 1. *Safety* and PersistentSafety* are closed under union; Safety* is also closed under intersection.*

Proof. The union and the intersection of prefix-closed properties is also prefix-closed. Also, the union of persistent prefix-closed properties is also persistent.

The intersection of persistent safety properties is not necessarily persistent:

Example 2. Let Σ be the set $\{0, 1\}$. Let $P = \{1^m \mid m \in \mathbb{N}\}$ and $P' = \{\epsilon\} \cup \{10^m \mid m \in \mathbb{N}\}$ be two persistent safety properties, where ϵ is the empty word (the word containing no letters). Then $P \cap P'$ is the finite safety property $\{\epsilon, 1\}$, which is not persistent. If one thinks that this happened because $P \cap P'$ does not contain any proper (i.e., non-empty) persistent property, then one can take instead the persistent safety properties $P = \{0^n \mid n \in \mathbb{N}\} \cdot \{1^m \mid m \in \mathbb{N}\}$ and $P' = \{0^n \mid n \in \mathbb{N}\} \cdot (\{\epsilon\} \cup \{10^m \mid m \in \mathbb{N}\})$, whose intersection is the safety property $\{0^n \mid n \in \mathbb{N}\} \cup \{0^n 1 \mid n \in \mathbb{N}\}$. This safety property is not persistent because its words ending in 1 cannot persist, but it contains the proper persistent safety property $\{0^n \mid n \in \mathbb{N}\}$.

Therefore, we can associate to any safety property in Safety^* a largest persistent safety property in $\text{PersistentSafety}^*$, by simply taking the union of all

persistent safety properties that are included in the original safety property (the empty property is one of them, the smallest):

Definition 3. For a safety property $P \in \text{Safety}^*$, let $P^\circ \in \text{PersistentSafety}^*$ be the largest persistent safety property with $P^\circ \subseteq P$.

The following example shows that one may need to eliminate infinitely many words from a safety property in order to obtain a persistent safety property:

Example 3. Let $\Sigma = \{0, 1\}$ and let P be the safety property $\{0^n \mid n \in \mathbb{N}\} \cup \{0^n 1 \mid n \in \mathbb{N}\}$. Then P° can contain no word ending with a 1 and can contain all the words of 0's. Therefore, $P^\circ = \{0^n \mid n \in \mathbb{N}\}$.

Finite safety properties obviously cannot contain any non-empty persistent safety property, that is, $P^\circ = \emptyset$ if P is finite. But what if P is infinite? Is it always the case that it contains a non-empty persistent safety property? Interestingly, it turns out that this is true if and only if Σ is finite:

Proposition 2. If Σ is finite and P is a safety property containing infinitely many words, then $P^\circ \neq \emptyset$.

Proof. For each letter $a \in \Sigma$, let us define the *derivative of P wrt a* , written $\delta_a(P)$, as the language $\{w \in \Sigma^* \mid aw \in P\}$. Since

$$P = \{\epsilon\} \cup \bigcup_{a \in \Sigma} \{a\} \cdot \delta_a(P)$$

since Σ is finite, and since P is infinite, it follows that there is some $a_1 \in \Sigma$ such that $\delta_{a_1}(P)$ is infinite; note that $a_1 \in P$ since P is prefix closed. Similarly, since $\delta_{a_1}(P)$ is infinite, there is some $a_2 \in \Sigma$ such that $\delta_{a_2}(\delta_{a_1}(P))$ is infinite and $a_1 a_2 \in P$. Iterating this reasoning, we can find some $a_n \in \Sigma$ for each $n \in \mathbb{N}$, such that $a_1 a_2 \dots a_n \in P$ and $\delta_{a_n}(\dots(\delta_{a_2}(\delta_{a_1}(P)))\dots)$ is infinite, that is, the set $\{w \in \Sigma^* \mid a_1 a_2 \dots a_n w \in P\}$ is infinite. It is now easy to see that the set $\{a_1 a_2 \dots a_n \mid n \in \mathbb{N}\} \subseteq P$ is persistent. Therefore, $P^\circ \neq \emptyset$.

The following example shows that Σ must indeed be finite in order for the result above to hold:

Example 4. Consider some infinite set of events or states Σ . Then we can label distinct elements in Σ with distinct labels in $\mathbb{N} \cup \{\infty\}$. We only need these elements from Σ ; therefore, without loss of generality, we can assume that $\Sigma = \mathbb{N} \cup \{\infty\}$. Let P be the safety property

$$\{\epsilon\} \cup \{\infty n(n-1) \dots (m+1)m \mid 0 \leq m \leq n+1\},$$

where ϵ is the empty word (the word containing no letters) and $(n+1) \dots n$ is also the empty word for any $n \in \mathbb{N}$. Then P° is the empty property. Indeed, note that any persistent safety property P' included in P cannot have traces ending in 0, because those cannot be continued into other traces in P ; since P'

cannot contain traces ending in 0, it cannot contain traces ending in 1 either, because such traces can only be continued with a 0 letter into traces in P , but those traces have already been decided that cannot be part of P' ; inductively, one can show that P' can contain no words ending in letters that are natural numbers in \mathbb{N} . Since the only trace in P ending in ∞ is ∞ itself and since ∞ can only be continued with a natural number letter into a trace in P but such trace cannot belong to P' , we deduce that P' can contain no word with letters in Σ . In particular, P° must be empty.

Even though we know that the largest persistent safety property P° included into a safety property P always exists because $\text{PersistentSafety}^*$ is closed under union, we would like to have a more constructive way to obtain it. A first and obvious thing to do is to eliminate from P all the “stuck” computations, that is, those which cannot be added any new state to obtain a trace that is also in P . This removal step does not destroy the prefix-closeness of P , but it may reveal new computations which are stuck. By iteratively eliminating all the computations that get stuck in a finite number of steps, one would expect to obtain a persistent safety property, namely precisely P° . It turns out that this is indeed true only if Σ is finite. If that is the case, then the following can also be used as an alternative definition of P° :

Proposition 3. *Given safety property $P \in \text{Safety}^*$, then let P° be the property $\{w \in P \mid (\exists a \in \Sigma) wa \in P\}$. Also, let $\{P_i \mid i \in \mathbb{N}\}$ be properties defined as $P_0 = P$ and $P_{i+1} = P_i^\circ$ for all $i \geq 0$. Then $P^\circ = \bigcap_{i \geq 0} P_i$ whenever Σ is finite.*

Proof. It is easy to see that if P is prefix-closed then $P^\circ \subseteq P$ is also prefix-closed, so P° is also a property in Safety^* . Therefore, the properties P_i form a sequence $P = P_0 \supseteq P_1 \supseteq P_2 \supseteq \dots$ of increasingly smaller safety properties.

Let us first prove that $\bigcap_{i \geq 0} P_i$ is a persistent safety property. Assume by contradiction that for some $w \in \bigcap_{i \geq 0} P_i$ there is no $a \in \Sigma$ such that $wa \in \bigcap_{i \geq 0} P_i$. In other words, we can find for each $a \in \Sigma$ some $i_a \geq 0$ such that $wa \notin P_{i_a}$. Since Σ is finite, we can let i be the largest among the natural numbers $i_a \in \mathbb{N}$ for all $a \in \Sigma$. Since $P_i \subseteq P_{i_a}$ for all $a \in \Sigma$, it should be clear that there is no $a \in \Sigma$ such that $wa \in P_i$, which means that $w \notin P_{i+1}$. This contradicts the fact that $w \in \bigcap_{i \geq 0} P_i$. Therefore, $\bigcap_{i \geq 0} P_i \in \text{PersistentSafety}^*$.

Let us now prove that $\bigcap_{i \geq 0} P_i$ is the largest persistent safety property included in P . Let P' be any persistent safety property included in P . We show by induction on i that $P' \subseteq P_i$ for all $i \in \mathbb{N}$. The base case, $P' \subseteq P_0$, is obvious. Suppose that $P' \subseteq P_i$ for some $i \in \mathbb{N}$ and let $w \in P'$. Since P' is persistent, there is some $a \in \Sigma$ such that $wa \in P' \subseteq P_i$, which means that $w \in P_{i+1}$. Since w was chosen arbitrarily, it follows that $P' \subseteq P_{i+1}$. Therefore, $P' \subseteq \bigcap_{i \geq 0} P_i$.

We next show that the finiteness of Σ was a necessary requirement in order for the result above to hold. In other words, we show that if Σ is allowed to be infinite then we can find a safety property $P \in \text{Safety}^*$ over Σ such that $P^\circ \in \text{PersistentSafety}^*$ and $\bigcap_{i \geq 0} P_i \in \text{Safety}^*$ are distinct. Since we showed in the proof of Proposition 3 that any persistent safety property P' is included in

$\bigcap_{i \geq 0} P_i$, it follows that $P^\circ \subseteq \bigcap_{i \geq 0} P_i$. Since P° is the largest persistent safety property included in P , one can easily show that $P^\circ = (\bigcap_{i \geq 0} P_i)^\circ$. However, if $P^\circ \neq \bigcap_{i \geq 0} P_i$ then $\bigcap_{i \geq 0} P_i$ cannot be persistent:

Example 5. Consider the safety property P over infinite $\Sigma = \mathbb{N} \cup \{\infty\}$ discussed in Example 4, namely $\{\epsilon\} \cup \{\infty n(n-1) \dots (m+1)m \mid 0 \leq m \leq n+1\}$. Then one can easily show by induction on $i \in \mathbb{N}$ that the properties P_i defined in Proposition 3 are the sets $\{\epsilon\} \cup \{\infty n(n-1) \dots (m+1)m \mid i \leq m \leq n+1\}$; in other words, each P_i excludes from P all the words whose last letters are smaller than i when regarded as natural numbers. Then the intersection $\bigcap_{i \geq 0} P_i$ contains no trace ending in a natural number; the only possibility left is then $\bigcap_{i \geq 0} P_i = \{\epsilon, \infty\}$, which is different from $P^\circ = \emptyset$ (see Example 4).

One may argue that $P^\circ \neq \bigcap_{i \geq 0} P_i$ above happened precisely because P° was empty. One can instead pick the safety property $Q = \{0^n \mid n \in \mathbb{N}\} \cdot P$. Then one can show following the same idea as in Example 4 that $Q^\circ = \{0^n \mid n \in \mathbb{N}\}$. Further, one can show that $Q_i = \{0^n \mid n \in \mathbb{N}\} \cdot P_i$, so $\bigcap_{i \geq 0} Q_i = \{0^n \mid n \in \mathbb{N}\} \cup \{0^n \infty \mid n \in \mathbb{N}\}$, which is different from Q° .

Persistency is reminiscent of “feasibility” introduced by Apt et al. [5] in the context of fairness, and of “machine closeness” introduced by Abadi and Lamport [1, 2] (see also Schneider [26]) in the context of refinement. Let us use the terminology “machine closeness”: a property L (typically a liveness or a fairness property) is *machine closed* for a property M (typically given as the language of some state machine) iff L does not prohibit any of the observable runtime behaviors of M , that is, iff $\text{prefixes}(M) = \text{prefixes}(M \cap L)$; for example, if M is the total property (i.e., every event is possible at any moment, i.e., $M = \Sigma^*$) and L is the property stating that “always eventually event a ”, then any prefix of M can be continued to obtain a property satisfying L . Persistency is related to machine closeness in that a safety property P is persistent if and only if P° is machine closed for P . In other words, there is nothing P can do in a finite amount of time that P° cannot do. However, there is a caveat here: since liveness and fairness are inherently infinite-trace notions, machine closeness (or feasibility) have been introduced in the context of infinite-traces. On the other hand, persistency makes sense only in the context of finite traces.

It is clear that $\text{PersistentSafety}^*$ is properly included in Safety^* . Yet, we next show that, surprisingly, there is a bijective correspondence between Safety^* and $\text{PersistentSafety}^*$, both having the cardinal of the continuum:

Theorem 1. $|\text{PersistentSafety}^*| = |\text{Safety}^*| = c$.

Proof. Since Σ^* is recursively enumerable and since $2^{\aleph_0} = c$, we can readily infer that $|\text{PersistentSafety}^*| \leq |\text{Safety}^*| \leq |\mathcal{P}(\Sigma^*)| = c$.

Let us now define an injective function φ from the open interval of real numbers $(0, 1)$ to $\text{PersistentSafety}^*$. Since $|\Sigma| \geq 2$, let us distinguish two different elements in Σ and let us label them $\bar{0}$ and $\bar{1}$. For a real $r \in (0, 1)$, let $\varphi(r)$ be the set $\{\bar{\alpha} \mid \alpha \in \{0, 1\}^* \text{ and } 0.\alpha < r\}$, where $0.\alpha$ is the (rational) number in $(0, 1)$ whose decimals in binary representation are α , and where $\bar{\alpha}$ is the word

in Σ^* corresponding to α . Note that the set $\varphi(r) \in \mathcal{P}(\Sigma^*)$ is prefix-closed for any $r \in (0, 1)$, and that if $w \in \varphi(r)$ then also $w\bar{0} \in \varphi(r)$ (the latter holds since, by real numbers conventions, $0.\alpha = 0.\alpha 0$), so $\varphi(r) \in \text{PersistentSafety}^*$. Since the set of rationals with finite number of decimals in binary representation is dense in \mathbb{R} and in particular in the interval $(0, 1)$, it follows that the function $\varphi : (0, 1) \rightarrow \text{PersistentSafety}^*$ is injective: indeed, if $r_1 \neq r_2 \in (0, 1)$, say $r_1 < r_2$, then there is some $\alpha \in \{0, 1\}^*$ such that $r_1 < 0.\alpha < r_2$, so $\varphi(r_1) \neq \varphi(r_2)$. Since the interval $(0, 1)$ has the cardinal of the continuum c , the existence of the injective function φ implies that $c \leq |\text{PersistentSafety}^*|$. By the Cantor-Bernstein-Schroeder theorem it follows that $|\text{PersistentSafety}^*| = |\text{Safety}^*| = c$.

The proof above could have been rearranged to avoid the need to use the set $\text{PersistentSafety}^*$. However, we prefer to keep it for two reasons:

1. For finite-traces, persistent safety properties appear to be more natural in the context of reactive systems than just prefix closed properties;
2. Persistent safety properties play a technical bridge role in the next section to show that the infinite-trace safety properties also have the cardinal c .

3.2 Safety Properties over Infinite Traces

The finite-trace safety properties defined above, persistent or not, rely on the intuition of a correct prefix: a safety property is identified with the set of all its finite prefixes. In the case of a persistent safety property, each “informal” infinite acceptable behavior is captured by its infinite set of finite prefixes. Even though persistent safety properties appear to capture well in a finite-trace setting the intuition of safety in the context of (infinite-trace) reactive systems, one could argue that it does not say anything about unacceptable infinite traces. Indeed, one may think that persistent safety properties do not capture the intuition that if an infinite trace is unacceptable then there must be some finite prefix of it which is already unacceptable. In this section we show that there is in fact a bijection between safety properties over infinite traces and persistent safety properties over finite traces as we defined them in the previous section.

The definition of safety properties over infinite traces below appears to be the most used definition of safety property in the literature; at our knowledge, it was formally introduced by Alpern and Schneider [4], but they credit the insights of their definition to Lamport [19]:

Definition 4. *Let Safety^ω be the set of infinite-trace properties $Q \in \mathcal{P}(\Sigma^\omega)$ s.t.: if $u \notin Q$ then there is a finite trace $w \in \text{prefixes}(u)$ s.t. $wv \notin Q$ for any $v \in \Sigma^\omega$.*

In other words, if an infinite behavior violates the safety property then there is some finite-trace “violation threshold”; once the violation threshold is reached, there is no chance to recover.

The following proposition can serve as an alternative and more compact definition of Safety^ω :

Proposition 4. $\text{Safety}^\omega = \{Q \in \mathcal{P}(\Sigma^\omega) \mid u \in Q \text{ iff } \text{prefixes}(u) \subseteq \text{prefixes}(Q)\}$.

Proof. Since $u \in Q$ implies $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$, the only thing left to show is that $Q \in \text{Safety}^\omega$ iff “ $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ implies $u \in Q$ ”; the latter is equivalent to “ $u \notin Q$ implies $\text{prefixes}(u) \not\subseteq \text{prefixes}(Q)$ ”, which is further equivalent to “ $u \notin Q$ implies there is some $w \in \text{prefixes}(u)$ s.t. $w \notin \text{prefixes}(Q)$ ”, which is indeed equivalent to $Q \in \text{Safety}^\omega$.

Another common intuition for safety properties over infinite traces is as *closed* sets in the topology corresponding to Σ^ω . Alpern and Schneider captured formally this intuition for the first time in [4]; then it was used as a convenient definition of safety by Abadi and Lamport [1, 2] among others:

Definition 5. *An infinite sequence $u^{(1)}, u^{(2)}, \dots$, of infinite traces in Σ^ω converges to $u \in \Sigma^\omega$, or u is a limit of $u^{(1)}, u^{(2)}, \dots$, written $u = \lim_i u^{(i)}$, iff for all $m \geq 0$ there is an $n \geq 0$ such that $u_1^{(i)} u_2^{(i)} \dots u_m^{(i)} = u_1 u_2 \dots u_m$ for all $i \geq n$. If $Q \in \mathcal{P}(\Sigma^\omega)$ then \overline{Q} , the closure of Q , is the set $\{\lim_i u^{(i)} \mid u^{(i)} \in Q \text{ for all } i \in \mathbb{N}\}$.*

It can be easily shown that the “overline” closure above is indeed a closure operator on Σ^ω , that is, it is extensive ($Q \subseteq \overline{Q}$), monotone ($Q \subseteq Q'$ implies $\overline{Q} \subseteq \overline{Q'}$), and idempotent ($\overline{\overline{Q}} = \overline{Q}$).

Definition 6. *Let $\text{Safety}_{\text{lim}}^\omega$ be the set of properties $\{Q \in \mathcal{P}(\Sigma^\omega) \mid Q = \overline{Q}\}$.*

As expected, the two infinite-trace safety property definitions are equivalent:

Proposition 5. *$\text{Safety}_{\text{lim}}^\omega = \text{Safety}^\omega$.*

Proof. All we need to prove is that for any $Q \in \mathcal{P}(\Sigma^\omega)$ and any $u \in \Sigma^\omega$, $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ iff $u = \lim_i u^{(i)}$ for some infinite sequence of infinite traces $u^{(1)}, u^{(2)}$. If $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ then one can find for each $i \geq 0$ some $u^{(i)} \in \Sigma^\omega$ such that $u_1 u_2 \dots u_i = u_1^{(i)} u_2^{(i)} \dots u_i^{(i)}$, so for each $m \geq 0$ one can pick $n = m$ such that $u_1 u_2 \dots u_m = u_1^{(i)} u_2^{(i)} \dots u_m^{(i)}$ for all $i \geq n$, so $u = \lim_i u^{(i)}$. Conversely, if $u = \lim_i u^{(i)}$ for some infinite sequence of infinite traces $u^{(1)}, u^{(2)}, \dots$ in Σ^ω , then for any $m \geq 0$ there is some $n \geq 0$ such that $u_1 u_2 \dots u_m = u_1^{(n)} u_2^{(n)} \dots u_m^{(n)}$, that is, for any prefix of u there is some $u' \in Q$ having the same prefix, that is, $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$.

The next result establishes the relationship between infinite-trace and finite-trace persistent safety properties, by proposing a concrete bijective mapping relating the two (as opposed to using cardinality arguments to indirectly show only the existence of such a mapping). Therefore, there is also a bijective correspondence between safety properties over infinite traces and the real numbers:

Theorem 2. $|\text{Safety}^\omega| = |\text{PersistentSafety}^*| = c$.

Proof. We show that there is a bijective function between the two sets of safety properties. Note that $\text{prefixes}(S) \in \text{PersistentSafety}^*$ for any $S \in \mathcal{P}(\Sigma^\omega)$, so prefixes is in fact a function $\mathcal{P}(\Sigma^\omega) \rightarrow \text{PersistentSafety}^*$. Let $\text{prefixes} : \text{Safety}^\omega \rightarrow \text{PersistentSafety}^*$ be the restriction of this prefix function to Safety^ω ; this function

is well-defined because $\text{prefixes}(S) \in \text{PersistentSafety}^*$ for any $S \in \mathcal{P}(\Sigma^\omega)$. Let us also define a function $\omega : \text{PersistentSafety}^* \rightarrow \text{Safety}^\omega$ as follows: $\omega(P) = \{u \in \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\}$. This function is well-defined: if $u \notin \omega(P)$ then by the definition of $\omega(P)$ there is some $w \in \text{prefixes}(u)$ such that $w \not\subseteq P$; since $w \in \text{prefixes}(wv)$ for any $v \in \Sigma^\omega$, it follows that $wv \notin \omega(P)$ for any $v \in \Sigma^\omega$.

We next show that prefixes and ω are inverse to each other. Let us first show that $\text{prefixes}(\omega(P)) = P$ for any $P \in \text{PersistentSafety}^*$. The inclusion $\text{prefixes}(\omega(P)) \subseteq P$ follows by the definition of $\omega(P)$: $\text{prefixes}(u) \subseteq P$ for any $u \in \omega(P)$. The inclusion $P \subseteq \text{prefixes}(\omega(P))$ follows from the fact that P is a persistent safety property: for any $w \in P$ one can iteratively build an infinite sequence v_1, v_2, \dots , such that $wv_1, wv_1v_2, \dots \in P$, so $wv_1v_2\dots \in \omega(P)$. Let us now show that $\omega(\text{prefixes}(Q)) = Q$ for any $Q \in \text{Safety}^\omega$. The inclusion $Q \subseteq \omega(\text{prefixes}(Q))$ is immediate. For the other inclusion, let $u \in \omega(\text{prefixes}(Q))$, that is, $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$. Suppose by contradiction that $u \notin Q$. Then there is some $w \in \text{prefixes}(u)$ such that $wv \notin Q$ for any $v \in \Sigma^\omega$. Since $w \in \text{prefixes}(u)$ and $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$, it follows that $w \in \text{prefixes}(Q)$, that is, that there is some $u' \in Q$ such that $u' = wv$ for some $v \in \Sigma^\omega$. This contradicts the fact that $wv \notin Q$ for any $v \in \Sigma^\omega$. Consequently, $u \in Q$.

The second part follows by Theorem 1.

3.3 Safety Properties over Finite and Infinite Traces

It is also common to define safety properties as properties over both finite and infinite traces, the intuition for the finite traces being that of unfinished computations. For example, Lamport [20] extends the notion of safety in Definition 4 to properties over both finite and infinite traces, while Schneider et al [27, 10] give an alternative definition of safety over finite and infinite traces. We define both approaches shortly and then show their equivalence and their 1-to-1 correspondence with real numbers. Before that, we argue that the mix of finite and infinite traces is less trivial than it may appear, by showing that there are significantly more prefix closed properties than in the case when only finite traces were considered.

Definition 7. Let $\text{PrefixClosed}^{*,\omega}$ be the set of prefix-closed sets of finite and infinite traces: for $Q \subseteq \Sigma^* \cup \Sigma^\omega$, $Q \in \text{PrefixClosed}^{*,\omega}$ iff $\text{prefixes}(Q) \subseteq Q$. Also, let $\text{PersistentPrefixClosed}^{*,\omega}$ be the set of persistent prefix-closed sets of finite and infinite traces: for $Q \in \text{PrefixClosed}^{*,\omega}$, $Q \in \text{PersistentPrefixClosed}^{*,\omega}$ iff $Q(w)$ for some $w \in \Sigma^*$ implies that there is some $a \in \Sigma$ such that $Q(wa)$.

The next result says that there is a bijective correspondence between prefix-closed and persistent prefix-closed properties also in the case of finite and infinite traces, but that there are exponentially more such properties than in the case of just finite traces:

Proposition 6. $|\text{PersistentPrefixClosed}^{*,\omega}| = |\text{PrefixClosed}^{*,\omega}| = 2^c$.

Proof. We show $2^c \leq |\text{PersistentPrefixClosed}^{*,\omega}| \leq |\text{PrefixClosed}^{*,\omega}| \leq 2^c$, where the middle inequality is immediate. To show $2^c \leq |\text{PersistentPrefixClosed}^{*,\omega}|$, let us define $\varphi: \mathcal{P}((0, 1)) \rightarrow \text{PersistentPrefixClosed}^{*,\omega}$ as

$$\varphi(R) = \bigcup_{0.\alpha \in R} \{\bar{\alpha}\} \cup \text{prefixes}(\bar{\alpha})$$

where we assume for any real number in the interval $(0, 1)$ its decimal binary representation $0.\alpha$ with $\alpha \in \{0, 1\}^\omega$ (if the number is rational then α may contain infinitely many ending 0's), and $\bar{\alpha}$ is the infinite trace in Σ^ω replacing each 0 and 1 in α by $\bar{0}$ and $\bar{1}$, respectively, where $\bar{0}$ and $\bar{1}$ are two arbitrary but fixed distinct elements in Σ (recall that $|\Sigma| \geq 2$). Note that $\varphi(R)$ is well-defined: it is clearly prefix-closed and it is also persistent because its finite traces are exactly prefixes of infinite traces, so they admit continuations in $\varphi(R)$. It is easy to see that φ is injective. Since $|(0, 1)| = c$, we conclude that $2^c \leq |\text{PersistentPrefixClosed}^{*,\omega}|$.

To show $|\text{PrefixClosed}^{*,\omega}| \leq 2^c$, note that any property in $\text{PrefixClosed}^{*,\omega}$ is a union of a subset in Σ^* and a subset in Σ^ω , so $|\text{PrefixClosed}^{*,\omega}| \leq 2^{|\Sigma^*|} \cdot 2^{|\Sigma^\omega|}$. Since $|\Sigma^*| = \aleph_0$, $|\Sigma^\omega| = c$, $2^{\aleph_0} = c$, and $c \cdot 2^c = 2^c$ (by absorption of transfinite cardinals), we get that $|\text{PrefixClosed}^{*,\omega}| \leq 2^c$.

The fact that properties in $\text{PersistentPrefixClosed}^{*,\omega}$ contain also infinite traces was crucial in showing the injectivity of φ in the proof above. A similar construction for the finite trace setting does *not* work. Indeed, if one tries to define a function $\varphi: \mathcal{P}((0, 1)) \rightarrow \text{PersistentSafety}^*$ as $\varphi(R) = \bigcup_{0.\alpha \in R} \text{prefixes}(\bar{\alpha})$, then one can show it well-defined but cannot show it injective: e.g., $\varphi((0, 0.5)) = \varphi((0, 0.5])$.

Since safety properties over finite and infinite traces are governed by the same intuitions as safety properties over only finite or over only infinite traces, the result above tells us that prefix closeness is not a sufficient condition to properly capture the safety properties. Schneider [27] proposes an additional condition in the context of his EM (execution monitoring) framework, namely that if an infinite trace is not in the property, then there is a finite prefix of it which is not in the property either. It is easy to see that this additional condition is equivalent to saying that an infinite trace is in the property whenever all its finite prefixes are in the property, which allows us to compactly define safety properties over finite and infinite traces in the EM style as follows:

Definition 8. *Safety_{EM}^{*,\omega} is the set $\{Q \subseteq \Sigma^* \cup \Sigma^\omega \mid u \in Q \text{ iff } \text{prefixes}(u) \subseteq Q\}$.*

Note that $\text{Safety}_{\text{EM}}^{*,\omega} \subset \text{PrefixClosed}^{*,\omega}$. We will shortly show that $\text{Safety}_{\text{EM}}^{*,\omega}$ is in fact exponentially smaller than $\text{PrefixClosed}^{*,\omega}$, by showing that $|\text{Safety}_{\text{EM}}^{*,\omega}| = c$.

The consecrated definition of a safety property in the context of both finite and infinite traces is perhaps the one proposed by Lamport in [20], which relaxes the one in Definition 4 by allowing u to range over both finite and infinite traces:

Definition 9. *Let $\text{Safety}^{*,\omega}$ be the set of finite- and infinite-trace properties $\{Q \subseteq \Sigma^* \cup \Sigma^\omega \mid u \notin Q \Rightarrow (\exists w \in \text{prefixes}(u)) (\forall v \in \Sigma^* \cup \Sigma^\omega) vw \notin Q\}$*

Schneider informally stated in [27] that the two definitions of safety above are equivalent. It is not hard to show it formally:

Proposition 7. $\text{Safety}_{\text{EM}}^{*,\omega} = \text{Safety}^{*,\omega}$.

Proof. First note that $\text{Safety}^{*,\omega} \subseteq \text{PrefixClosed}^{*,\omega}$: if $wu \in Q \in \text{Safety}^{*,\omega}$ and $w \notin Q$ then there is some $w' \in \text{prefixes}(w)$, say $w = w'w''$, such that $w'v \notin Q$ for any v , in particular $w'w''u \notin Q$, which contradicts $wu \in Q$.

$\text{Safety}_{\text{EM}}^{*,\omega} \subseteq \text{Safety}^{*,\omega}$: let $Q \in \text{Safety}_{\text{EM}}^{*,\omega}$ and $u \in \Sigma^* \cup \Sigma^\omega$ s.t. $\text{prefixes}(u) \subseteq Q$; if $u \notin Q$ then there is some $w \in \text{prefixes}(u)$ s.t. $wv \notin Q$ for any v , in particular for v the empty word, that is, $w \notin Q$, which contradicts $\text{prefixes}(u) \subseteq Q$.

$\text{Safety}_{\text{EM}}^{*,\omega} \subseteq \text{Safety}^{*,\omega}$: let $u \notin Q \in \text{Safety}_{\text{EM}}^{*,\omega}$; then $\text{prefixes}(u) \not\subseteq Q$, that is, there is some $w \in \text{prefixes}(u)$ s.t. $w \notin Q$; since Q is prefix-closed, it follows that $wv \notin Q$ for any $v \in \Sigma^* \cup \Sigma^\omega$.

We next show that there is a bijective correspondence between the safety properties over finite or infinite traces above and the finite trace safety properties in Section 3.1:

Theorem 3. $|\text{Safety}^{*,\omega}| = |\text{Safety}_{\text{EM}}^{*,\omega}| = |\text{Safety}^*| = c$.

Proof. $\text{Safety}^* \subset \text{Safety}_{\text{EM}}^{*,\omega}$ since the properties in $\text{Safety}_{\text{EM}}^{*,\omega}$ are prefix-closed, so $|\text{Safety}^*| \leq |\text{Safety}_{\text{EM}}^{*,\omega}|$.

Since the functions $\text{prefixes}: \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$ and $\text{prefixes}: \mathcal{P}(\Sigma^\omega) \rightarrow \mathcal{P}(\Sigma^*)$ have actual co-domains Safety^* and $\text{PersistentSafety}^*$, respectively, they can be organized as a function $\text{prefixes}: \text{Safety}_{\text{EM}}^{*,\omega} \rightarrow \text{Safety}^*$. Let us show that this function is injective. Let us assume $Q \neq Q' \in \text{Safety}_{\text{EM}}^{*,\omega}$, say $u \in Q$ and $u \notin Q'$, s.t. $\text{prefixes}(Q) = \text{prefixes}(Q')$. Since $u \in Q \in \text{Safety}_{\text{EM}}^{*,\omega}$ it follows that $\text{prefixes}(u) \subseteq \text{prefixes}(Q) \subseteq Q$, which implies that $\text{prefixes}(u) \subseteq \text{prefixes}(Q') \subseteq Q'$; since $Q' \in \text{Safety}_{\text{EM}}^{*,\omega}$, it follows that $u \in Q'$, contradiction. Therefore, $\text{prefixes}: \text{Safety}_{\text{EM}}^{*,\omega} \rightarrow \text{Safety}^*$ is injective, which proves that $\text{Safety}_{\text{EM}}^{*,\omega} \leq \text{Safety}^*$.

The rest follows by Proposition 7 and Theorem 1.

3.4 “Always Past” Characterization of Safety Properties

Another common way to specify safety properties is by giving an arbitrary property on finite traces, not necessarily prefix closed, and then to require that any acceptable behavior must have all its finite prefixes in the given property. A particularly frequent case is when one specifies the property of the finite-prefixes using the past-time fragment of linear temporal logics (LTL). For example, Manna and Pnueli [21] call the resulting “always (past LTL)” properties *safety formulae*; many other authors, including ourselves, adopted the terminology “safety formula” from Manna and Pnueli, though some qualify it as “LTL safety formula”. An example of an LTL safety formula is “always (b implies eventually in the past a)”, written using LTL notation as “ $\Box(b \rightarrow \diamond a)$ ”; here the past time formula “ $b \rightarrow \diamond a$ ” compactly specifies all the finite-traces

$$\{ws w' s' \mid w, w' \in \Sigma^*, s, s' \in \Sigma, a(s) \text{ and } b(s') \text{ hold}\} \cup \{ws \mid w \in \Sigma^*, s \in \Sigma, b(s) \text{ does not hold}\}.$$

We will investigate the case when safety properties are expressed as LTL safety formulae, as well as optimal monitoring techniques for such safety properties, in Section 8. In the remaining of this section we assume that the past time prefix properties are given as ordinary sets of finite-traces (so we make abstraction of how these properties are expressed) and show not only that the resulting “always past” properties are safety properties, but also that any safety properties can be expressed as an “always past” property. This holds for all the variants of safety properties (i.e., over finite traces, over infinite traces, or over both finite and infinite traces).

Definition 10. *Let $P \subseteq \Sigma^*$ be any property over finite traces. Then we define the “always past” property $\Box P$ as follows:*

- (finite traces) $\{w \in \Sigma^* \mid \text{prefixes}(w) \subseteq P\}$; and*
- (infinite traces) $\{u \in \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\}$; and*
- (finite and infinite traces) $\{u \in \Sigma^* \cup \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\}$.*

Let Safety_{\Box}^ , $\text{Safety}_{\Box}^\omega$ and $\text{Safety}_{\Box}^{*,\omega}$ be the corresponding sets of properties.*

In Section 8 we show that the language $\mathcal{L}(\Box\varphi)$, that corresponds to the LTL safety formula $\Box\varphi$ for φ some past-time LTL formula, is a property in $\text{Safety}_{\Box}^\omega$. If one was interested in a finite-trace or in a both finite and infinite trace semantics of LTL, then one could have shown that $\mathcal{L}(\Box\varphi) \in \text{Safety}_{\Box}^*$ or that $\mathcal{L}(\Box\varphi) \in \text{Safety}_{\Box}^{*,\omega}$.

Intuitively, one can regard the square “ \Box ” as a closure operator. Technically, it is not precisely a closure operator because it does not operate on the same set: it takes finite-trace properties to any of the three types of properties considered. Since prefixes takes properties back to finite-trace properties, we can show the following result saying that the square is a “closure operator via prefixes ”, and that safety properties are precisely the sets of words which are closed this way:

Proposition 8. *The following hold for all three types of safety properties:*

- $\Box(\text{prefixes}(\Box P)) = \Box P$ for any $P \subseteq \Sigma^*$;
- Q is a safety property iff $\Box(\text{prefixes}(Q)) = Q$.

Proof. ...

We next show that the “always past” properties are all safety properties and, moreover, that any safety property can be expressed as an “always past” property:

Theorem 4. *The following hold:*

- $\text{Safety}_{\Box}^* = \text{Safety}^*$,
- $\text{Safety}_{\Box}^\omega = \text{Safety}^\omega$, and
- $\text{Safety}_{\Box}^{*,\omega} = \text{Safety}^{*,\omega}$.

Therefore, each of the “always past” safety properties have the cardinal c.

Proof. We prove each of the equalities by double inclusion.

$\text{Safety}_{\square}^* \subseteq \text{Safety}^*$. It is true because any property $\square P$ in $\text{Safety}_{\square}^*$ is prefix-close.

$\text{Safety}^* \subseteq \text{Safety}_{\square}^*$. If $P \in \text{Safety}^*$ then we claim that $P = \square P$, so $P \in \text{Safety}_{\square}^*$.

Indeed, since P is prefix-close, $\text{prefixes}(w) \subseteq P$ for any $w \in P$, so $w \in \square P$; also, since $w \in \text{prefixes}(w)$, it follows that for any $w \in \square P$, $w \in P$.

$\text{Safety}_{\square}^{\omega} \subseteq \text{Safety}^{\omega}$. Let $\square P$ be an “always past” property in $\text{Safety}_{\square}^{\omega}$, and let u be an infinite trace in Σ^{ω} such that $u \notin \square P$. Then it follows that $\text{prefixes}(u) \not\subseteq P$,

that is, there is some $w \in \text{prefixes}(u)$ such that $w \notin P$. Since $w \in \text{prefixes}(wv)$ for any $v \in \Sigma^{\omega}$, it means that there is no $v \in \square P$ such that $\text{prefixes}(wv) \subseteq P$, that is, there is no $v \in \Sigma^{\omega}$ such that $wv \in \square P$. Therefore, $\square P \in \text{Safety}^{\omega}$.

$\text{Safety}^{\omega} \subseteq \text{Safety}_{\square}^{\omega}$. If $Q \in \text{Safety}^{\omega}$ then we claim that $Q = \square \text{prefixes}(Q)$. The

inclusion $Q \subseteq \square \text{prefixes}(Q)$ is clear, because $u \in Q$ implies $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$. For the other inclusion, note that if $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ for some $u \in \Sigma^{\omega}$, then u must be in Q : if $u \notin Q$ then by the definition of $Q \in \text{Safety}^{\omega}$, there is some $w \in \text{prefixes}(u)$ which cannot be completed into an infinite trace in Q , which contradicts $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$.

$\text{Safety}_{\square}^{*,\omega} \subseteq \text{Safety}^{*,\omega}$. By Proposition 7, it suffices to show that $\text{Safety}_{\square}^{*,\omega} \subseteq \text{Safety}_{\text{EM}}^{*,\omega}$.

Let $\square P$ be an “always past” property in $\text{Safety}_{\square}^{*,\omega}$, and let $u \in \Sigma^* \cup \Sigma^{\omega}$ such that $\text{prefixes}(u) \subseteq \text{prefixes}(\square P)$. Since $\text{prefixes}(\square P) \subseteq P$, it follows that $u \in \square P$; therefore, $\square P \in \text{Safety}_{\text{EM}}^{*,\omega}$.

$\text{Safety}^{*,\omega} \subseteq \text{Safety}_{\square}^{*,\omega}$. It is straightforward to see that $Q \in \text{Safety}_{\text{EM}}^{*,\omega}$ implies $Q = \square \text{prefixes}(Q)$.

The cardinality part follows by Theorems 1, 2, and 3.

Proposition 8 and Theorem 4 give yet another characterization for safety properties over any of the three combinations of traces, namely one in the style of the equivalent formulation of safety over infinite traces in Proposition 4: Q is a safety property iff it contains precisely the words whose prefixes are in $\text{prefixes}(Q)$.

4 On Monitoring Safety Properties

safety properties are exactly the properties that can be monitored
persistent monitors = synchronous monitors
discuss pitfalls and principles, and then define the complexity of monitoring a safety property
show that monitoring of safety properties can be arbitrarily hard

4.1 Specifying Safety Properties as Monitors

Safety properties are difficult to work with as flat sets of finite or infinite words, not only because they can contain infinitely many words, but also because such a flat representation is inconvenient for further analysis. It is important therefore to *specify* safety properties using formalisms that are easier to represent

and reason about. The next sections in this paper investigate several dedicated formalisms that proved to be convenient in specifying safety, such as finite state machines, regular expressions and temporal logics, together with corresponding limitations and efficient monitor synthesis techniques. In this section we formalize the intuitive notion of a *monitor* as a special state machine and give yet another characterization of safety properties, namely as *monitorable properties*. Since monitorable properties are completely defined by their monitors, it follows that *all* safety properties can be specified by their corresponding monitors.

Recall that we work under the assumption that Σ is a set of events or program states such that $|\Sigma| \leq \aleph_0$.

Definition 11. *A Σ -monitor, or just a monitor (when Σ is understood), is a triple $\mathcal{M} = (S, s_0, M: S \times \Sigma \rightarrow S)$, where S is a set of states, $s_0 \in S$ is the initial state, and M is a deterministic partial transition function.*

Therefore, a monitor as defined above is nothing but a deterministic state machine without final states. One could have defined monitors to be precisely standard state machines, but the subsequent technical developments would have been slightly more involved. In fact, we aim at shortly giving an even more compact definition of a monitor, that we will call *canonical monitor*, which appears to be sufficient to capture any safety property. The intuition for a monitor is the expected one: the monitor is driven by events generated by the observed program (the letters in Σ) – each newly received event drives the monitor from its current state to some other state, as indicated by the transition function M ; if the monitor ever gets stuck, that is, if the transition function M is undefined on the current state and the current event, then the monitored property is declared violated at that point by the monitor.

For any partial function $M: S \times \Sigma \rightarrow S$, we obey the following common notational convention. If $s \in S$ and $w = w_1 w_2 \dots w_k \in \Sigma^*$, we write “ $M(s, w) \downarrow$ ” whenever $M(s, w)$ is defined, that is, whenever $M(s, w_1)$ and $M(M(s, w_1), w_2)$ and ... and $M(\dots(M(s, w_1), w_2)\dots, w_k)$ are all defined, which is nothing but only saying that $M(\dots(M(s, w_1), w_2)\dots, w_k)$ is defined. If we write $M(s, w) = s'$ for some $s' \in S$, then, as expected, we mean that $M(\dots(M(s, w_1), w_2)\dots, w_k)$ is defined and equal to s' .

A monitor specifies a finite-trace property, an infinite-trace property, as well as a finite- and infinite-trace property:

Definition 12. *Given a monitor $\mathcal{M} = (S, s_0, M: S \times \Sigma \rightarrow S)$, we define the following properties:*

- $\mathcal{L}^*(\mathcal{M}) = \{w \in \Sigma^* \mid M(s_0, w) \downarrow\}$,
- $\mathcal{L}^\omega(\mathcal{M}) = \{u \in \Sigma^\omega \mid M(s_0, w) \downarrow \text{ for all } w \in \text{prefixes}(u)\}$, and
- $\mathcal{L}^{*,\omega}(\mathcal{M}) = \mathcal{L}^*(\mathcal{M}) \cup \mathcal{L}^\omega(\mathcal{M})$.

We call $\mathcal{L}^(\mathcal{M})$ the finite-trace property specified by \mathcal{M} , call $\mathcal{L}^\omega(\mathcal{M})$ the infinite-trace property specified by \mathcal{M} , and call $\mathcal{L}^{*,\omega}(\mathcal{M})$ the finite- and infinite-trace property specified by \mathcal{M} .*

Proposition 9. *For any monitor \mathcal{M} , the following hold:*

- $\mathcal{L}^*(\mathcal{M}) \in \text{Safety}^*$,
- $\mathcal{L}^\omega(\mathcal{M}) \in \text{Safety}^\omega$,
- $\mathcal{L}^{*,\omega}(\mathcal{M}) \in \text{Safety}^{*,\omega}$, and

Proof. They all follow by Theorem 4: taking P in Definition 10 to be the property $\{w \in \Sigma^* \mid M(s_0, w) \downarrow\}$, then $\Box P$ over finite traces is precisely $\mathcal{L}^*(\mathcal{M})$, over infinite traces is precisely $\mathcal{L}^\omega(\mathcal{M})$, and over finite and infinite traces is precisely $\mathcal{L}^{*,\omega}(\mathcal{M})$, so the three languages are in Safety_{\Box}^* , $\text{Safety}_{\Box}^\omega$, and $\text{Safety}_{\Box}^{*,\omega}$, respectively.

Definition 13. *For a property $P \subseteq \Sigma^* \cup \Sigma^\omega$, we let $\text{Monitors}(P)$ be the set of monitors $\{\mathcal{M} \mid \mathcal{L}^{*,\omega}(\mathcal{M}) = P\}$. If $\text{Monitors}(P) \neq \emptyset$ then P is called monitorable and the elements of $\text{Monitors}(P)$ are called monitors of P . We define the following classes of properties:*

- $\text{Monitorable}^* = \{P \subseteq \Sigma^* \mid P \text{ monitorable}\}$,
- $\text{Monitorable}^\omega = \{P \subseteq \Sigma^\omega \mid P \text{ monitorable}\}$, and
- $\text{Monitorable}^{*,\omega} = \{P \subseteq \Sigma^* \cup \Sigma^\omega \mid P \text{ monitorable}\}$.

we don't require that M should be decidable; for example, it can be an oracle, or, in the context of social domains, a committee deciding/voting whether to let it go or not

Definition 14. *A canonical Σ -monitor is a partial function $\mathcal{N} : \mathbb{N} \times \Sigma \rightarrow \mathbb{N}$. Let $\mathcal{S}_{\mathcal{N}} = \{n \mid (\exists w) \mathcal{N}(0, w) = n\}$ be the states of \mathcal{N} . As before, let*

- $\mathcal{L}^*(\mathcal{N}) = \{w \in \Sigma^* \mid \mathcal{N}(0, w) \downarrow\}$,
- $\mathcal{L}^\omega(\mathcal{N}) = \{u \in \Sigma^\omega \mid \mathcal{N}(0, w) \downarrow \text{ for all } w \in \text{prefixes}(u)\}$, and
- $\mathcal{L}^{*,\omega}(\mathcal{N}) = \mathcal{L}^*(\mathcal{N}) \cup \mathcal{L}^\omega(\mathcal{N})$.

Proposition 10. *A property $P \subseteq \Sigma^* \cup \Sigma^\omega$ is monitorable iff there is some canonical monitor \mathcal{N} such that $P = \mathcal{L}^{*,\omega}(\mathcal{N})$.*

Proof. ...

Definition 15. *A canonical monitor $\mathcal{K} : \mathbb{N} \times \Sigma \rightarrow \mathbb{N}$ is persistent iff for any $s \in \mathcal{S}_{\mathcal{K}}$, there is an $a \in \Sigma$ such that $\mathcal{K}(s, a) \downarrow$. Let $\text{PersistentMonitorable}^*$ be the set of finite-trace properties monitorable by persistent monitors, $\{\mathcal{L}^*(\mathcal{K}) \mid \mathcal{K} \text{ persistent}\}$.*

Definition 16. *If $\mathcal{N}_1, \mathcal{N}_2 : \mathbb{N} \times \Sigma \rightarrow \mathbb{N}$ are monitors, then let $\mathcal{N}_1 \sqsubseteq \mathcal{N}_2$ iff for any $s \in \mathbb{N}$ and any $a \in \Sigma$, if $\mathcal{N}_1(s, a)$ defined then $\mathcal{N}_2(s, a)$ also defined and $\mathcal{N}_2(s, a) = \mathcal{N}_1(s, a)$.*

Proposition 11. *$(\{\mathcal{K} \mid \mathcal{K} \sqsubseteq \mathcal{N} \text{ and } \mathcal{K} \text{ persistent}\}, \sqsubseteq)$ is a complete lattice for any canonical monitor \mathcal{N} .*

Proof. ...

Definition 17. For any canonical monitor \mathcal{N} , let \mathcal{N}° be the \sqsubseteq -maximal element of the complete lattice $(\{\mathcal{K} \mid \mathcal{K} \sqsubseteq \mathcal{N} \text{ and } \mathcal{K} \text{ persistent}\}, \sqsubseteq)$.

Proposition 12. The following hold for any canonical monitor \mathcal{N} :

- $\mathcal{L}^\omega(\mathcal{N}) = \mathcal{L}^\omega(\mathcal{N}^\circ)$,
- $\mathcal{L}^*(\mathcal{N}^\circ) = \mathcal{L}^*(\mathcal{N})^\circ$, and
- \mathcal{N} persistent iff $\mathcal{L}^*(\mathcal{N})$ persistent.

Proof. ...

Theorem 5. The following hold:

- $\text{Monitorable}^* = \text{Safety}^*$,
- $\text{PersistentMonitorable}^* = \text{PersistentSafety}^*$,
- $\text{Monitorable}^\omega = \text{Safety}^\omega$, and
- $\text{Monitorable}^{*,\omega} = \text{Safety}^{*,\omega}$.

Proof. ...

Definition 18. A canonical monitor \mathcal{N} such that $\text{prefixes}(\mathcal{L}^\omega(\mathcal{N})) = \mathcal{L}^*(\mathcal{N})$ is called synchronous.

Proposition 13. A monitor is synchronous iff it is persistent.

Proof. ...

4.2 The Complexity of Monitoring a Safety Property

We here address the problem of defining the complexity of monitoring. Before we give our definition, let us first discuss some pitfalls in defining this notion. Our definition for the complexity of monitoring resulted as a consequence of trying to avoid these pitfalls. Let P be a safety property.

Pitfall 1.

The complexity of monitoring P is nothing but the complexity of checking, for an input word $w \in \Sigma^$, whether $w \in \text{prefixes}(P)$.*

This would be an easy to formulate decision problem, but, unfortunately, does not capture well the intuition of monitoring, because it does not require that the word w be processed incrementally, as its letters become available from the observed system. Incremental processing of letters can make a huge difference in both how complex monitoring is and how monitoring complexity can be defined. For example, it is well-known that the membership problem of a finite word to the language of an extended regular expression (ERE), i.e., a regular expression extended with complement operators, is a polynomial problem (the classic algorithm by Hopcroft and Ullman [13] runs in space $O(m^2 \cdot n)$ and time $O(m^3 \cdot n)$, where m is the size of the word and n that of the expression). However, as shown in Section 6, there are EREs defining safety properties whose monitoring requires non-elementary space and time. Of course, this non-elementary lower-bound is expressed only as a function of the size of the ERE representing the safety property; it does not take into account the size of the monitored trace. This leads us to our first guiding principle:

Principle 1.

The complexity of monitoring a safety property P should depend only upon P , not upon the trace being monitored.

Indeed, since monitoring is a process that involves potentially unbounded traces, if the complexity of monitoring a property P were expressed as a function of the execution trace as well, then that complexity measure would be close to meaningless in practice, because monitoring reactive systems would have unbounded complexity. For example, consider an operating system monitoring some safety property on how its resources are being used by the various running processes; what one would like to know here is what is the runtime overhead of monitoring that safety property at each relevant event, and not the obvious fact that the more the operating system runs the larger the total runtime overhead is.

Nevertheless, one can admittedly argue that it would still be useful to know how complex the monitoring of P against a given finite trace w is, in terms of both the size of (some representation of) P and the size of w ; however, this is nothing but a conventional membership test decision problem, that has nothing to do with monitoring. If one picks some arbitrary off-the-shelf efficient algorithm for membership testing and uses that at each newly received event on the existing finite execution trace, then one may obtain a “monitoring” algorithm whose complexity to process each event grows in time, as events are processed. In the context of monitoring a reactive system, that means that eventually the monitoring process may become unfeasible, regardless of how many resources are initially available and regardless of how efficient the membership testing algorithm is. What one needs in order for the monitoring process to stay feasible regardless of how many events are observed, is a special membership algorithm that processes each event as received and whose state or processing time does not increase potentially unbounded as events are received. Therefore, one needs an algorithm which, if it takes resources R to check w , then it takes at most $R + \Delta$ to check a one-event continuation wa of w , where Δ *does not depend on* w . In other words, one needs a *monitor for P of complexity Δ* .

Pitfall 2.

P is typically infinite, so the complexity of monitoring P should be a function of the size of some finite specification, or representation, of P .

Indeed, since Principle 1 tells us that the complexity of monitoring P is a function of P only and not of the monitored trace, one may be tempted to conclude that it is a function of the *size* of some convenient encoding of P . There are at least two problems with this approach, that we discuss below.

- One problem is that the same property P can be specified many different ways as a structure of finite size; for example, it can be specified as a regular expression, as an extended regular expression, as a temporal logic formula, as an ordinary automaton, as a push-down automaton, etc. These formalisms may represent P as specifications of quite different sizes. Which is the most appropriate? It is, nevertheless, interesting and important to study the complexity of monitoring safety properties expressed using different specification

formalisms, as a function of the property representation size, because that can give us an idea of the amount of resources needed to monitor a particular specification; the rest of the paper is, in fact, dedicated to this problem. However, one should be aware that such a complexity measure is an attribute of the corresponding specification formalisms, not of the specified property itself. Indeed, the higher this complexity measure for a particular formalism, the higher the encoding strength of safety properties in that formalism: for example, as shown in Section 6, the complexity of monitoring safety properties expressed as EREs is non-elementary in the size of the original ERE, while the complexity of monitoring the same property expressed as an ordinary regular expression is linear in the size of the regular expression. Does that mean that one can monitor safety properties expressed as regular expressions non-elementarily more efficiently than one can monitor safety properties expressed as EREs? Of course not, because EREs and regular expressions have the same expressivity, so they specify exactly the same safety properties. All it means is that EREs can express safety-properties non-elementarily more compactly than ordinary regular expressions.

- Another problem with this approach is that apparently appropriate representations of P may be significantly larger than it takes to monitor P . One may say, for example, that, whenever possible, a natural way to specify a particular safety property is as a finite-state machine, e.g., as a monitor like in Definition 11. To be more concrete, consider that the safety property P_n saying “every 2^n -th event is a ” is specified as a monitor of 2^n states that transits with any event from each state to the next one, except for the 2^n -th state, which has only one transition, with event a , back to state 1. Therefore, the size of this representation of P_n is $\Omega(2^n)$. Assuming that each state takes n bits of storage (for example, assume that states are exactly the binary encodings of the numbers $1, 2, 3, \dots, 2^n$) and that the next state can be calculated from the current state in linear complexity with the size of the state (which is true in our scenario), then it is clear that the actual complexity of monitoring P_n is $O(n)$. If the complexity of monitoring P_n were a function of the size of the specification of P_n , then one could wrongly conclude that the complexity of monitoring “every 2^n -th event is a ” is $O(\log n)$.

Therefore, a safety property P has an inherent complexity w.r.t. monitoring, complexity which has nothing to do with how P is represented, or encoded, or specified. It is that inherent complexity attribute of safety properties that we are after here. From the discussion above, we draw our second guiding principle:

Principle 2.

The complexity of monitoring a safety property P is an attribute of P alone, not a function of the size of some adhoc representation of P .

By Theorem 5, safety properties are precisely those properties that are monitorable, that is, those properties P for which there are (finite-state or not) monitors $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ whose (finite-trace, infinite-trace, or finite- and infinite-trace – this depends upon the type of P) language is precisely P . Any

algorithm, program or system that one may come up with to be used as a monitor for P , can be organized as a monitor of the form $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ for P . Consequently, the complexity of monitoring P cannot be smaller than the functional complexity of the partial function $M : S \times \Sigma \rightarrow S$ corresponding to some “best” monitor \mathcal{M} for P ; if there are no additional restrictions, then by “best” monitor we mean the one whose functional complexity of M is smallest. In particular, if there is no monitor for P whose transition partial function M is decidable, then we can safely say that the problem of monitoring P is undecidable. This discussion leads to the following:

Pitfall 3.

The complexity of monitoring P is the functional complexity of function M , where $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ is the “best” monitor for P .

Since safety properties are precisely the monitorable properties, this appears to be a very natural definition for the complexity of monitoring. While the functional complexity of the monitor function is indeed important because it directly influences the efficiency of monitoring, it is *not* a sufficient measure for the complexity of monitoring. That is because the functional complexity of M only says how complex M is in terms of the size of *its input*; it does not say anything about how large the state of the monitor can grow in time. For example, in Section 6 we recall a rewriting-based monitoring algorithm for EREs from [23], whose states are EREs and whose transition is a derivative operation of functional complexity $O(n^2)$ taking an ERE of size n into an ERE of size $O(n^2)$. It would be very misleading to say that the complexity of monitoring EREs is $O(n^2)$, because it may sound much better than it actually is: the n^2 factor accumulates as events are processed. As shown in Section 6, any monitor for EREs, including the one based on derivatives, eventually requires non-elementary resources (in the size of the ERE) to process a new event.

Therefore, while the complexity of the function M being executed at each newly received event by a monitor \mathcal{M} is definitely a necessary and important factor to be considered when defining the complexity of monitoring using \mathcal{M} , it is not sufficient. One also needs to take into account the size of the input that is being passed to the monitoring function, that is, the size of the monitor state together with the size of the received event. In particular, a monitor storing all the observed trace has unbounded complexity, say ∞ , even though its monitoring function has trivial complexity (e.g., the “event storing” function has linear complexity). More generally, if a property admits no finite-state monitor, then we’d like to say that its monitoring complexity is ∞ : indeed, for any monitor for such a property and for any amount of resources R , there is some sequence of events that would lead the monitor to a state that needs more than R resources to be stored or computed. These observations lead us to the following:

Principle 3. The complexity of monitoring P is a function of both the functional complexity of M and of the size of the states in S , where $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ is an appropriately chosen (“best”) monitor for P .

We next follow the three principles above and derive our definition for the complexity of monitoring a safety property P . Before that, let us first define the complexity of monitoring a safety property using a particular monitor for that property, or in other words, let us first define the complexity of a monitor.

During a monitoring session using a monitor, at any moment in time one needs to store at least one state, namely the state that the monitor is currently in. When receiving a new event, the monitor launches its transition function on the current state and the received input. Therefore, the (worst-case) complexity of monitoring with $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ could be defined as

$$\max\{FC(M(s, a)) \mid s \in S, a \in \Sigma\},$$

where $FC(M(s, a))$ is the functional complexity of evaluating M on state s and event a , as a function of the sizes of s and a . In other words, the worst-case monitoring complexity of a particular monitor is the maximal functional complexity that its transition function has on any state and any input; this functional complexity is expressed as a function of the size of the pair (state,event). In order for such a definition to make sense formally, one would need to define or axiomatize the size of monitor states and the size of events.

Definition 19. *Given a monitor $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$, we define the complexity of monitoring \mathcal{M} , written $C_{Mon}(\mathcal{M})$, as the function*

$$FC(M)(\log |S|) : \mathbb{N} \rightarrow \mathbb{N},$$

which is the “uncurried” version applied on $\log |S|$ of the worst-case functional complexity $FC(M) : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ of the partial function M as a function of the size of the pair (state,event) being passed to it.

we assume that the complexity of monitoring a safety property P is the worst-case complexity of monitoring it using some appropriate, “best” monitor for P .

a safety property P would be

$$\min\{\max\{FC(M(s, a)) \mid s \in S, a \in \Sigma\} \mid \mathcal{M} = (S, s_0, M) \in \text{Monitors}(P)\},$$

where $FC(M(s, a))$ is the functional complexity of evaluating M on state s and event a , as a function of the sizes of s and a . In other words, we assume that the complexity of monitoring a safety property P is the worst-case complexity of monitoring it using some appropriate, “best” monitor for P . The worst-case monitoring complexity of a particular monitor is the maximal functional complexity that its transition function has on any state and any input; this functional complexity is expressed as a function of the size of the pair (state,event). Therefore, in order for such a definition to make sense formally, one would need to define or axiomatize the size of monitor states and the size of events.

Since in order to distinguish N elements one needs $\log(N)$ space, we deduce that one needs at least $\log(|S_M|)$ space to store the state of the monitor in its worst-case monitoring scenario (recall that each state in S_M is reachable).

Definition 20. We define the complexity of monitoring a safety property P , written $C_{Mon}(P)$, as the

$$\min\{FC(M) \circ \langle \log(|S|), 1_\Sigma \rangle \mid \mathcal{M} = (S, s_0, M) \in \text{Monitors}(P)\}$$

4.3 Monitoring Safety Properties is Arbitrarily Hard

We show that the problem of monitoring a safety property can be arbitrarily complex. The previous section tells us that there are as many safety properties as real numbers. Therefore, it is not surprising that some of them can be very hard or impossible to monitor. In this section we formalize this intuitive argument. Our approach is to show that we can associate a safety property P_S to any set of natural numbers S , such that monitoring that safety property is as hard as checking membership of arbitrary natural numbers to S . The result then follows from the fact that checking memberships of natural numbers to sets of natural numbers is a problem that can be arbitrarily complex.

Theorem 1 indirectly says that we can associate a persistent safety property to any set of natural numbers (sets of natural numbers are in a bijective correspondence with the real numbers). However, it is not clear how that safety property looks and neither how to monitor it. We next give a more concrete mapping from sets of natural numbers to (persistent) safety properties and show that monitoring the property is equivalent to testing membership to the set. It suffices to assume that Σ contains only two elements, say $\Sigma = \{0, 1\}$.

Definition 21. Let $P_- : \mathcal{P}(\mathbb{N}) \rightarrow \text{PersistentSafety}^*$ be the mapping defined as follows: for any $S \subseteq \mathbb{N}$, let P_S be the set $1^* \cup \{1^k 0 \mid k \in S\} \cdot \{0, 1\}^*$.

It is easy to see that P_S is a persistent safety property over finite traces. Also, it is easy to see that the bijection in the proof of theorem 2 associates to P_S the safety property over infinite traces $1^\omega \cup \{1^k 0 \mid k \in S\} \cdot \{0, 1\}^\omega$.

Let us now investigate the problem of monitoring P_S .

Proposition 14. For any $S \subseteq \mathbb{N}$, monitoring P_S is equivalent to deciding membership of natural numbers to S .

Proof. If M_S is an oracle deciding membership of natural numbers to S , that is, if $M_S(n)$ is true iff $n \in S$, then one can build a monitor for P_S as follows: for a given trace, incrementally read and count the number of prefix 1's; if no 0 is ever observed then monitor indefinitely without reporting any violation; when a first 0 is observed, if any, ask if $M(k)$, where k is the number of 1's observed; if $M(k)$ is false, then report violation; if $M(k)$ is true, then continue monitoring indefinitely and never report violation. It is clear that this is indeed an online monitor for P_S .

Conversely, if we had any monitor for P_S then we could build a decision procedure for membership to S as follows: given $k \in \mathbb{N}$, send to the monitor a sequence of k ones followed by a 0; if the monitor reports violation then deduce

that $k \notin S$; if the monitor does not report violation, then deduce that $k \in S$. It is clear that this is a decision procedure for membership to S .

The proof works for both persistent safety properties over finite traces and for safety properties over infinite traces.

The claim in the title of this section follows now from the fact that the set S of natural numbers can be chosen so that its membership problem is arbitrarily complex. For example, since there are as many subsets of natural numbers as real numbers while there are only as many Turing machines as natural numbers, it follows that there are many (exponentially) more sets of natural numbers that are not recognized by Turing machines than those that are. In particular, there are sets of natural numbers corresponding to any degree in the arithmetic hierarchy, i.e., to predicates $A(k)$ of the form $(Q_1 k_1)(Q_2 k_2) \cdots (Q_n k_n) R(k, k_1, k_2, \dots, k_n)$, where Q_1, Q_2, \dots, Q_n are alternating (universal or existential) quantifiers and R is a recursive/decidable relation: for A such a predicate, let S_A be the set of natural numbers $\{k \mid A(k)\}$. Recall that if Q_1 is \forall then A is called a Π_n property, while if Q_1 is \exists then A is called a Σ_n property. In particular, $\Sigma_0 = \Pi_0$ and they contain precisely the recursive/decidable properties, Σ_1 contains precisely the class of recursively enumerable problems, Π_1 contains precisely the co-recursively enumerable problems, etc.; a standard Π_2 problem is TOTALITY: given $k \in \mathbb{N}$, is it true that Turing machine with Gödel number k terminates on all inputs? Since each level in the arithmetic hierarchy contains problems strictly harder than problems on the previous layer (because $\Sigma_n \cup \Pi_n \subsetneq \Sigma_{n+1} \cap \Pi_{n+1}$), the arithmetic hierarchy gives us a universe of safety properties whose monitoring can be arbitrarily hard.

Within the decidable fragment, as expected, monitoring safety properties can also have any complexity. Indeed, pick for example any NP-complete problem and let S be the set of inputs (coded as natural numbers) for which the problem has a positive answer; then, as explained in the proof of Proposition 14, monitoring P_S against input $1^k 0$ is equivalent to deciding membership of k to S , which is further equivalent to answering the NP-complete problem on input k . Of course, in practice a particular (implementation of a) monitor can be more complex than the corresponding membership problem; for example, monitors corresponding to NP-complete problems are most likely exponential. Also, note that a monitor for P_S needs not necessarily do its complex computation on an input $1^k 0$ when it encounters the 0. It can perform intermediate computations as it reads the prefix 1's and thus pay a lesser computational price when the 0 is encountered. What Proposition 14 says is that the *total* complexity to process the input $1^k 0$ can be no lower than the complexity of checking whether $k \in S$.

5 Monitoring FSM Safety Properties

...

6 Monitoring ERE Safety Needs Non-Elementary Space

Extended regular expressions (EREs) add complementation (\neg) to regular expressions (REs). Complementation can be handy when defining safety properties, because it allows one to say both what should never happen as well as what could happen. In particular, in the context of monitoring EREs, one can switch between the expression of bad prefixes of a safety property and that of good prefixes by just applying a complement operator.

In this section we show that any monitor for safety properties expressed as EREs requires non-elementary space. More precisely, for a given $n \in \mathbb{N}$, we build an ERE of size $O(n^3)$ such that its language is prefix closed and any monitor for its language requires $2^{2^{\cdot 2^n}}$ space, with n nested power of 2 operations.

6.1 Discussion and Relevance of the Lower-Bound Result

Since regular expressions (REs) and deterministic (DFA) and non-deterministic (NFA) finite-state automata are enumerable structures while the set of safety properties is in bijection with the set of real numbers, there are many safety properties that cannot be expressed using REs or automata (or any other formalism whose objects are enumerable). Nevertheless, there are many safety properties of interest that can be expressed as REs or automata. Safety properties over finite or infinite traces can be expressed as REs in at least two different ways:

1. Use an RE to express the language of its bad prefixes; or
2. Use an RE to express the language of its good prefixes.

In the first case, the RE captures the finite-trace behaviors that should never happen, while in the second the RE captures the ones that could possibly happen.

Obviously, not all REs correspond to safety properties in one or the other of the two cases above. For example, the RE $(0 \cdot 1)^+$ cannot express the bad prefixes of a safety property, because “01” is a bad prefix while “010” is not. In order for an RE to express the bad prefixes of a safety property, it should have the property that once w is in its language, then all ww' for any w' should also be in its language. The RE $(0 \cdot 1)^+$ cannot express the good prefixes of a safety property either: “0101” is a good prefix while 010 is not. The language of an RE must be prefix closed in order to express the good prefixes of a safety property.

In both cases above, monitoring the safety property can be done very efficiently (linearly in the size of the RE, both space-wise and time-wise) by first translating its corresponding RE into an NFA, for example using a technique such as Thompson’s [30], and then doing one of the following:

1. In the first case, simulate the NFA-to-DFA construction on the fly as events are received. The state of the monitor is therefore a set of states of the NFA. At start, that set contains only the initial state of the NFA. For each new event, construct the next set by collecting all the NFA states that can be reached via the received event from any of the existing states in the set. If a

final state is reached then report violation of the property: bad prefix found. Since the final states in the NFA symbolize the reach of a bad prefix and since bad prefixes have “no future” in a safety property, the NFA associated to the original RA can be optimized (in case it is not already optimal directly from its construction) by removing any edges out of its final states.

2. In the second case, the monitor works the same way as in the first case, but checking at each time that the monitor state (also a set of NFA states) contains at least one final state of the NFA; if that is not the case, then report violation: prefix found which is not good. If one is willing to pay the exponential price and determinize the NFA of good prefixes, then one can further optimize the resulting DFA (in case it is not already optimized by construction) by collapsing all its non-final states into a “dead-end” state: indeed, the reachability of a non-final state signifies the reachability of a bad prefix, which has “no future”. It is not clear whether or how to optimize the NFA of good prefixes using the additional info that it is a safety property.

Extended regular expressions (EREs) add complementation (\neg) to REs. Meyer and Stockmeyer [29, 28] showed that EREs can express languages non-elementarily more compactly than REs. In other words, for any constant $k \geq 1$, one can find EREs of large enough size $n \in \mathbb{N}$ for which there is no RE having the same language of size less than $2^{2^{\dots^{2^n}}}$, with k nested power operations. Meyer and Stockmeyer [29, 28] showed that several other problems concerning EREs are also non-elementary, including: the equivalence of EREs, the emptiness of the language of an ERE (and implicitly the emptiness of the complement of the language of an ERE), the automata generation (NFA or DFA) from an ERE, etc. Note that it is straightforward to generate potentially non-elementarily large automata from EREs. All one needs to do is to iteratively apply NFA-to-DFA transformations for EREs under complement operators and then complement the resulting DFAs (by complementing their final states). Since each NFA-to-DFA transformation may lead to an exponential explosion on the number of states and since the ERE can have arbitrarily many nested complement operators, the resulting NFA or DFA can be non-elementarily larger than the original ERE.

As already mentioned above, if we allow complementation then we can easily switch from an expression defining the bad prefixes of a safety property to one defining its good prefixes, and backwards, by applying a complement operator. Therefore, from here on, when we say that an ERE expresses a safety property, without any loss of generality we assume that it defines the good prefixes of the safety property; in particular, we assume that its language is prefix closed. Clearly, if one can afford to generate an automaton from the ERE expressing a safety property, then one can and probably should use that automaton as a monitor for the safety property. However, since such an automaton can be enormous compared to the the size of the original ERE, a natural question to ask is whether one can generate monitors for safety properties expressed as EREs that need less than non-elementary space in the size of the original ERE. We next give a negative answer to this question.

Let us first discuss our subsequent lower bound result from a more conceptual perspective. Notice that *synchronous monitoring*, that is, the monitoring process where an error is detected as soon as it appears, is harder than checking for satisfiability (or non-emptiness); indeed, if a formula in a particular formalism is not satisfiable (or it has an empty language), then a synchronous monitor should detect that before the first event is observed. We refer the interested reader to [22] for a discussion on various types of monitoring, including synchronous versus asynchronous monitoring. Synchronous monitors need to either directly (e.g., by accumulating logical constraints while verifying their consistency) or indirectly (e.g., by generating statically an automaton or a structure containing all possible future behaviors) check for satisfiability (or non-emptiness) of the remaining requirements as events are observed online. Unfortunately, this is an expensive process that may be desired to be avoided, at the expense of delaying the detection of violations. For example, the rewriting based monitoring approach for LTL in [22] delays the detection of violations of LTL formulae of the form “(next φ) and (next $\neg\varphi$)” for one more event, to avoid invoking an expensive satisfiability checker for LTL but to instead invoke a propositional satisfiability checker which is less expensive in practice; this is closely related to the notion of “informative prefixes” in [17] that “tell all the story”. Our subsequent lower-bound result states that any monitor for safety properties expressed using EREs, *synchronous or asynchronous*, requires non-elementary space.

Let us now clarify that our lower bound result is not a consequence of the lower-bound result by Meyer and Stockmeyer in [29] (see [28] for a proof of that result). A first reason is that neither the ERE constructed in [28, 29] nor its complement is prefix closed. Indeed, we here focus on a subset of EREs, rather than arbitrary EREs, namely those whose languages are prefix closed, so they express good prefixes of safety properties. Supposing that one could modify the “hard” ERE in [28, 29] whose complement non-emptiness requires non-elementary space into one whose language is prefix-closed and whose size is linear in the size of the original one, the fact that synchronous monitoring of EREs is harder than checking for emptiness does not necessarily imply that monitoring that hard ERE requires non-elementary space. In fact, monitoring that particular ERE requires constant time to process each event, because it is equivalent with an automaton of one state – it takes, however, non-elementary space to compute that one state automaton. What it says is therefore that the *initialization step of ERE-safety synchronous monitoring* requires, in the worst case, non-elementary space.

Interestingly, if one could modify the results in [28, 29] to hold for prefix-closed EREs, including especially the result stating that finding an RE equivalent to an ERE is a non-elementary problem, then one could show that *synchronous monitoring* of safety properties expressed as EREs requires non-elementary space! Indeed, supposing that one had for any ERE a monitor that takes only elementary space in its worst-case monitoring scenario, then one could use that monitor to generate a DFA for the ERE as follows: start with the initial state of the monitor and then discover and store new states of the monitor by feeding arbitrary

(but finite in number) events to each state of the monitor until no new state is discovered. This closure operation takes as much time and space as the number of states the monitor reach; since by assumption the monitor needs “only” elementary space to store its state in any scenario, we deduce that the obtained automaton has size elementary in the size of the ERE (and it also takes elementary time and space to generate it). In other words, we could find an elementary algorithm to associate to any (prefix closed) ERE an equivalent RE, contradicting the non-elementary lower bound in [28, 29] (again, supposing that the lower-bound results in [28, 29] could be modified for prefix-closed EREs).

Unfortunately, it is not that clear how to reduce the non-elementary problems in [28, 29] to *asynchronous monitoring*, and thus to conclude that asynchronous monitoring also requires non-elementary space. That is because a “smart” asynchronous monitor may in principle collapse states (when regarded as an automaton as in the construction above) in rather unexpected ways, just because it “knows” that eventually an error may be reported anyway if observation continues indefinitely from that state on; in other words, states with the property that “eventually violation detected in the future” may be collapsed as equivalent by an asynchronous monitor. This way, the DFA extracted from a monitor for a safety property expressed as an ERE may be significantly smaller than the DFA corresponding to the ERE (and obviously, it may have a different language).

Our next result shows that asynchronous monitoring of ERE-safety also requires non-elementary space, which is a more general lower-bound result than the space non-elementarity of synchronous monitoring. Moreover, it gives an alternative proof of the non-elementary lower-bounds by Stockmeyer and Meyer [28, 29], because automata corresponding to safety-defining EREs are just special cases of monitors, so they must also take non-elementary space. Moreover, we improve the results in [28, 29] by showing that their lower-bounds also hold for a *subset of EREs*, namely those corresponding to safety properties.

Summarizing the discussion above, we believe that the main contributions of our subsequent lower-bound result are the following:

1. We show that asynchronous monitoring already requires non-elementary space, same as synchronous monitoring. For example, an ERE monitoring algorithm was presented by Roşu and Viswanathan in [23], which “rewrites” or “derives” the ERE by each letter in the input word; the derivation process consists of some straightforward rewrite rules, some for expanding the ERE others for contracting it via simplifications. No comprehensive and expensive check for emptiness on the resulting ERE is performed, except what is done by the simplification rules (for example $\emptyset \cdot R \rightarrow \emptyset$ and $\epsilon \cdot R \rightarrow R$, etc.). A check for emptiness can and should be eventually performed (for example, a check for emptiness can be done periodically, say every 10^x events for some convenient x). This gives us an asynchronous ERE monitoring algorithm, which, unlike the simplistic NFA/DFA generation algorithm, does not pay upfront the potentially non-elementary worst-case penalty! However, our subsequent lower bound result tells that there is a worst-case scenario in which one cannot avoid the non-elementary space required to store the continuously

changing ERE if one wants to correctly eventually detect violations of the original ERE. And that is the case for any synchronous or asynchronous monitoring algorithm for safety properties expressed as EREs.

2. We propose a different technique to prove non-elementary lower-bounds, fundamentally different from the one in [28]. The technique in [28] is based on diagonalization arguments and encodings of accepting Turing machine computations as finite trace words. Our technique is inspired from an idea by Chandra, Kozen and Stockmeyer [6] introduced to show the power of alternation and then used by several authors to prove exponential lower bounds [15, 16, 23, 22]. At our knowledge, the use of such a technique to show non-elementary lower bounds is novel. The idea of the technique in [6] is to define expressions having as languages $L_n = \{w^{(1)}\#w^{(2)}\#\dots\#w^{(k)}\$w \mid w^{(1)}, w^{(2)}, \dots, w^{(k)}, w \in \{0, 1\}^n, (\exists 1 \leq i \leq k) w^{(i)} = w\}$. Our idea is to define, using EREs, words of the form $X_n\$X'_n$, where X_n and X'_n are *n-deep nested sets* starting with elements in $\{0, 1\}^n$ (i.e., sets of sets of ... of sets of elements in $\{0, 1\}^n$, with n power set operations), such that X'_n is *n-nested included* in X_n , where *i-nested inclusion* is standard inclusion when $i = 1$ and, if $i > 1$, then it is defined inductively as: X'_i is *i-nested included* in X_i iff for each $X'_{i-1} \in X'_i$, there is some $X_{i-1} \in X_i$ such that X'_{i-1} is $(i - 1)$ -nested included in X_{i-1} .

One more observation is in place before we move on to the technical details. It is known that the *membership problem* for EREs, testing whether a word w of size m is in the language of an ERE of size n , is polynomial in m and n . For example, the classic algorithm by Hopcroft and Ullman [13] runs in space $O(m^2 \cdot n)$ and time $O(m^3 \cdot n)$; slightly improved algorithms have been proposed by several authors [12, 31–33, 18, 14], reducing space requirements to $O(m^2 \cdot k + m \cdot n)$ and time to $O(m^3 \cdot k + m^2 \cdot n)$ or worse, where k is the number of complement operators in the ERE; a recent ERE membership algorithm was proposed by the author in [25], which runs in space $O(m \cdot \log m \cdot 2^n)$ and time $O(m^2 \cdot \log m \cdot 2^n)$ when $m > 2^n$. These algorithms appear to be efficient, because they are polynomial or simply exponential in the ERE, so one may think that one may devise an elementary ERE-safety monitoring algorithm as follows: store the trace of events and at each newly received event invoke one of these “efficient” ERE membership algorithms. While this algorithm will indeed be elementary in the size of the ERE *and* the size of the trace, our lower bound result says that it will, in fact, be *non-elementary in the size of only the ERE!* In other words, for a carefully chosen “hard” ERE of size n , there are finite traces of large enough size m so that checking their membership is a problem which is non-elementary in n ; this will indeed happen when m is non-elementarily larger than n .

6.2 The Lower-Bound Result

EREs define languages by inductively applying *union* ($+$), *concatenation* (\cdot), *Kleene Closure* (\star), *intersection* (\cap), and *complementation* (\neg). The language of an ERE R , say $\mathcal{L}(R)$, is defined inductively as follows, where $s \in \Sigma$:

- $\mathcal{L}(\emptyset) = \emptyset$,

- $\mathcal{L}(\epsilon) = \{\epsilon\}$,
- $\mathcal{L}(s) = \{s\}$,
- $\mathcal{L}(R_1 + R_2) = \mathcal{L}(R_1) \cup \mathcal{L}(R_2)$,
- $\mathcal{L}(R_1 \cdot R_2) = \mathcal{L}(R_1) \cdot \mathcal{L}(R_2)$,
- $\mathcal{L}(R^*) = (\mathcal{L}(R))^*$,
- $\mathcal{L}(R_1 \cap R_2) = \mathcal{L}(R_1) \cap \mathcal{L}(R_2)$,
- $\mathcal{L}(\neg R) = \neg \mathcal{L}(R)$.

If R does not contain \neg and \cap then it is a *regular expression* (RE). By applying De Morgan's law $R_1 \cap R_2 \equiv \neg(\neg R_1 + \neg R_2)$, EREs can be linearly (in both time and size) translated into equivalent EREs without intersection; therefore, intersection of EREs is just syntactic sugar. The *size* of an ERE is the total number of occurrences of letters and composition operators ($+$, \cdot , \star , and \neg) that it contains. In what follows we assume that Σ is finite. For notational simplicity, in what follows we let Σ also denote the RE $s_1 + s_2 + \dots + s_n$ where $\Sigma = \{s_1, s_2, \dots, s_n\}$ and let Σ^* denote both the language $\{s_1, s_2, \dots, s_n\}^*$ and the RE $(s_1 + s_2 + \dots + s_n)^*$.

For $n \in \mathbb{N}$, let us define inductively the following alphabets and languages:

- $\Sigma_0 = \{0, 1\}$ and $\Xi_0 = \{0, 1\}^n$, and
- $\Sigma_i = \Sigma_{i-1} \cup \{\#_i\}$ and $\Xi_i = \{\#_i \#_i\} \cup (\{\#_i\} \cdot \Xi_{i-1})^+ \cdot \{\#_i\}$, for all $1 \leq i \leq n$.

In the above, $\#_i$ are n fresh letters. The intuition for the languages Ξ_i above is to encode nested sets of depth i that contain sets of words of n bits at their deepest level. The symbols $\#_i$ play the role of markers separating the elements of such sets. For example, the word $\#_2 \#_1 \#_1 \#_2 \#_1 01 \#_1 10 \#_1 \#_2 \#_1 \#_1 \#_2$ encodes $\{\{\}, \{01, 10\}, \{\}\}$, that is, the set $\{\{\}, \{01, 10\}\}$; since the multiplicity and order of elements in sets are irrelevant, the same set can have (infinitely) many different encodings. Formally, let us define the following *set* functions associating to encodings in Ξ_i corresponding nested sets:

- $set_0: \Xi_0 \rightarrow \{0, 1\}^n$ is the identity function on Ξ_0 , i.e., $set_0(w) = w$;
- $set_i: \Xi_i \rightarrow \mathcal{P}^i(\{0, 1\}^n)$, where \mathcal{P}^i is the power set operator applied i times, $set_i(\#_i \#_i) = \{\}$, $set_i(\#_i X_{i-1} \#_i) = \{set_{i-1}(X_{i-1})\}$, and $set_i(\#_i X_{i-1} X_i) = \{set_{i-1}(X_{i-1})\} \cup set_i(X_i)$, for all $1 \leq i \leq n$, $X_{i-1} \in \Xi_{i-1}$, and $X_i \in \Xi_i$.

Note that $|set_0(\Xi_0)| = 2^n$ and $set_i(\Xi_i) = \mathcal{P}(set_{i-1}(\Xi_{i-1}))$ for all $1 \leq i \leq n$; therefore, $|set_i(\Xi_i)| = 2^{2^{\cdot^{2^n}}}$ for all $1 \leq i \leq n$, with $i+1$ nested power operations.

Let us define *nested-inclusion* relations $_ \underline{\Xi}_i _ : \mathcal{P}^i(\{0, 1\}^n) \times \mathcal{P}^i(\{0, 1\}^n)$ for $0 \leq i \leq n$ and *nested-membership* relations $_ \in_i _ : \mathcal{P}^{i-1}(\{0, 1\}^n) \times \mathcal{P}^i(\{0, 1\}^n)$ for $1 \leq i \leq n$ as follows:

- $_ \underline{\Xi}_0 _$ is the identity on $\{0, 1\}^n$ and $_ \underline{\Xi}_1 _$ is $_ \subseteq _ : \mathcal{P}(\{0, 1\}^n) \times \mathcal{P}(\{0, 1\}^n)$,
- $_ \in_1 _$ is $_ \in _ : \{0, 1\}^n \times \mathcal{P}(\{0, 1\}^n)$, and for $1 < i$,
- $S_{i-1} \in S_i$ iff there is some $S'_{i-1} \in S_i$ such that $S_{i-1} \underline{\Xi}_i S'_{i-1}$, and
- $S_i \underline{\Xi}_i S'_i$ iff $S_{i-1} \in S'_i$ for each $S_{i-1} \in S_i$.

For example, if $n = 2$ then $\{\{00, 01\}, \{01, 10\}, \{11\}\} \subseteq_2 \{\{00, 01, 10\}, \{00, 11\}\}$ because $\{00, 01\}, \{01, 10\} \subseteq_1 \{00, 01, 10\}$ and $\{11\} \subseteq_1 \{00, 11\}$.

We can now define Σ as $\Sigma_n \cup \{\$\}$ and the infinite trace property P_n^ω over Σ :

$$(\epsilon \cup (\Sigma_n^* \cup \{X_n \$ X'_n \mid X_n, X'_n \in \Xi_n \text{ and } \text{set}_n(X'_n) \subseteq_n \text{set}_n(X_n)\}) \cdot \{\$\}) \cdot \Sigma_n^\omega.$$

An infinite trace in P_n^ω therefore contains at most two \$ letters and infinitely many letters in Σ_n . There are no restrictions on the appearance of the letters in Σ_n when there is no \$ letter or when there is precisely one \$ letter. However, if the infinite trace contains precisely two \$ letters, that is, if it has the form $w \$ w' \$ u$ for some $w, w' \in \Sigma_n^*$ and some $u \in \Sigma_n^\omega$, then w and w' must be in Ξ_n and the nested set corresponding to w must nested-include the nested set corresponding to w' ; there are no restrictions on u .

Proposition 15. $P_n^\omega \in \text{Safety}^\omega$.

Proof. There are two cases to analyze for an infinite trace that is not in P_n^ω : when it contains more than two \$ letter, or when it has the form $w \$ w' \$ u$ with $w, w' \in \Sigma_n^*$ and $u \in \Sigma_n^\omega$, but it is not the case that $w, w' \in \Xi_n$ and $\text{set}_n(w') \subseteq_n \text{set}_n(w)$. In the first case, we can pick the first prefix of the infinite trace containing three \$ letters in total; clearly, this finite trace prefix cannot be continued into any acceptable infinite trace. In the second case, since there are no restrictions on the letters in u , we can easily see that the prefix $w \$ w' \$$ is already a violation threshold: there is no $u' \in \Sigma_n^\omega$ such that $w \$ w' \$ u' \in P_n^\omega$.

The bijection in the proof of Theorem 2 associates to each infinite-trace safety property a persistent finite-trace safety property by taking its prefixes. Let P_n be the persistent finite-trace safety property $\text{prefixes}(P_n^\omega)$ corresponding to P_n^ω . It is easy to see that P_n is the property

$$\Sigma_n^* \cup \Sigma_n^* \cdot \{\$\} \cdot \Sigma_n^* \cup \{X_n \$ X'_n \mid X_n, X'_n \in \Xi_n \text{ and } \text{set}_n(X'_n) \subseteq_n \text{set}_n(X_n)\} \cdot \{\$\} \cdot \Sigma_n^*.$$

Note that monitoring P_n^ω is the same as monitoring P_n : in both cases, besides the capability to checking whether there are more than two \$ letters, which is trivial, the monitor needs to store sufficient information about the nested set corresponding to X_n , so that, when the first \$ is seen, to be able to check whether it nested-includes the set corresponding to the upcoming, yet unknown X'_n .

Theorem 6. *Any synchronous or asynchronous monitor for P_n or P_n^ω needs space non-elementary in n , namely $\Omega(2^{2^{\dots^{2^n}}})$, with n nested power operations.*

Proof. Suppose that M is a monitor for P_n or P_n^ω and suppose that, during a monitoring session, it reads the prefix $X_n \in \Xi_n$. Regardless of how M is defined or implemented, in particular regardless of whether it reports violations synchronously or asynchronously, when the first \$ event is encountered, the state of M must contain enough information to sooner or later be able to decide whether the set $\text{set}_n(X'_n)$ corresponding to *any* upcoming (unknown at the time

the \$ is observed) sequence X'_n is nested-included in $set_n(X_n)$. Since $set(X'_n)$ can in particular be equal to $set_n(X_n)$, and since once the second \$ event is observed there is no further event that may bring new knowledge to the monitor, we deduce that M must be able to distinguish any two different sets in $set_n(\Xi_n)$ when the first \$ event is encountered, that is, M 's states must be different after reading words in Ξ_n whose corresponding nested sets are different. Therefore, M must be able to distinguish $|set_n(\Xi_n)|$ different possibilities. Since one needs $\Omega(\log N)$ space to distinguish among N different situations (one label for each), we conclude that M needs space $\Omega(\log(|set_n(\Xi_n)|))$, that is, $\Omega(2^{2^n})$ with n nested power operations. Hence, any monitor for P_n needs non-elementarily large space in n .

We next show how to construct an ERE polynomial in size with n whose language is precisely P_n .

Theorem 7. *There is an ERE of size $O(n^3)$ whose language is P_n .*

Proof. Note that P_n is a union of three languages, the first two being trivial to express as languages of corresponding REs. As expected, the difficult part is to associate an ERE to the language

$$\{X_n \$ X'_n \mid X_n, X'_n \in \Xi_n \text{ and } set_n(X'_n) \subseteq_n set_n(X_n)\}.$$

Note that so far we did not need complementation. The property above can, however, be expressed as an ERE of size $O(n^3)$ using $O(n)$ nested complement operators. The idea is to define iteratively a sequence of EREs \mathbb{K}_i for $0 \leq i \leq n$ whose languages contain words $X_i w \$ w' X'_i$ with $set_i(X'_i) \subseteq_i set(X_i)$, which are contiguous fragments of desired words $X_n \$ X'_n$. Then \mathbb{K}_n would be the language that we are looking for. To define \mathbb{K}_i , we observe that the nested-inclusion $S'_i \subseteq_i S_i$ is equivalent to: there is no $S'_{i-1} \in S'_i$ such that it is not the case that we can find some $S_{i-1} \in S_i$ such that $S'_{i-1} \subseteq_{i-1} S_{i-1}$. This crucial observation will allow us to define \mathbb{K}_i in terms of \mathbb{K}_{i-1} . We next develop the technical details.

Let us first define regular patterns corresponding to each of the languages Ξ_i for $0 \leq i \leq n$; to avoid introducing new names, we ambiguously let the corresponding regular expressions have the same names as their languages:

- Let $\Xi_0 = (0 + 1)^n$, where for an RE, r^n is $r \cdot r \cdot \dots \cdot r$ (n times); and
- Let $\Xi_i = \#_i \cdot \#_i + (\#_i \cdot \Xi_{i-1})^+ \cdot \#_i$ for all $1 \leq i \leq n$.

Iteratively eliminating the Ξ_{i-1} regular expressions from the right-hand-sides, we eventually obtain $n + 1$ regular patters, each of size $O(n)$ (the size of Ξ_0 as a regular expression is $O(n)$ and each Ξ_i adds a constant size to that of Ξ_{i-1}).

Next we define REs for the languages $\text{prefixes}(\Xi_i)$ and $\text{suffixes}(\Xi_i)$ of prefixes and respectively suffixes of words in Ξ_i , for all $0 \leq i \leq n$. We only discuss the prefix closure languages; the suffix closures are dual. The prefix closures can be defined relatively easily inductively as follows:

- $\text{prefixes}(\Xi_0) = \bigcup_{k=0}^n \{0, 1\}^k$, and

$$\begin{aligned}
 - \text{prefixes}(\Xi_i) &= \{\epsilon, \#_i \#_i\} \cup \{\#_i\} \cdot \text{prefixes}(\Xi_{i-1}) \cup (\{\#_i\} \cdot \Xi_{i-1})^+ \cup (\{\#_i\} \cdot \\
 &\quad \Xi_{i-1})^+ \cdot \{\#_i\} \cdot \text{prefixes}(\Xi_{i-1}) \\
 &= \{\#_i \#_i\} \cup (\{\#_i\} \cdot \Xi_{i-1})^* \cdot (\{\epsilon\} \cup \{\#_i\} \cdot \text{prefixes}(\Xi_{i-1})).
 \end{aligned}$$

These languages can be expressed with the following REs; as before, we ambiguously use the same names for the corresponding REs:

$$\begin{aligned}
 - \text{Let } \text{prefixes}(\Xi_0) &= \epsilon + (0 + 1) + (0 + 1)^2 + \dots + (0 + 1)^n = (\epsilon + 0 + 1)^n; \text{ and} \\
 - \text{Let } \text{prefixes}(\Xi_i) &= \#_i \cdot \#_i + (\#_i \cdot \Xi_{i-1})^* \cdot (\epsilon + \#_i \cdot \text{prefixes}(\Xi_{i-1})).
 \end{aligned}$$

Iteratively eliminating the REs $\text{prefixes}(\Xi_{i-1})$ from the right-hand-sides, we eventually obtain $n + 1$ REs, each of size $O(i^2 + n)$ (the size of $\text{prefixes}(\Xi_0)$ as an RE is $O(n)$ and each $\text{prefixes}(\Xi_i)$ adds size $O(i)$ to that of $\text{prefixes}(\Xi_{i-1})$). Dually,

$$\begin{aligned}
 - \text{Let } \text{suffixes}(\Xi_0) &= \epsilon + (0 + 1) + (0 + 1)^2 + \dots + (0 + 1)^n = (\epsilon + 0 + 1)^n; \text{ and} \\
 - \text{Let } \text{suffixes}(\Xi_i) &= \#_i \cdot \#_i + (\epsilon + \text{suffixes}(\Xi_{i-1}) \cdot \#_i) \cdot (\Xi_{i-1} \cdot \#_i)^*.
 \end{aligned}$$

We next define REs L_i and R_i for $0 \leq i \leq n$ whose languages contain the contiguous fragments of words in Ξ_n that are allowed to appear to the left and to the right of $\$,$ respectively, so that words in $\mathcal{L}(L_i)$ start with $\#_i$ and words in $\mathcal{L}(R_i)$ end with $\#_i$. Let us also assume by convention that $L_{n+1} = R_{n+1} = \epsilon$ (the RE whose language contains only the empty word). It is easy to see that L_i and R_i can be defined as follows:

$$\begin{aligned}
 - \text{Let } L_i &= \#_i \cdot \Sigma_n^* \cap \text{suffixes}(\Xi_n), \text{ and} \\
 - \text{Let } R_i &= \Sigma_n^* \cdot \#_i \cap \text{prefixes}(\Xi_n).
 \end{aligned}$$

Note that words in L_i and R_i may not necessarily start or end with a word in Ξ_i : indeed, the $\#_i$ that may start or end L_i or R_i could very well be followed or preceded, respectively, by a $\#_{i+1}$ or, if $i = n$, by $\$$.

Let us also define the EREs \bar{L}_i and \bar{R}_i whose languages are included in those of L_i and R_i , respectively, and whose words start or end with words in Ξ_i :

$$\begin{aligned}
 - \text{Let } \bar{L}_i &= \Xi_i \cdot \Sigma_n^* \cap \text{suffixes}(\Xi_n), \text{ and} \\
 - \text{Let } \bar{R}_i &= \Sigma_n^* \cdot \Xi_i \cap \text{prefixes}(\Xi_n).
 \end{aligned}$$

It is not difficult to see that $\bar{L}_i = \Xi_i \cdot L_{i+1}$ and $\bar{R}_i = R_{i+1} \cdot \Xi_i$. Note that the sizes of L_i, R_i, \bar{L}_i and \bar{R}_i are $O(n^2)$.

Let us now define the EREs \mathbb{K}_i for $0 \leq i \leq n$ as follows:

$$\begin{aligned}
 - \mathbb{K}_0 &= \bar{L}_0 \cdot \$ \cdot \bar{R}_0 \cap \bigcap_{k=0}^{n-1} (\Sigma_0^k \cdot 0 \cdot \Sigma^* \cdot 0 \cdot \Sigma_0^{n-k-1} + \Sigma_0^k \cdot 1 \cdot \Sigma^* \cdot 1 \cdot \Sigma_0^{n-k-1}), \text{ and} \\
 - \mathbb{K}_i &= \bar{L}_i \cdot \$ \cdot \bar{R}_i \cap \neg((\neg(\#_i \cdot \Xi_{i-1})^* \cdot \#_i \cdot \mathbb{K}_{i-1}) \cap L_i \cdot \$ \cdot R_{i-1}) \cdot (\#_i \cdot \Xi_{i-1})^* \cdot \#_i.
 \end{aligned}$$

We next show by induction on i that $\mathcal{L}(\mathbb{K}_i)$ is the language

$$\{X_i w X'_i \mid X_i, X'_i \in \Xi_i, w \in \mathcal{L}(L_{i+1} \cdot \$ \cdot R_{i+1}), \text{set}_i(X'_i) \subseteq_i \text{set}_i(X_i)\}.$$

It is easy to see that $\mathcal{L}(\mathbb{K}_0) = \{X_0 w X_0 \mid X_0 \in \Xi_0, w \in \mathcal{L}(L_1 \cdot \$ \cdot R_1)\}$, because the large conjunct in \mathbb{K}_0 states that the words formed with the first n letters and with the last ones, respectively, are equal and in Ξ_0 , and because $\bar{L}_0 \cdot \$ \cdot \bar{R}_0 =$

8 Optimal Monitoring of “Always Past” Temporal Safety

A monitor synthesis algorithm from linear temporal logic (LTL) safety formulae of the form $\Box\varphi$ where φ is a past time LTL formula was presented in [11]. The generated monitors implemented the recursive semantics of past-time LTL using a dynamic programming technique, and needed $O(|\varphi|)$ time to process each new event and $O(|\varphi|)$ total space. Some compiler-like optimizations of the generated monitors were also proposed in [11], which would further reduce the required space. It is not clear how much the required space could be reduced by applying those optimizations.

We here show how to generate using a divide-and-conquer technique directly monitors that need $O(k)$ space and still $O(|\varphi|)$ time, where k is the number of temporal operators in φ .

8.1 The Monitor Synthesis Algorithm

For simplicity, we assume only two past operators, namely \circ (previously) and \mathcal{S} (since). Let us first note that one cannot asymptotically reduce the space requirements below $\Omega(k)$, where k is the number of temporal operators appearing in the formula to monitor φ . Indeed, one can take $\varphi = (\#_1 \rightarrow t_1) \wedge \cdots \wedge (\#_k \rightarrow t_k)$, where for each $1 \leq i \leq k$, $\#_i$ is some event and t_i is some temporal formula containing precisely one past temporal operator, i.e., a \circ or a \mathcal{S} . Any monitor for φ must directly or indirectly store the status of each t_i at every event, to be able to react accordingly in case the next event is some $\#_i$. Assuming that the events $\#_i$ are distinct and that the formulae t_i are unrelated, then the monitor needs to distinguish among 2^k possible states, so it needs $\Omega(k)$ space.

In what follows, we assume the usual recursive semantics of LTL, also presented below, restricted to safety formulae of the form $\Box\varphi$, where φ is a past-time LTL. We adopt the simplifying assumption that the empty trace invalidates any atomic proposition and any past temporal operator; as argued in [11], this may not always be the best choice, but other semantic variations regarding the empty trace present no difficulties for monitoring.

Definition 22. (adapted from [21]) *LTL formulae of the form $\Box\varphi$ (read “always φ ”), where φ is a past-time LTL formula, are called LTL safety formulae; we may call them just safety formulae when LTL is understood from the context. An infinite trace $u \in \Sigma^\omega$ satisfies $\Box\varphi$, written $u \models \Box\varphi$, iff each $w \in \text{prefixes}(u)$ satisfies the past-time LTL formula φ , written also $w \models \varphi$ and defined inductively as follows:*

$$\begin{array}{ll}
 w \models \text{true} & \text{is always true,} \\
 ws \models a & \text{iff } a(s) \text{ holds,} \\
 w \models \neg\varphi & \text{iff } w \not\models \varphi, \\
 w \models \varphi_1 \wedge \varphi_2 & \text{iff } w \models \varphi_1 \text{ and } w \models \varphi_2, \\
 ws \models \circ\varphi & \text{iff } w \models \varphi, \\
 ws \models \varphi_1 \mathcal{S} \varphi_2 & \text{iff } ws \models \varphi_2 \text{ or } ws \models F_\varphi \text{ and } w \models \varphi_1 \mathcal{S} \varphi_2 \\
 \epsilon \models \varphi & \text{is false otherwise}
 \end{array}$$

Given safety formula $\Box\varphi$, we let $\mathcal{L}(\Box\varphi) \subseteq \Sigma^\omega$ be the set $\{u \in \Sigma^\omega \mid u \models \Box\varphi\}$.

Proposition 16. $\mathcal{L}(\Box\varphi) \in \text{Safety}^\omega$ for any past-time LTL formula φ .

Proof. By the definition of $\mathcal{L}(\Box\varphi)$ in Definition 22 and the definition of $\Box P$ in Definition 10, one can easily note that $\mathcal{L}(\Box\varphi) = \Box\mathcal{L}(\varphi)$, where $\mathcal{L}(\varphi) = \{w \in \Sigma^* \mid w \models \varphi\}$. Therefore, $\mathcal{L}(\Box\varphi) \in \text{Safety}_\Box^\omega$. The rest follows by Theorem 4.

Let us next investigate the problem of monitoring safety properties $P \in \text{Safety}^\omega$ expressed as languages of safety formulae, that is, $P = \mathcal{L}(\Box\varphi)$ for some past-time LTL formula φ . Because of the recursive nature of the satisfaction relation, a first important observation is that the generated monitor only needs to store information regarding the status of temporal operators from the previous state. More precisely, the monitor needs one bit per temporal operator, keeping the satisfaction status of the subformula corresponding to that temporal operator; when a new state is received, the satisfaction status of the subformula is recalculated according to the recursive semantics above and then the bit is updated. The order in which the temporal operators are processed when a new state is received is important: the nested operators must be processed first.

We next present the actual monitor synthesis algorithm at a high-level. We refrain from giving detailed pseudocode as we did in [11], because different applications may choose different implementation paradigms. For example, we are currently using rewriting techniques to implement the monitor synthesis algorithms in MOP [7]; Section 8.2 shows our complete Maude rewriting implementation of the subsequent monitor synthesis algorithm.

Step 1 Let $\varphi_1, \dots, \varphi_k$ be the k subformulae of φ corresponding to temporal operators, such that, if φ_i is a subformula of φ_j , then $i < j$; this can be easily achieved by a DFS traversal of φ .

Step 2 Let $bit[1..k]$ be a vector of k bits initialized to 0 (or false); $bit[i]$ will store information related to φ_i from the previous state:

- if $\varphi_i = \circ\psi$ then $bit[i]$ says if ψ was satisfied at the previous state;
- if $\varphi_i = \psi \mathcal{S} \psi'$ then $bit[i]$ says if φ_i was satisfied at the previous step.

Step 3 Let $bit'[1..k]$ be another vector of k bits; this will be used to store temporary results, which will be moved eventually into the vector $bit[1..k]$.

Step 4 Generate a loop that executes whenever a new state s is available; the body of the loop executes the following code:

Step 4.1 For each i from 1 to k execute a bit assignment as follows, where for a subformula ψ of φ , $\overline{\psi}$ is the boolean expression replacing in ψ each non-nested temporal subformula φ_j by $bit[j]$ if φ_j is a “previously” formula or by $bit'[j]$ if φ_j is a “since” formula, and each remaining atomic proposition a by its satisfaction in the current state, $a(s)$:

- if $\varphi_i = \circ\psi$ then generate the assignment $bit'[i] := \overline{\psi}$
- if $\varphi_i = \psi \mathcal{S} \psi'$ then generate the assignment $bit'[i] := \overline{\psi'} \vee \overline{\psi} \wedge bit[i]$

Step 4.2 Generate the conditional: if $\overline{\varphi}$ is false then error (formula violated)

Step 4.3 Generate code to move the contents of $bit'[1..k]$ into $bit[1..k]$.

Note that the generated monitors are well-defined, because each time a $\overline{\psi}$ boolean expression is generated, all the bits in $bit'[1..k]$ that are needed are already calculated. One can also perform boolean simplifications when calculating $\overline{\psi}$ to reduce runtime overhead even further. For example, in our implementation that also generated the code below (see Section 8.2), we used the simplification $\neg\neg\psi = \psi$. To illustrate the monitor generation algorithm above, let us consider the past time LTL formula: $\varphi = \neg(a \wedge \neg(\circ b \wedge (c \mathcal{S} (d \wedge (\neg e \mathcal{S} f))))))$. Step 1 produces the following enumeration of φ 's subformulae: $\varphi_1 = \circ b$, $\varphi_2 = \neg e \mathcal{S} f$, and $\varphi_3 = c \mathcal{S} (d \wedge (\neg e \mathcal{S} f))$. The other steps eventually generate the code:

```

bit[1..3] := false;      // three global bits
foreach new state s do {
    // first update the bits in a consistent order
    bit'[1] := b(s);
    bit'[2] := f(s) ∨ (¬e(s) ∧ bit[2]);
    bit'[3] := d(s) ∧ bit'[2] ∨ (c(s) ∧ bit[3]);
    // then check whether the formula is violated
    if a(s) ∧ ¬(bit[1] ∧ bit'[3]) then Error;
    // finally, update the state of the monitor
    bit[1..3] := bit'[1..3]
}

```

It is easy to see that for any past LTL formula φ of k temporal operators, the state of the generated monitor is encoded on k bits, namely the vector $bit[1..k]$. The runtime of the generated monitor is still $O(|\varphi|)$, because each temporal operator in φ results in an assignment and a read operation in the monitor, while each boolean operator in φ is “executed” by the monitor.

8.2 A Maude Implementation of the Monitor Synthesizer

We here show a term rewriting implementation of the algorithm above, using the Maude system [8]. Implementations in other languages are obviously also possible; however, rewriting proved to be an elegant means to generate monitors from logical formulae in several other contexts, and so seems to be here. In what follows we show the complete Maude code that takes as input a formula, parses it, generates the monitor, and then pretty prints it. We use the **K** technique here [24], which is a rewriting-based language and/or logic definitional technique; to use **K**, one needs to first upload the generic, i.e., application-independent, module discussed at the end of this section.

Atomic Predicates We start by defining the atomic state predicates that one can use in formulae. These can be either identifiers (of the form `'a`, `'abc`, `'a123`, etc.; these are provided by the Maude builtin module **QID**):

```

fmod PREDICATE is
--- atomic predicates can be quoted identifiers

```

```
protecting QID .
sort Predicate .
subsort Qid < Predicate .
endfm
```

Syntax of Formulae Let us next define the syntax of formulae. We here use Maude’s mixfix notation for defining syntax as algebraic operators, where underscores stay for arguments. Also, note that operators are assigned precedences (declared as operator attributes), to relieve the user from writing parentheses (the lower the precedence the tighter the binding):

```
fmod SYNTAX is
protecting PREDICATE .
sort Formula .
subsort Predicate < Formula .
op !_ : Formula -> Formula [prec 20] .
op _/\_ : Formula Formula -> Formula [prec 23] .
op 0_ : Formula -> Formula [prec 21] .
op _S_ : Formula Formula -> Formula [prec 22] .
endfm
```

Target Language We are done with the input language. Let us now define the output language. We need a very simple language for implementing the generated monitors, namely one with limited assignment, conditional and looping. The generated code, as well as the target language, play no role in this paper; one is expected to change the language below to one’s desired target language (Java, C, C#, assembler, etc.). Our chosen language below has bits, expressions, statements and code. Bits are also expressions; code is a list of statements composed sequentially using “;” or just concatenation. The syntax below is also making use of precedence attributes. The `format` attributes are useful solely for pretty-printing reasons (see Maude’s manual [8] for details on formatting):

```
fmod CODE is
--- syntax for the generated code
protecting PREDICATE + INT + STRING .
sorts Bit Exp Statement Code .
subsorts Bit < Exp .
subsort Statement < Code .
ops (bit[_]) (bit'[_]) : Nat -> Bit .
ops (bit[1 .. _]) (bit'[1 .. _]) : Int -> Bit .
op _(s) : Predicate -> Exp [prec 0] .
ops true false : -> Exp .
op !_ : Exp -> Exp [prec 20] .
op _/\_ : Exp Exp -> Exp [prec 23] .
op _\/_ : Exp Exp -> Exp [prec 24] .
op _:=_ : Exp Exp -> Statement [prec 27 format(ni d d d)] .
op if_then_ : Exp Statement -> Statement
[format(ni d d ++ --) prec 30] .
```

```

op foreach new state s do _ : Code -> Statement
    [format(n d d d s++ --n)] .
op Error : -> Statement [format(ni d)] .
op //_ : String -> Statement [format(ni d d)] .
op nil : -> Code .
op _;_ : Code Code -> Code [assoc id: nil prec 40] .
op __ : Code Code -> Code [assoc id: nil prec 40] .
op {_} : Code -> Statement [format(d d --ni ++)] .
--- code simplification rules
var B : Exp .
eq !! B = B .
endfm

```

The following module defines the actual monitor synthesis algorithm. We use the K definitional technique here, because it yields a very compact implementation. K is centered on the basic intuition of *computation*; computations are encoded as first-order data-structures that “evolve”, via rewriting, to *results*. Computations are sequentialized using the list constructor “ $_ \rightarrow _$ ”; thus, if K and K' are computations, then $K \rightarrow K'$ is the computation consisting of K followed by K' . Computations may eventually yield results; for example, $K \rightarrow K'$ may rewrite (in context) to $R \rightarrow K'$, meaning that R is the result that K reduces to. An important feature of K is that one can schedule lists of tasks for reduction; for example, $[K1, K2, K3] \rightarrow K$ may eventually reduce to $[R1, R2, R3] \rightarrow K$, where $R1$, $R2$, and $R3$ are the results that $K1$, $K2$, and $K3$ reduce to, in this order. To use K , one needs to import the module K discussed at the end of this section. The equations of the module K (three in total) are all about reducing a list of computations to a list of results, supposing that one knows how to reduce one computation to one result.

K is a definitional framework that is generic in computations and results. More precisely, it provides sorts `KComputation` and `KResult`, and expects its user to define the desired computations and results, as well as rules to reduce a computation to a result. Computations typically can be reduced to results only in context; to facilitate this, K provides a sort `KConfiguration`, which is also supposed to be populated accordingly. The sort `KConfiguration` is a multi-set sort over a sort `KConfigurationItem`, where the multi-set constructor is just concatenation; also, the sort `KComputation` is a list sort over `KComputationItem`, where the list constructor is $_ \rightarrow _$. To make use of K , one needs to first define constructors for the sorts `KConfigurationItem`, `KComputationItem` and `KResult`, and then to define how each computation item reduces to a result.

In our case, the computations are the formulae or subformulae that still need to be processed, and the results are the corresponding boolean expressions that need to be checked in the current (generated code) context to see whether the formula has been violated or not. We define the following additional constructors: we add four constructors for configurations, namely “`k`” that wraps the current computation, “`code`” that wraps the current generated code, and “`nextBit`” that wraps the next available bit; we add one main constructor for computations, “`form`”, that wraps a formula, and one constant computation item per operator

in the input language (the later is needed to know how to combine back the results of the corresponding subexpressions; finally, we add one constructor for results, “exp”, that wraps a boolean expression.

The formula is processed in a depth-first-order, following a divide-and-conquer philosophy. Each subformula is decomposed into a list of computation subtasks consisting of its subformulae, then the corresponding results are composed back into a result corresponding to the original subformula. Recall that equations/rules apply wherever they match, not only at the top. Let us only discuss the two equations defining the “since” ($_S$), the last two in the module below. The first one is straightforward: it decomposes the task of processing $F1 \ S \ F2$ to the subtasks of processing $F1$ and $F2$; the computation item S is placed in the computation structure to prepare the terrain for the next equation. The next equation applies after $F1$ and $F2$ have been processed, say to expressions $B1$ and $B2$, respectively; if C is the code generated so far and if $I+1$ is the next bit available, then the boolean expression corresponding to the current since formula is indeed $\text{bit}'(I+1)$, provided that one adds the corresponding code capturing the recursive semantics of since to the generated code.

```
fmod MONITOR-GENERATION is
  protecting K + SYNTAX + CODE .
  op k : KComputation -> KConfigurationItem .
  op code : Code -> KConfigurationItem .
  op nextBit : Nat -> KConfigurationItem .
  op process : Formula -> KConfigurationItem .
  op form : Formula -> KComputationItem .
  op exp : Exp -> KResult .
  ops ! /\ 0 S : -> KComputationItem .
  var P : Predicate . vars F F1 F2 : Formula . var C : Code .
  var I : Nat . vars B B1 B2 : Exp . var K : KComputation .
  eq process(F) = k(form(F)) code(nil) nextBit(0) .
  eq k(form(P) -> K) = k(exp(P(s)) -> K) .
  eq form(! F) = form(F) -> ! .
  eq exp(B) -> ! = exp(! B) .
  eq form(F1 /\ F2) = [form(F1),form(F2)] -> /\ .
  eq [exp(B1),exp(B2)] -> /\ = exp(B1 /\ B2) .
  eq form(0 F) = form(F) -> 0 .
  eq k(exp(B) -> 0 -> K) code(C) nextBit(I)
    = k(exp(bit[I + 1]) -> K) code(C ; bit'[I + 1] := B)
      nextBit(I + 1) .
  eq form(F1 S F2) = [form(F1), form(F2)] -> S .
  eq k([exp(B1),exp(B2)] -> S -> K) code(C) nextBit(I)
    = k(exp(bit'[I + 1]) -> K)
      code(C ; bit'[I + 1] := B2 \/ B1 /\ bit[I + 1]) nextBit(I + 1) .
endfm
```

Putting It All together The following module plugs the code generated above into the general pattern:


```

=====
reduce in PRETTY-PRINT :
  genMonitor(! ('a /\ ! (0 'b /\ 'c S ('d /\ ! 'e S 'f)))) .
rewrites: 46 in -93406740ms cpu (1ms real) (~ rewrites/second)
result Code:
bit[1 .. 3] := false ;
foreach new state s do {
  // "first update the bits in a consistent order"
  bit'[1] := 'b(s) ;
  bit'[2] := 'f(s) \/ ! 'e(s) /\ bit[2] ;
  bit'[3] := 'd(s) /\ bit'[2] \/ 'c(s) /\ bit[3] ;
  // "then check whether the formula is violated"
  if 'a(s) /\ ! (bit[1] /\ bit'[3]) then
    Error ;
  // "finally, update the state of the monitor"
  bit[1 .. 3] := bit'[1 .. 3]
}

Maude>

```

The K Module One should upload the next module whenever one wants to use the K technique to define a language, logic or tool. Note that the module below has nothing to do with our particular logic under consideration in this paper; that is the reason for which we exiled it here.

```

fmod K is
  sorts KConfigurationItem KConfiguration .
  subsort KConfigurationItem < KConfiguration .
  op empty : -> KConfiguration .
  op _ _ : KConfiguration KConfiguration -> KConfiguration [assoc comm id: empty] .

  sorts KComputationItem KNeComputation KComputation .
  subsort KComputationItem < KNeComputation < KComputation .
  op nil : -> KComputation .
  op _->_ : KComputation KComputation -> KComputation [assoc id: nil] .
  op _->_ : KNeComputation KNeComputation -> KNeComputation [ditto] .

  sort KComputationList .
  subsort KComputation < KComputationList .
  op nil : -> KComputationList .
  op _ , _ : KComputationList KComputationList -> KComputationList [assoc id: nil] .

  sort KResult KResultList .
  subsorts KResult < KResultList < KComputation .
  op nil : -> KResultList .
  op _ , _ : KResultList KResultList -> KResultList [assoc id: nil] .

  op [_] : KComputationList -> KComputationItem .
  op [_] : KResultList -> KComputationItem .

```

```
op [_|_] : KComputationList KResultList -> KComputationItem .

var K : KNeComputation . var Kl : KComputationList .
var R : KResult . var Rl : KResultList .
eq [K,Kl] = K -> [Kl | nil] .
eq R -> [K,Kl | Rl] = K -> [Kl | Rl,R] .
eq R -> [nil | Rl] = [Rl,R] .
endfm
```

To use `K`, after importing the module above, one should define one's own constructors for configuration items (sort `KConfigurationItem`), for computation items (sort `KComputationItem`), and for results (sort `KResult`). For our example, we defined all these at the beginning of the module `MONITOR-GENERATION`.

9 Conclusion

In this paper we made the following contributions:

1. by showing that their sets have the same cardinal, we showed that various definitions of safety properties over finite, infinite, or both finite and infinite traces considered in the literature are equivalent;
2. we introduced persistent (finite-trace) safety properties and showed their one-to-one correspondence with the infinite trace safety properties;
3. we showed that all safety properties can be characterized as “always past” properties, that is, as properties that their prefixes must always satisfy;
4. we showed that monitoring safety properties is an arbitrarily hard problem;
5. we showed that the problem of monitoring safety properties expressed as extended regular expressions requires non-elementary space; and
6. we gave an optimal monitor synthesis algorithm for past LTL.

References

1. M. Abadi and L. Lamport. The existence of refinement mappings. In *LICS*, pages 165–175. IEEE Computer Society, 1988.
2. M. Abadi and L. Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, 1991.
3. C. Allan, P. Avgustinov, A. S. Christensen, L. Hendren, S. Kuzins, O. Lhoták, O. de Moor, D. Sereni, G. Sittampalam, and J. Tibble. Adding trace matching with free variables to AspectJ. In R. P. Gabriel, editor, *ACM Conference on Object-Oriented Programming, Systems and Languages (OOPSLA)*, pages 345–364. ACM Press, 2005.
4. B. Alpern and F. B. Schneider. Defining liveness. *IPL*, 21(4):181–185, 1985.
5. K. R. Apt, N. Francez, and S. Katz. Appraising fairness in languages for distributed programming. In *POPL*, pages 189–198, 1987.
6. A. K. Chandra, D. Kozen, and L. J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.

7. F. Chen, M. D'Amorim, and G. Roşu. Checking and correcting behaviors of Java programs at runtime with Java-MOP. In *RV'05*, volume 144(4) of *ENTCS*, 2005.
8. M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott. Maude Manual. <http://maude.cs.uiuc.edu>.
9. M. d'Amorim and G. Roşu. Efficient monitoring of ω -languages. In K. Etessami and S. K. Rajamani, editors, *CAV*, volume 3576 of *LNCS*, pages 364–378, 2005.
10. K. W. Hamlen, J. G. Morrisett, and F. B. Schneider. Computability classes for enforcement mechanisms. *ACM Trans. Program. Lang. Syst.*, 28(1):175–205, 2006.
11. K. Havelund and G. Roşu. Efficient monitoring of safety properties. *Software Tools and Technology Transfer*, 6(2):158–173, 2004. (also TACAS'02, LNCS 2280).
12. S. Hirst. A new algorithm solving membership of extended regular expressions. Technical report, The University of Sydney, 1989.
13. J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
14. L. Ilie, B. Shan, and S. Yu. Fast algorithms for extended regular expression matching and searching. In *Proceedings of STACS'03*, volume 2607 of *LNCS*, pages 179–190, 2003.
15. O. Kupferman and M. Y. Vardi. Freedom, Weakness, and Determinism: From Linear-Time to Branching-Time. In *Proc. of the IEEE Symposium on Logic in Computer Science*, pages 81–92, 1998.
16. O. Kupferman and M. Y. Vardi. Model Checking of Safety Properties. In *Proc. of CAV'99: Conference on Computer-Aided Verification*, Trento, Italy, 1999.
17. O. Kupferman and M. Y. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 19(3):291–314, 2001.
18. O. Kupferman and S. Zuhovitzky. An improved algorithm for the membership problem for extended regular expressions. In *Proc. of MFCS'02*, volume 2420 of *LNCS*, pages 446–458, 2002.
19. L. Lamport. Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.*, 3(2):125–143, 1977.
20. L. Lamport. Logical foundation. In M. W. Alford, J. P. Ansart, G. Hommel, L. Lamport, B. Liskov, G. P. Mullery, F. B. Schneider, M. Paul, and H. J. Siegert, editors, *Distributed systems: Methods and tools for specification. An advanced course*, volume 190 of *LNCS*, pages 119–130. Springer-Verlag, 1985.
21. Z. Manna and A. Pnueli. *Temporal verification of reactive systems: safety*. Springer-Verlag New York, Inc., New York, NY, USA, 1995.
22. G. Roşu and K. Havelund. Rewriting-based techniques for runtime verification. *Automated Software Engineering*, 12(2):151–197, 2005.
23. G. Roşu and M. Viswanathan. Testing extended regular language membership incrementally by rewriting. In *RTA'03*, volume 2706 of *LNCS*. Springer, 2003.
24. G. Roşu. K: a rewrite-based framework for modular lang. design, semantics, analysis and implementation (V2). Technical Report UIUCDCS-R-2006-2802, 2006.
25. G. Rosu. An effective algorithm for the membership problem for extended regular expressions. In *Proceedings of the 10th International Conference on Foundations of Software Science and Computation Structures (FOSSACS'07)*, volume 4423 of *LNCS*, pages 332–345. Springer-Verlag, 2007.
26. F. B. Schneider. *On Concurrent Programming*. Springer, 1997.
27. F. B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 3(1):30–50, 2000.
28. L. J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, Massachusetts Institute of Technology, 1974.

29. L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time (preliminary report). In *STOC*, pages 1–9. ACM Press, 1973.
30. K. Thompson. Regular expression search algorithm. *CACM*, 11(6):419–422, 1968.
31. H. Yamamoto. An automata-based recognition algorithm for semi-extended regular expressions. In *Proc. of MFCS'00*, volume 1893 of *LNCS*, pages 699–708, 2000.
32. H. Yamamoto. A new recognition algorithm for extended regular expressions. In *Proceedings of ISAAC'01*, volume 2223 of *LNCS*, pages 257–267, 2001.
33. H. Yamamoto and T. Miyazaki. A fast bit-parallel algorithm for matching extended regular expressions. In *Proc. of COCOON'03*, volume 2697 of *LNCS*, pages 222–231, 2003.