

UC Berkeley

UC Berkeley Previously Published Works

Title

From LTL to rLTL monitoring: improved monitorability through robust semantics

Permalink

<https://escholarship.org/uc/item/3zw1858p>

Journal

Formal Methods in System Design, 59(1-3)

ISSN

0925-9856

ISBN

978-1-4503-7018-9

Authors

Masclé, Corto
Neider, Daniel
Schwenger, Maximilian
[et al.](#)

Publication Date

2021-12-01

DOI

10.1007/s10703-022-00398-4

Peer reviewed



From LTL to rLTL monitoring: improved monitorability through robust semantics

Corto Mascle¹ · Daniel Neider²  · Maximilian Schwenger³ · Paulo Tabuada⁴ · Alexander Weinert⁵ · Martin Zimmermann^{6,7}

Received: 28 June 2021 / Accepted: 13 September 2022 / Published online: 21 October 2022
© The Author(s) 2022

Abstract

Runtime monitoring is commonly used to detect the violation of desired properties in safety critical cyber-physical systems by observing its executions. Bauer et al. introduced an influential framework for monitoring Linear Temporal Logic (LTL) properties based on a three-valued semantics for a finite execution: the formula is already satisfied by the given execution, it is already violated, or it is still undetermined, i.e., it can still be satisfied and violated by appropriate extensions of the given execution. However, a wide range of formulas are not monitorable under this approach, meaning that there are executions for which satisfaction and violation will always remain undetermined no matter how it is extended. In particular, Bauer et al. report that 44% of the formulas they consider in their experiments fall into this category. Recently, a robust semantics for LTL was introduced to capture different degrees by which a property can be violated. In this paper we introduce a robust semantics for finite strings and show its potential in monitoring: every formula considered by Bauer et al. is monitorable under our approach. Furthermore, we discuss which properties that come naturally in LTL monitoring—such as the realizability of all truth values—can be transferred to the robust setting. We show that LTL formulas with robust semantics can be monitored by deterministic automata, and provide tight bounds on the size of the constructed automaton. Lastly, we report on a prototype implementation and compare it to the LTL monitor of Bauer et al. on a sample of examples.

Keywords Runtime monitoring · Robust Linear Temporal Logic

1 Introduction

Runtime monitoring is nowadays routinely used to assess the satisfaction of properties of systems during their execution. To this end, a monitor, a finite-state device that runs in

This work has partly been conducted at the Max Planck Institute for Software Systems, Kaiserslautern, Germany.

✉ Daniel Neider
daniel.neider@uol.de

Extended author information available on the last page of the article

parallel to the system during deployment, evaluates it with respect to a fixed property. This is especially useful for systems that cannot be verified prior to deployment and, for this reason, can contain hidden bugs. While it is useful to catch and document these bugs during an execution of a system, we find that the current approach to runtime verification based on Linear Temporal Logic (LTL) [14] is not sufficiently informative, especially in what regards a system's robustness. Imagine that we are monitoring a property φ and that this property is violated during an execution. In addition to be alerted to the presence of a bug, there are several other questions we would like to have answered such as: although φ was falsified, was there a *weaker* version of φ that was still satisfied or did the system fail catastrophically? Similarly, if we consider a property of the form $\varphi \rightarrow \psi$, where φ is an environment assumption and ψ is a system guarantee, and the environment violates φ *slightly* along an execution can we still guarantee that ψ is only *slightly* violated?

Answering these questions requires a logical formalism for specifying properties that provides meaning to terms such as *weaker* and *slightly*. Formalizing these notions within temporal logic, so as to be able to reason about the robustness of a system, was the main impetus behind the definition of *robust Linear-time Temporal Logic* (rLTL) [58]. While reasoning in LTL yields a binary result, rLTL adopts a five-valued semantics representing different *shades of violation*. Consider, for example, the specification $\Box a \rightarrow \Box b$ requiring that b is always satisfied provided a is always satisfied. In LTL, if the premise a is violated in a single position of the trace, then the specification is satisfied vacuously, eliminating all requirements on the system regarding $\Box b$. In this case, rLTL detects a mild violation of the premise and thus allows for a mild violation of the conclusion.

While recent work covers the synthesis [58] and verification problem [5, 6, 58] for rLTL, the runtime verification problem is yet to be addressed, except for a preliminary version of the results in this paper presented in the 2020 International Conference on Hybrid Systems: Computation and Control [47]. Since runtime verification can only rely on finite traces by its nature, interesting theoretical questions open up for rLTL with finite semantics. On the practical side, the very same reasons that make runtime verification for LTL so useful also motivate the need for developing a finite semantics suitable for rLTL runtime verification. To this end, we tackle the problem of evaluating a property over infinite traces based on a finite prefix similarly to Bauer et al. [14]. If the available information is insufficient to declare a specification violated or satisfied, the monitor reports a ?. This concept is applied to each degree of violation of the rLTL semantics. Thus, the rLTL monitor's verdict consists of four three-valued bits, as the rLTL semantics is based on four two-valued bits. Each bit represents a degree of violation of the specification in increasing order of severity.

As an example, consider an autonomous drone that may or may not be in a stable state.¹ The specification requires that it remains stable throughout the entire mission. However, if the take-off is shaky due to bad weather, the drone is unstable for the first couple of minutes. An LTL monitor thus jumps to the conclusion that the specification is violated whereas an rLTL monitor only reports a *partial* violation. As soon as the drone stabilizes, the LTL monitor does not indicate any improvement while the rLTL monitor refines its verdict to also report a partial satisfaction.

Some interesting properties that come naturally with LTL monitoring cannot be seamlessly lifted to rLTL monitoring. While it is obvious that all three truth values for finite trace LTL, i.e., satisfied, violated, and unknown, can be realized for some prefix and formula, the same does not hold for rLTL. Intuitively, the second and third bit of the rLTL monitor's four-bit output for the property $\Box a$ represent whether a eventually holds forever or whether it holds

¹ By this we mean, e.g., that the error in tracking a desired trajectory is below a certain threshold.

infinitely often, respectively. Based on a prefix, a monitor cannot distinguish between these two shades of violation, rendering some monitor outputs unrealizable.

In addition to that, we investigate how the level of informedness of an LTL monitor relates to the one of an rLTL monitor. The first observation is that a verdict of an LTL monitor can be refined at most once, from an unknown to either true or false. With rLTL semantics, however, a monitor can refine its output for a given formula up to four times. Secondly, an LTL monitor can only deliver meaningful verdicts for *monitorable* [13] properties. Intuitively, a property is monitorable if every prefix can be extended by a finite continuation that gives a definite verdict. We adapt the definition to robust monitoring and show that neither does LTL monitorability imply rLTL monitorability, nor vice versa.

Notwithstanding the above, empirical data suggests that rLTL monitoring indeed provides more information than LTL monitoring: This paper presents an algorithm synthesizing monitors for rLTL specifications. An implementation thereof allows us to validate the approach by replicating the experiments of Bauer et al. [13]. As performance metric, we use LTL and rLTL monitorability. While 44% of the formulas considered by Bauer et al. [13] are not LTL-monitorable, we show all of them to be rLTL-monitorable. This indicates that rLTL monitoring is an improvement over LTL monitoring in terms of monitorability and complements the theoretical results with a practical validation.

This paper is an extended version of the work presented in the 2020 International Conference on Hybrid Systems: computation and control [47]. The main research contributions are a finite trace semantics for rLTL coupled with an investigation of its properties when compared to LTL, as well as an algorithm to synthesize monitors for rLTL specifications. Our construction is doubly-exponential in the size of the formula, showing that rLTL monitoring is no more costly than LTL monitoring. In addition to the original work [47], this article features (i) a more detailed discussion of the properties of our finite trace semantics for rLTL, (ii) a new running example detailing each step of the monitor construction, (iii) a new example illustrating the nesting of rLTL operators, (iv) refined complexity bounds on our monitor construction, and (v) all proofs omitted from the conference paper, which provide important additional insight into the problem of monitoring rLTL properties.

Related work

In runtime verification [22, 35, 42, 49] the specification is often given in LTL [46]. While properties arguing about the past or current state of a system are always monitorable [34], LTL can also express assumptions on the future that cannot be validated using only a finite prefix of a word. Thus, adaptations of LTL have been proposed which include different notions of a next step on finite words [24, 45], lifting LTL to a three- or four-valued domain [13, 14], or applying predictive measures to rule out impossible extensions of words [60].

Non-binary monitoring has also been addressed by adding quantitative measures such as counting events [9, 48]. Most notably, Bartocci et al. [10] evaluate the “likelihood” that a satisfying or violating continuation will occur. To this end, for a given prefix, they count how long a continuation needs to be such that the specification is satisfied/violated; these numbers are then compared against each other. The resulting verdict is quinary: satisfying/violating, presumably satisfying/violating, or inconclusive. This approach is similar in nature to our work as it assesses the degree of satisfaction or violation of a given prefix. However, the motivation and niche of both approaches differs: Bartocci et al.’s approach computes—intuitively speaking—the amount of work that is required to satisfy or violate a specification,

which allows for estimating the likelihood of satisfaction. Our approach, however, focuses on measuring the extent to which a specification was satisfied or violated.

Apart from that, monitoring tools collecting statistics [1, 4, 30] become increasingly popular: Snort [55] is a commercial tool for rule-based network monitoring and computing efficient statistics, Beep Beep 3 [33] is a tool based on a query language allowing for powerful aggregation functions and statistical measures. On the downside, these tools impose the overhead of running a heavy-weight application on the monitored system. In contrast, we generate monitor automata out of an rLTL formula. Such an automaton can easily and automatically be implemented on almost any system with statically determined memory requirements and negligible performance overhead. Similarly, the Copilot [52] framework based on synchronous languages [16, 19] transforms a specification in a declarative data-flow language into a C implementation of a monitor with constant space and time requirements. Lola [2, 19] allows for more involved computations, also incorporating parametrization [27] and real-time capabilities [28] while retaining constant space and time requirements.

Another approach is to enrich temporal logics with quantitative measures such as taking either the edit distance [37], counting the number of possible infinite models for LTL [31, 59], incorporating aggregation expressions into metric first-order temporal logic [11], or using averaging temporal operators that quantify the degree of satisfaction of a signal for a specification by integrating the signal w.r.t. a constant reference signal [3].

Rather than enriching temporal logics with such strong quantitative measures, we consider a robust version of LTL: rLTL [5–7, 58]. Robust semantics yields information about to which degree a trace violates a property. We adapt the semantics to work with finite traces by allowing for intermediate verdicts. Here, a certain degree of violation can be classified as “indefinite” and refined when more information becomes available to the monitor. Similarly, for Signal Temporal Logic [43, 44], Fainekos et al. [25] introduced a notion of spatial robustness based on interpreting atomic propositions over the real numbers. The sign of the real number provides information about satisfaction/violation while its absolute value provides information about robustness, i.e., how much can this value be altered without changing satisfaction/violation. This approach is complementary to ours since the notion of robustness in rLTL is related to the temporal evolution of atomic propositions which are interpreted classically, i.e., over the Booleans. Donze et al. [21] introduced a notion of robustness closer to rLTL in the sense that it measures how long we need to wait for the truth value of a formula to change. For this, Cralley et al. [18] presented a convenient toolbox, achieving high efficiency through parallel evaluation. While the semantics of rLTL does not allow for quantifying the exact delay needed to change the truth value of a formula, it allows for distinguishing between the influence that different temporal evolutions, e.g., delays, persistence, and recurrence, have on the truth value of an LTL formula. Closer to rLTL is the work of Radionova et al. [54] (see also [57]) that established an unexpected connection between LTL and filtering through a quantitative semantics based on convolution with a kernel. By using different kernels, one can express weaker or stronger interpretations of the same formula. However, this requires the user to choose multiple kernels and to use multiple semantics to reason about how the degradation of assumptions leads to the degradation of guarantees. In contrast, no such choices are required in rLTL. Finally, it is worth mentioning that extensions similar to rLTL have been proposed for other temporal logics, such as prompt LTL and linear dynamic logic [50, 51].

Another venue for robust monitoring is machine learning. Cheng [17] presents an algorithm for generating monitors evaluating the distance between the input of a neural net and its training data. While neural nets are prone to fragility, the monitor is provably robust in the sense that minor input deviations invariably lead to minor changes in the output. Similarly,

Finkbeiner et al. [29] generate monitors for medical cyber-physical systems controlled by machine learned components. Due to the complexity of the underlying specification language, they opt for the simpler task of analyzing the robustness of the specification instead. If the specification is robust, then so will be the generated monitors.

2 Robust Linear Temporal Logic

Throughout this work, we assume basic familiarity with classical LTL and refer the reader to a textbook for more details on the logic (see, e.g., [8]). Moreover, let us fix some finite set P of atomic propositions throughout the paper and define $\Sigma = 2^P$. We denote the set of finite and infinite words over Σ by Σ^* and Σ^ω , respectively. The empty word is denoted by ε and \sqsubseteq and \sqsubset denote the non-strict and the strict prefix relation, respectively. Moreover, we denote the set of Booleans by $\mathbb{B} = \{0, 1\}$.

The logics LTL and rLTL share the same syntax save for a dot superimposed on temporal operators. More precisely, the syntax of rLTL is given by the grammar

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \varphi \rightarrow \varphi \mid \odot\varphi \mid \varphi \mathbf{U} \varphi \mid \varphi \mathbf{R} \varphi \mid \diamond\varphi \mid \square\varphi,$$

where p ranges over atomic propositions in P and the temporal operators \odot , \mathbf{U} , \mathbf{R} , \diamond and \square correspond to “next”, “until”, “release”, “eventually”, and “always”, respectively.² The size $|\varphi|$ of a formula φ is the number of its distinct subformulas. Furthermore, we denote the set of all LTL and rLTL formulas over P by Φ_{LTL} and Φ_{rLTL} , respectively.

The development of rLTL was motivated by the observation that the difference between “minor” and “major” violations of a formula cannot be adequately described in a two-valued semantics. If an LTL formula φ , for example, demands that the property p holds at all positions of a word $\sigma \in \Sigma^\omega$, then σ violates φ even if p does not hold at only a single position, a very minor violation. The semantics of LTL, however, does not differentiate between the σ above and a σ' in which the property p never holds, a major violation of the property φ .

In order to alleviate this shortcoming, Tabuada and Neider introduced Robust Linear-time Temporal Logic (rLTL) [58], whose semantics allows for distinguishing various “degrees” to which a word violates a formula. More precisely, the semantics of rLTL are defined over the set $\mathbb{B}_4 = \{0000, 0001, 0011, 0111, 1111\}$ of five *truth values*, each of which is a monotonically increasing sequence of four bits. We order the truth values in \mathbb{B}_4 by $0000 < 0001 < 0011 < 0111 < 1111$.

Intuitively, this order reflects increasingly desirable outcomes. If the specification is $\square p$, the least desirable outcome, represented by 0000, is that p never holds on the entire trace. A slightly more desirable outcome is that p at least holds *sometime* but not infinitely often, which results in the value 0001. An even more desirable outcome would be if p holds infinitely often, while also being violated infinitely often, represented by 0011. Climbing up the ladder of desirable outcomes, the next best one requires p to hold infinitely often while being violated only finitely often, represented by the value 0111. Lastly, the optimal outcome fully satisfies $\square p$, so p holds the entire time, represented by 1111. Thus, the first bit states whether $\square p$ is satisfied, the second one stands for $\diamond \square p$, the third one for $\square \diamond p$, and the fourth one for $\diamond p$. If all of them are 0, $\square \neg p$ holds. The robust release is defined analogously.

² Note that we include the operators \wedge , \rightarrow , and \mathbf{R} explicitly in the syntax as they cannot be derived from other operators due to the many-valued nature of rLTL. Following the original work on rLTL [58], we also include the operators \diamond and \square explicitly (which can be derived from \mathbf{U} and \mathbf{R} , respectively).

The robust eventually-operator considers future positions in the trace and returns the truth value with the least degree of violation, which is a maximization with respect to the order defined above. This closely resembles the LTL definition. The robust until is defined analogously.

Based on this, the boolean conjunction and disjunction are defined as min and max, respectively, w.r.t. the order defined above, which generalizes the classical definition thereof. For the implication, consider a specification $\Box a \rightarrow \Box g$, where $\Box a$ is an assumption on the environment and $\Box g$ is a system guarantee. If the truth value of $\Box g$ is greater or equal to the one of $\Box a$, the implication is fully satisfied. Thus, the rLTL semantics takes the violation of the assumption into account and lowers the requirements on the guarantees. However, if the guarantee exhibits a greater violation than the assumptions, the truth value of the implication is the same as the one of the guarantee. Lastly, the intuition behind the negation is that every truth value that is not 1111 constitutes a violation of the specification. Thus, the negation thereof is a full satisfaction (1111). The negation of the truth value representing a perfect satisfaction (1111) is a full violation (0000).

To introduce the semantics, we need some additional notation: for a word $\sigma = \sigma(0)\sigma(1)\sigma(2)\dots \in \Sigma^\omega$ and a natural number n , define $\sigma[n, \infty) = \sigma(n)\sigma(n+1)\sigma(n+2)\dots$, (i.e., as the suffix of σ obtained by removing the first n letters of σ). To be able to refer to individual bits of an rLTL truth value $\beta \in \mathbb{B}_4$, we use $\beta[i]$ with $i \in \{1, \dots, 4\}$ as to denote the i -th bit of β .

For the sake of a simpler presentation, we denote the semantics of both LTL and rLTL not in terms of satisfaction relations but by means of *valuation functions*. For LTL, the valuation function $V : \Sigma^\omega \times \Phi_{LTL} \rightarrow \mathbb{B}$ assigns to each infinite word $\sigma \in \Sigma^\omega$ and each LTL formula $\varphi \in \Phi_{LTL}$ the value 1 if σ satisfies φ and the value 0 if σ does not satisfy φ , and is defined as usual (see, e.g., [8]). The semantics of rLTL, on the other hand, is more complex and formalized next by an valuation function $V_r : \Sigma^\omega \times \Phi_{rLTL} \rightarrow \mathbb{B}_4$ mapping an infinite word $\sigma \in \Sigma^\omega$ and an rLTL formula φ to a truth value in \mathbb{B}_4 .

- $V_r(\sigma, p) = \begin{cases} 1111 & \text{if } p \in \sigma(0), \\ 0000 & \text{if } p \notin \sigma(0), \end{cases}$
- $V_r(\sigma, \neg\varphi) = \begin{cases} 1111 & \text{if } V_r(\sigma, \varphi) \neq 1111, \\ 0000 & \text{if } V_r(\sigma, \varphi) = 1111, \end{cases}$
- $V_r(\sigma, \varphi_1 \wedge \varphi_2) = \min\{V_r(\sigma, \varphi_1), V_r(\sigma, \varphi_2)\}$,
- $V_r(\sigma, \varphi_1 \vee \varphi_2) = \max\{V_r(\sigma, \varphi_1), V_r(\sigma, \varphi_2)\}$,
- $V_r(\sigma, \varphi_1 \rightarrow \varphi_2) = \begin{cases} 1111 & \text{if } V_r(\sigma, \varphi_1) \leq V_r(\sigma, \varphi_2), \\ V_r(\sigma, \varphi_2) & \text{if } V_r(\sigma, \varphi_1) > V_r(\sigma, \varphi_2), \end{cases}$
- $V_r(\sigma, \odot \varphi) = V_r(\sigma[1, \infty), \varphi)$,
- $V_r(\sigma, \diamond \varphi) = \beta$ with $\beta[i] = \max_{n \geq 0} V_r(\sigma[n, \infty), \varphi)[i]$ for $i \in \{1, \dots, 4\}$,
- $V_r(\sigma, \Box \varphi) = \beta$ with

$$\begin{aligned} \beta[1] &= \min_{n \geq 0} V_r(\sigma[n, \infty), \varphi)[1], \\ \beta[2] &= \max_{m \geq 0} \min_{n \geq m} V_r(\sigma[n, \infty), \varphi)[2], \\ \beta[3] &= \min_{m \geq 0} \max_{n \geq m} V_r(\sigma[n, \infty), \varphi)[3], \\ \beta[4] &= \max_{n \geq 0} V_r(\sigma[n, \infty), \varphi)[4], \end{aligned}$$

– $V_r(\sigma, \varphi_1 \mathbf{U} \varphi_2) = \beta$ with

$$\beta[i] = \max_{n \geq 0} \min\{V_r(\sigma[n, \infty), \varphi_2)[i], \min_{0 \leq n' < n} V_r(\sigma[n', \infty), \varphi_1)[i]\},$$

for $i \in \{1, \dots, 4\}$,

– $V_r(\sigma, \varphi_1 \mathbf{R} \varphi_2) = \beta$ with

$$\beta[1] = \min_{n \geq 0} \max\{V_r(\sigma[n, \infty), \varphi_2)[1], \max_{0 \leq n' < n} V_r(\sigma[n', \infty), \varphi_1)[1]\},$$

$$\beta[2] = \max_{m \geq 0} \min_{n \geq m} \max\{V_r(\sigma[n, \infty), \varphi_2)[2], \max_{0 \leq n' < n} V_r(\sigma[n', \infty), \varphi_1)[2]\},$$

$$\beta[3] = \min_{m \geq 0} \max_{n \geq m} \max\{V_r(\sigma[n, \infty), \varphi_2)[3], \max_{0 \leq n' < n} V_r(\sigma[n', \infty), \varphi_1)[3]\}, \text{ and}$$

$$\beta[4] = \max_{n \geq 0} \max\{V_r(\sigma[n, \infty), \varphi_2)[4], \max_{0 \leq n' < n} V_r(\sigma[n', \infty), \varphi_1)[4]\}.$$

So as to not clutter this section too much, we refer the reader to the original work by Tabuada and Neider [58] for a thorough introduction and motivation to the preceding semantics. However, we here want to illustrate the definition above and briefly argue that it indeed captures the intuition described at the beginning of this section. To this end, we reconsider the formulas $\Box p$, $\Box a \rightarrow \Box g$, $\Box (q \rightarrow \Diamond p)$ in Examples 1, 2, and 3 respectively.

Example 1 Consider the formula $\Box p$ and the following five infinite words over the set $P = \{p\}$ of atomic propositions:

$\sigma_1 = \{p\}^\omega$	(“ p holds always”)
$\sigma_2 = \emptyset\{p\}^\omega$	(“ p holds almost always”)
$\sigma_3 = (\emptyset\{p\})^\omega$	(“ p holds infinitely often”)
$\sigma_4 = \{p\}\emptyset^\omega$	(“ p holds finitely often”)
$\sigma_5 = \emptyset^\omega$	(“ p holds never”)

Let us begin the example with the word $\sigma_1 = \{p\}^\omega$. It is not hard to verify that $V_r(\sigma_1, \Box p)[1] = 1$ because p always holds in σ_1 , i.e., $\min_{n \geq 0} V_r(\sigma[n, \infty), p)[1] = 1$ for $n \geq 0$. Using the same argument, we also have $V_r(\sigma_1, \Box p)[2] = V_r(\sigma_1, \Box p)[3] = V_r(\sigma_1, \Box p)[4] = 1$. Thus, $V_r(\sigma_1, \Box p) = 1111$.

As another example, consider the word $\sigma_2 = \emptyset\{p\}^\omega$. In this case, we have $V_r(\sigma_1, \Box p)[1] = 0$ because $V_r(\sigma[0, \infty), p)[1] = 0$ (p does not hold in the first symbol of σ_2). However, $V_r(\sigma_1, \Box p)[2] = 1$ because p holds almost always, i.e., $\max_{m \geq 0} \min_{n \geq m} V_r(\sigma[n, \infty), p)[2] = 1$. Moreover, $V_r(\sigma_1, \Box p)[3] = V_r(\sigma_1, \Box a)[4] = 1$ and, therefore, $V_r(\sigma_2, \Box p) = 0111$. Similarly, we obtain $V_r(\sigma_3, \Box p) = 0011$, $V_r(\sigma_4, \Box p) = 0001$, and $V_r(\sigma_5, \Box p) = 0000$.

In conclusion, this indeed illustrates that the semantics of the robust always is in accordance with the intuition provided at the beginning of this section. □

Example 2 Let us now consider the more complex formula $\Box a \rightarrow \Box g$, where we interpret a to be an assumption on the environment of a cyber-physical system and g one of its guarantees. Moreover, let σ be an infinite word over $P = \{a, g\}$ such that $V_r(\sigma, \Box a \rightarrow \Box g) = 1111$. We now distinguish various cases.

First, let us assume that σ is such that $V_r(\sigma, \Box a) = 1111$, i.e., a always holds. By definition of the robust implication and since $V_r(\sigma, \Box a \rightarrow \Box g) = 1111$, this can only be the case if $V_r(\sigma, \Box g) = 1111$. Thus, the formula $\Box a \rightarrow \Box g$ ensures that if the environment assumption a always holds, so does the system guarantee g .

Next, assume that σ is such that $V_r(\sigma, \Box a) = 0111$, i.e., a does not always hold but almost always. By definition of the robust implication and since $V_r(\sigma, \Box a \rightarrow \Box g) = 1111$, this can only be the case if $V_r(\sigma, \Box g) \geq 0111$. In this case, the formula $\Box a \rightarrow \Box g$ ensures that if the environment assumption a holds almost always, then the system guarantee g holds almost always or—even better—always.

It is not hard to verify that we obtain similar results for the cases $V_r(\sigma, \Box a) \in \{0011, 0001, 0000\}$. In other words, the semantics of rLTL ensures that the violation of the system guarantee g is always proportional to the violation of the environment assumption a (given that $V_r(\sigma, \Box a \rightarrow \Box g)$ evaluates to 1111). Again, this illustrates that the semantics of the implication is in accordance with the intuition provided at the beginning of this section. \square

Example 3 As a last example, let us discuss the nesting of temporal operators. Consider the formula $\varphi = \Box (q \rightarrow \Diamond p)$ where we interpret q as a request and p as a response.

We have $V_r(\sigma[n, \infty), \Diamond p) = 1111$ if $\sigma[n, \infty)$ contains a response, otherwise we have $V_r(\sigma[n, \infty), \Diamond p) = 0000$. Similarly, we have $V_r(\sigma[n, \infty), q \rightarrow \Diamond p) = 1111$ if $q \in \sigma(n)$ implies that $\sigma[n, \infty)$ contains a response. On the other hand, if $q \in \sigma(n)$ and $\sigma[n, \infty)$ does not contain a response then we have $V_r(\sigma[n, \infty), q \rightarrow \Diamond p) = 0000$.

From these observations, we can deduce $V_r(\sigma, \varphi) = 1111$ if every request in σ is followed by a response, which is equivalent to the LTL formula $\varphi_1 = \Box (q \rightarrow \Diamond p)$ that expresses a request-response property. Further, we have $V_r(\sigma, \varphi) = 0111$ if and only if σ violates φ_1 and if from some point onwards, every request in σ is followed by a response. This is equivalent to the LTL formula $\neg\varphi_1 \wedge \varphi_2$ with $\varphi_2 = (\Box \Diamond q) \rightarrow (\Box \Diamond p)$, which expresses strong fairness. Similarly, we have $V_r(\sigma, \varphi) = 0011$ if and only if σ violates φ_2 and if for infinitely many positions, if there is a request in σ at that position, then it is followed by a response. This is equivalent to the LTL formula $\neg\varphi_2 \wedge \varphi_3$ with $\varphi_3 = (\Diamond \Box q) \rightarrow (\Box \Diamond p)$, which expresses weak fairness. Moreover, we have $V_r(\sigma, \varphi) = 0001$ if and only if σ violates φ_3 and if there is some position such that if there is a request in σ at that position, then it is followed by a response. This is equivalent to the LTL formula $\neg\varphi_3 \wedge \varphi_4$ with $\varphi_4 = (\Box q) \rightarrow (\Diamond p)$, which expresses a very weak notion of fairness. Finally, we have $V_r(\sigma, \varphi) = 0000$ if and only if σ violates φ_4 .

For $i \in \{1, 2, 3\}$, the LTL formula φ_i implies φ_{i+1} . Thus, if a trace σ violates φ_{i+1} , it also violates φ_i . This further illustrates the monotonicity of rLTL. This monotonicity also allows us to only require that φ_{i+1} violates φ_i in the intuitive explanations above, instead of having to require violations of all $\varphi_{i'}$ with $i' \leq i$. \square

It is important to note that rLTL is an extension of LTL. In fact, the LTL semantics can be recovered from the first bit of the rLTL semantics (after every implication $\varphi \rightarrow \psi$ has been replaced with $\neg\varphi \vee \psi$).³

Lemma 1 ([58], Proposition 5) *Let φ be an LTL formula without implications, and let φ' be the corresponding rLTL formula (obtained by dotting all temporal operators). Then, we have $V_r(\sigma, \varphi')[1] = V(\sigma, \varphi)$ for every trace σ .*

³ It turns out that Tabuada and Neider’s original proof [58, Proposition 5] has a minor mistake. Although the first bit of the rLTL semantics coincides with the original LTL semantics for all formulas that do not contain implications, the formula $\Box \neg a \rightarrow \Box a$ is an example witnessing this claim is no longer correct in the presence of implications, e.g., for $\{a\} \emptyset^\omega$. However, this issue can be fixed by replacing every implication $\varphi \rightarrow \psi$ with $\neg\varphi \vee \psi$. This substitution results in an equivalent LTL formula for which the first bit of the rLTL semantics indeed coincides with the LTL semantics.

Table 1 The function $\text{ltl}: \{1, \dots, 4\} \times \Phi_{rLTL} \rightarrow \Phi_{LTL}$

Operator	Symbol	Semantics ($\varphi, \psi \in \Phi_{rLTL}$)
Atomic proposition	$p \in P$	$1 \leq i \leq 4: \text{ltl}(i, p) = p$
Negation	\neg	$1 \leq i \leq 4: \text{ltl}(i, \neg\varphi) := \neg\text{ltl}(1, \varphi)$
Disjunction	\vee	$1 \leq i \leq 4: \text{ltl}(i, \varphi \vee \psi) := \text{ltl}(i, \varphi) \vee \text{ltl}(i, \psi)$
Conjunction	\wedge	$1 \leq i \leq 4: \text{ltl}(i, \varphi \wedge \psi) := \text{ltl}(i, \varphi) \wedge \text{ltl}(i, \psi)$
Implication	\rightarrow	$1 \leq i \leq 3: \text{ltl}(i, \varphi \rightarrow \psi) := (\text{ltl}(i, \varphi) \rightarrow \text{ltl}(i, \psi)) \wedge \text{ltl}(i + 1, \varphi \rightarrow \psi);$ $\text{ltl}(4, \varphi \rightarrow \psi) := \text{ltl}(4, \varphi) \rightarrow \text{ltl}(4, \psi)$
Robust next	\odot	$1 \leq i \leq 4: \text{ltl}(i, \odot\varphi) := \odot \text{ltl}(i, \varphi)$
Robust eventually	\diamond	$1 \leq i \leq 4: \text{ltl}(i, \diamond\varphi) := \diamond \text{ltl}(i, \varphi)$
Robust always	\square	$\text{ltl}(1, \square\varphi) := \square \text{ltl}(1, \varphi); \text{ltl}(2, \square\varphi) := \square \square \text{ltl}(2, \varphi);$ $\text{ltl}(3, \square\varphi) := \square \diamond \text{ltl}(3, \varphi); \text{ltl}(4, \square\varphi) := \diamond \text{ltl}(4, \varphi)$
Robust until	U	$1 \leq i \leq 4: \text{ltl}(i, \varphi \text{ U } \psi) := \text{ltl}(i, \varphi) \text{ U } \text{ltl}(i, \psi)$
Robust release	\mathbf{R}	$\text{ltl}(1, \varphi \mathbf{R} \psi) := \text{ltl}(1, \varphi) \mathbf{R} \text{ltl}(1, \psi);$ $\text{ltl}(2, \varphi \mathbf{R} \psi) := \diamond \square \text{ltl}(2, \psi) \vee \diamond \text{ltl}(2, \varphi);$ $\text{ltl}(3, \varphi \mathbf{R} \psi) := \square \diamond \text{ltl}(3, \psi) \vee \diamond \text{ltl}(3, \varphi);$ $\text{ltl}(4, \varphi \mathbf{R} \psi) := \diamond \text{ltl}(4, \psi) \vee \diamond \text{ltl}(4, \varphi)$

To reduce the number of cases we have to consider in our inductive proofs (for instance the one for Lemma 3), we note that the robust eventually and the robust always operator are syntactic sugar. Formally, we say that two rLTL formulas φ_1, φ_2 are equivalent if $V_r(\sigma, \varphi_1) = V_r(\sigma, \varphi_2)$ for every $\sigma \in \Sigma^\omega$. Now, let $\top = p \vee \neg p$ and $\perp = p \wedge \neg p$ for some atomic proposition p . Then, the robust eventually and the robust always are, as usual, expressible in terms of the robust until and the robust release, respectively.

Remark 1 1. $\diamond\varphi$ and $\top \text{ U } \varphi$ are equivalent.
 2. $\square\varphi$ and $\perp \mathbf{R} \varphi$ are equivalent.

2.1 An alternative definition of robust semantics for LTL

Before we introduce rLTL monitoring, we need to introduce an alternative definition of the semantics of rLTL, which is more convenient to prove some of the results from Sect. 3. This alternative definition has been introduced in later works on rLTL [5, 6].

Definition 1 Let the function $\text{ltl}: \{1, \dots, 4\} \times \Phi_{rLTL} \rightarrow \Phi_{LTL}$ be inductively defined as in Table 1. The rLTL semantics is then given as the valuation function $V_r: \Sigma^\omega \times \Phi_{rLTL} \rightarrow \mathbb{B}_4$, where for every $\sigma \in \Sigma^\omega$, every rLTL formula φ , and every $i \in \{1, \dots, 4\}$, the i -th bit of $V_r(\sigma, \varphi)$ is defined as $V_r(\sigma, \varphi)[i] = V(\sigma, \text{ltl}(i, \varphi))$ (i.e., via the semantics of the LTL formulas $\text{ltl}(i, \varphi)$).

As a consequence of Lemma 1 (cf. [58], Proposition 5), we know that rLTL is at least as expressive as LTL. The latter definition of the semantics of rLTL shows that it is not more expressive than LTL, in the sense that for all rLTL formulas there exist LTL formulas giving the truth values of each of the four bits. However, it is more convenient to work with one formula of rLTL than to work with the four LTL formulas capturing it.

A useful feature of the alternative semantics is the following property: to determine the truth value of an rLTL formula φ on σ , it suffices to determine the truth values of the LTL

formulas $\text{ltl}(i, \varphi)$ on σ . For certain formulas, $\text{ltl}(i, \varphi)$ is obtained from φ by a very simple rewriting, as shown below.

Remark 2 Let φ be an rLTL formula that has no always in the scope of a negation and only uses negation, conjunction, disjunction, next, eventually, and always. Then,

- $\text{ltl}(1, \varphi)$ is equivalent to the formula obtained from φ by replacing every \odot by \circ , every \diamond by \diamond , and every \square by \square ,
- $\text{ltl}(2, \varphi)$ is equivalent to the formula obtained from φ by replacing every \odot by \circ , every \diamond by \diamond , and every \square by $\diamond \square$,
- $\text{ltl}(3, \varphi)$ is equivalent to the formula obtained from φ by replacing every \odot by \circ , every \diamond by \diamond , and every \square by $\square \diamond$, and
- $\text{ltl}(4, \varphi)$ is equivalent to the formula obtained from φ by replacing every \odot by \circ , every \diamond by \diamond , and every \square by $\diamond \square$.

3 Monitoring robust LTL

In their work on LTL monitoring, Bauer et al. [14] define the problem of runtime monitoring as “*check[ing] LTL properties given finite prefixes of infinite [words]*”. More formally, given some prefix $u \in \Sigma^*$ and some LTL formula φ , it asks whether all, some, or no infinite extension $u\sigma \in \Sigma^\omega$ of u by some $\sigma \in \Sigma^\omega$ satisfies φ . To reflect these three possible results, the authors use the set $\mathbb{B}^? = \{0, ?, 1\}$ to define a three-valued logic that is syntactically identical to LTL, but equipped with a semantics in form of an evaluation function $V^m: \Sigma^* \times \Phi_{LTL} \rightarrow \mathbb{B}^?$ over finite prefixes. This semantics is defined such that $V^m(u, \varphi)$ is equal to 0 (is equal to 1) if no (if every) extension $u\sigma$ of u satisfies φ . If neither is the case, i.e., if there is an extension of u that satisfies φ and there is an extension of u that does not satisfy φ , then $V^m(u, \varphi)$ is equal to ?.

We aim to extend the approach of Bauer et al. to rLTL, whose semantics is based on truth values from the set \mathbb{B}_4 (containing the sequences of length four in 0^*1^*). As a motivating example, let us consider the formula $\varphi = \square s$ for some atomic proposition s and study which situations can arise when monitoring this formula. Note that the truth value of φ can be obtained by concatenating the truth values of the LTL formulas $\varphi_1 = \square s$, $\varphi_2 = \diamond \square s$, $\varphi_3 = \square \diamond s$, and $\varphi_4 = \diamond s$.

First, consider the empty prefix and its two extensions \emptyset^ω and $\{s\}^\omega$. We have $V_r(\emptyset^\omega, \varphi) = 0000$ and $V_r(\{s\}^\omega, \varphi) = 1111$. Thus, all four bits can both be equal to 0 and 1. This situation is captured by the sequence $????$ which signifies that for every position i and every bit $b \in \mathbb{B}$, there exists an extension of ε that has bit b in the i -th position of the truth value with respect to φ .

Now, consider the prefix $\{s\}$ for which we have $V_r(\{s\}\sigma, \varphi)[4] = 1$ for every $\sigma \in \Sigma^\omega$ as $\varphi_4 = \diamond s$ is satisfied on each extension of $\{s\}$ (s has already occurred). On the other hand, $V_r(\{s\}\emptyset^\omega, \varphi) = 0001$ and $V_r(\{s\}\{s\}^\omega, \varphi) = 1111$, i.e., the first three bits can both be 0 and 1 by picking an appropriate extension. Hence, the situation is captured by the sequence $???1$, signifying that the last bit is determined by the prefix, but the first three are not. Using dual arguments, the sequence $0???$ is used for the prefix \emptyset , signifying that the first bit is determined by the prefix as every extension violates $\varphi_1 = \square s$. However, the last three bits are not yet determined by the prefix, hence the trailing $?s$.

Finally, consider the prefix $\{s\}\emptyset$. Using the same arguments as for the previous two prefixes, we obtain $V_r(\{s\}\emptyset\sigma, \varphi)[1] = 0$ and $V_r(\{s\}\emptyset\sigma, \varphi)[4] = 1$ for every $\sigma \in \Sigma^\omega$. Also, as before, we have $V_r(\{s\}\emptyset\emptyset^\omega, \varphi) = 0001$ and $V_r(\{s\}\emptyset\{s\}^\omega, \varphi) = 0111$. Hence, here we obtain the

sequence $0??1$ signifying that the first and last bit are determined by the prefix, but the middle two are not.

In general, we use truth values of the form $0^*?^*1^*$, which follows from the fact that the truth values of rLTL are in 0^*1^* . Hence, let $\mathbb{B}_4^?$ denote the set of sequences of length four in $0^*?^*1^*$. Based on $\mathbb{B}_4^?$, we now formally define the rLTL monitoring semantics as a bitwise generalization of the LTL definition.

Definition 2 The semantics of the robust monitor $V_r^m: \Sigma^* \times \Phi_{rLTL} \rightarrow \mathbb{B}_4^?$ is defined as $V_r^m(u, \varphi) = \beta$ with

$$\beta[i] = \begin{cases} 0 & \text{if } V_r(u\sigma, \varphi)[i] = 0 \text{ for all } \sigma \in \Sigma^\omega; \\ 1 & \text{if } V_r(u\sigma, \varphi)[i] = 1 \text{ for all } \sigma \in \Sigma^\omega; \text{ and} \\ ? & \text{otherwise,} \end{cases}$$

for every $i \in \{1, \dots, 4\}$, every rLTL formula φ , and every $u \in \Sigma^*$.

First, let us remark that our notion of rLTL monitoring indeed refines the notion of LTL monitoring, which follows immediately from Lemma 1.

Remark 3 Let φ be an LTL formula without implications, and let φ' be the corresponding rLTL formula (obtained by dotting all temporal operators). Then, we have $V_r^m(u, \varphi')[1] = V^m(u, \varphi)$ for every $u \in \Sigma^*$.

Using rLTL monitoring semantics, we are able to infer information about the infinite run of a system after having read only a finite prefix thereof. In fact, this robust semantics provides far more information about the degree of violation of the specification than classical LTL monitoring as each bit of the monitoring output represents a degree of violation of the specification: a ? turning into a 0 or 1 indicates a deterioration or improvement in the system’s state, respectively. Consider, for instance, an autonomous drone with specification $\varphi = \Box s$ where s denotes a state of stable flight (recall the motivating example on Page 11). Initially, the monitor would output $????$ due to a lack of information. If taking off under windy conditions, the state s is not reached initially, hence the monitor issues a warning by producing $V_r^m(\emptyset^n, \varphi) = 0???$ for every $n > 0$. Thus, the safety condition is violated temporarily, but not irrecoverably. Hence, mitigation measures can be initiated. Upon success, the monitoring output turns into $V_r^m(\emptyset^n \{s\}, \varphi) = 0??1$ for every $n > 0$, signaling that flight was stable for some time.

Before we continue, let us first state that the new semantics is well-defined, i.e., that the sequence $\beta[1]\beta[2]\beta[3]\beta[4]$ in Definition 2 is indeed in $\mathbb{B}_4^?$.

Lemma 2 $V_r^m(u, \varphi) \in \mathbb{B}_4^?$ for every rLTL formula φ and every $u \in \Sigma^*$.

Proof Let $V_r^m(u, \varphi)[i] = 0$ and $j < i$. By definition of V_r^m , we have $V_r(u\sigma, \varphi)[i] = 0$ for every $\sigma \in \Sigma^\omega$. Hence, due to the monotonicity of the truth values from \mathbb{B}_4 used to define V_r , we obtain $V_r(u\sigma, \varphi)[j] = 0$ for every such σ . Hence, $V_r^m(u, \varphi)[j] = 0$.

A dual argument shows that $V_r^m(u, \varphi)[i] = 1$ and $j > i$ implies $V_r^m(u, \varphi)[j] = 1$. Combining both properties yields $V_r^m(u, \varphi) \in 0^*?^*1^*$, i.e., $V_r^m(u, \varphi) \in \mathbb{B}_4^?$. \square

After having shown that every possible output of V_r^m is in $\mathbb{B}_4^?$, the next obvious question is whether V_r^m is surjective, i.e., whether every truth value $\beta \in \mathbb{B}_4^?$ is realized by some prefix $u \in \Sigma^*$ and some rLTL formula φ in the sense that $V_r^m(u, \varphi) = \beta$. Recall the motivating example above: The formula $\Box s$ realizes at least the following four truth values:

Table 2 Realizable truth values. For every truth value β , the next two columns show prefixes u and formulas φ such that $V_r^m(u, \varphi) = \beta$, or that β is unrealizable

Value	Prefix	Formula	Value	Prefix	Formula
0000	ε	$a \wedge \neg a$	0?11	$\emptyset\{a\}$	$\Box a \vee \Box \neg a$
000?	ε	$\Diamond \Box a \wedge \Diamond \neg \Diamond a$	0111	$\emptyset\{a\}$	$a \mathbf{R} a$
0001	unrealizable		????	ε	$\Box a$
00??	ε	$\Box a \wedge \Box \neg a$???1	$\{a\}$	$\Box a$
00?1	$\emptyset\{a\}$	$\Box a \wedge \Box \neg a$??11	ε	$\Box a \vee \Diamond \neg \Diamond a$
0011	unrealizable		?111	ε	$\Box a \vee \neg \Diamond \neg \Diamond \neg a$
0???	\emptyset	$\Box a$	1111	ε	$a \vee \neg a$
0??1	$\emptyset\{a\}$	$\Box a$			

???? (on ε), ???1 (on $\{s\}$), 0??? (on \emptyset), and 0??1 (on $\{s\}\emptyset$). It is not hard to convince oneself that these are all truth values realized by $\Box s$ as they represent the following four types of prefixes that can be distinguished: the prefix is empty (truth value ????), the prefix is in $\{s\}^+$ (truth value ???1), the prefix is in \emptyset^+ (truth value 0???), or the prefix contains both an $\{s\}$ and an \emptyset (truth value 0??1).

For most other truth values, it is straightforward to come up with rLTL formulas and prefixes that realize them. See Table 2 for an overview and recall Remark 2, which is applicable to all these formulas.

For others, such as 0011, it is much harder. Intuitively, to realize 0011, one needs to find an rLTL formula φ and a prefix $u \in \Sigma^*$ such that the formula obtained by replacing all \Box in φ by $\Diamond \Box$ is not satisfied by any extension of u , but the formula obtained by replacing all \Box in φ by $\Box \Diamond$ is satisfied by every extension of u .⁴ Thus, intuitively, the prefix has to differentiate between a property holding almost always and holding infinitely often. It turns out that no such u and φ exist. A similar argument is true for 0001, leading to the following theorem.

Theorem 1 *All truth values except for 0011 and 0001 are realizable.*

The unrealizability results for the truth values 0011 and 0001 are based on the following technical lemma (the reader might want to skip the proof for now and consult it at a later time).

Lemma 3 *Let φ be an rLTL formula. Then, the following holds:*

1. $V_r(u\emptyset^\omega, \varphi)[2] = V_r(u\emptyset^\omega, \varphi)[3]$ for all $u \in \Sigma^*$.
2. $V_r(u^\omega, \varphi)[3] = V_r(u^\omega, \varphi)[4]$ for all non-empty $u \in \Sigma^*$.
3. If φ does not contain the release operator, then $V_r(u^\omega, \varphi)[1] = V_r(u^\omega, \varphi)[2]$ for all non-empty $u \in \Sigma^*$.

Proof The proofs of all three items proceed by induction over the construction of φ . The induction start and the induction steps for Boolean connectives can be abstracted into the following closure property, which follows easily from the original definition of V_r in Sect. 2:

⁴ Note that this intuition breaks down in the presence of implications and negation, due to their non-standard definitions.

Let $T \subseteq \mathbb{B}_4$ contain 0000 and 1111. If $V_r(\sigma, \varphi_1)$ and $V_r(\sigma, \varphi_2)$ are in T , then so are $V_r(\sigma, p)$ for atomic propositions p , $V_r(\sigma, \neg\varphi_1)$, $V_r(\sigma, \varphi_1 \wedge \varphi_2)$, $V_r(\sigma, \varphi_1 \vee \varphi_2)$, and $V_r(\sigma, \varphi_1 \rightarrow \varphi_2)$.

Claim 1 The induction start and the induction step for the Boolean operators follow from the closure property, where we pick T to be the set of truth values from \mathbb{B}_4 whose second and third bit coincide. Furthermore, due to Remark 1, we only have to consider the inductive steps for the next, until, and release operator. All three cases rely on the following simple fact: A suffix $u\emptyset^\omega[n, \infty)$ for some n is again of the form $u'\emptyset^\omega$, i.e., the induction hypothesis is applicable to suffixes. Also, if $n \geq |u|$, then $u\emptyset^\omega[n, \infty) = \emptyset^\omega$. In particular, $u\emptyset^\omega$ has only finitely many distinct suffixes.

So, first consider a formula of the form $\varphi = \bigcirc \varphi_1$. Then, we have, for an arbitrary $u \in \Sigma^*$,

$$\begin{aligned} V_r(u\emptyset^\omega, \varphi)[2] &= V_r(u\emptyset^\omega[1, \infty), \varphi_1)[2] \\ &= V_r(u\emptyset^\omega[1, \infty), \varphi_1)[3] = V_r(u\emptyset^\omega, \varphi)[3], \end{aligned}$$

where the second equality is due to the induction hypothesis being applied to the suffix $u\emptyset^\omega[1, \infty)$.

Next, consider a formula of the form $\varphi = \varphi_1 \text{ U } \varphi_2$. Then, we have, for an arbitrary $u \in \Sigma^*$,

$$\begin{aligned} V_r(u\emptyset^\omega, \varphi)[2] &= \max_{n \geq 0} \min\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[2], \min_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[2]\} \\ &= \max_{n \geq 0} \min\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[3], \min_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\} \\ &= V_r(u\emptyset^\omega, \varphi)[3], \end{aligned}$$

where the second equality follows from an application of the induction hypothesis to the suffixes $u\emptyset^\omega[n, \infty)$ and $u\emptyset^\omega[n', \infty)$.

It remains to consider a formula of the form $\varphi = \varphi_1 \text{ R } \varphi_2$. Then, we have, for an arbitrary $u \in \Sigma^*$, that $V_r(u\emptyset^\omega, \varphi)[2]$ is by definition equal to

$$\begin{aligned} &\max_{m \geq 0} \min_{n \geq m} \max\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[2], \max_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[2]\} \\ &= \max_{m \geq 0} \min_{n \geq m} \max\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[3], \max_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\} \\ &= \max_{m \geq |u|} \min_{n \geq m} \max\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[3], \max_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\} \\ &= \max_{m \geq |u|} \min_{n \geq m} \max\{V_r(\emptyset^\omega, \varphi_2)[3], \max_{0 \leq n' \leq |u|} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\} \\ &= \max\{V_r(\emptyset^\omega, \varphi_2)[3], \max_{0 \leq n' \leq |u|} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\}, \end{aligned}$$

The first equality follows from twice applying the induction hypothesis. For the second one, observe that

$$\min_{n \geq m} \max\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[3], \max_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\}$$

is increasing in m . For the third one, note that for all $n \geq |u|$, $u\emptyset^\omega[n, \infty) = \emptyset^\omega$, which means that we have eliminated every occurrence of m and n . This explains the last equality. Similarly, $V_r(u\emptyset^\omega, \varphi)[3]$ is by definition equal to

$$\min_{m \geq 0} \max_{n \geq m} \max\{V_r(u\emptyset^\omega[n, \infty), \varphi_2)[3], \max_{0 \leq n' < n} V_r(u\emptyset^\omega[n', \infty), \varphi_1)[3]\}$$

$$= \max\{V_r(\vartheta^\omega, \varphi_2)[3], \max_{0 \leq n' \leq |u|} V_r(u\vartheta^\omega[n', \infty), \varphi_1)[3]\},$$

where the equality again follows from all suffixes $u\vartheta^\omega[n, \infty)$ with $n \geq |u|$ being equal to ϑ^ω . Thus, we have derived the desired equality between $V_r(u\vartheta^\omega, \varphi)[2]$ and $V_r(u\vartheta^\omega, \varphi)[3]$.

Claim 2 The induction start and the induction steps for Boolean operators follow from the closure property, where we here pick T to be the set of truth values from \mathbb{B}_4 whose third and fourth bit coincide. For $u = u(0) \cdots u(|u| - 1)$ and $n < |u|$, we define $\rho(u, n) = u(n) \cdots u(|u| - 1)u(0) \cdots u(n - 1)$, i.e., $\rho(u, n)$ is obtained by “rotating” u n times. The induction steps for the temporal operators are based on the following simple fact: The suffix $u^\omega[n, \infty)$ is equal to $(\rho(u, n \bmod |u|))^\omega$, i.e., the induction hypothesis is applicable to the suffixes. In particular, u^ω has only finitely many distinct suffixes, which all appear infinitely often in a cyclic order.

Now, the induction steps for the next and until operator are analogous to their counterparts in Item 1, as the only property we require there is that the induction hypothesis is applicable to suffixes. Hence, due to Remark 1, it only remains to consider the inductive step for the release operator.

So consider a formula of the form $\varphi = \varphi_1 \mathbf{R} \varphi_2$. Then, we have, for an arbitrary $u \in \Sigma^*$, that $V_r(u^\omega, \varphi)[3]$ is by definition equal to

$$\begin{aligned} & \min_{m \geq 0} \max_{n \geq m} \max\{V_r(u^\omega[n, \infty), \varphi_2)[3], \max_{0 \leq n' < n} V_r(u^\omega[n', \infty), \varphi_1)[3]\} \\ &= \min_{m \geq 0} \max_{n \geq m} \max\{V_r(u^\omega[n, \infty), \varphi_2)[4], \max_{0 \leq n' < n} V_r(u^\omega[n', \infty), \varphi_1)[4]\} \\ &= \max_{0 \leq n < |u|} \max\{V_r((\rho(u, n))^\omega, \varphi_2)[4], \max_{0 \leq n' < n} V_r((\rho(u, n'))^\omega, \varphi_1)[4]\}, \end{aligned}$$

where the first equality follows from twice applying the induction hypothesis and the second one is due to all suffixes $u^\omega[n, \infty)$ being equal to $\rho(u, n \bmod |u|)^\omega$, and that there are only finitely many, which all appear infinitely often in a cyclic order among the $(\rho(u, n))^\omega$ for $0 \leq n < |u|$.

Similarly, $V_r(u^\omega, \varphi)[4]$ is by definition equal to

$$\begin{aligned} & \max_{n \geq 0} \max\{V_r(u^\omega[n, \infty), \varphi_2)[4], \max_{0 \leq n' < n} V_r(u^\omega[n', \infty), \varphi_1)[4]\} \\ &= \max_{0 \leq n < |u|} \max\{V_r((\rho(u, n))^\omega, \varphi_2)[4], \max_{0 \leq n' < n} V_r((\rho(u, n'))^\omega, \varphi_1)[4]\}, \end{aligned}$$

where the equality again follows from all suffixes $u^\omega[n, \infty)$ being equal to $\rho(u, n \bmod |u|)^\omega$, and that there are only finitely many, which appear in a cyclic order: In particular, after the first $|u|$ suffixes, we have seen all of them. Thus, we have derived the desired equality between $V_r(u^\omega, \varphi)[3]$ and $V_r(u^\omega, \varphi)[4]$.

Claim 3 The induction start and the induction steps for Boolean operators are covered by the closure property, where we here pick T to be the set of truth values from \mathbb{B}_4 whose first and second bit coincide. The cases of the next and until operator are again analogous to the first and second item. Hence, we only have to consider the inductive step for the always operator, as we here only consider formulas without release.

So, consider a formula of the form $\varphi = \Box \varphi_1$. Here, we again rely on the fact that the suffix $u^\omega[n, \infty)$ is equal to $(\rho(u, n \bmod |u|))^\omega$. By definition, $V_r(u^\omega, \varphi)[1]$ is equal to

$$\min_{n \geq 0} V_r(u^\omega[n, \infty), \varphi_1)[1] = \min_{n \geq 0} V_r(u^\omega[n, \infty), \varphi_1)[2] = \min_{0 \leq n < |u|} V_r((\rho(u, n))^\omega, \varphi_1)[2],$$

where the first equality is due to the induction hypothesis and the second one due to the fact that u^ω has only finitely many suffixes, which are all already realized by some $u^\omega[n, \infty)$ for $0 \leq n < |u|$.

Similarly, $V_r(u^\omega, \varphi)[2]$ is by definition equal to

$$\begin{aligned} \max_{m \geq 0} \min_{n \geq m} V_r(u^\omega[n, \infty), \varphi_1)[1] &= \max_{m \geq 0} \min_{n \geq m} V_r(u^\omega[n, \infty), \varphi_1)[2] \\ &= \min_{0 \leq n < |u|} V_r((\rho(u, n))^\omega, \varphi_1)[2], \end{aligned}$$

where the two equalities follow as before: the first by induction hypothesis and the second one by the fact that u^ω has only finitely many suffixes, which all appear infinitely often in a cyclic order and which are all already realized by some $u^\omega[n, \infty)$ for $0 \leq n < |u|$. Thus, we have derived the desired equality between $V_r(u^\omega, \varphi)[1]$ and $V_r(u^\omega, \varphi)[2]$. \square

Now, we are able to prove Theorem 1.

Proof We begin by showing that 0011 and 0001 are not realizable.

First, towards a contradiction, assume there is an rLTL formula φ and a prefix u such that $V_r^m(u, \varphi) = 0011$, i.e., for every extension $u\sigma$, we have $V_r(u\sigma)[2] = 0$ and $V_r(u\sigma)[3] = 1$. However, by picking $\sigma = \emptyset^\omega$ we obtain the desired contradiction to Lemma 3.1.

The proof for 0001 is similar. Assume there is an rLTL formula φ and a prefix u such that $V_r^m(u, \varphi) = 0001$. Due to Lemma 4, we can assume that u is non-empty. Thus, we have $V_r(u^\omega, \varphi) = 0001$ by definition of V_r^m , which contradicts Lemma 3.2.

Finally, applying Lemma 3.3, one can show that no rLTL formula without the release operator realizes 0111. However, we show below that it is realizable by a formula with the release operator.

Next, we show that every other truth value $\beta \notin \{0011, 0001\}$ is indeed realizable. The witnessing pairs of prefixes and formulas are presented in Table 2.

First, consider $\beta = 0111$ with prefix $u = \emptyset\{a\}$ and formula $\varphi = a \mathbf{R} a$. We have $\text{ltl}(1, \varphi) = a \mathbf{R} a$ and $\text{ltl}(2, \varphi) = \diamond \square a \vee \diamond a$. Note that $a \mathbf{R} a$ is violated by $u\sigma$, for every $\sigma \in \Sigma^\omega$. Dually, $\diamond \square a \vee \diamond a$ is satisfied by $u\sigma$, for every $\sigma \in \Sigma^\omega$. Hence, for arbitrary $\sigma \in \Sigma^\omega$, we have $V_r(u\sigma, \varphi)[1] = 0$ and $V_r(u\sigma, \varphi)[2] = 1$. Hence, we have $V_r(u\sigma, \varphi) = 0111$ for every σ , as this is the only truth value that matches this pattern. Hence, by definition, we obtain $V_r^m(u, \varphi) = 0111$.

The verification for all other truth values is based on Remark 2, which is applicable to all formulas φ in the third column witnessing the realization of a truth value $\beta \neq 0111$. Now, for every such truth value β and corresponding pair (u, φ) , one can easily verify the following:

- If $\beta[i] = 0$, then no $u\sigma$ satisfies $\text{ltl}(i, \varphi)$.
- If $\beta[i] = 1$, then every $u\sigma$ satisfies $\text{ltl}(i, \varphi)$.
- If $\beta[i] = ?$, then there are σ, σ' such that $u\sigma$ satisfies $\text{ltl}(i, \varphi)$ and such that $u\sigma'$ violates $\text{ltl}(i, \varphi)$. In all such cases, $\sigma, \sigma' \in \{\emptyset^\omega, \{a\}^\omega, \{a\}\emptyset^\omega, \emptyset\{a\}^\omega, (\{a\}\emptyset)^\omega\}$ suffice.

We leave the details of this slightly tedious, but trivial, verification to the reader. \square

As shown in Table 2, all of the realizable truth values except for 0111 are realized by formulas using only conjunction, disjunction, negation, eventually, and always. Further, 0111 can only be realized by a formula with the release operator while the truth values 0011 and 0001 are indeed not realizable at all.

Note that the two unrealizable truth values 0011 and 0001 both contain a 0 that is directly followed by a 1. The proof of unrealizability formalizes the intuition that such an “abrupt”

transition from definitive violation of a property to definitive satisfaction of the property cannot be witnessed by any finite prefix. Finally, the only other truth value of this form, 0111, is only realizable by using a formula with the release operator.

Going again back to the motivating example $\square s$, consider the evolution of the truth values on the sequence $\varepsilon, \{s\}, \{s\}\emptyset$: They are $????, ???1$, and $0???1$, i.e., 0’s and 1’s are stable when extending a prefix, only a ? may be replaced by a 0 or a 1. This property holds in general. To formalize this, say that $\beta' \in \mathbb{B}_4^2$ is more specific than $\beta \in \mathbb{B}_4^2$, written as $\beta \leq \beta'$, if, for all $i, \beta[i] \neq ?$ implies $\beta'[i] = \beta[i]$.

Lemma 4 *Let φ be an rLTL formula and $u, u' \in \Sigma^*$. If $u \sqsubseteq u'$, then $V_r^m(u, \varphi) \leq V_r^m(u', \varphi)$.*

Proof Let $u \sqsubseteq u'$ and assume we have $V_r^m(u, \varphi)[i] \in \{0, 1\}$. Thus, by definition, $V_r(u\sigma, \varphi)[i] = V_r^m(u, \varphi)[i]$ for every $\sigma \in \Sigma^\omega$. Now, as u is a prefix of u' , we can decompose u' into $u' = uv$ for some $v \in \Sigma^*$ and every extension $u'\sigma'$ of u' is the extension $uv\sigma'$ of u . Hence, we have $V_r(u'\sigma', \varphi)[i] = V_r(uv\sigma', \varphi)[i] = V_r^m(u, \varphi)[i]$ for every $\sigma' \in \Sigma^\omega$. Thus, $V_r^m(u', \varphi)[i] = V_r^m(u, \varphi)[i]$.

As this property holds for every i , we obtain $V_r^m(u, \varphi) \leq V_r^m(u', \varphi)$. □

Let us discuss two properties of the semantics: *impartiality* and *anticipation* [20]. Impartiality states that a definitive verdict will never be revoked: If $V_r^m(u, \varphi)[i] \neq ?$, then for all finite extensions $v \in \Sigma^*$, the verdict will not change, so $V_r^m(uv, \varphi)[i] = V_r^m(u, \varphi)[i]$. This property follows immediately from Lemma 4. Anticipation requires that a definitive verdict is decided as soon as possible, i.e., if $V_r^m(u, \varphi)[i] = ?$, then u can still be extended to satisfy and to violate φ with the i -th bit. Formally, there have to exist infinite extensions σ_0 and σ_1 such that $V_r(u\sigma_0, \varphi)[i] = 0$ and $V_r(u\sigma_1, \varphi)[i] = 1$. Anticipation holds by definition of $V_r^m(u, \varphi)$.

Due to Lemma 4, for a fixed formula, the prefixes of every infinite word can assume at most five different truth values, which are all of increasing specificity. It turns out that this upper bound is tight. To formalize this claim, we denote the strict version of \leq by $<$, i.e., $\beta < \beta'$ if and only if $\beta \leq \beta'$ and $\beta \neq \beta'$.

Lemma 5 *There is an rLTL formula φ and prefixes $u_0 \sqsubset u_1 \sqsubset u_2 \sqsubset u_3 \sqsubset u_4$ such that $V_r^m(u_0, \varphi) < V_r^m(u_1, \varphi) < V_r^m(u_2, \varphi) < V_r^m(u_3, \varphi) < V_r^m(u_4, \varphi)$.*

Proof Consider the sequence β_0, \dots, β_4 with $\beta_j = 0^j ?^{4-j}$ and note that we have $\beta_j < \beta_{j+1}$ for every $j < 4$. Furthermore, let $u_j = \emptyset^j$ for $j \in \{0, \dots, 4\}$. We construct a formula φ such that $V_r^m(u_j, \varphi) = \beta_j$ for every $j \in \{0, \dots, 4\}$.

To this end, let

- $\psi_{\beta_1} = \diamond (a \wedge \square \neg \diamond a)$,
- $\psi_{\beta_2} = \square (a \wedge \odot \neg a) \wedge \neg \diamond \neg \diamond a$, and
- $\psi_{\beta_3} = \diamond \square a \wedge \diamond \neg \diamond a$.

Later, we rely on the following fact about these formulas, which can easily be shown by applying Remark 2: we have $V_r^m(u, \psi_{\beta_j}) = \beta_j$ for every prefix u .

Further, for $j \in \{0, 1, 2, 3\}$, let ψ_j be a formula that requires the proposition a to be violated at the first $j - 1$ positions, but to hold at the j -th position (recall that we start counting at zero), i.e., $\psi_j = (\bigwedge_{0 \leq j' < j} \odot^{j'} \neg a) \wedge \odot^j a$. Here, we define the nesting of next operators as usual: $\odot^0 \xi = \xi$ and $\odot^{j+1} \xi = \odot \odot^j \xi$. By definition, we have $V_r(\emptyset^{j+1}\sigma, \psi_j) = 0000$ for every $\sigma \in \Sigma^\omega$ (\dagger).

Now, we define

$$\varphi = \psi_0 \vee \bigvee_{j=1}^3 (\psi_{\beta_j} \wedge \psi_j)$$

and claim that it has the desired properties. To this end, we note that property (\dagger) implies $V_r(\emptyset^4\sigma, \varphi) = 0000$ for every $\sigma \in \Sigma^\omega$ $(\dagger\dagger)$, as every disjunct of φ contains a conjunct of the form ψ_j for some $j \leq 3$. Also, let us mention that Remark 2 is applicable to φ .

It remains to prove $V_r^m(u_j, \varphi) = \beta_j$ for every $j \in \{0, \dots, 4\}$.

- For $j = 0$, we have $u_0 = \varepsilon$ and $\beta_0 = 0000$. Hence, it suffices to present $\sigma_0, \sigma_1 \in \Sigma^\omega$ such that $V_r(\sigma_0, \varphi) = 0000$ and $V_r(\sigma_1, \varphi) = 1111$.

Due to property $(\dagger\dagger)$, we can pick $\sigma_0 = \emptyset^\omega$. To conclude, we pick $\sigma_1 = \{a\}^\omega$, as we have

$$V_r(\sigma_1, \varphi) \geq V_r(\sigma_1, \psi_0) = V_r(\{a\}^\omega, a) = 1111,$$

where the first inequality follows from ψ_0 being a disjunct of φ .

- For $j = 1$, we have $u_1 = \emptyset$ and $\beta_1 = 0111$. To show $V_r^m(u_1, \varphi) = \beta_1$, it suffices to present $\sigma_0, \sigma_1 \in \Sigma^\omega$ such that $V_r(u_1\sigma_0, \varphi) = 0000$, $V_r(u_1\sigma_1, \varphi) = 0111$, and show that $V_r(u_1\sigma, \varphi)[1] = 0$ for every $\sigma \in \Sigma^\omega$. First, we again pick $\sigma_0 = \emptyset^\omega$ due to property $(\dagger\dagger)$. Now, consider $\sigma_1 = \{a\}^\omega$. Then,

$$\begin{aligned} V_r(u_1\{a\}^\omega, \psi_{\beta_j} \wedge \psi_j) &= \min\{V_r(u_1\{a\}^\omega, \psi_{\beta_j}), V_r(u_1\{a\}^\omega, \psi_j)\} \\ &= \min\{0111, 1111\} = 0111, \end{aligned}$$

where $V_r(u_1\{a\}^\omega, \psi_{\beta_j}) = 0111$ can easily be verified using Remark 2. To conclude, using Remark 2, one can easily verify that $\text{ltl}(1, \varphi)$ is not satisfied by $u_1\sigma$ for any $\sigma \in \Sigma^\omega$.

- The reasoning for $j = 2, 3$ is along the same lines as the one for $j = 1$ and is left to the reader.
- For $j = 4$, we have $u_4 = \emptyset\emptyset\emptyset\emptyset$ and $\beta_4 = 0000$. Hence, our claim follows directly from property $(\dagger\dagger)$, which shows $V_r(u_4\sigma, \varphi) = 0000$ for every $\sigma \in \Sigma^\omega$.

□

After determining how many different truth values can be assumed by prefixes of a single infinite word, an obvious question is how many truth values can be realized by a fixed formula on *different* prefixes. It is not hard to combine the formulas in Table 2 to a formula that realizes all truth values not ruled out by Theorem 1.⁵

Lemma 6 *There is an rLTL formula φ such that for every $\beta \in \mathbb{B}_4^2 \setminus \{0011, 0001\}$ there is a prefix u_β with $V_r^m(u_\beta, \varphi) = \beta$.*

Proof For every $\beta \in \mathbb{B}_4^2 \setminus \{0011, 0001\}$ let φ_β be an rLTL formula and u'_β be a prefix, both over $\{a\}$, with $V_r^m(u'_\beta, \varphi_\beta) = \beta$. Such formulas and prefixes exist as shown in Table 2.

Now, consider the formula

$$\varphi = \bigvee_{\beta \in \mathbb{B}_4^2 \setminus \{0011, 0001\}} a_\beta \wedge \varphi_\beta$$

⁵ Note that there are formulas in publicly available repositories that assume *many* truth values. One example is the formula

$$(((a \wedge d) \vee (\neg a \wedge \neg d)) \wedge \square(\neg b \vee (\neg a \wedge d))) \vee (((\neg a \wedge d) \vee (a \wedge \neg d)) \wedge \diamond(b \wedge (a \vee \neg d))) \vee (a \wedge \square b),$$

which is taken from the LTLStore [38] and assumes ten different truth values.

over the propositions $\{a\} \cup \{a_\beta \mid \beta \in \mathbb{B}_4^2 \setminus \{0011, 0001\}\}$.

By construction, we have $V_r^m(u_\beta, \varphi) = \beta$ for every β , where

$$u_\beta = (u'_\beta(0) \cup \{a_\beta\})u'_\beta(1) \cdots u'_\beta(|u'_\beta| - 1),$$

i.e., we obtain u_β from u'_β by adding the proposition a_β to the first letter. Hence, φ has the desired properties. □

Finally, let us consider the notion of *monitorability* [53], an important concept in the theory of runtime monitoring. As a motivation, consider the LTL formula $\psi = \Box \Diamond s$ and an arbitrary prefix $u \in \Sigma^*$. Then, the extension $u\{s\}^\omega$ satisfies ψ while the extension $u\emptyset^\omega$ does not satisfy ψ , i.e., satisfaction of ψ is independent of any prefix u . Hence, we have $V^m(u, \psi) = ?$ for every prefix u , i.e., monitoring the formula ψ does not generate any information.

In general, for a fixed LTL formula φ , a prefix $u \in \Sigma^*$ is called *ugly* if we have $V^m(uv, \varphi) = ?$ for every finite $v \in \Sigma^*$, i.e., every finite extension of u yields an indefinite verdict.⁶ Now, φ is *LTL-monitorable* if there is no ugly prefix with respect to φ . A wide range of LTL formulas (e.g., $\psi = \Box \Diamond s$ as above) are unmonitorable in that sense. In particular, 44% of the LTL formulas considered in the experiments of Bauer et al. are not LTL-monitorable.

We next generalize the notion of monitorability to rLTL. In particular, we answer whether there are unmonitorable rLTL formulas. Then, in Sect. 5, we exhibit that all LTL formulas considered by Bauer et al.’s experimental evaluation, even the unmonitorable ones, are monitorable under rLTL semantics. To conclude the motivating example, note that the rLTL analogue $\Box \Diamond s$ of the LTL formula ψ induces two truth values from \mathbb{B}_4^2 indicating whether s has been true at least once (truth value $???1$) or not (truth value $????$). Even more so, every prefix inducing the truth value $????$ can be extended to one inducing the truth value $???1$.

Definition 3 Let φ be an rLTL formula. A prefix $u \in \Sigma^*$ is called *ugly* if we have $V^m(uv, \varphi) = ???? for every finite $v \in \Sigma^*$. Further, φ is *rLTL-monitorable* if it has no ugly prefix.$

As we have argued above, the formula $\Box \Diamond s$ has no ugly prefix, i.e., it is rLTL-monitorable. Thus, we have found an unmonitorable LTL formula whose rLTL analogue (the formula obtained by adding dots to all temporal operators) is monitorable. The converse statement is also true. There is a monitorable LTL formula whose rLTL analogue is unmonitorable. To this end, consider the LTL formula

$$(\Box s \wedge \Box \neg s) \rightarrow (\Diamond \Box s \wedge \Diamond \neg \Diamond s),$$

which is a tautology and therefore monitorable. On the other hand, we claim that $\emptyset\{s\}$ is an ugly prefix for the rLTL analogue φ obtained by adding dots to the temporal operators. To this end note that we have both $V_r(\emptyset\{s\}v\emptyset^\omega, \varphi) = 1111$ and $V_r(\emptyset\{s\}v\{s\}^\omega, \varphi) = 0000$ for every $v \in \Sigma^*$. Hence, $V_r^m(\emptyset\{s\}v, \varphi) = ???? for every such v , i.e., $\emptyset\{s\}$ is indeed ugly and φ therefore not rLTL-monitorable.$

Thus, there are formulas that are unmonitorable under LTL semantics, but monitorable under rLTL semantics and there are formulas that are unmonitorable under rLTL semantics, but monitorable under LTL semantics. Using these formulas one can also construct a formula that is unmonitorable under both semantics.

⁶ Note that the good/bad prefixes introduced by Kupfermann and Vardi [40] can only be extended into infinite words satisfying/unsatisfying the formula, respectively, and thus provide a verdict immediately. On the other hand, no finite extension of an ugly prefix [14] allows to conclude on the satisfaction of the formula.

To this end, fix LTL formulas φ_ℓ and φ_r over disjoint sets of propositions and a fresh proposition p not used in either formula such that

- φ_ℓ has an ugly prefix u_ℓ under LTL semantics, and
- φ_r (with dotted operators) has an ugly prefix u_r under rLTL semantics.

We can assume both prefixes to be non-empty, as ugliness is closed under finite extensions. Let $\varphi = (p \wedge \varphi_\ell) \vee (\neg p \wedge \varphi_r)$. Then, the prefix obtained from u_ℓ by adding the proposition p to the first letter is ugly for φ under LTL semantics and u_r is ugly for φ (with dotted operators) under rLTL semantics.

As a final example, recall that we have shown that $\Box \Diamond s$ is rLTL-monitorable and consider its negation $\neg \Box \Diamond s$. It is not hard to see that $V_r^m(u, \varphi) = \text{????}$ holds for every prefix u . Hence, ε is an ugly prefix for the formula, i.e., we have found another unmonitorable rLTL formula. In particular, the example shows that, unlike for LTL, rLTL-monitorability is not preserved under negation.

After having studied properties of rLTL monitorability, we next show our main result: The robust monitoring semantics V_r^m can be implemented by finite-state machines.

4 Construction of rLTL monitors

An rLTL monitor is an implementation of the robust monitoring semantics V_r^m in form of a finite-state machine with output. More precisely, an *rLTL monitor* for an rLTL formula φ is a finite-state machine \mathcal{M}_φ that on reading an input $u \in \Sigma^*$ outputs $V_r^m(u, \varphi)$. In this section, we show how to construct rLTL monitors and that this construction is asymptotically not more expensive than the construction of LTL monitors. Let us fix an rLTL formula φ for the remainder of this section.

Our rLTL monitor construction is inspired by Bauer et al. [14] and generates a sequence of finite-state machines (i.e., Büchi automata over infinite words, (non)deterministic automata over finite words, and Moore machines). Underlying these machines are *transition structures* $\mathcal{T} = (Q, q_I, \Delta)$ consisting of a nonempty, finite set Q of states, an initial state $q_I \in Q$, and a transition relation $\Delta \subseteq Q \times \Sigma \times Q$. An (infinite) run of \mathcal{T} on a word $\sigma = a_0a_1a_2 \dots \in \Sigma^\omega$ is a sequence $\rho = q_0q_1 \dots$ of states such that $q_0 = q_I$ and $(q_j, a_j, q_{j+1}) \in \Delta$ for $j \in \mathbb{N}$. Finite runs on finite words are defined analogously. The transition structure \mathcal{T} is *deterministic* if (a) $(q, a, q') \in \Delta$ and $(q, a, q'') \in \Delta$ imply $q' = q''$ and (b) for each $q \in Q$ and $a \in \Sigma$ there exists a $(q, a, q') \in \Delta$.

We then replace the transition relation Δ by a function $\delta: Q \times \Sigma \rightarrow Q$. Finally, we define the *size* of a transition structure \mathcal{T} as $|\mathcal{T}| = |Q|$ in order to measure its complexity.

Our construction then proceeds in three steps:

1. We bring φ into an operational form by constructing Büchi automata $\mathcal{A}_\beta^\varphi$ for each truth value $\beta \in \mathbb{B}_4$ that can decide the valuation $V_r(\sigma, \varphi)$ of infinite words $\sigma \in \Sigma^\omega$.
2. Based on these Büchi automata, we then construct nondeterministic automata $\mathcal{B}_\beta^\varphi$ that can decide whether a finite word $u \in \Sigma^*$ can still be extended to an infinite word $u\sigma \in \Sigma^\omega$ with $V_r(u\sigma, \varphi) = \beta$.
3. We determinize the nondeterministic automata obtained in Step 2 and combine them into a single Moore machine that computes $V_r^m(u, \varphi)$.

Let us now describe each of these steps in detail.

Step 1 We first translate the rLTL formula φ into several Büchi automata using a construction by Tabuada and Neider [58], summarized in Theorem 2 below. A (*nondeterministic*) *Büchi*

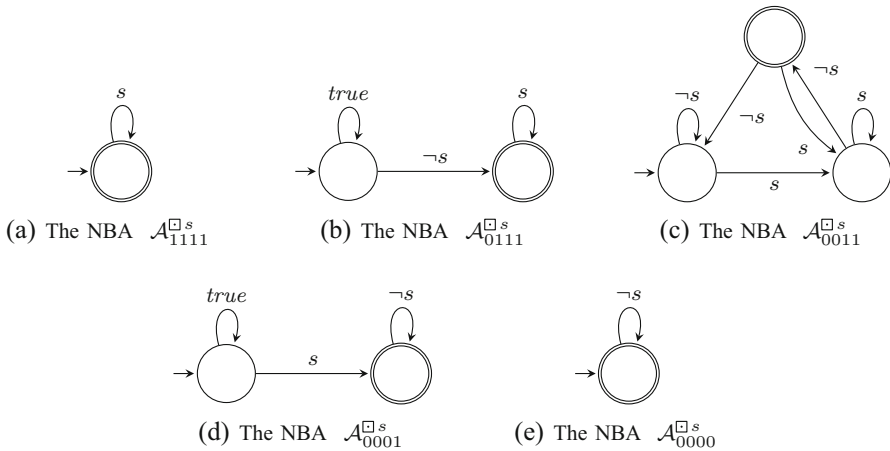


Fig. 1 The Büchi automata $\mathcal{A}_\beta^{\Box s}$ constructed in Step 1 of our monitor construction

automaton (NBA) is a four-tuple $\mathcal{A} = (Q, q_I, \Delta, F)$ where $\mathcal{T} = (Q, q_I, \Delta)$ is a transition structure and $F \subseteq Q$ is a set of accepting states. A run π of \mathcal{A} on $\sigma \in \Sigma^\omega$ is a run of \mathcal{T} on σ , and we say that π is accepting if it contains infinitely many states from F . The automaton \mathcal{A} accepts a word σ if there exists an accepting run of \mathcal{A} on σ . The language $\mathcal{L}(\mathcal{A})$ is the set of all words accepted by \mathcal{A} , and the size of \mathcal{A} is defined as $|\mathcal{A}| = |\mathcal{T}|$.

Theorem 2 (Tabuada and Neider [58]) *Given a truth value $\beta \in \mathbb{B}_4$, one can construct a Büchi automaton $\mathcal{A}_\beta^\varphi$ with $2^{\mathcal{O}(|\varphi|)}$ states such that $\mathcal{L}(\mathcal{A}_\beta^\varphi) = \{\sigma \in \Sigma^\omega \mid V_r(\sigma, \varphi) = \beta\}$. This construction can be performed in $2^{\mathcal{O}(|\varphi|)}$ time.*

The Büchi automata $\mathcal{A}_\beta^\varphi$ for $\beta \in \mathbb{B}_4$ serve as building blocks for the next steps. However, before we proceed, let us illustrate this step with an example.

Example 4 Let us consider the formula $\varphi = \Box s$, which already served as a running example in Sect. 3. Applying Theorem 2 results in the five nondeterministic Büchi automata $\mathcal{A}_\beta^\varphi$, one for each $\beta \in \mathbb{B}_4$, shown in Fig. 1. We here use the standard way to represent finite-state machines graphically. States are drawn as circles and transitions are drawn as arrows. Moreover, the initial state has an incoming arrow, while accepting states are indicated by double circles. Finally, note that we use propositional formulas to symbolically define sets of transitions. For instance, a transition labeled with s in Fig. 1a represents all transitions labeled with a symbol from the set $\{A \subseteq P \mid s \in A\} \subseteq \Sigma$. In particular, *true* represents all symbols in Σ . \square

Step 2 For each Büchi automaton $\mathcal{A}_\beta^\varphi$ obtained in the previous step, we now construct a nondeterministic automaton $\mathcal{B}_\beta^\varphi$ over finite words. This automaton determines whether a finite word $u \in \Sigma^*$ can be continued to an infinite word $u\sigma \in \mathcal{L}(\mathcal{A}_\beta^\varphi)$ (i.e., $V_r(u\sigma, \varphi) = \beta$) and is used later to construct the rLTL monitor.

A *nondeterministic finite automaton (NFA)* is a four-tuple $\mathcal{A} = (Q, q_I, \Delta, F)$ that is syntactically identical to a Büchi automaton. The size of \mathcal{A} is defined analogously to Büchi automata. In contrast to Büchi automata, however, NFAs only admit finite runs on finite words, i.e., a run of \mathcal{A} on $u = a_0 \cdots a_{n-1} \in \Sigma^*$ is a sequence $q_0 \cdots q_n$ such that $q_0 = q_I$




Fig. 2 The NFAs $\mathcal{B}_\beta^{\square s}$ constructed in Step 2 of our monitor construction

and $(q_j, a_j, q_{j+1}) \in \Delta$ for every $j < n$. A run $q_0 \dots q_n$ is called *accepting* if $q_n \in F$. Accepted words as well as the language of \mathcal{A} are again defined analogously to the Büchi case. If (Q, q_I, Δ) is deterministic, \mathcal{A} is a *deterministic finite automaton (DFA)*. It is well-known that for each NFA \mathcal{A} one can construct a DFA \mathcal{A}' with $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$ and $|\mathcal{A}'| \in \mathcal{O}(2^{|\mathcal{A}|})$.

Given the Büchi automaton $\mathcal{A}_\beta^\varphi = (Q_\beta, q_{I,\beta}, \Delta_\beta, F_\beta)$, we first compute the set $F_\beta^* = \{q \in Q_\beta \mid \mathcal{L}(\mathcal{A}_\beta^\varphi(q)) \neq \emptyset\}$, where $\mathcal{A}_\beta^\varphi(q)$ denotes the Büchi automaton $\mathcal{A}_\beta^\varphi$ but with initial state q instead of q_I . Intuitively, the set F_β^* contains all states $q \in Q_\beta$ from which there exists an accepting run in $\mathcal{A}_\beta^\varphi$ and, hence, indicates whether a finite word $u \in \Sigma^*$ reaching a state of F_β^* can be extended to an infinite word $u\sigma' \in \mathcal{L}(\mathcal{A}_\beta^\varphi)$. The set F_β^* can be computed, for instance, using a nested depth-first search [56] for each state $q \in Q_\beta$. Since each such search requires time quadratic in $|\mathcal{A}_\beta^\varphi|$, the set F_β^* can be computed in time $\mathcal{O}(|\mathcal{A}_\beta^\varphi|^3)$.

Using F_β^* , we define the NFA $\mathcal{B}_\beta^\varphi = (Q_\beta, q_{I,\beta}, \Delta_\beta, F_\beta^*)$. It shares the transition structure of $\mathcal{A}_\beta^\varphi$ and uses F_β^* as the set of accepting states. Let us illustrate this construction using our running example.

Example 5 Given the NBAs $\mathcal{A}_\beta^\varphi$ from Step 1 of our construction, we now compute the corresponding NFAs $\mathcal{B}_\beta^\varphi$, which are depicted in Fig. 2. Note that the transition structure has remained the same as compared to the preceding step (see Fig. 1). By contrast, the accepting states have changed according to the definition of F_β^* , causing all states to be accepting. Note, however, that this does not mean that the resulting NFAs accept any finite word. For instance, the NFA $\mathcal{B}_{1111}^{\square s}$ in Fig. 2a is a counterexample to this claim. \square

The next lemma now states that $\mathcal{B}_\beta^\varphi$ indeed recognizes prefixes of words in $\mathcal{L}(\mathcal{A}_\beta^\varphi)$.

Lemma 7 *Let $\beta \in \mathbb{B}_4$ and $u \in \Sigma^*$. Then, $u \in \mathcal{L}(\mathcal{B}_\beta^\varphi)$ if and only if there exists an infinite word $\sigma \in \Sigma^\omega$ with $V_r(u\sigma, \varphi) = \beta$.*

Proof We show both directions separately.

From left to right Assume $u \in \mathcal{L}(\mathcal{B}_\beta^\varphi)$. Moreover, let $q \in F_\beta^*$ be the accepting state reached by $\mathcal{B}_\beta^\varphi$ on an accepting run on u (which exists since $u \in \mathcal{L}(\mathcal{B}_\beta^\varphi)$). By definition of F_β^* , this

means that $\mathcal{L}(\mathcal{A}_\beta^\varphi(q)) \neq \emptyset$, say $\sigma \in \mathcal{L}(\mathcal{A}_\beta^\varphi(q))$. Since $\mathcal{A}_\beta^\varphi$ and $\mathcal{B}_\beta^\varphi$ share the same transition structures, the run of $\mathcal{B}_\beta^\varphi$ on u is also a run of $\mathcal{A}_\beta^\varphi$ on u , which both lead to state q . Therefore, $u\sigma \in \mathcal{L}(\mathcal{A}_\beta^\varphi)$. By Theorem 2, this is equivalent to $V_r(u\sigma, \varphi) = \beta$.

From right to left Let $u \in \Sigma^*$ and $\sigma \in \Sigma^\omega$ such that $V(u\sigma, \varphi) = \beta$. By Theorem 2, we have $u\sigma \in \mathcal{L}(\mathcal{A}_\beta^\varphi)$. Consider an accepting run of $\mathcal{A}_\beta^\varphi$ on $u\sigma$, and let q be the state that $\mathcal{A}_\beta^\varphi$ reaches after reading the finite prefix u . Since $u\sigma \in \mathcal{L}(\mathcal{A}_\beta^\varphi)$, this means that $\sigma \in \mathcal{L}(\mathcal{A}_\beta^\varphi(q))$. Thus, $q \in F_\beta^*$ because $\mathcal{L}(\mathcal{A}_\beta^\varphi(q)) \neq \emptyset$. Moreover, since the run of $\mathcal{A}_\beta^\varphi$ on u is also a run of $\mathcal{B}_\beta^\varphi$ on u , the NFA $\mathcal{B}_\beta^\varphi$ can also reach state q after reading u . Therefore, $u \in \mathcal{L}(\mathcal{B}_\beta^\varphi)$ since $q \in F_\beta^*$. \square

Before we continue to the last step in our construction, let us briefly comment on the complexity of computing the NFAs $\mathcal{B}_\beta^\varphi$. Since $\mathcal{B}_\beta^\varphi$ and $\mathcal{A}_\beta^\varphi$ share the same underlying transition structure, we immediately obtain $|\mathcal{B}_\beta^\varphi| \in 2^{O(|\varphi|)}$. Moreover, the construction of $\mathcal{B}_\beta^\varphi$ is dominated by the computation of the set F_β^* and, hence, can be done in time $2^{O(|\varphi|)}$.

Step 3 In the final step, we construct a Moore machine implementing an rLTL monitor for φ . Formally, a *Moore machine* is a five-tuple $\mathcal{M} = (Q, q_I, \delta, \Gamma, \lambda)$ consisting of a deterministic transition structure (Q, q_I, δ) , an output alphabet Γ , and an output function $\lambda: Q \rightarrow \Gamma$. The size of \mathcal{M} as well of runs of \mathcal{M} are defined as for DFAs. In contrast to a DFA, however, a Moore machine \mathcal{M} computes a function $\lambda_{\mathcal{M}}: \Sigma^* \rightarrow \Gamma$ that is defined by $\lambda_{\mathcal{M}}(u) = \lambda(q_n)$ where q_n is the last state reached on the unique finite run $q_0 \cdots q_n$ of \mathcal{M} on its input $u \in \Sigma^*$.

The first step in the construction of the Moore machine is to determinize the NFAs $\mathcal{B}_\beta^\varphi$, obtaining equivalent DFAs $\mathcal{C}_\beta^\varphi = (Q'_\beta, q'_{I,\beta}, \delta'_\beta, F'_\beta)$ of at most exponential size in $|\mathcal{B}_\beta^\varphi|$. Subsequently, we combine these DFAs into a single Moore machine \mathcal{M}_φ implementing the desired rLTL monitor. Intuitively, this Moore machine is the product of the DFAs $\mathcal{C}_\beta^\varphi$ for each $\beta \in \mathbb{B}_4$ and tracks the run of each individual DFA on the given input. Formally, \mathcal{M}_φ is defined as follows.

Definition 4 Let $\mathbb{B}_4 = \{\beta_1, \beta_2, \beta_3, \beta_4, \beta_5\}$. We define $\mathcal{M}_\varphi = (Q, q_I, \Gamma, \delta, \lambda)$ by

- $Q = Q'_{\beta_1} \times Q'_{\beta_2} \times Q'_{\beta_3} \times Q'_{\beta_4} \times Q'_{\beta_5}$;
- $q_I = (q'_{1,\beta_1}, q'_{1,\beta_2}, q'_{1,\beta_3}, q'_{1,\beta_4}, q'_{1,\beta_5})$;
- $\delta((q_1, q_2, q_3, q_4, q_5), a) = (q'_1, q'_2, q'_3, q'_4, q'_5)$ where $q'_j = \delta'_{\beta_j}(q_j, a)$ for each $j \in \{1, \dots, 5\}$;
- $\Gamma = \mathbb{B}_4^2$; and
- $\lambda((q_1, q_2, q_3, q_4, q_5)) = \xi(\{\beta_j \in \mathbb{B}_4 \mid q_j \in F'_{\beta_j}, j \in \{1, \dots, 5\}\})$,

where the surjective function $\xi: 2^{\mathbb{B}_4} \rightarrow \mathbb{B}_4^2$ translates sets $B \subseteq \mathbb{B}_4$ of truth values to the robust monitoring semantics as follows: $\xi(B) = \beta^? \in \mathbb{B}_4^2$ with

$$\beta^?[j] = \begin{cases} 0 & \text{if } \beta[j] = 0 \text{ for each } \beta \in B; \\ 1 & \text{if } \beta[j] = 1 \text{ for each } \beta \in B; \text{ and} \\ ? & \text{otherwise.} \end{cases}$$

Let us illustrate this last step of our construction by means of our running example.

Example 6 Given the NFAs $\mathcal{B}_\beta^{\square s}$ from Step 2 of our construction, we first apply a standard determinization step. This process results in equivalent DFAs $\mathcal{C}_\beta^{\square s}$, which are shown in Fig. 3.

The final, minimized monitor $\mathcal{M}_{\square s}$, which results from the Cartesian product of all DFAs, is shown in Fig. 4. Note that this monitor has four different verdicts, shown as labels




Fig. 3 The DFAs $C_{\beta}^{\square s}$ constructed in Step 3 of our monitor construction




Fig. 4 The final monitor $\mathcal{M}_{\square s}$

next to each state. These are four of the verdicts used to prove results in Table 2 (on Page 13). □

The main result of this paper now shows that the Moore machine \mathcal{M}_{φ} implements V_r^m , i.e., we have $\lambda_{\mathcal{M}_{\varphi}}(u) = V_r^m(u, \varphi)$ for every prefix u .

Theorem 3 For every rLTL formula φ , one can construct an rLTL monitor of size $2^{2^{\mathcal{O}(|\varphi|)}}$.

Proof First observe that ξ indeed produces a valid value of \mathbb{B}_4^2 (i.e., a truth value of the form $0^*?^*1^*$). This follows immediately from the definition of ξ and the fact that the truth values of rLTL are sequences in 0^*1^* .

Next, we observe that \mathcal{M}_{φ} reaches state $(q_1, q_2, q_3, q_4, q_5)$ after reading a word $u \in \Sigma^*$ if and only if for each $\beta_j \in \mathbb{B}_4$ the DFA $C_{\beta_j}^{\varphi}$ reaches state q_j after reading u . A simple induction over the length of inputs fed to \mathcal{M}_{φ} proves this.

Now, let us fix a word $u \in \Sigma^*$ and assume that $(q_1, q_2, q_3, q_4, q_5)$ is the state reached by \mathcal{M}_{φ} after reading u . This means that each individual DFA $C_{\beta_j}^{\varphi} = (Q'_{\beta_j}, q'_{1,\beta_j}, \delta'_{\beta_j}, F'_{\beta_j})$ reaches state q_j after reading u . Let now

$$B = \{\beta_j \in \mathbb{B}_4 \mid q_j \in F'_{\beta_j}, j \in \{1, \dots, 5\}\}$$

as in the definition of the output function λ of \mathcal{M}_φ . By applying Lemma 7, we then obtain

$$\beta_j \in B \Leftrightarrow q_j \in F'_{\beta_j} \Leftrightarrow u \in L(C_{\beta_j}^\varphi) \Leftrightarrow u \in L(\mathcal{B}_{\beta_j}^\varphi) \Leftrightarrow \exists \sigma \in \Sigma^\omega : V_r(u\sigma, \varphi) = \beta_j.$$

To conclude the proof, it is left to show that $\xi(B) = V_r^m(u, \varphi)$. We show this for each bit individually using a case distinction over the elements of $\mathbb{B}^? = \{0, ?, 1\}$. So as to clutter this proof not too much, however, we only discuss the case of ? here, while noting that the remaining two cases can be proven analogously. Thus, let $i \in \{1, \dots, 4\}$. Then,

$$\begin{aligned} \xi(B)[i] = ? &\Leftrightarrow \exists \beta, \beta' \in B : \beta[i] = 0 \text{ and } \beta'[i] = 1 \\ &\Leftrightarrow \exists \sigma_0, \sigma_1 \in \Sigma^\omega : V_r(u\sigma_0, \varphi)[i] = 0 \text{ and } V_r(u\sigma_1, \varphi)[i] = 1 \\ &\Leftrightarrow V_r^m(u, \varphi)[i] = ?. \end{aligned}$$

Since $\lambda((q_1, q_2, q_3, q_4, q_5)) = \xi(B)$, the Moore machine \mathcal{M}_φ indeed outputs $V_r^m(u, \varphi)$ for every word $u \in \Sigma^*$. Moreover, \mathcal{M}_φ has $2^{2^{\mathcal{O}(|\varphi|)}}$ states because the DFAs $C_\beta^\varphi = (Q'_\beta, q'_{1,\beta}, \delta'_\beta, F'_\beta)$ are of at most exponential size in $|\mathcal{B}_\beta^\varphi|$, which in turn is at most exponential in $|\varphi|$. In total, this proves Theorem 3. \square

In a final post-processing step, we minimize \mathcal{M}_φ (e.g., using one of the standard algorithms for deterministic automata). As a result, we obtain the unique minimal monitor for the given rLTL formula.

It is left to determine the complexity of our rLTL monitor construction. Since each DFA C_β^φ is in the worst case exponential in the size of the NFA $\mathcal{B}_\beta^\varphi$, we immediately obtain that C_β^φ is at most of size $2^{2^{\mathcal{O}(|\varphi|)}}$. Thus, the Moore machine \mathcal{M}_φ is at most of size $2^{2^{\mathcal{O}(|\varphi|)}}$ as well and can be effectively computed in doubly-exponential time in $|\varphi|$. Note that this matches the complexity bound of Bauer et al.’s approach for LTL runtime monitoring [14]. Moreover, this bound is tight since rLTL subsumes LTL (see Remark 3): Every monitor for an rLTL formula (without implications) can be turned into a monitor for the corresponding LTL formula by projecting every output to its first bit. Thus, the doubly-exponential bound, which is tight for LTL [14, 40], is also tight for rLTL. Hence, robust runtime monitoring asymptotically incurs no extra cost compared to classical LTL runtime monitoring. However, it provides more useful information as we demonstrate next in our experimental evaluation.

5 Experimental evaluation

Besides incorporating a notion of robustness into classical LTL monitoring, our rLTL monitoring approach also promises to provide richer information than its LTL counterpart. In this section, we evaluate empirically whether this promise is actually fulfilled. More precisely, we answer the following two questions on a comprehensive suite of benchmarks:

1. How does rLTL monitoring compare to classical LTL monitoring in terms of monitorability?
2. For formulas that are both LTL-monitorable and rLTL-monitorable, how do both approaches compare in terms of the size of the resulting monitors and the time required to construct them?

To answer these research questions, we have implemented a prototype, which we named `rLTL-mon`. Our prototype is written in Java and builds on top of two libraries: Owl [39], a library for LTL and automata over infinite words, as well as AutomataLib (part of LearnLib

[36]), a library for automata over finite words and Moore machines. For technical reasons (partly due to limitations of the Owl library and partly to simplify the implementation), `rLTL-mon` uses a monitor construction that is slightly different from the one described in Sect. 4: Instead of translating an rLTL formula into nondeterministic Büchi automata, `rLTL-mon` constructs deterministic parity automata. These parity automata are then directly converted into DFAs, thus skirting the need for a detour over NFAs and a subsequent determinization step. Note, however, that this alternative construction produces the same rLTL monitors than the one described in Sect. 4. Moreover, it has the same asymptotic complexity. The sources of our prototype are available online under the MIT license.⁷

Benchmarks and experimental setup

The starting point of our evaluation was the original benchmark suite of Bauer et al. [14], which is based on a survey by Dwyer on frequently used software specification patterns [23]. This benchmark suite consists of 97 LTL formulas and covers a wide range of patterns, including safety, scoping, precedence, and response patterns. For our rLTL monitor construction, we interpreted each LTL formula in the benchmark suite as an rLTL formula (by treating every operator as a robust operator).

We compared `rLTL-mon` to Bauer et al.'s implementation of their LTL monitoring approach, which the authors named `LTL3 tools`. This tool uses `LTL2BA` [32] to translate LTL formulas into Büchi automata and AT&T's `fsmlib` as a means to manipulate finite-state machines. Since `LTL2BA`'s and Owl's input format for LTL formulas do not match exactly, we have translated all benchmarks into a suitable format using a python script.

We conducted all experiments on an Intel Core i5-6600 @ 3.3 GHz in a virtual machine with 4 GB of RAM running Ubuntu 18.04 LTS. As no monitor construction took longer than 600 s, we did not impose any time limit.

Results

Our evaluation shows that `LTL3 tools` and `rLTL-mon` are both able to generate monitors for all 97 formulas in Bauer et al.'s benchmark suite.⁸ Aggregated statistics of this evaluation are visualized in Fig. 5.⁹


The histogram in Fig. 5a shows the aggregate number of LTL and rLTL monitors with respect to their number of states. As Bauer et al. already noted in their original work, the resulting LTL monitors are quite small (none had more than six states), which they attribute to Dwyer et al.'s specific selection of formulas [23]. A similar observation is also true for the rLTL monitors: None had more than eight states.

To determine which formulas are monitorable and which are not, we used a different, though equivalent definition, which is easy to check on the monitor itself: an LTL formula (rLTL formula) is monitorable if and only if the unique minimized LTL monitor (rLTL monitor) does not contain a sink-state with universal self-loop that outputs “?” (that outputs “????”). In other words, even if a finite word does not allow us to infer anything about the satisfaction of the LTL (rLTL) formula by infinite words extending it, it can always be


⁷ <https://github.com/logic-and-learning/rtl-mon>.

⁸ Note that the tools disagreed on one monitor where `LTL3 tools` constructed a monitor with 1 state whereas `rLTL-mon` constructed an LTL monitor with 8 states. The respective formula was removed from the reported results.

⁹ Detailed results can be found in Tables 3 and 4 in the appendix.



(a) Histogram of the number of monitors with respect to their size



(b) Analysis of the monitor construction for the 54 formulas that are both LTL-monitorable and rLTL-monitorable

Fig. 5 Comparison of rLTL-mon and LTL₃ tools on Bauer et al.’s benchmarks [14]

extended into another finite word that does. Bauer et al. report that 44.3% of the LTL monitors (43 out of 97) have this property (in fact, exactly the 43 LTL monitors with a single state), which means that 44.3% of all formulas in their benchmark suite are not LTL-monitorable. By contrast, all these formulas are rLTL-monitorable. Moreover, in 78.4% of the cases (76 out of 97), the rLTL monitor has more distinct outputs than the LTL monitor, indicating that the rLTL monitor provides more fine-grained information of the property being monitored; in the remaining 21.6%, both monitors have the same number of distinct outputs. These results answer our first research question strongly in favor of rLTL monitoring: *rLTL monitoring did in fact provide more information than its classical LTL counterpart. In particular, only 55.7% of the benchmarks are LTL-monitorable, whereas 100% are rLTL-monitorable.*

Let us now turn to our second research question and compare both monitoring approaches on the 54 formulas that are both LTL-monitorable and rLTL-monitorable. For these formulas, Fig. 5b further provides statistical analysis of the generated monitors in terms of their size (left diagram) as well as the time required to generate them (right diagram). Each box in the diagrams shows the lower and upper quartile (left and right border of the box, respectively), the median (line within the box), and minimum and maximum (left and right whisker, respectively).

Let us first consider the size of the monitors (left diagram of Fig. 5b). The majority of LTL monitors (52) has between two and four states, while the majority of rLTL monitors (45) has between two and five states. For 21 benchmarks, the LTL and rLTL monitors are of equal size, while the rLTL monitor is larger for the remaining 33 benchmarks. On average, rLTL monitors are about 1.5 times larger than the corresponding LTL monitors.

Let us now discuss the time taken to construct the monitors. As the diagram on the right-hand-side of Fig. 5b shows, LTL₃ tools was considerably faster than rLTL-mon on a majority of benchmarks (around 0.1 s and 2.6 s per benchmark, respectively). For all

54 benchmarks, the rLTL monitor construction took longer than the construction of the corresponding LTL monitor (although there are two non-LTL-monitorable formulas for which the construction of the rLTL monitor was faster). However, we attribute this large runtime gap partly to the overhead caused by repeatedly starting the Java virtual machine, which is not required in the case of `LTL3 TOOLS`. Note that this is not a concern in practice as a monitor is only constructed once before it is deployed.

Finally, our analysis answers our second question: *rLTL monitors are only slightly larger than the corresponding LTL monitors and although they require considerably more time to construct, the overall construction time was negligible for almost all benchmarks.*

6 Conclusion

We adapted the three-valued LTL monitoring semantics of Bauer et al. to rLTL, proved that the construction of monitors is asymptotically no more expensive than the one for LTL, and validated our approach on the benchmark of Bauer et al.: All formulas are rLTL-monitorable and the rLTL monitor is strictly more informative than its LTL counterpart for 77% of their formulas.

Recall Theorem 1, which states that the truth values 0011 and 0001 are not realizable. This points to a drawback regarding the two middle bits: When considering the formula $\Box a$, the second bit represents $\Diamond \Box a$ and the third bit $\Box \Diamond a$. A prefix cannot possibly provide enough information to distinguish these two formulas. On the other hand, the truth value $??11$ is realizable, which shows that the middle bits can be relevant. In further work, we will investigate the role of the middle bits in rLTL monitoring.

Moreover, the informedness of a monitor can be increased further when attributing a special role to the last position(s) of a prefix. Even though a prefix of the form $\emptyset^+ \{a\}^+$ does not fully satisfy $\Diamond \Box a$, neither does it fully violate it. If the system just now reached a state in which $\{a\}$ always holds, an infinite continuation of the execution would satisfy the specification. So rather than reporting an undetermined result, the monitor could indicate that an infinite repetition of the last position of the prefix would satisfy the specification. Similarly, for a prefix $\{a\}^+ \emptyset$, the specification $\Box \Diamond a$ is undetermined. While an infinite repetition of the last position ($\{a\}^+ \emptyset^\omega$) does not satisfy the specification, an infinite repetition of the last two positions ($\{a\}^+ (\emptyset \{a\})^\omega$) would. We plan to investigate an extension of rLTL which takes this observation into account.

Bauer et al. [12] proposed an orthogonal approach with the logic RV-LTL. Here, the specification can contain the strong (weak) next-operator whose operand is consider violated (satisfied) at the last position of the trace. A formula that is undetermined under the strong semantics and satisfied (violated) under the weak semantics is considered *potentially true* (*potentially false*). Incorporating one of these approaches into rLTL monitoring could refine its output and thus increase its level of informedness.

Moreover, desired properties for cyber-physical systems often include real-time components such as “touch the ground at most 15 s after receiving a landing command”. Monitors for logics taking real-time into account [15], such as STL [43, 44], induce high computational overhead at runtime when compared to LTL and rLTL monitors. Thus, a real-time extension for rLTL retaining its low runtime cost would greatly increase its viability as specification language.

Acknowledgements The authors would like to thank Li Bingchen for discovering the formula mentioned in Footnote 5. The work of Daniel Neider was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Grant No. 434592664. The work of Maximilian Schwenger was supported by the European Research Council (ERC) Grant OSARES (No. 683300) and the German Research Foundation (DFG) as part of the Collaborative Research Center “Center for Perspicuous Computing” (TRR 248, 389792660). The work of Paulo Tabuada was partially supported by the NSF project 1645824. The work of Alexander Weinert was supported by the Saarbrücken Graduate School of Computer Science. The work of Martin Zimmermann was supported by the Engineering and Physical Sciences Research Council (EPSRC) project EP/S032207/1.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data availability All data generated or analyzed during this study can be obtained through the repository at <https://github.com/logic-and-learning/rtdl-monitoring>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Experimental results

The following two tables provide detailed results of our experimental evaluation.

Table 3 Summary of the result when comparing the monitor construction of rLTL against LTL; continued in Table 4

Property	# States		# Outputs		Monitorable		Time in s	
	rLTL	LTL	rLTL	LTL	rLTL	LTL	rLTL	LTL
Constrained response-chain 2-1	3	1	2	1	✓	×	2.06	0.07
Constrained 3-2 response chain	3	1	2	1	✓	×	1.79	0.03
Constrained 3-2 response chain	4	1	2	1	✓	×	24.21	284.53
Constrained 2-1 response chain	3	1	2	1	✓	×	1.94	0.03
Existence	6	3	4	2	✓	✓	1.71	0.02
2 Bounded Existence	8	1	2	1	✓	×	3.44	0.02
Response	2	1	2	1	✓	×	2.17	0.02
Existence	6	3	4	2	✓	✓	1.59	0.02
Existence	3	3	3	3	✓	✓	1.23	0.01
Existence	2	1	2	1	✓	×	1.11	0.02
Existence	2	1	2	1	✓	×	1.12	0.01
Existence	2	1	2	1	✓	×	1.12	0.02
Response	2	1	2	1	✓	×	2.15	0.01
Response	2	1	2	1	✓	×	2.15	0.02
Existence	5	3	3	2	✓	✓	1.14	0.01
Absence	4	4	3	3	✓	✓	2.34	0.02
Absence	4	4	2	2	✓	✓	3.10	0.02
Absence	5	3	4	2	✓	✓	5.07	0.02
Absence	4	4	2	2	✓	✓	2.08	0.02
Response	2	1	2	1	✓	×	1.64	0.01
Response	2	1	2	1	✓	×	1.61	0.01
GlobalResponse	2	1	2	1	✓	×	6.67	0.04
Precedence	5	3	5	3	✓	✓	1.35	0.01

Table 3 continued

Property	# States		# Outputs		Monitorable		Time in s	
	rLTL	LTL	rLTL	LTL	rLTL	LTL	rLTL	LTL
Absence	4	2	4	2	✓	✓	1.53	0.03
Response	2	1	2	1	✓	×	1.34	0.02
Universal	4	2	4	2	✓	✓	1.67	0.02
Absence	4	4	2	2	✓	✓	2.12	0.03
Absence	4	4	2	2	✓	✓	2.10	0.03
Absence	4	4	3	3	✓	✓	2.67	0.86
Absence	3	3	2	2	✓	✓	1.74	0.01
Absence	3	3	2	2	✓	✓	3.20	0.42
Universal	4	2	4	2	✓	✓	1.44	0.01
Response	5	5	2	2	✓	✓	8.33	0.34
Precedence	5	3	5	3	✓	✓	1.50	0.01
Absence	5	3	4	2	✓	✓	6.30	0.02
Absence	5	3	4	2	✓	✓	6.59	0.02
Absence	4	4	3	3	✓	✓	1.44	0.02
Absence	5	3	4	2	✓	✓	4.17	0.02
Constrained Response-Chain 3-1	3	1	2	1	✓	×	35.40	319.74
Absence	7	4	4	2	✓	✓	35.57	2.11
Absence	4	4	3	3	✓	✓	1.42	0.01
Absence	5	3	4	2	✓	✓	4.04	0.02
Universal	4	2	4	2	✓	✓	1.38	0.02
Response	4	4	3	3	✓	✓	1.81	0.02
Response	2	1	2	1	✓	×	1.45	0.02
Response	2	1	2	1	✓	×	1.59	0.02
Response	2	1	2	1	✓	×	1.46	0.01
Response	2	1	2	1	✓	×	1.79	0.02

References

1. Abbas H, Rodionova A, Bartocci E, Smolka SA, Grosu R (2017) Quantitative regular expressions for arrhythmia detection algorithms. In: Feret J, Koepl H (eds) CMSB 2017, LNCS, vol 10545. Springer, Berlin, pp 23–39. https://doi.org/10.1007/978-3-319-67471-1_2
2. Adolf F, Faymonville P, Finkbeiner B, Schirmer S, Torens C (2017) Stream runtime monitoring on UAS. In: Lahiri SK, Reger G (eds) RV 2017, LNCS, vol 10548. Springer, Berlin, pp 33–49. https://doi.org/10.1007/978-3-319-67531-2_3
3. Akazaki T, Hasuo I (2015) Time robustness in MTL and expressivity in hybrid system falsification. In: Kroening D, Pasareanu CS (eds) CAV 2015, LNCS, vol 9207. Springer, Berlin, pp 356–374. https://doi.org/10.1007/978-3-319-21668-3_21
4. Alur R, Fisman D, Raghthaman M (2016) Regular programming for quantitative properties of data streams. In: Thiemann P (ed) ESOP 2016, LNCS, vol 9632. Springer, Berlin, pp 15–40. https://doi.org/10.1007/978-3-662-49498-1_2
5. Anevlavis T, Neider D, Phillippe M, Tabuada P (2019) Evrostos: the rLTL verifier. In: Ozay N, Prabhakar P (eds) HSCC 2019. ACM, New York, pp 218–223. <https://doi.org/10.1145/3302504.3311812>
6. Anevlavis T, Phillippe M, Neider D, Tabuada P (2018) Verifying rLTL formulas: now faster than ever before! In: CDC 2018, pp 1556–1561. IEEE. <https://doi.org/10.1109/CDC.2018.8619014>

Table 4 Summary of the result when comparing the monitor construction of rLTL against LTL; continuation of Table 3

Property	# States		# Outputs		Monitorable		Time in s	
	rLTL	LTL	rLTL	LTL	rLTL	LTL	rLTL	LTL
Constrained response	6	3	4	2	✓	✓	1.52	0.02
Absence	4	2	4	2	✓	✓	1.26	0.01
Response	2	1	2	1	✓	×	1.35	0.02
Response	2	1	2	1	✓	×	1.35	0.01
Unknown	8	3	4	2	✓	✓	2.54	0.02
Existence	2	2	2	2	✓	✓	0.95	0.01
Unknown	2	1	2	1	✓	×	1.34	0.02
Unknown	2	1	2	1	✓	×	1.31	0.01
Response	2	1	2	1	✓	×	550.84	2.63
Unknown	6	6	3	3	✓	✓	0.88	0.02
Unknown	3	3	2	2	✓	✓	0.95	0.02
Existence	2	2	2	2	✓	✓	0.85	0.01
Existence	2	1	2	1	✓	×	0.88	0.02
Always	2	1	2	1	✓	×	1.03	0.01
Universal	2	1	2	1	✓	×	0.89	0.01
Universal	3	3	2	2	✓	✓	1.31	0.01
Universal	2	1	2	1	✓	×	0.88	0.01
Existence	2	2	2	2	✓	✓	0.82	0.02
Absence	6	3	4	2	✓	✓	1.45	0.01
Response	2	1	2	1	✓	×	1.16	0.01
Existence	3	3	2	2	✓	✓	1.11	0.01
GlobalUniversal GlobalAbsence	4	1	2	1	✓	×	1.29	0.01
Response	2	1	2	1	✓	×	1.39	0.01
Universal	4	2	4	2	✓	✓	1.49	0.01
Response	2	1	2	1	✓	×	1.23	0.01
Response	2	1	2	1	✓	×	1.17	0.01
Response Chain 1-2	2	1	2	1	✓	×	2.40	0.01
Universal	4	2	4	2	✓	✓	1.11	0.01
Absence	4	2	4	2	✓	✓	2.63	0.01
Absence	4	2	4	2	✓	✓	1.01	0.01
Universal	4	2	4	2	✓	✓	1.12	0.01
Universal	4	2	4	2	✓	✓	1.10	0.01
Universal	4	2	4	2	✓	✓	1.12	0.01
Response	2	1	2	1	✓	×	1.17	0.01
Absence	6	3	4	2	✓	✓	1.50	0.01
Universal	4	2	4	2	✓	✓	1.01	0.01

Table 4 continued

Property	# States		# Outputs		Monitorable		Time in s	
	rLTL	LTL	rLTL	LTL	rLTL	LTL	rLTL	LTL
Absence	4	4	2	2	✓	✓	1.63	0.02
Response	2	1	2	1	✓	×	1.30	0.01
Universal	4	2	4	2	✓	✓	1.11	0.01
Response	2	1	2	1	✓	×	1.17	0.01
Response	2	1	2	1	✓	×	1.27	0.02
Response	2	1	2	1	✓	×	1.17	0.01
Existence	6	3	4	2	✓	✓	1.32	0.01
Response	2	1	2	1	✓	×	1.18	0.01
Response	2	1	2	1	✓	×	1.17	0.01
Response	2	1	2	1	✓	×	1.17	0.01
Unknown	4	2	4	2	✓	✓	1.25	0.01
Universal	4	2	4	2	✓	✓	0.91	0.01

7. Anevlavis T, Philippe M, Neider D, Tabuada P (2022) Being correct is not enough: efficient verification using robust linear temporal logic. *ACM Trans Comput Log* 23(2):8:1-8:39. <https://doi.org/10.1145/3491216>
8. Baier C, Katoen J (2008) Principles of model checking. MIT Press, Cambridge
9. Barringer H, Falcone Y, Havelund K, Reger G, Rydeheard DE (2012) Quantified event automata: towards expressive and efficient runtime monitors. In: Giannakopoulou D, Méry D (eds) FM 2012, LNCS, vol 7436, pp 68–84. Springer. https://doi.org/10.1007/978-3-642-32759-9_9
10. Bartocci E, Bloem R, Nickovic D, Röck F (2018) A counting semantics for monitoring LTL specifications over finite traces. In: Chockler H, Weissenbacher G (eds) CAV 2018, LNCS, vol 10981. Springer, Berlin, pp 547–564. https://doi.org/10.1007/978-3-319-96145-3_29
11. Basin DA, Klaedtke F, Marinovic S, Zalinescu E (2015) Monitoring of temporal first-order properties with aggregations. *Form Methods Syst Des* 46(3):262–285. <https://doi.org/10.1007/s10703-015-0222-7>
12. Bauer A, Leucker M, Schallhart C (2007) The good, the bad, and the ugly, but how ugly is ugly? In: Sokolsky O, Tasiran S (eds) RV 2007, LNCS, vol 4839. Springer, Berlin, pp 126–138. https://doi.org/10.1007/978-3-540-77395-5_11
13. Bauer A, Leucker M, Schallhart C (2010) Comparing LTL semantics for runtime verification. *J Log Comput* 20(3):651–674. <https://doi.org/10.1093/logcom/exn075>
14. Bauer A, Leucker M, Schallhart C (2011) Runtime verification for LTL and TLTL. *ACM Trans Softw Eng Methodol* 20(4):14:1-14:64. <https://doi.org/10.1145/2000799.2000800>
15. Bernstein A, Harter PK Jr (1981) Proving real-time properties of programs with temporal logic. In: Howard J, Reed DP (ed) SOSP 1981. ACM, New York, pp 1–11. <https://doi.org/10.1145/800216.806585>
16. Caspi P, Pilaud D, Halbwach N, Plaice J (1987) Lustre: a declarative language for programming synchronous systems. In: POPL 1987, pp 178–188. ACM Press, New York. <https://doi.org/10.1145/41625.41641>
17. Cheng C (2021) Provably-robust runtime monitoring of neuron activation patterns. In: DATE 2021, pp 1310–1313. IEEE. <https://doi.org/10.23919/DATES1398.2021.9473957>
18. Cralley J, Spantidi O, Hoxha B, Fainekos G (2020) Tltk: a toolbox for parallel robustness computation of temporal logic specifications. In: Deshmukh J, Nickovic D (eds) RV 2020, LNCS, vol 12399. Springer, Berlin, pp 404–416. https://doi.org/10.1007/978-3-030-60508-7_22
19. D’Angelo B, Sankaranarayanan S, Sánchez C, Robinson W, Finkbeiner B, Sipma HB, Mehrotra S, Manna Z (2005) LOLA: runtime monitoring of synchronous systems. In: TIME 2005, pp 166–174. IEEE Computer Society. <https://doi.org/10.1109/TIME.2005.26>
20. Decker N, Leucker M, Thoma D (2013) Impartiality and anticipation for monitoring of visibly context-free properties. In: Legay A, Bensalem S (ed) RV 2013, LNCS, vol 8174, pp 183–200. Springer, Berlin. https://doi.org/10.1007/978-3-642-40787-1_11

21. Donzé A, Ferrère T, Maler O (2013) Efficient robust monitoring for STL. In: Sharygina N, Veith H (ed) CAV 2013, LNCS, vol 8044, pp 264–279. Springer, Berlin. https://doi.org/10.1007/978-3-642-39799-8_19
22. Drusinsky D (2000) The temporal rover and the ATG rover. In: Havelund K, Penix J, Visser W (ed) SPIN 2000, LNCS, vol 1885, pp 323–330. Springer, Berlin. https://doi.org/10.1007/10722468_19
23. Dwyer MB, Avrunin GS, Corbett JC (1999) Patterns in property specifications for finite-state verification. In: Boehm BW, Garland D, Kramer J (ed) ICSE 1999, pp 411–420. ACM, New York. <https://doi.org/10.1145/302405.302672>
24. Eisner C, Fisman D, Havlicek J, Lustig Y, McIsaac A, Campenhout DV (2003) Reasoning with temporal logic on truncated paths. In: Hunt WA, Somenzi F (ed) CAV 2003, LNCS, vol 2725, pp 27–39. Springer, Berlin. https://doi.org/10.1007/978-3-540-45069-6_3
25. Fainekos GE, Pappas GJ (2006) Robustness of temporal logic specifications for continuous-time signals. *Theor Comput Sci* 410(42):4262–4291. <https://doi.org/10.1016/j.tcs.2009.06.021>
26. Falcone Y, Sánchez C (eds) (2016) RV 2016, LNCS, vol 10012. Springer, New York. <https://doi.org/10.1007/978-3-319-46982-9>
27. Faymonville P, Finkbeiner B, Schirmer S, Torfah H A stream-based specification language for network monitoring. In: Falcone and Sánchez [26], pp 152–168. https://doi.org/10.1007/978-3-319-46982-9_10
28. Faymonville P, Finkbeiner B, Schledjewski M, Schwenger M, Tenstrup L, Stenger M, Torfah H (2019) Streamlab: stream-based monitoring of cyber-physical systems. In: CAV 2019. To appear
29. Finkbeiner B, Keller A, Schmidt J, Schwenger M (2021) Robust monitoring for medical cyber-physical systems. In: MCPS 2021, pp 17–22. Association for computing machinery, New York, USA. <https://doi.org/10.1145/3446913.3460318>
30. Finkbeiner B, Sankaranarayanan S, Sipma H (2005) Collecting statistics over runtime executions. *Form Methods Syst Des* 27(3):253–274. <https://doi.org/10.1007/s10703-005-3399-3>
31. Finkbeiner B, Torfah H (2017) The density of linear-time properties. In: D’Souza D, Kumar KN (ed) ATVA 2017, LNCS, vol 10482, pp 139–155. Springer, New York. https://doi.org/10.1007/978-3-319-68167-2_10
32. Gastin P, Oddoux D (2001) Fast LTL to Büchi automata translation. In: Berry G, Comon H, Finkel A (ed) CAV 2001, LNCS, vol 2102, pp 53–65. Springer, New York. https://doi.org/10.1007/3-540-44585-4_6
33. Hallé S When RV meets CEP. In: Falcone and Sánchez [26], pp 68–91. https://doi.org/10.1007/978-3-319-46982-9_6
34. Havelund K, Rosu G (2002) Synthesizing monitors for safety properties. In: Katoen J, Stevens P (ed) TACAS 2002, LNCS, vol 2280, pp 342–356. Springer, New York. https://doi.org/10.1007/3-540-46002-0_24
35. Havelund K, Rosu G (2004) An overview of the runtime verification tool Java PathExplorer. *Form Methods Syst Des* 24(2):189–215. <https://doi.org/10.1023/B:FORM.0000017721.39909.4b>
36. Isberner M, Howar F, Steffen B (2015) The open-source learnlib - A framework for active automata learning. In: Kroening D, Pasareanu CS (ed) CAV 2015 (Part I), LNCS, vol 9206, pp 487–495. Springer, New York. https://doi.org/10.1007/978-3-319-21690-4_32
37. Jaksic S, Bartocci E, Grosu R, Nguyen T, Nickovic D (2018) Quantitative monitoring of STL with edit distance. *Form Methods Syst Des* 53(1):83–112. <https://doi.org/10.1007/s10703-018-0319-x>
38. Kretínský J, Meggendorfer T, Sickert S (2018) LTL store: repository of LTL formulae from literature and case studies. [arXiv:1807.03296](https://arxiv.org/abs/1807.03296)
39. Kretínský J, Meggendorfer T, Sickert S (2018) Owl: A library for ω -words, automata, and LTL. In: Lahiri, Wang [41], pp 543–550. https://doi.org/10.1007/978-3-030-01090-4_34
40. Kupferman O, Vardi MY (2001) Model checking of safety properties. *Form Methods Syst Des* 19(3):291–314. <https://doi.org/10.1023/A:1011254632723>
41. Lahiri SK, Wang C (eds) (2018) ATVA 2018, LNCS, vol 11138. Springer, Cham. <https://doi.org/10.1007/978-3-030-01090-4>
42. Lee I, Kannan S, Kim M, Sokolsky O, Viswanathan M (1999) Runtime assurance based on formal specifications. In: Arabnia HR (ed) PDPTA 1999. CSREA Press, Las Vegas, pp 279–287
43. Maler O, Nickovic D (2004) Monitoring temporal properties of continuous signals. In: Lakhnech Y, Yovine S (ed) FORMATS and FTRTFT 2004, LNCS, vol 3253, pp 152–166. Springer, Cham. https://doi.org/10.1007/978-3-540-30206-3_12
44. Maler O, Nickovic D, Pnueli A (2008) Checking temporal properties of discrete, timed and continuous behaviors. In: Avron A, Dershowitz N, Rabinovich A (ed) Pillars of computer science, essays dedicated to Boris (Boaz) Trakhtenbrot on the occasion of his 85th birthday, LNCS, vol 4800, pp 475–505. Springer, Cham. https://doi.org/10.1007/978-3-540-78127-1_26

45. Maler O, Pnueli A (1995) Timing analysis of asynchronous circuits using timed automata. In: Camurati P, Evekling H (ed) CHARME 1995, LNCS, vol 987, pp 189–205. Springer, Cham. https://doi.org/10.1007/3-540-60385-9_12
46. Manna Z, Pnueli A (1995) Temporal verification of reactive systems-safety. Springer, Berlin
47. Mascle C, Neider D, Schwenger M, Tabuada P, Weinert A, Zimmermann M (2020) From LTL to rltl monitoring: improved monitorability through robust semantics. In: Ames AD, Seshia SA, Deshmukh J (ed) HSCC 2020, pp 7:1–7:12. ACM, New York. <https://doi.org/10.1145/3365365.3382197>
48. Medhat R, Bonakdarpour B, Fischmeister S, Joshi Y (2016) Accelerated runtime verification of LTL specifications with counting semantics. In: Falcone and Sánchez [26], pp 251–267. https://doi.org/10.1007/978-3-319-46982-9_16
49. Moosbrugger P, Rozier KY, Schumann J (2017) R2U2: monitoring and diagnosis of security threats for unmanned aerial systems. *Form Methods Syst Des* 51(1):31–61. <https://doi.org/10.1007/s10703-017-0275-x>
50. Neider D, Weinert A, Zimmermann M (2019) Robust, expressive, and quantitative linear temporal logics: pick any two for free. In: Leroux J, Raskin J (ed) Proceedings tenth international symposium on games, automata, logics, and formal verification, GandALF 2019, Bordeaux, 2–3rd Sept 2019, EPTCS, vol 305, pp 1–16. <https://doi.org/10.4204/EPTCS.305.1>
51. Neider D, Weinert A, Zimmermann M (2021) Robust, expressive, and quantitative linear temporal logics: pick any two for free. *Inf Comput*. <https://doi.org/10.1016/j.ic.2021.104810>
52. Pike L, Goodloe A, Morisset R, Niller S (2010) Copilot: a hard real-time runtime monitor. In: Barringer H, Falcone Y, Finkbeiner B, Havelund K, Lee I, Pace GJ, Rosu G, Sokolsky O, Tillmann N (ed) RV 2010, LNCS, vol 6418, pp 345–359. Springer, Cham. https://doi.org/10.1007/978-3-642-16612-9_26
53. Pnueli A, Zaks A (2006) PSL model checking and run-time verification via testers. In: Misra J, Nipkow T, Sekerinski E (ed) FM 2006, LNCS, vol 4085, pp 573–586. Springer, Cham. https://doi.org/10.1007/11813040_38
54. Rodionova A, Bartocci E, Nickovic D, Grosu R (2016) Temporal logic as filtering. In: Proceedings of the 19th international conference on hybrid systems: computation and control, HSCC '16, pp 11–20. ACM, New York. <https://doi.org/10.1145/2883817.2883839>
55. Roesch M (1999) Snort: lightweight intrusion detection for networks. In: Parter DW (ed) LISA 1999, pp 229–238. USENIX, Berkeley
56. Schwoon S, Esparza J (2005) A note on on-the-fly verification algorithms. In: Halbwachs N, Zuck LD (ed) TACAS 2005, LNCS, vol 3440, pp 174–190. Springer, Cham. https://doi.org/10.1007/978-3-540-31980-1_12
57. Silvetti S, Nenzi L, Bartocci E, Bortolussi L (2018) Signal convolution logic. In: Lahiri and Wang [41], pp 267–283. https://doi.org/10.1007/978-3-030-01090-4_16
58. Tabuada P, Neider D (2016) Robust linear temporal logic. In: Talbot J, Regnier L (ed) CSL 2016, LIPIcs, vol 62, pp 10:1–10:21. Schloss Dagstuhl-LZI. <https://doi.org/10.4230/LIPIcs.CSL.2016.10>
59. Torfah H, Zimmermann M (2018) The complexity of counting models of linear-time temporal logic. *Acta Inf* 55(3):191–212. <https://doi.org/10.1007/s00236-016-0284-z>
60. Zhang X, Leucker M, Dong W (2012) Runtime verification with predictive semantics. In: Goodloe A, Person S (ed) NFM 2012, LNCS, vol 7226, pp 418–432. Springer, Cham. https://doi.org/10.1007/978-3-642-28891-3_37

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Corto Mascle¹ · Daniel Neider²  · Maximilian Schwenger³ · Paulo Tabuada⁴ · Alexander Weinert⁵ · Martin Zimmermann^{6,7}

Corto Mascle
corto.mascle@labri.fr

Maximilian Schwenger
schwenger@react.uni-saarland.de

Paulo Tabuada
tabuada@ee.ucla.edu

Alexander Weinert
alexander.weinert@dlr.de

Martin Zimmermann
mzi@cs.aau.dk

- ¹ LaBRI, University of Bordeaux, Talence, France
- ² Safety and Explainability of Learning Systems Group, Carl von Ossietzky University of Oldenburg, Oldenburg, Germany
- ³ Reactive Systems Group, Saarland University, Saarbrücken, Germany
- ⁴ Department of Electrical and Computer Engineering, UCLA, Los Angeles, USA
- ⁵ German Aerospace Center (DLR), Cologne, Germany
- ⁶ University of Liverpool, Liverpool, UK
- ⁷ Present Address: Aalborg University, Aalborg, Denmark