

Robust, Expressive, and Quantitative Linear Temporal Logics: Pick any Two for Free

Daniel Neider

Max Planck Institute for Software Systems, 67663 Kaiserslautern, Germany

Alexander Weinert

German Aerospace Center (DLR), Institute for Software Technology, 51147 Cologne, Germany

Martin Zimmermann

*University of Liverpool, Liverpool L69 3BX, United Kingdom
Present address: Aalborg University, 9220 Aalborg, Denmark*

Abstract

Linear Temporal Logic (LTL) is the standard specification language for reactive systems and is successfully applied in industrial settings. However, many shortcomings of LTL have been identified in the literature, among them the limited expressiveness, the lack of quantitative features, and the inability to express robustness. There is work on overcoming these shortcomings, but each of these is typically addressed in isolation. This is insufficient for applications where all shortcomings manifest themselves simultaneously.

Here, we tackle this issue by introducing logics that address more than one shortcoming. To this end, we combine the logics Linear Dynamic Logic, Prompt-LTL, and robust LTL, each addressing one aspect, to new logics. For all combinations of two aspects, the resulting logic has the same desirable algorithmic properties as plain LTL. In particular, the highly efficient algorithmic backends that have been developed for LTL are also applicable to these new logics. Finally, we discuss how to address all three aspects simultaneously.

*This work was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) project no. 434592664 and the EPSRC project “quantMD” (EP/S032207/1).

**This work is an extended version of an earlier conference paper [1]. We detail the extensions in Section 1.1.

Email addresses: neider@mpi-sws.org (Daniel Neider), alexander.weinert@dlr.de (Alexander Weinert), mzi@cs.aau.dk (Martin Zimmermann)

1. Introduction

Linear Temporal Logic (LTL) [2] is amongst the most prominent and most important specification languages for reactive systems, e.g., non-terminating controllers interacting with an antagonistic environment. Verification of such systems against LTL specifications is routinely applied in industrial settings nowadays [3, 4]. Underlying this success story is the exponential compilation property [5]: every LTL formula can be effectively translated into an equivalent Büchi automaton of exponential size (and it turns out that this upper bound is tight). In fact, almost all verification algorithms for LTL are based on this property, which is in particular true for the popular polynomial space model checking algorithm and the doubly-exponential time synthesis algorithms. Other desirable properties of LTL include its compact and variable-free syntax and its intuitive semantics.

Despite the success of LTL, a plethora of extensions of LTL have been studied, all addressing individual and specific shortcomings of LTL, e.g., its limited expressiveness, its lack of quantitative features, and its inability to express robustness. Commonly, extensions of LTL as described above are only studied in isolation—the logics are either more expressive, or quantitative, or robust. One notable exception is Parametric LDL (PLDL) [6], which adds quantitative operators and increased expressiveness while maintaining the exponential compilation property and intuitive syntax and semantics. In practical settings, however, it does not suffice to address one shortcoming of LTL while ignoring the others. Instead, one needs a logic that combines multiple extensions while still maintaining the desirable properties of LTL. The overall goal of this paper is, hence, to bridge this gap, thereby enabling expressive, quantitative, and robust verification and synthesis.

It is a well-known fact that LTL is strictly weaker than Büchi automata, i.e., it does not harness the full expressive power of the algorithmic backends. Thus, increasing the expressiveness of LTL has generated much attention [7, 8, 5, 9] as it can be easily exploited: as long as the new logic also has the exponential compilation property, the same optimized backends as for LTL can be used. A prominent and recent example of such an extension that yields the full expressive power of Büchi automata is Linear Dynamic Logic (LDL) [8], which adds to LTL temporal operators guarded by regular expressions. In fact, the guarded operators can express all temporal operators of LTL, i.e., we discard the temporal operators and only allow guarded operators.

As an example, consider the specification “ p holds at every even time point, but may or may not hold at odd time points”. It is well-known that this property is not expressible in LTL, as LTL, intuitively, is unable to count modulo a fixed number. However, the specification is easily expressible in LDL as $[r]p$, where r is the regular expression $(\mathbf{tt} \cdot \mathbf{tt})^*$. The formula requires p to be satisfied at every position j such that the prefix up to position j matches the regular expression r (which is equivalent to j being even), i.e., \mathbf{tt} is an atomic regular expression that matches every letter. In this work, we consider LDL instead of the alternatives cited above for its conceptual simplicity: LDL has a simple and

variable-free syntax based on regular expressions as well as intuitive semantics (assuming some familiarity with regular expressions).

Another serious shortcoming of LTL (and LDL) is its inability to adequately express timing bounds. For example, consider the specification “every request q is eventually answered by a response p ”, which is expressed in LTL as $\Box(q \rightarrow \Diamond p)$. It is satisfied, even if the waiting time between requests q and responses p diverges to infinity, although such a behavior is typically undesired. Again, a long line of research has addressed this second shortcoming of LTL [10, 6, 11, 12, 13]. The most basic one is Prompt-LTL [12], which adds the prompt-eventually operator $\Diamond_{\mathbf{p}}$ to LTL. To retain decidability [10], one has to give up negation and implication when adding the prompt-eventually operator. This is no restriction for LTL, as every formula has an equivalent one in negation normal form.

The semantics is now defined with an additional parameter k , which bounds the scope of $\Diamond_{\mathbf{p}}$: $\Box(q \rightarrow \Diamond_{\mathbf{p}} p)$ requires every request q to be answered within k steps, when evaluated with respect to k . The resulting logic is a quantitative one: either one quantifies the parameter k existentially and obtains a boundedness problem, e.g., “does there exist a bound k such that every request can be answered within k steps”, or one even aims to determine the optimal bound k . Again, Prompt-LTL retains the desirable properties of LTL, i.e., the exponential compilation property as well as intuitive syntax and semantics. Furthermore, Prompt-LTL captures the technical core of the alternatives cited above, e.g., decision problems for the more general logic PLTL [10] can be reduced to those for Prompt-LTL. For these reasons, we study Prompt-LTL in this work.

Finally, a third line of extensions of LTL is concerned with the concept of robustness, which is much harder to formalize. This is reflected by a multitude of incomparable notions of robustness in verification [14, 15, 16, 17, 18, 19, 20, 21, 22]. Here, we are interested in robust LTL (rLTL) [22], which equips LTL with a five-valued semantics that captures different degrees of violations of universal specifications. As an example, consider the specification “if property φ always holds true, then property ψ also always holds true”, which is expressed in LTL as $\Box\varphi \rightarrow \Box\psi$ and is typical for systems that have to interact with an antagonistic environment. In classical semantics, the whole formula is satisfied as soon as the assumption φ is violated once, even if the guarantee ψ is violated as well. By contrast, the semantics of robust LTL ensures that the degree of the violation of $\Box\psi$ is always proportional to the degree of the violation of $\Box\varphi$. To this end, the degree of a violation of a property $\Box\varphi$ is expressed by five different truth values: either φ always holds, or φ is violated only finitely often, violated infinitely often, violated almost always, or violated always. Again, robust LTL has the exponential compilation property and an intuitive syntax (though its semantics is more intricate). In this work, we consider robust LTL, as it is the first logic that intrinsically captures the notion of robustness in LTL. In particular, formulas of robust LTL are evaluated over traces with Boolean truth values for atomic propositions and do not require non-Boolean assignments, which are often hard to determine in real-life applications.

We consider here the fragment $\text{rLTL}(\Box, \Diamond)$ of rLTL that only contains the temporal operators eventually and always, as it already captures the most in-

approach by exhibiting a direct translation of rLDL into Büchi automata via alternating automata. In contrast, to prove the exponential compilation property for rPrompt-LTL, we present a generalization of the reduction-based approach translating rPrompt-LTL into equivalent Prompt-LTL formulas of linear size, which have the exponential compilation property.

Finally, in Section 5, we discuss the combination of all three aspects. Recall that we present a direct translation to automata for rLDL to prove the exponential compilation property. For reasons explained in Section 5, it is challenging to prove the property for rLDL by a reduction to LDL. In contrast, we present a reduction from rPrompt-LTL to Prompt-LTL to prove the exponential compilation property for rPrompt-LTL. Here, it is challenging to prove the property by a direct translation to automata. Hence, both approaches seem inadequate to deal with the combination of all three extensions. Ultimately, we leave the question of whether the logic combining all three aspects has the exponential compilation property for future work.

2. Preliminaries

We denote the non-negative integers by \mathbb{N} , the set $\{0, 1\}$ of Boolean truth values by \mathbb{B} , and the power set of S by 2^S . By convention, we have $\min \emptyset = 1$ and $\max \emptyset = 0$ when the operators range over subsets of \mathbb{B} . Following Tabuada and Neider [22], the set of truth values for robust semantics is $\mathbb{B}_4 = \{0000, 0001, 0011, 0111, 1111\}$, which are ordered by $0000 \prec 0001 \prec 0011 \prec 0111 \prec 1111$. We write \preceq for the non-strict variant of \prec and define $\min \emptyset = 1111$ and $\max \emptyset = 0000$ when the operators range over subsets of \mathbb{B}_4 .

Throughout this work, we fix a finite non-empty set P of atomic propositions and define the shorthands $\mathbf{tt} = p \vee \neg p$ and $\mathbf{ff} = p \wedge \neg p$ for some atomic proposition p . For a set $A \subseteq P$ and a propositional formula ϕ over P , we write $A \models \phi$ if the variable valuation mapping elements in A to 1 and elements in $P \setminus A$ to 0 satisfies ϕ . A trace (over P) is an infinite sequence $w \in (2^P)^\omega$. Given a trace $w = w(0)w(1)w(2) \cdots$ and a position $j \in \mathbb{N}$, we define $w[0, j] = w(0) \cdots w(j-1)$ and $w[j, \infty) = w(j)w(j+1)w(j+2) \cdots$, i.e., $w[0, j]$ is the prefix of length j of w and $w[j, \infty)$ the remaining suffix. In particular, $w[0, 0]$ is empty and $w[0, \infty)$ is w .

In the remainder of this section, we introduce the logics we generalize in this work, namely Robust Linear Temporal Logic (rLTL(\square, \diamond)) [22], Linear Dynamic Logic (LDL) [8], and Prompt Linear Temporal Logic (Prompt-LTL) [12]. Also, we introduce Prompt Linear Dynamic Logic (Prompt-LDL) [6], which we use in some proofs. See Table 1 for an overview. References for the results mentioned in the table are given in the following subsections introducing the logics.

We define the semantics of all these logics by evaluation functions V mapping a trace, a formula, and a bound (in the case of a quantitative logic) to a truth value. This is prudent for robust semantics, as it allows us to introduce useful notation naturally. For the sake of consistency, we also use this approach for the other logics, whose semantics is typically defined via satisfaction relations. In particular, V^R , V^P , and V^V denote the evaluation functions of rLTL(\square, \diamond), LDL,

Logic	Operators	Complexity	
		Model Checking	Synthesis
rLTL(\Box, \Diamond)	$\neg, \wedge, \vee, \rightarrow, \Box, \Diamond$	NP-hard/in PSPACE	2EXPTIME-compl.
LDL	$\neg, \wedge, \vee, \rightarrow, \langle r \rangle, [r]$	PSPACE-compl.	2EXPTIME-compl.
Prompt-LTL	$\wedge, \vee, \mathcal{O}, \mathbf{U}, \mathbf{R}, \Diamond_{\mathbf{P}}$	PSPACE-compl.	2EXPTIME-compl.
Prompt-LDL	$\wedge, \vee, \langle r \rangle, [r], \langle r \rangle_{\mathbf{P}}$	PSPACE-compl.	2EXPTIME-compl.

Table 1: The logics our work is based on.

and Prompt-LTL, respectively. Nevertheless, our definitions here are equivalent to the original definitions.

165 2.1. Robust Linear Temporal Logic

The main impetus behind the introduction of robust LTL was the need to capture the concept of robustness in temporal logics. As a first motivating example consider the LTL formula $\Box p$, stating that p holds at every position. Consequently, the formula is violated if there is a single position where p does
170 not hold. However, this is a very “mild” violation of the property and there are much more “severe” violations. As exhibited by Tabuada and Neider, there are four canonical *degrees* of violation of $\Box p$: (i) p is violated at finitely many positions, (ii) p is violated at infinitely many positions, (iii) p is violated at all but finitely many positions, and (iv) p is violated at all positions. These first
175 three degrees are captured by the LTL formulas $\Diamond \Box p$, $\Box \Diamond p$, and $\Diamond p$, which are all weakenings of $\Box p$. All five possibilities, satisfaction and four degrees of violation, are captured in robust LTL by the truth values

$$1111 \succ 0111 \succ 0011 \succ 0001 \succ 0000$$

introduced above. By design, the formula $\Box p$ of robust LTL¹ has

- truth value 1111 on all traces where p holds at all positions,
- 180 • truth value 0111 on all traces where p holds at all but finitely many positions,
- truth value 0011 on all traces where p holds at infinitely many positions and does not hold at infinitely many positions,
- truth value 0001 on all traces where p only holds at finitely many positions,
185 and
- truth value 0000 on all traces where p holds at no position.

¹Following the precedent for robust LTL, we use dots to distinguish operators of robust logics from those of classical logics throughout the paper.

As a further example, consider the formula $\Box p \rightarrow \Box q$. For this formula, the robust semantics captures the intuition described in the introduction: the implication is satisfied (i.e., has truth value 1111), if the degree of violation of the property “always q” is at most the degree of violation of the property “always p”. Thus, if p is violated finitely often, then q may also be violated finitely often (but not infinitely often) while still satisfying the implication.

Formally, the formulas of $\text{rLTL}(\Box, \Diamond)$ are given by the grammar

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \Diamond\varphi \mid \Box\varphi,$$

where p ranges over the atomic propositions in P . Note that the syntax of $\text{rLTL}(\Box, \Diamond)$ explicitly contains implication and conjunction; due to the many-valued semantics of $\text{rLTL}(\Box, \Diamond)$ introduced below, these two operators cannot be recovered from disjunction and negation. We define the size $|\varphi|$ of a formula as the number of distinct subformulas of φ .

Intuitively, conjunction and disjunction are defined as usual using minimization and maximization relying on the order of truth values indicated above while negation is based on the intuition that 1111 represents satisfaction and all other truth values represent degrees of violation. Hence, a negation $\neg\varphi$ is satisfied (i.e., has truth value 1111), if φ has truth value less than 1111, and it is violated (i.e., has truth value 0000) if φ has truth value 1111. Finally, the semantics of the eventually operator is defined as usual, i.e., the truth value of $\Diamond\varphi$ on w is the maximal truth value that is assumed by φ on some suffix of w .

This intuition is formalized in the evaluation function V^R , which maps a trace $w \in (2^P)^\omega$ and an $\text{rLTL}(\Box, \Diamond)$ formula φ to a truth value $V^R(w, \varphi)$ in \mathbb{B}_4 and which is defined as follows [22]:

- $V^R(w, p) = \begin{cases} 1111 & \text{if } p \in w(0), \\ 0000 & \text{if } p \notin w(0), \end{cases}$
- $V^R(w, \neg\varphi) = \begin{cases} 1111 & \text{if } V^R(w, \varphi) \neq 1111, \\ 0000 & \text{if } V^R(w, \varphi) = 1111, \end{cases}$
- $V^R(w, \varphi_0 \wedge \varphi_1) = \min\{V^R(w, \varphi_0), V^R(w, \varphi_1)\},$
- $V^R(w, \varphi_0 \vee \varphi_1) = \max\{V^R(w, \varphi_0), V^R(w, \varphi_1)\},$
- $V^R(w, \varphi_0 \rightarrow \varphi_1) = \begin{cases} 1111 & \text{if } V^R(w, \varphi_0) \preceq V^R(w, \varphi_1), \\ V^R(w, \varphi_1) & \text{if } V^R(w, \varphi_0) \succ V^R(w, \varphi_1), \end{cases}$
- $V^R(w, \Diamond\varphi) = b_1 b_2 b_3 b_4$ with $b_i = \max_{j \geq 0} V_i^R(w[j, \infty), \varphi)$, and
- $V^R(w, \Box\varphi) = b_1 b_2 b_3 b_4$ with
 - $b_1 = \min_{j \geq 0} V_1^R(w[j, \infty), \varphi),$
 - $b_2 = \max_{j' \in \mathbb{N}} \min_{j \geq j'} V_2^R(w[j, \infty), \varphi),$
 - $b_3 = \min_{j' \in \mathbb{N}} \max_{j \geq j'} V_3^R(w[j, \infty), \varphi),$ and
 - $b_4 = \max_{j \geq 0} V_4^R(w[j, \infty), \varphi).$

220 Here, $V_i^R(w, \varphi)$ denotes the projection of $V^R(w, \varphi)$ to its i -th component, i.e., we have

$$V^R(w, \varphi) = V_1^R(w, \varphi)V_2^R(w, \varphi)V_3^R(w, \varphi)V_4^R(w, \varphi).$$

The first bit of the semantics captures the classical semantics of LTL, i.e., we have $V_1^R(w, \varphi) = 1$ if and only if w satisfies φ classically. Intuitively, the next three bits are obtained by weakening the semantics of the subformulas of the form $\Box\varphi$: instead of (classically) requiring every position to satisfy φ , the second bit is one if almost all positions satisfy φ (i.e., $\Diamond\Box\varphi$ holds), the third bit is one if infinitely many positions satisfy φ (i.e., $\Box\Diamond\varphi$ holds), and the fourth bit is one if at least one position satisfies φ (i.e., $\Diamond\varphi$ holds). Note that the semantics of negation and implication are also non-classical and break the intuition given above, e.g., for formulas of the form $\neg\Box\varphi$. For a full motivation and explanation of the semantics, we refer to the original work introducing $\text{rLTL}(\Box, \Diamond)$ [22] as well as follow-up work [23, 24, 25].

Verification problems with $\text{rLTL}(\Box, \Diamond)$ specifications have a threshold $\beta \in \mathbb{B}_4$ as an additional input and ask every trace to evaluate to at least β . Tabuada and Neider showed that the model checking problem with $\text{rLTL}(\Box, \Diamond)$ specifications can be solved in polynomial space² and that infinite games with $\text{rLTL}(\Box, \Diamond)$ specifications can be solved in doubly-exponential time [22]. The lower bounds presented in Table 1 are derived from the special case of $\text{LTL}(\Box, \Diamond)$ [26, 27], which is a fragment of $\text{rLTL}(\Box, \Diamond)$.

240 When Tabuada and Neider introduced robust LTL, they first considered the fragment $\text{rLTL}(\Box, \Diamond)$ without the next, until, and release operators, which already captures the most interesting problems arising from adding robustness. Then, they added the missing operators and studied the full logic [22]. Here, we follow their approach and only consider generalizations of the fragment $\text{rLTL}(\Box, \Diamond)$ that only contains the temporal operators \Box and \Diamond . We comment on the effect of this restriction when defining the combinations of logics.

2.2. Linear Dynamic Logic

250 The logic LDL has only two temporal operators, $\langle r \rangle$ and $[r]$, which can be understood as guarded variants of the classical eventually and always operators from LTL, respectively. Both are guarded by regular expressions r over the atomic propositions that may contain tests, which are again LDL formulas. These two operators together with Boolean connectives capture the full expressive power of the ω -regular expressions, i.e., LDL exceeds the expressiveness of LTL.

The formulas of LDL are given by the grammar

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \langle r \rangle \varphi \mid [r] \varphi \\ r &::= \phi \mid \varphi? \mid r + r \mid r ; r \mid r^* \end{aligned}$$

²Tabuada and Neider only showed that their algorithm runs in exponential time, but using standard on-the-fly techniques [5] it can also be implemented in polynomial space.

where p ranges over the atomic propositions in P and where ϕ ranges over arbitrary propositional formulas over P . The regular expressions have two types of atoms: propositional formulas ϕ over the atomic propositions and tests $\varphi?$, where φ is again an LDL formula. As we will see later, the semantics of these two kinds of atoms differ significantly. We refer to formulas of the form $\langle r \rangle \varphi$ and $[r] \varphi$ as diamond formulas and box formulas, respectively. In both cases, we call r the guard of the operator.

We denote the set of subformulas of φ by $\text{cl}(\varphi)$. Guards are not subformulas, but the formulas appearing in the tests are, e.g., we have

$$\text{cl}(\langle p? ; q \rangle p') = \{p, p', \langle p? ; q \rangle p'\}.$$

The size $|\varphi|$ of φ is the sum of $|\text{cl}(\varphi)|$ and the sum of the lengths of the guards appearing in φ (measured in the number of operators), taking each occurrence of a guard in the syntax tree (not the syntax DAG like for subformulas) into account.

Formally, a formula $\langle r \rangle \varphi$ is satisfied by a trace w , if there is some j such that the prefix $w[0, j)$ matches the regular expression r and the corresponding suffix $w[j, \infty)$ satisfies φ . Dually, a formula $[r] \varphi$ is satisfied by a trace w if for every j with $w[0, j)$ matching r , $w[j, \infty)$ satisfies φ . Thus, while the classical eventually and always operator range over all positions, the operators of LDL range only over those positions whose induced prefix matches the guard of the operator.

Analogously to the definition for $\text{rLTL}(\Box, \Diamond)$, and slightly non-standard, we define the semantics of LDL by specifying an evaluation function V^{D} mapping a trace w and a formula φ to a truth value from \mathbb{B} denoting whether w satisfies φ or not. Also, our presentation of the semantics here is slightly cumbersome, in particular the definition for the implication, again to align with the definition for $\text{rLTL}(\Box, \Diamond)$. Nevertheless, our definition below is equivalent to the classical semantics of LDL (cf. [8, 28, 6]) via a satisfaction relation \models in the following sense: we have $V^{\text{D}}(w, \varphi) = 1$ if and only if $w \models \varphi$.

- $V^{\text{D}}(w, p) = \begin{cases} 1 & \text{if } p \in w(0), \\ 0 & \text{if } p \notin w(0), \end{cases}$
- $V^{\text{D}}(w, \neg\varphi) = \begin{cases} 1 & \text{if } V^{\text{D}}(w, \varphi) = 0, \\ 0 & \text{if } V^{\text{D}}(w, \varphi) = 1, \end{cases}$
- $V^{\text{D}}(w, \varphi_0 \wedge \varphi_1) = \min\{V^{\text{D}}(w, \varphi_0), V^{\text{D}}(w, \varphi_1)\},$
- $V^{\text{D}}(w, \varphi_0 \vee \varphi_1) = \max\{V^{\text{D}}(w, \varphi_0), V^{\text{D}}(w, \varphi_1)\},$
- $V^{\text{D}}(w, \varphi_0 \rightarrow \varphi_1) = \begin{cases} 1 & \text{if } V^{\text{D}}(w, \varphi_0) \leq V^{\text{D}}(w, \varphi_1), \\ V^{\text{D}}(w, \varphi_1) & \text{if } V^{\text{D}}(w, \varphi_0) > V^{\text{D}}(w, \varphi_1), \end{cases}$
- $V^{\text{D}}(w, \langle r \rangle \varphi) = \max_{j \in \mathcal{R}(w, r)} V^{\text{D}}(w[j, \infty), \varphi),$ and
- $V^{\text{D}}(w, [r] \varphi) = \min_{j \in \mathcal{R}(w, r)} V^{\text{D}}(w[j, \infty), \varphi).$

Here, the match set $\mathcal{R}(w, r) \subseteq \mathbb{N}$ contains all positions j such that $w[0, j)$ matches r . Recall that $w[0, j)$ denotes the prefix of w of length j , i.e., $w[0, j) =$

$w(0) \cdots w(j-1)$. In particular, $w[0, 0]$ is empty and $w[0, \infty)$ is w . Now, $\mathcal{R}(w, r)$ is defined inductively as follows:

- 295 • $\mathcal{R}(w, \phi) = \{1\}$ if $w(0) \models \phi$ (i.e., we evaluate ϕ in standard Boolean semantics) and $\mathcal{R}(w, \phi) = \emptyset$ otherwise, for propositional ϕ .
- $\mathcal{R}(w, \varphi?) = \{0\}$ if $V^D(w, \varphi) = 1$ and $\mathcal{R}(w, \varphi?) = \emptyset$ otherwise.
- $\mathcal{R}(w, r_0 + r_1) = \mathcal{R}(w, r_0) \cup \mathcal{R}(w, r_1)$.
- $\mathcal{R}(w, r_0; r_1) = \{j_0 + j_1 \mid j_0, j_1 \geq 0 \text{ and } j_0 \in \mathcal{R}(w, r_0) \text{ and } j_1 \in \mathcal{R}(w[j_0, \infty), r_1)\}$. Thus, for j to be in $\mathcal{R}(w, r_0; r_1)$, it has to be the sum of natural numbers j_0 and j_1 such that w has a prefix of length j_0 that matches r_0 and $w[j_0, \infty)$ has a prefix of length j_1 that matches r_1 .
- 300 • $\mathcal{R}(w, r^*) = \{0\} \cup \{j_1 + \cdots + j_\ell \mid 0 \leq j_{\ell'} \in \mathcal{R}(w[j_1 + \cdots + j_{\ell'-1}, \infty), r) \text{ for all } \ell' \in \{1, \dots, \ell\}\}$, where we use $j_1 + \cdots + j_0 = 0$. Thus, for j to be in $\mathcal{R}(w, r^*)$, it has to be expressible as $j = j_1 + \cdots + j_\ell$ with non-negative $j_{\ell'}$ such that the prefix of w of length j_1 matches r , the prefix of length j_2 of $w[j_1, \infty)$ matches r , and in general, the prefix of length $j_{\ell'}$ of $w[j_1 + \cdots + j_{\ell'-1}, \infty)$ matches r , for every $\ell' \in \{1, \dots, \ell\}$.
- 305

Due to tests, membership of some j in $\mathcal{R}(w, r)$ does, in general, not only depend on the prefix $w[0, j]$, but on the complete trace w . Also, the semantics of the propositional atom ϕ differ from the semantics of the test $\phi?$: the former consumes an input letter, while tests do not. Hence, the guards of LDL feature both kinds of atoms.

Fix some trace w , a formula φ , and a guard r . We say that a position j of w is an r -match if $j \in \mathcal{R}(w, r)$. Further, j is a φ -satisfying position of w if $V^D(w[j, \infty), \varphi) = 1$. Thus, the formula $\langle r \rangle \varphi$ requires some φ -satisfying r -match to exist. Dually, $[r] \varphi$ requires every r -match of w to be φ -satisfying (in particular, this is the case if there is no r -match in w). Thus, the diamond operator generalizes the eventually operator and the box operator generalizes the always operator, which are the respective special cases for a trivial guard that matches every position, e.g., \mathbf{tt}^* . Similarly, the next, until, and release operator of LTL can be expressed in LDL (the latter two use tests in the guards). Thus, LTL is a fragment of LDL. Furthermore, it is known that LDL captures the ω -regular languages [8].

Model checking against LDL specifications is PSPACE-complete and solving LDL games is 2EXPTIME-complete [6, 8].

2.3. Prompt Linear Temporal Logic

To express timing constraints, the logic Prompt-LTL adds the prompt-eventually operator $\diamond_{\mathbf{p}}$ to LTL. Intuitively, the new operator requires its argument to be satisfied within a bounded number of steps.

330 The formulas of Prompt-LTL are given by the grammar

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \bigcirc \varphi \mid \varphi \mathbf{U} \varphi \mid \varphi \mathbf{R} \varphi \mid \diamond_{\mathbf{p}} \varphi$$

where p ranges over the atomic propositions in P . The size $|\varphi|$ of a formula φ is defined as the number of its distinct subformulas.

In Prompt-LTL, formulas are in negation normal form and implication is disallowed. Both requirements are necessary to preserve monotonicity of the prompt-eventually $\diamond_{\mathbf{p}}$ with respect to the parameter k bounding it (Alur et al. [10] provide a detailed discussion). We define the shorthands $\diamond\varphi = \mathbf{tt}\mathbf{U}\varphi$ and $\square\varphi = \mathbf{ff}\mathbf{R}\varphi$.

Again, we define the semantics by an evaluation function $V^{\mathbf{P}}$ mapping a trace $w \in (2^P)^\omega$, a bound $k \in \mathbb{N}$ for the prompt operators, and a formula φ to a truth value in \mathbb{B} (which is again equivalent to the standard definition). This function is defined as usual for all Boolean and standard temporal operators (ignoring the bound k), while a formula $\diamond_{\mathbf{p}}\varphi$ is satisfied with respect to the bound k if φ holds within the next k steps, i.e., the prompt-eventually behaves like the classical eventually with a bounded scope [12]:

$$\begin{aligned}
345 \quad & \bullet V^{\mathbf{P}}(w, k, p) = \begin{cases} 1 & \text{if } p \in w(0), \\ 0 & \text{if } p \notin w(0), \end{cases} & \bullet V^{\mathbf{P}}(w, k, \neg p) = \begin{cases} 1 & \text{if } p \notin w(0), \\ 0 & \text{if } p \in w(0), \end{cases} \\
& \bullet V^{\mathbf{P}}(w, k, \varphi_0 \wedge \varphi_1) = \min\{V^{\mathbf{P}}(w, k, \varphi_0), V^{\mathbf{P}}(w, k, \varphi_1)\}, \\
& \bullet V^{\mathbf{P}}(w, k, \varphi_0 \vee \varphi_1) = \max\{V^{\mathbf{P}}(w, k, \varphi_0), V^{\mathbf{P}}(w, k, \varphi_1)\}, \\
& \bullet V^{\mathbf{P}}(w, k, \mathbf{O}\varphi) = V^{\mathbf{P}}(w[1, \infty), k, \varphi), \\
& \bullet V^{\mathbf{P}}(w, k, \varphi_0 \mathbf{U}\varphi_1) = \max_{j \in \mathbb{N}} \min\{V^{\mathbf{P}}(w[j, \infty), k, \varphi_1), \min_{0 \leq j' < j} V^{\mathbf{P}}(w[j', \infty)), k, \varphi_0\}, \\
350 \quad & \bullet V^{\mathbf{P}}(w, k, \varphi_0 \mathbf{R}\varphi_1) = \min_{j \in \mathbb{N}} \max\{V^{\mathbf{P}}(w[j, \infty), k, \varphi_1), \max_{0 \leq j' < j} V^{\mathbf{P}}(w[j', \infty)), k, \varphi_0\}, \\
& \bullet V^{\mathbf{P}}(w, k, \diamond_{\mathbf{p}}\varphi) = \max_{0 \leq j \leq k} V^{\mathbf{P}}(w[j, \infty), k, \varphi).
\end{aligned}$$

In verification problems for Prompt-LTL, the bound k on the prompt-eventualities is existentially quantified. Kupferman et al. proved that Prompt-LTL model checking is PSPACE-complete and that solving games with Prompt-LTL winning conditions is 2EXPTIME-complete [12].³

2.4. Prompt Linear Dynamic Logic

In our proofs, we also use Prompt-LDL, which can be seen as a combination of LDL and Prompt-LTL. This logic has been studied by Faymonville and Zimmermann [6] as a fragment of Parametric LTL [10] (although the logic has never been explicitly named).

The formulas of Prompt-LDL are given by the grammar

$$\begin{aligned}
\varphi & ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle r \rangle \varphi \mid [r] \varphi \mid \langle r \rangle_{\mathbf{p}} \varphi \\
r & ::= \phi \mid \varphi? \mid r + r \mid r ; r \mid r^*
\end{aligned}$$

³Instead of games, they actually considered the related framework of realizability, an abstract type of game without underlying graph. However, realizability and graph-based games are interreducible (also, see [29]).

where p again ranges over the atomic propositions in P and ϕ ranges over propositional formulas over P . As in Prompt-LTL, we have to disallow arbitrary negations and implications. The size of a formula is defined as for LDL.

Furthermore, the semantics of Prompt-LDL is obtained by combining the one of LDL and the one of Prompt-LTL: Again, we define an evaluation function V^{PD} mapping a trace w , a bound k , and a formula φ to a truth value.

- $V^{\text{PD}}(w, k, p) = \begin{cases} 1 & \text{if } p \in w(0), \\ 0 & \text{if } p \notin w(0), \end{cases}$
- $V^{\text{PD}}(w, k, \neg p) = \begin{cases} 1 & \text{if } p \notin w(0), \\ 0 & \text{if } p \in w(0), \end{cases}$
- $V^{\text{PD}}(w, k, \varphi_0 \wedge \varphi_1) = \min\{V^{\text{PD}}(w, k, \varphi_0), V^{\text{PD}}(w, k, \varphi_1)\},$
- $V^{\text{PD}}(w, k, \varphi_0 \vee \varphi_1) = \max\{V^{\text{PD}}(w, k, \varphi_0), V^{\text{PD}}(w, k, \varphi_1)\},$
- $V^{\text{PD}}(w, k, \langle r \rangle \varphi) = \max_{j \in \mathcal{R}(w, k, r)} V^{\text{PD}}(w[j, \infty), k, \varphi),$
- $V^{\text{PD}}(w, k, [r] \varphi) = \min_{j \in \mathcal{R}(w, k, r)} V^{\text{PD}}(w[j, \infty), k, \varphi),$ and
- $V^{\text{PD}}(w, k, \langle r \rangle_{\mathbf{p}} \varphi) = \max_{j \in \mathcal{R}(w, k, r) \cap \{0, \dots, k\}} V^{\text{PD}}(w[j, \infty), k, \varphi).$

Here, $\mathcal{R}(w, k, r)$ is defined as $\mathcal{R}(w, r)$, but propagates the bound k to evaluate tests. Hence, we define $\mathcal{R}(w, k, \varphi?) = \{0\}$ if $V^{\text{D}}(w, k, \varphi) = 1$ and $\mathcal{R}(w, k, \varphi?) = \emptyset$ otherwise. All other cases are defined as for LDL, but propagate the bound k .

Prompt-LDL as defined here is a syntactic fragment of Parametric LDL [6] and subsumes LTL. Hence, its model checking problem is PSPACE-complete and the synthesis problem is 2EXPTIME-complete. Here, the bound k is again uniformly existentially quantified in verification problems.

3. Robust and Prompt Linear Temporal Logic

We begin our treatment of combinations of the three basic logics by introducing robust semantics for Prompt-LTL, obtaining the logic rPrompt-LTL. To this end, we add the prompt-eventually operator to rLTL(\Box, \Diamond) while disallowing implications and restricting negation to retain decidability (cf. [10]). The formulas of rPrompt-LTL are given by

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Diamond \varphi \mid \Box \varphi \mid \Diamond_{\mathbf{p}} \varphi,$$

where p ranges over the set P of atomic propositions. The size $|\varphi|$ of a formula φ is the number of its distinct subformulas.

The semantics of rPrompt-LTL is given by an evaluation function V^{RP} mapping a trace w , a bound k for the prompt-eventuallys, and a formula φ to a truth value in \mathbb{B}_4 . To simplify our notation, we write $V_i^{\text{RP}}(w, k, \varphi)$ for $i \in \{1, 2, 3, 4\}$ to denote the i -th bit of $V^{\text{RP}}(w, k, \varphi)$, i.e.,

$$V^{\text{RP}}(w, k, \varphi) = V_1^{\text{RP}}(w, k, \varphi) V_2^{\text{RP}}(w, k, \varphi) V_3^{\text{RP}}(w, k, \varphi) V_4^{\text{RP}}(w, k, \varphi).$$

The semantics of Boolean connectives as well as of the eventually and always operators is defined as for robust LTL. The motivation behind these definitions

is carefully and convincingly discussed by Tabuada and Neider [22]. The semantics of the prompt-eventually operator bounds its scope to the next k positions as in classical Prompt-LTL [12].

- $V^{\text{RD}}(w, k, p) = \begin{cases} 1111 & \text{if } p \in w(0), \\ 0000 & \text{if } p \notin w(0), \end{cases}$
- $V^{\text{RD}}(w, k, \neg p) = \begin{cases} 1111 & \text{if } p \notin w(0), \\ 0000 & \text{if } p \in w(0), \end{cases}$
- $V^{\text{RD}}(w, k, \varphi_0 \wedge \varphi_1) = \min\{V^{\text{RD}}(w, k, \varphi_0), V^{\text{RD}}(w, k, \varphi_1)\},$
- $V^{\text{RD}}(w, k, \varphi_0 \vee \varphi_1) = \max\{V^{\text{RD}}(w, k, \varphi_0), V^{\text{RD}}(w, k, \varphi_1)\},$
- $V^{\text{RP}}(w, k, \diamond \varphi) = b_1 b_2 b_3 b_4$ where $b_i = \max_{j \in \mathbb{N}} V_i^{\text{RP}}(w[j, \infty), k, \varphi)$,⁴ and
- $V^{\text{RP}}(w, k, \square \varphi) = b_1 b_2 b_3 b_4$ where
 - $b_1 = \min_{j \in \mathbb{N}} V_1^{\text{RP}}(w[j, \infty), k, \varphi)$, i.e., $b_1 = 1$ iff φ holds always,
 - $b_2 = \max_{j' \in \mathbb{N}} \min_{j' \leq j} V_2^{\text{RP}}(w[j, \infty), k, \varphi)$, i.e., $b_2 = 1$ iff φ holds almost always,
 - $b_3 = \min_{j' \in \mathbb{N}} \max_{j' \leq j} V_3^{\text{RP}}(w[j, \infty), k, \varphi)$, i.e., $b_3 = 1$ iff φ holds infinitely often, and
 - $b_4 = \max_{j \in \mathbb{N}} V_4^{\text{RP}}(w[j, \infty), k, \varphi)$ i.e., $b_4 = 1$ iff φ holds at least once.
- $V^{\text{RP}}(w, k, \diamond_{\mathbf{p}} \varphi) = b_1 b_2 b_3 b_4$ where $b_i = \max_{0 \leq j \leq k} V_i^{\text{RP}}(w[j, \infty), k, \varphi)$.

It is easy to verify that $V^{\text{RP}}(w, k, \varphi)$ is well-defined, i.e., $V^{\text{RP}}(w, k, \varphi) \in \mathbb{B}_4$ for all w, k , and φ .

Example 1. Consider the formula $\varphi = \square \diamond_{\mathbf{p}} s$, where we interpret occurrences of the atomic proposition s as synchronizations. Then, the different degrees of satisfaction of the formula express the following possibilities, when evaluating it with respect to $k \in \mathbb{N}$: (i) the distance between synchronizations is bounded by k , (ii) from some point onwards, the distance between synchronizations is bounded by k , (iii) there are infinitely many synchronizations, and (iv) there is at least one synchronization. Note that the last two possibilities are independent of k , which is explained by simple logical equivalences, e.g., the third possibility reads actually as follows: there are infinitely many positions such that a synchronization occurs within distance k . However, it is easy to see that is equivalent to the property stated above.

In the next two sections, we solve the model checking problem and the synthesis problem for rPrompt-LTL. To this end, we translate every rPrompt-LTL formula into a sequence of five Prompt-LTL formulas that capture the five degrees of satisfaction and violation by making the semantics of the robust always operator explicit. This is a straightforward generalization of the, in the terms of the introduction, reduction-based approach to robust LTL [22].

⁴This definition is equivalent to $V^{\text{RP}}(w, k, \diamond \varphi) = \max_{j \in \mathbb{N}} V^{\text{RP}}(w[j, \infty), k, \varphi)$ due to monotonicity of the truth values, which is closer to the classical semantics of the eventually operator. A similar equivalence holds for $\diamond_{\mathbf{p}} \varphi$.

Lemma 1. For every rPrompt-LTL formula φ and every $\beta \in \mathbb{B}_4$, there is a Prompt-LTL formula φ_β of size $\mathcal{O}(|\varphi|)$ such that $V^{\text{RP}}(w, k, \varphi) \succeq \beta$ if and only if $V^{\text{P}}(w, k, \varphi_\beta) = 1$.

Proof. If $\beta = 0000$, then we can pick $\varphi_\beta = \mathbf{tt}$, independently of φ . Otherwise, we obtain the result by induction over the construction of φ implementing the intuition behind the robust semantics:

- $p_\beta = p$ and $(\neg p)_\beta = \neg p$ for all atomic propositions $p \in P$ and all $\beta \succ 0000$.
- $(\varphi_0 \wedge \varphi_1)_\beta = (\varphi_0)_\beta \wedge (\varphi_1)_\beta$ for all $\beta \succ 0000$.
- $(\varphi_0 \vee \varphi_1)_\beta = (\varphi_0)_\beta \vee (\varphi_1)_\beta$ for all $\beta \succ 0000$.
- $(\diamond \varphi)_\beta = \diamond(\varphi_\beta)$ for all $\beta \succ 0000$.
- $(\square \varphi)_{1111} = \square(\varphi_{1111})$.
- $(\square \varphi)_{0111} = \diamond \square(\varphi_{0111})$.
- $(\square \varphi)_{0011} = \square \diamond(\varphi_{0011})$.
- $(\square \varphi)_{0001} = \diamond(\varphi_{0001})$.
- $(\diamond_{\mathbf{p}} \varphi)_\beta = \diamond_{\mathbf{p}}(\varphi_\beta)$ for all $\beta \succ 0000$.

A straightforward induction shows that the resulting formula has the desired properties. \square

Note that the logic rLTL(\square, \diamond) is not a fragment of rPrompt-LTL as we have to disallow negation and implication to retain decidability [10]. Conversely, Prompt-LTL is also not a fragment of rPrompt-LTL as we omitted the next, until, and release operator. However, we present a reduction-based approach from rPrompt-LTL to Prompt-LTL. Thus, one could easily add the additional temporal operators to rPrompt-LTL while maintaining the result of Lemma 1. We prefer not to do so for the sake of accessibility and brevity.

3.1. Model Checking

Let us now consider the rPrompt-LTL model checking problem, which asks whether all executions of a given finite transition system satisfy a given specification expressed as an rPrompt-LTL formula with truth value at least $\beta \in \mathbb{B}_4$. More formally, we assume the system under consideration to be modeled as a (labeled and initialized) transition system $\mathcal{S} = (S, s_I, E, \lambda)$ over P consisting of a finite set S of states containing the initial state s_I , a directed edge relation $E \subseteq S \times S$, and a state labeling $\lambda: S \rightarrow 2^P$ that maps each state to the set of atomic propositions that hold true in this state. A path through \mathcal{S} is a sequence $\rho = s_0 s_1 s_2 \dots$ satisfying $s_0 = s_I$ and $(s_j, s_{j+1}) \in E$ for every $j \in \mathbb{N}$, and $\Pi_{\mathcal{S}}$ denotes the set of all paths through \mathcal{S} . Finally, the trace of a path $\rho = s_0 s_1 s_2 \dots \in \Pi_{\mathcal{S}}$ is the sequence $\lambda(\rho) = \lambda(s_0)\lambda(s_1)\lambda(s_2)\dots$ of labels induced by ρ .

465 **Problem 1.** Let φ be an rPrompt-LTL formula, \mathcal{S} a transition system, and $\beta \in \mathbb{B}_4$. Is there a natural number k such that $V^{\text{RP}}(\lambda(\rho), k, \varphi) \succeq \beta$ holds true for all paths $\rho \in \Pi_{\mathcal{S}}$?

Our solution relies on Lemma 1 and on Prompt-LTL model checking being in PSPACE [12].

470 **Theorem 1.** rPrompt-LTL model checking is in PSPACE.

Proof. By Lemma 1, there exists a $k \in \mathbb{N}$ such that $V^{\text{RP}}(\lambda(\rho), k, \varphi) \succeq \beta$ holds true for all paths $\rho \in \Pi_{\mathcal{S}}$ if and only if there exists a $k \in \mathbb{N}$ such that $V^{\text{P}}(\lambda(\rho), k, \varphi_{\beta}) = 1$ for all paths $\rho \in \Pi_{\mathcal{S}}$. The latter is an instance of the Prompt-LTL model checking problem, which is in PSPACE [12]. \square

475 We do not claim PSPACE-hardness because model checking the fragment of LTL with disjunction, conjunction, always, and eventually operators only (and classical semantics) is NP-complete [27]. Since this fragment can be embedded into rPrompt-LTL (via a translation of this LTL fragment into rPrompt-LTL using techniques similar to those presented by Tabuada and Neider [22] for translating LTL into rLTL(\square, \diamond)), we obtain at least NP-hardness for Problem 1. As we have no next, until, and release operators (by our own volition), we cannot easily claim PSPACE-hardness. In contrast, the solution of the Prompt-LTL model checking problem consists of a reduction to LTL model checking that introduces until operators (see [12]). Hence, we leave the fragment mentioned above, for which NP membership is known. However, adding next, until, and release to rPrompt-LTL yields a PSPACE-hard model checking problem.

3.2. Synthesis

Next, we consider the problem of synthesizing reactive controllers from rPrompt-LTL specifications. In this context, we rely on the classical reduction from reactive synthesis to infinite-duration two-player games over finite graphs. In particular, we show how to construct a finite-state winning strategy for games with rPrompt-LTL winning conditions, which immediately correspond to implementations of reactive controllers. Throughout this section, we assume familiarity with games over finite graphs (see, e.g., [30, Chapter 2]).

495 We consider rPrompt-LTL games over P , which are triples $\mathcal{G} = (G, \varphi, \beta)$ consisting of a labeled game graph G , an rPrompt-LTL formula φ , and a truth value $\beta \in \mathbb{B}_4$. A labeled game graph $G = (V_0, V_1, E, \lambda)$ consists of a directed graph $(V_0 \cup V_1, E)$, two finite, disjoint sets of vertices V_0 and V_1 , and a function $\lambda: V_0 \cup V_1 \rightarrow 2^P$ mapping each vertex v to the set $\lambda(v)$ of atomic propositions that hold true in v . We denote the set of all vertices by $V = V_0 \cup V_1$ and assume that game graphs do not have terminal vertices, i.e., $\{v\} \times V \cap E \neq \emptyset$ for each $v \in V$.

505 As in the classical setting, rPrompt-LTL games are played by two players, Player 0 and Player 1, who move a token along the edges of the game graph ad infinitum (if the token is currently placed on a vertex $v \in V_i$, $i \in \{0, 1\}$, then Player i decides the next move). The resulting infinite sequence $\rho = v_0 v_1 v_2 \dots \in$

V^ω of vertices is called a play and induces a trace $\lambda(\rho) = \lambda(v_0)\lambda(v_1)\lambda(v_2)\cdots \in (2^P)^\omega$.

A strategy of Player 0 is a mapping $f: V^*V_0 \rightarrow V$ that prescribes where
510 to move the token depending on the finite play prefix constructed so far. A
play $v_0v_1v_2\cdots$ is played according to f if $v_{j+1} = f(v_0\cdots v_j)$ for every j with
 $v_j \in V_0$. A strategy f of Player 0 is winning from a vertex $v \in V$ if there is
a $k \in \mathbb{N}$ such that all plays ρ that start in v and that are played according
to f satisfy $V^{\text{RP}}(\lambda(\rho), k, \varphi) \succeq \beta$, i.e., the evaluation of φ with respect to k on
515 $\lambda(\rho)$ determines the winner of the play ρ . Further, a (winning) strategy is a
finite-state strategy if there exists a finite-state machine computing it in the
usual sense (see [30, Chapter 2] for details).

We are interested in solving rPrompt-LTL games, i.e., in solving the following
problem.

520 **Problem 2.** *Let \mathcal{G} be an rPrompt-LTL game and v a vertex. Determine whether
Player 0 has a winning strategy for \mathcal{G} from v and compute a finite-state winning
strategy if so.*

Again, our solution to this problem relies on Lemma 1 and the fact that
solving Prompt-LTL games is in 2EXPTIME [12, 29].

525 **Theorem 2.** *Solving rPrompt-LTL games is 2EXPTIME-complete.*

Proof. The lower bound follows from the special case of LTL(\square, \diamond) [26], which
is a fragment of rPrompt-LTL. On the other hand, the upper bound is again
proven by a reduction to Prompt-LTL: Player 0 having a winning strategy
for (G, φ, β) from v is equivalent to her having a winning strategy for the
530 Prompt-LTL game (G, φ_β) from v . The latter problem can be solved in doubly-
exponential time and a finite-state strategy can effectively be computed [29]. \square

Here we have a matching lower bound, as solving games with LTL conditions
without next, until, and release is already 2EXPTIME-hard [26].

4. Robust Linear Dynamic Logic

Next, we “robustify” LDL by generalizing the ideas underlying robust LTL
to LDL, obtaining the logic rLDL. Again, following the precedent of robust
LTL, we equip robust operators with dots to distinguish them from non-robust
ones. The formulas of rLDL are given by the grammar

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \langle r \cdot \rangle \varphi \mid [\cdot r \cdot] \varphi \\ r &::= \phi \mid \varphi? \mid r + r \mid r ; r \mid r^* \end{aligned}$$

535 where p ranges over the atomic propositions in P and ϕ over propositional for-
mulas over P . We refer to formulas of the form $\langle r \cdot \rangle \varphi$ and $[\cdot r \cdot] \varphi$ as diamond
formulas and box formulas, respectively. In both cases, r is the guard of the
operator. An atom $\varphi?$ of a regular expression is a test. We use the abbrevi-
ations $\mathbf{tt} = p \vee \neg p$ and $\mathbf{ff} = p \wedge \neg p$ for some $p \in P$ and note that both are

540 formulas and guards. We denote the set of subformulas of φ by $\text{cl}(\varphi)$. Guards are not subformulas, but the formulas appearing in the tests are, e.g., we have

$$\text{cl}(\langle \cdot p? ; q \cdot \rangle p') = \{p, p', \langle \cdot p? ; q \cdot \rangle p'\}.$$

The size $|\varphi|$ of φ is the sum of $|\text{cl}(\varphi)|$ and the sum of the lengths of the guards appearing in φ (measured in the number of operators), taking each occurrence of a guard in the syntax tree (not the syntax DAG like for subformulas) into
545 account.

Before we introduce the semantics of rLDL we first recall the semantics of the robust always operator $\Box\varphi$ in robust LTL. To this end, call a position j of a trace φ -satisfying if the suffix starting at position j satisfies φ . Now, the robust semantics are based on the following five cases, where the latter four distinguish
550 various degrees of violating the formula $\Box\varphi$: either all positions are φ -satisfying (\Box), almost all positions are φ -satisfying ($\Diamond\Box$), infinitely many positions are φ -satisfying ($\Box\Diamond$), some position is φ -satisfying (\Diamond), or no position is φ -satisfying.

A similar approach for a formula $[\cdot r \cdot]\varphi$ would be to consider the following possibilities, where a position j of a trace w is an r -match if the prefix of w
555 up to and including position $j - 1$ is in the language of r : all r -matches are φ -satisfying, almost all r -matches are φ -satisfying, infinitely many r -matches are φ -satisfying, some r -match is φ -satisfying, or no r -match is φ -satisfying. On a trace w with infinitely many r -matches, this is the natural generalization of the robust semantics. A trace, however, may only contain finitely many r -matches,
560 or none at all. In the former case, there are not infinitely many φ -satisfying r -matches, but all r -matches could satisfy φ . Thus, the monotonicity of the cases is violated. We overcome this by interpreting “almost all” as “all” and “infinitely many” as “some” if there are only finitely many r -matches.⁵

Also, the guard r may contain tests, which have to be evaluated to determine
565 whether a position is an r -match. For this, we have to use the appropriate semantics for the robust box operator. For example, if we interpret $[\cdot r \cdot]\varphi$ to mean “almost all r -matches satisfy φ ”, then the robust box operators in tests of r are evaluated with this interpretation as well. This may, however, violate monotonicity (see Example 3), which we therefore hardcode in the semantics.

570 We now formalize the informal description above and subsequently show that this formalization satisfies all desired properties. To this end, we again define an evaluation function V^{RD} mapping a trace w and a formula φ to a truth value. Also, we again denote the projection of $V^{\text{RD}}(w, \varphi)$ to its i -th bit by $V_i^{\text{RD}}(w, \varphi)$. For atomic propositions, conjunction, disjunction, negation, and implication,
575 the definition is the same as for robust LTL on Page 7.

To define the semantics of the diamond and the box operator, we need to first define the semantics of the guards: The match set $\mathcal{R}_i^{\text{RD}}(w, r) \subseteq \mathbb{N}$ for

⁵There is an alternative definition inspired by the semantics of LTL on finite traces: Here, both $\Diamond\Box\varphi$ and $\Box\Diamond\varphi$ are equivalent to “ φ holds at the last position”. This suggests interpreting “almost all r -matches are φ -satisfying” and “infinitely many r -matches are φ -satisfying” as “the last r -match is φ -satisfying” in case there are only finitely many r -matches. Arguably, this definition is less intuitive than the one we propose to pursue.

$i \in \{1, 2, 3, 4\}$ contains all positions j of w such that $w[0, j]$ matches r (with tests in r being evaluated depending on the value of i) and is defined inductively as follows:

- $\mathcal{R}_i^{\text{RD}}(w, \phi) = \{1\}$ if $w(0) \models \phi$ and $\mathcal{R}_i^{\text{RD}}(\phi, w) = \emptyset$ otherwise, for propositional ϕ .
- $\mathcal{R}_i^{\text{RD}}(w, \varphi?) = \{0\}$ if $V_i^{\text{RD}}(w, \varphi) = 1$ and $\mathcal{R}_i^{\text{RD}}(w, \varphi?) = \emptyset$ otherwise.
- $\mathcal{R}_i^{\text{RD}}(w, r_0 + r_1) = \mathcal{R}_i^{\text{RD}}(w, r_0) \cup \mathcal{R}_i^{\text{RD}}(w, r_1)$.
- $\mathcal{R}_i^{\text{RD}}(w, r_0; r_1) = \{j_0 + j_1 \mid j_0, j_1 \geq 0 \text{ and } j_0 \in \mathcal{R}_i^{\text{RD}}(w, r_0) \text{ and } j_1 \in \mathcal{R}_i^{\text{RD}}(w[j_0, \infty), r_1)\}$, i.e., for j to be in $\mathcal{R}_i^{\text{RD}}(w, r_0; r_1)$, it has to be the sum of natural numbers j_0 and j_1 such that w has a prefix of length j_0 that matches r_0 and $w[j_0, \infty)$ has a prefix of length j_1 that matches r_1 (where in both cases the tests are again evaluated depending on the value of i).
- $\mathcal{R}_i^{\text{RD}}(w, r^*) = \{0\} \cup \{j_1 + \dots + j_\ell \mid 0 \leq j_{\ell'} \in \mathcal{R}_i^{\text{RD}}(w[j_1 + \dots + j_{\ell'-1}, \infty), r) \text{ for all } \ell' \in \{1, \dots, \ell\}\}$, where we use $j_1 + \dots + j_0 = 0$. Thus, for j to be in $\mathcal{R}_i^{\text{RD}}(w, r^*)$, it has to be expressible as $j = j_1 + \dots + j_\ell$ with non-negative $j_{\ell'}$ such that the prefix of w of length j_1 matches r , the prefix of length j_2 of $w[j_1, \infty)$ matches r , and in general, the prefix of length $j_{\ell'}$ of $w[j_1 + \dots + j_{\ell'-1}, \infty)$ matches r , for every $\ell' \in \{1, \dots, \ell\}$ (where the tests are evaluated depending on i).

Due to tests, membership of j in $\mathcal{R}_i^{\text{RD}}(w, r)$ does, in general, not only depend on the prefix $w[0, j]$, but on the complete trace w . Also, the semantics of the propositional atom ϕ differs from the semantics of the test $\phi?$: the former consumes an input letter, while the latter one does not. Thus, rLDL (as LDL) features both kinds of atoms. We define the intuition given above via

- $V^{\text{RD}}(w, \langle r \cdot \rangle \varphi) = b_1 b_2 b_3 b_4$ where $b_i = \max_{j \in \mathcal{R}_i^{\text{RD}}(w, r)} V_i^{\text{RD}}(w[j, \infty), \varphi)$ and
- $V^{\text{RD}}(w, [r \cdot] \varphi) = b_1 b_2 b_3 b_4$ with $b_i = \max\{b'_1, \dots, b'_i\}$ for every $i \in \{1, 2, 3, 4\}$, where

$$\begin{aligned}
& - b'_1 = \min_{j \in \mathcal{R}_1^{\text{RD}}(w, r)} V_1^{\text{RD}}(w[j, \infty), \varphi), \\
& - b'_2 = \begin{cases} \max_{j' \in \mathbb{N}} \min_{j \in \mathcal{R}_2^{\text{RD}}(w, r) \cap \{j', j'+1, j'+2, \dots\}} V_2^{\text{RD}}(w[j, \infty), \varphi) & \text{if } |\mathcal{R}_2^{\text{RD}}(w, r)| = \infty, \\ \min_{j \in \mathcal{R}_2^{\text{RD}}(w, r)} V_2^{\text{RD}}(w[j, \infty), \varphi) & \text{if } 0 < |\mathcal{R}_2^{\text{RD}}(w, r)| < \infty, \\ 1 & \text{if } |\mathcal{R}_2^{\text{RD}}(w, r)| = 0, \end{cases} \\
& - b'_3 = \begin{cases} \min_{j' \in \mathbb{N}} \max_{j \in \mathcal{R}_3^{\text{RD}}(w, r) \cap \{j', j'+1, j'+2, \dots\}} V_3^{\text{RD}}(w[j, \infty), \varphi) & \text{if } |\mathcal{R}_3^{\text{RD}}(w, r)| = \infty, \\ \max_{j \in \mathcal{R}_3^{\text{RD}}(w, r)} V_3^{\text{RD}}(w[j, \infty), \varphi) & \text{if } 0 < |\mathcal{R}_3^{\text{RD}}(w, r)| < \infty, \\ 1 & \text{if } |\mathcal{R}_3^{\text{RD}}(w, r)| = 0, \end{cases} \\
& - b'_4 = \begin{cases} \max_{j \in \mathcal{R}_4^{\text{RD}}(w, r)} V_4^{\text{RD}}(w[j, \infty), \varphi) & \text{if } |\mathcal{R}_4^{\text{RD}}(w, r)| > 0, \\ 1 & \text{if } |\mathcal{R}_4^{\text{RD}}(w, r)| = 0. \end{cases}
\end{aligned}$$

610 To give an intuitive description of the semantics, let us first generalize the notion of r -matches and φ -satisfiability. We say that a position j of w is an r -match of degree β if $j \in \mathcal{R}_i^{\text{RD}}(w, r)$ for the unique i with $\beta = 0^{i-1}1^{5-i}$, which requires all tests in r to be evaluated w.r.t. V_i^{RD} (i.e., to some truth value at least β). Similarly, we say that a position j of w is φ -satisfying of degree β if
615 $V^{\text{RD}}(w[j, \infty], \varphi) \succeq \beta$, or if, equivalently, $V_i^{\text{RD}}(w[j, \infty], \varphi) = 1$ for the unique i with $\beta = 0^{i-1}1^{5-i}$.

Now, consider the b'_i defining the semantics of the robust box operator: We have $b'_1 = 1$ if all r -matches of degree 1111 are φ -satisfying of degree 1111. This is in particular satisfied if there is no such match. Further, if there are infinitely
620 (finitely) many r -matches of degree 0111, then $b'_2 = 1$ if almost all (if all) those matches are φ -satisfying of degree 0111. Dually, if there are infinitely (finitely) many r -matches of degree 0011, then $b'_3 = 1$ if infinitely many (at least one) of those matches are (is) φ -satisfying of degree 0011. Finally, if there is at least one r -match of degree 0001, then $b'_4 = 1$ if at least one of those matches is
625 φ -satisfying of degree 0001. The cases where there is no r -match are irrelevant due to monotonicity, so we hardcode them to 1.

Example 2. Consider the formula $[\cdot r \cdot] q \rightarrow [\cdot \text{tt} ; r \cdot] p$ with $r = (\text{tt}; \text{tt})^*$, which expresses that the degree of violation of q at even positions should at most be the degree of violation of p at odd positions. Such a property cannot be expressed
630 in $r\text{LTL}(\Box, \Diamond)$, as even $[\cdot r \cdot] q$ is known to be inexpressible in LTL [31].

First, we state that the semantics is well-defined. This is not obvious due to the case distinctions and the use of the matching sets $\mathcal{R}_i^{\text{RD}}$ for different i .

Lemma 2. We have $V^{\text{RD}}(w, \varphi) \in \mathbb{B}_4$ for every trace w and every formula φ .

Proof. We proceed by induction over the structure of φ . The cases of atomic
635 propositions and Boolean connectives are trivial, as they return values from \mathbb{B}_4 by definition, provided the arguments are from \mathbb{B}_4 . Similarly, we have $V^{\text{RD}}(w, [\cdot r \cdot] \varphi) = b_1 b_2 b_3 b_4 \in \mathbb{B}_4$, as the maximization “ $b_i = \max\{b'_1, \dots, b'_i\}$ ” in the definition enforces the desired monotonicity of the bits b_i .

To conclude, consider a diamond formula $\langle \cdot r \cdot \rangle \varphi$. Applying the induction
640 hypothesis to the tests of r and an induction over the construction of r shows

$$\mathcal{R}_1^{\text{RD}}(w, r) \subseteq \mathcal{R}_2^{\text{RD}}(w, r) \subseteq \mathcal{R}_3^{\text{RD}}(w, r) \subseteq \mathcal{R}_4^{\text{RD}}(w, r)$$

for every trace w . Hence, an application of the induction hypothesis for φ yields

$$\begin{aligned} \max_{j \in \mathcal{R}_1^{\text{RD}}(w, r)} V_1^{\text{RD}}(w[j, \infty], \varphi) &\leq \max_{j \in \mathcal{R}_2^{\text{RD}}(w, r)} V_2^{\text{RD}}(w[j, \infty], \varphi) \\ &\leq \max_{j \in \mathcal{R}_3^{\text{RD}}(w, r)} V_3^{\text{RD}}(w[j, \infty], \varphi) \leq \max_{j \in \mathcal{R}_4^{\text{RD}}(w, r)} V_4^{\text{RD}}(w[j, \infty], \varphi) \end{aligned}$$

for every trace w . Hence, $V^{\text{RD}}(w, \langle \cdot r \cdot \rangle \varphi) \in \mathbb{B}_4$. \square

To conclude the definition of the semantics, we give an example witnessing that the maximization over the b'_i in the semantics of the box operator is indeed necessary to obtain monotonicity.

645 **Example 3.** Let $\varphi = [r] \mathbf{ff}$ with $r = ([\mathbf{tt}^*] p)?$. Moreover, consider the
trace $w = \emptyset\{p\}^\omega$. Then, we have $V^{\text{RD}}(w, [\mathbf{tt}^*] p) = 0111$ and consequently
 $\mathcal{R}_1^{\text{RD}}(w, r) = \emptyset$ and $\mathcal{R}_2^{\text{RD}}(w, r) = \{0\}$. Therefore, $\min_{j \in \mathcal{R}_1^{\text{RD}}(w, r)} V_1^{\text{RD}}(w[j, \infty), \mathbf{ff}) =$
 $\min \emptyset = 1$, but $\min_{j \in \mathcal{R}_2^{\text{RD}}(w, r)} V_2^{\text{RD}}(w[j, \infty), \mathbf{ff}) = \min\{0\} = 0$. Thus, the bits b'_1
and b'_2 inducing $V^{\text{RD}}(w, [r] \mathbf{ff})$ are not monotonic, which explains the need to
650 maximize over the b'_i to obtain the semantics of the robust box operator. The
traces $(\emptyset\{p\})^\omega$ and $\{p\}\emptyset^\omega$ witness that monotonicity can also be violated for the
pairs b'_2, b'_3 and b'_3, b'_4 .

We prove that rLDL has the exponential compilation property. This allows
us to solve the model checking and the synthesis problem using well-known and
655 efficient automata-based algorithms. Furthermore, we are able to show that the
complexity of these algorithms is asymptotically the same as the complexity of
the algorithms for plain LDL and LTL. In the terminology introduced in the
introduction, we present a direct translation, i.e., we translate rLDL directly
into automata.

660 **Theorem 3.** Let φ be an rLDL formula, $n = |\varphi|$, and $\beta \in \mathbb{B}_4$. There is
a non-deterministic Büchi automaton $\mathfrak{B}_{\varphi, \beta}$ with $2^{\mathcal{O}(n)}$ states recognizing the
language $\{w \in (2^P)^\omega \mid V^{\text{RD}}(w, \varphi) \succeq \beta\}$.

In order to prove this theorem, we first recall in Section 4.1 how to translate
guards r into finite non-deterministic automata with special features to account
665 for tests. Then, in Section 4.2, we present the translation of rLDL into weak
alternating Büchi automata of linear size⁶, which can then be further transformed
into non-deterministic Büchi automata of exponential size and deterministic parity
automata of doubly-exponential size. Such automata are needed for solving
the model checking problem and the synthesis problem, respectively.

670 4.1. Translating Guards into Automata

Recall that P is the (finite) set of atomic propositions. An automaton with
tests $\mathfrak{G} = (Q, 2^P, q_I, \delta, F, t)$ consists of a finite set Q of states, the alphabet 2^P ,
an initial state $q_I \in Q$, a transition function $\delta: Q \times (2^P \cup \{\varepsilon\}) \rightarrow 2^Q$, a set F
of final states, and a partial function t , which assigns to states $q \in Q$ an rLDL
675 formula $t(q)$. These should be thought of as the analogue of tests, i.e., if $t(q)$
is defined, then a run visiting q is only successful if the word that remains to be
processed from q onwards satisfies the formula $t(q)$.

We write $q \xrightarrow{a} q'$ if $q' \in \delta(q, a)$ for $a \in 2^P \cup \{\varepsilon\}$. An ε -path π from q to q' in
 \mathfrak{G} is a sequence $\pi = q_1 \cdots q_k$ of $k \geq 1$ states with $q = q_1 \xrightarrow{\varepsilon} \cdots \xrightarrow{\varepsilon} q_k = q'$. Let
680 $t(\pi) = \{t(q_i) \mid 1 \leq i \leq k\}$ denote the set of tests visited by π and let $\Pi(q, q')$
denote the set of all ε -paths from q to q' .

A run of \mathfrak{G} on $w(0) \cdots w(j-1) \in (2^P)^*$ is a sequence $q_0 q_1 \cdots q_j$ of states
such that $q_0 = q_I$ and for every j' in the range $0 \leq j' \leq j-1$ there is a state $q'_{j'}$,

⁶In contrast to the original conference version [1], we here use weak alternating Büchi
automata instead of weak parity automata, which allows us to simplify the construction.

reachable from $q_{j'}$ via an ε -path $\pi_{j'}$ and such that $q_{j'+1} \in \delta(q_{j'}, w(j'))$. The run is accepting if there is a $q'_j \in F$ reachable from q_j via an ε -path π_j . This slightly unusual definition of runs (but equivalent to the standard one) simplifies our reasoning below. Also, the definition is oblivious to the tests assigned by t . To take them into account, we define for $i \in \{1, 2, 3, 4\}$

$$\mathcal{R}_i^{\text{RD}}(w, \mathfrak{G}) = \{j \mid \mathfrak{G} \text{ has an accepting run on } w[0, j] \text{ with } \varepsilon\text{-paths } \pi_0, \dots, \pi_j \text{ s.t.}$$

$$V_i^{\text{RD}}(w[j', \infty), \bigwedge t(\pi_{j'})) = 1 \text{ for every } j' \text{ in the range } 0 \leq j' \leq j\}.$$

Here, $\bigwedge t(\pi_{j'})$ is the conjunction of all formulas in $t(\pi_{j'})$.

Every guard (which is just a regular expression with tests) can be turned into an equivalent automaton with tests via a straightforward generalization of the classical Thompson construction turning classical regular expressions into ε -NFA, to which one adds a rule turning a test into a one-state automaton whose state is labeled with this test (see Figure 2 of Faymonville and Zimmermann [32] for details).

Lemma 3. *Every guard r can be translated into an automaton with tests \mathfrak{G}_r such that $\mathcal{R}_i^{\text{RD}}(w, r) = \mathcal{R}_i^{\text{RD}}(w, \mathfrak{G}_r)$ for every $i \in \{1, 2, 3, 4\}$ and with $|\mathfrak{G}_r| \in \mathcal{O}(|r|)$. Furthermore, all final states of \mathfrak{G}_r are terminal, i.e., they have no outgoing transitions.*

The automaton \mathfrak{G}_r is independent of i , as this value only determines how tests are evaluated. These are handled “externally” in the definition of the semantics. Having thus demonstrated how to turn guards into automata, we now demonstrate how to do the same for rLDL formulas.

4.2. Translating rLDL into Alternating Automata

In this subsection, we translate rLDL formulas into weak alternating Büchi automata, which are known to be translatable into non-deterministic Büchi automata of exponential size [33]. Hence, the linear translation from rLDL to weak alternating Büchi automata we are about to present implies the exponential compilation property for rLDL.

An alternating Büchi automaton $\mathfrak{A} = (Q, \Sigma, q_I, \delta, F)$ consists of a finite set Q of states, an alphabet Σ , an initial state $q_I \in Q$, a transition function $\delta: Q \times \Sigma \rightarrow \mathcal{B}^+(Q)$, and a set $F \subseteq Q$ of accepting states. Here, $\mathcal{B}^+(Q)$ denotes the set of positive Boolean combinations over Q , which contains in particular the formulas **tt** (true) and **ff** (false).

A run of \mathfrak{A} on $w = w(0)w(1)w(2)\dots \in \Sigma^\omega$ is a directed graph $\rho = (V, E)$ where $V \subseteq Q \times \mathbb{N}$ and $((q, n), (q', n')) \in E$ implies $n' = n + 1$, such that $(q_I, 0) \in V$, and such that for all $(q, n) \in V$ we have $\text{Succ}_\rho(q, n) \models \delta(q, w(n))$. Here $\text{Succ}_\rho(q, n)$ denotes the set of successors of (q, n) in ρ projected to Q . A run ρ is accepting if all infinite paths (projected to Q) through ρ satisfy the Büchi condition, i.e., an accepting state is visited infinitely often on the path. The language $L(\mathfrak{A})$ contains all $w \in \Sigma^\omega$ that have an accepting run of \mathfrak{A} .

715 Given \mathfrak{A} as above, its transition graph (Q, E) is the directed graph such that
 $(q, q') \in E$ if and only if q' appears in $\delta(q, a)$ for some $a \in \Sigma$. The automaton \mathfrak{A}
is weak, if every strongly connected component C of (Q, E) is either a subset of F
or a subset of $Q \setminus F$, i.e., every cycle has either only accepting states or only
rejecting states.

720 Using standard constructions, weak alternating Büchi automata are easily
seen to be closed under all Boolean operations. Fix automata $\mathfrak{A}_0 = (Q_0, \Sigma, q_I^0, \delta_0, F_0)$
and $\mathfrak{A}_1 = (Q_1, \Sigma, q_I^1, \delta_1, F_1)$.

- 725 • $(Q_0, \Sigma, q_I^0, \overline{\delta_0}, \overline{F})$ recognizes $\Sigma^\omega \setminus L(\mathfrak{A}_0)$, where $\overline{F} = Q_0 \setminus F$ and where $\overline{\delta_0}$
is the dual of δ_0 , i.e., $\overline{\delta_0}(q, A)$ is obtained from $\delta_0(q, A)$ by replacing each
disjunction by a conjunction, each conjunction by a disjunction, each **tt**
by **ff**, and each **ff** by **tt**.
- The disjoint union of \mathfrak{A}_0 and \mathfrak{A}_1 with a fresh initial state q_I and $\delta(q_I, A) =$
 $\delta_0(q_I^0, A) \wedge \delta_1(q_I^1, A)$ recognizes $L(\mathfrak{A}_0) \cap L(\mathfrak{A}_1)$.
- 730 • The disjoint union of \mathfrak{A}_0 and \mathfrak{A}_1 with a fresh initial state q_I and $\delta(q_I, A) =$
 $\delta_0(q_I^0, A) \vee \delta_1(q_I^1, A)$ recognizes $L(\mathfrak{A}_0) \cup L(\mathfrak{A}_1)$.

The latter two constructions can obviously be generalized to unions and inter-
sections of arbitrary arity while still only requiring a single fresh state.

We prove the following lemma, which implies Theorem 3, as alternating
Büchi automata can be translated into non-deterministic Büchi automata of
size $2^{\mathcal{O}(n)}$ [33].

Lemma 4. *For every rLDL formula φ and every $\beta \in \mathbb{B}_4$, there is a weak alter-
nating Büchi automaton $\mathfrak{A}_{\varphi, \beta}$ with $\mathcal{O}(|\varphi|)$ states recognizing the language $\{w \in$
 $(2^P)^\omega \mid V^{\text{RD}}(w, \varphi) \succeq \beta\}$.*

Proof. We first construct the desired automaton by induction over the structure
of φ . Then, we estimate its size. We begin by noting that $\mathfrak{A}_{\varphi, 0000}$ is trivial for
every formula φ , as it has to accept every input. Hence, we only consider
 $\beta \succ 0000$ in the remainder of the proof.

For an atomic proposition $p \in P$, $\mathfrak{A}_{p, \beta}$ for $\beta \succ 0000$ is an automaton that
accepts exactly those w with $p \in w(0)$. Such an automaton can easily be
constructed.

Now, consider a negation $\varphi = \neg\varphi'$: by definition, we have $V^{\text{RD}}(w, \varphi) = 0000$
if $V^{\text{RD}}(w, \varphi') = 1111$, and $V^{\text{RD}}(w, \varphi) = 1111$ if $V^{\text{RD}}(w, \varphi') \neq 1111$. Thus, $\mathfrak{A}_{\varphi, \beta}$
for $\beta \succ 0000$ has to accept the language $\{w \mid V^{\text{RD}}(w, \varphi') \neq 1111\}$, which is the
complement of the language of $\mathfrak{A}_{\varphi', 1111}$. Hence, we obtain the desired automaton
by applying the closure properties.

Next, consider a conjunction $\varphi = \varphi_0 \wedge \varphi_1$. To this end, recall that $V^{\text{RD}}(w, \varphi) =$
 $\min\{V^{\text{RD}}(w, \varphi_0), V^{\text{RD}}(w, \varphi_1)\}$. Hence, $\mathfrak{A}_{\varphi, \beta}$ has to recognize the language $L(\mathfrak{A}_{\varphi_0, \beta}) \cap$
 $L(\mathfrak{A}_{\varphi_1, \beta})$. Again, we obtain the desired automaton by applying the closure prop-
erties.

755 The construction for a disjunction $\varphi = \varphi_0 \vee \varphi_1$ is dual to the conjunction:
we have $V^{\text{RD}}(w, \varphi) = \max\{V^{\text{RD}}(w, \varphi_0), V^{\text{RD}}(w, \varphi_1)\}$ and thus construct $\mathfrak{A}_{\varphi, \beta}$

such that it recognizes the language $L(\mathfrak{A}_{\varphi_0, \beta}) \cap L(\mathfrak{A}_{\varphi_1, \beta})$. Hence, we obtain the desired automaton by applying the closure properties.

For an implication $\varphi = \varphi_0 \rightarrow \varphi_1$, we again implement the semantics via Boolean combinations of automata. Recall that $V^{\text{RD}}(w, \varphi_0 \rightarrow \varphi_1)$ is equal to 1111 if $V^{\text{RD}}(w, \varphi_0) \preceq V^{\text{RD}}(w, \varphi_1)$. Otherwise, it is equal to $V^{\text{RD}}(w, \varphi_1)$.

Here, we need to construct auxiliary automata $\mathfrak{A}_{\varphi_i, \beta}^{\equiv}$ that accept the traces w with $V^{\text{RD}}(w, \varphi_i) = \beta$. For $\beta = 1111$, this automaton is equal to $\mathfrak{A}_{\varphi_i, \beta}$ and for $\beta \prec 1111$ it is obtained by constructing the automaton recognizing $L(\mathfrak{A}_{\varphi_i, \beta}) \setminus \mathfrak{A}_{\varphi_i, \beta'}$, where β' is the next-larger truth value after β . Here, the set difference is implemented by taking the intersection with the complement. Note that, unlike $\mathfrak{A}_{\varphi_i, \beta}$, which accepts w if $V^{\text{RD}}(w, \varphi_i)$ is greater or equal than β , the automaton $\mathfrak{A}_{\varphi_i, \beta}^{\equiv}$ only accepts w if $V^{\text{RD}}(w, \varphi_i)$ is equal to β .

Now, we combine the auxiliary automata to construct $\mathfrak{A}_{\varphi, \beta}$ so that it recognizes the language

$$\left(\bigcup_{\substack{\beta_0, \beta_1 \in \mathbb{B}_4 \\ \beta_0 \preceq \beta_1}} L(\mathfrak{A}_{\varphi_0, \beta_0}^{\equiv}) \cap L(\mathfrak{A}_{\varphi_1, \beta_1}^{\equiv}) \right) \cup \left(\bigcup_{\substack{\beta_0, \beta_1 \in \mathbb{B}_4 \\ \beta_0 \succ \beta_1 \preceq \beta}} L(\mathfrak{A}_{\varphi_0, \beta_0}^{\equiv}) \cap L(\mathfrak{A}_{\varphi_1, \beta_1}^{\equiv}) \right).$$

The left part covers all cases in which the implication evaluates to 1111. Due to $1111 \succeq \beta$ for every β , this part is equal for all automata. The right part covers all other cases, which depend on β . So, we obtain the desired automaton by applying the closure properties.

Now, we turn to the constructions for the guarded temporal operators, which are more involved as we have to combine automata for guards, for the tests occurring in them, and for formulas. We follow the general construction presented by Faymonville and Zimmermann [6], but generalize it to deal with the richer truth values underlying the robust semantics. Intuitively, for a diamond formula $\langle \cdot r \cdot \rangle \varphi'$ we construct an automaton that checks whether its input has a prefix that matches r (with the required degree) such that the corresponding suffix satisfies φ' (again with the required degree). While checking for the prefix matching r , the automaton we construct also has to check that the tests in r are satisfied (with the required degree). To this end, we use alternation to spawn copies of the automata we have already constructed for the tests. Finally, the construction for a box formula will be dual, but technically slightly more involved due to the robust semantics of the box-operator and the case distinctions involved in its definition.

First, we consider a diamond formula $\varphi = \langle \cdot r \cdot \rangle \varphi'$ with tests $\theta_1, \dots, \theta_n$ in r . Recall that we have $V^{\text{RD}}(w, \varphi) = b_1 b_2 b_3 b_4$ where $b_i = \max_{j \in \mathcal{R}_i^{\text{RD}}(w, r)} V_i^{\text{RD}}(w[j, \infty), \varphi')$ for all $i \in \{1, 2, 3, 4\}$. Thus, $\mathfrak{A}_{\varphi, \beta}$ has to accept w if and only if w has an r -match of degree β that is φ' -satisfying of degree β .

By induction hypothesis, we have automata $\mathfrak{A}_{\varphi', \beta}$ and $\mathfrak{A}_{\theta_j, \beta}$ for every test θ_j in r . Also, we have an ε -NFA with tests \mathfrak{G}_r equivalent to r due to Lemma 3. We combine these automata to the alternating automaton $\mathfrak{A}_{\varphi, \beta}$ by non-deterministically guessing a (finite) run of \mathfrak{G}_r . Whenever the run encounters a final state, the

automaton may jump to the initial state of $\mathfrak{A}_{\varphi',\beta}$ and then behave like that automaton. Furthermore, while simulating \mathfrak{G}_r , $\mathfrak{A}_{\varphi,\beta}$ also has to verify that the tests occurring along the guessed run of \mathfrak{G}_r hold true by universally spawning

800 copies of $\mathfrak{A}_{\theta_j,\beta}$ each time a state labeled with θ_j is traversed. Since we do not allow for ε -transitions in alternating automata, we have to eliminate the ε -transitions of \mathfrak{G}_r during the construction of $\mathfrak{A}_{\varphi,\beta}$. Finally, in order to prevent $\mathfrak{A}_{\varphi,\beta}$ from simulating \mathfrak{G}_r ad infinitum, the states copied from \mathfrak{G}_r are all rejecting, which forces the jump to $\mathfrak{A}_{\varphi,\beta}$ to be executed eventually.

805 Formally, we define $\mathfrak{A}_{\varphi,\beta} = (Q, 2^P, q_I, \delta, F)$ where

- Q is the disjoint union of the sets of states of the automata \mathfrak{G}_r , $\mathfrak{A}_{\theta_j,\beta}$ for $j \in \{1, \dots, n\}$, and $\mathfrak{A}_{\varphi',\beta}$,
 - q_I is the initial state of \mathfrak{G}_r ,
 - F is the union of the accepting states of $\mathfrak{A}_{\theta_j,\beta}$ for $j \in \{1, \dots, n\}$, and of
- 810 $\mathfrak{A}_{\varphi',\beta}$,

and where δ is defined as follows: if q is a state of \mathfrak{G}_r , then

$$\delta(q, A) = \begin{cases} \bigvee_{q' \in Q^r} \bigvee_{\pi \in \Pi(q, q')} \bigvee_{p \in \delta^r(q', A)} (p \wedge \bigwedge_{\theta_j \in t(\pi)} \delta^j(q_I^j, A)) \\ \bigvee_{q' \in F^r} \bigvee_{\pi \in \Pi(q, q')} (\delta'(q', A) \wedge \bigwedge_{\theta_j \in t(\pi)} \delta^j(q_I^j, A)) \end{cases}$$

where q_I^j and q'_I are the initial states of $\mathfrak{A}_{\theta_j,\beta}$ and $\mathfrak{A}_{\varphi',\beta}$, respectively, where Q^r (F^r) is the set of (final) states of \mathfrak{G}_r , where δ_r , δ' , and δ_j are the transition functions of \mathfrak{G}_r , $\mathfrak{A}_{\varphi',\beta}$, and $\mathfrak{A}_{\theta_j,\beta}$ respectively, and where the sets $\Pi(q, q')$ of ε -

815 paths are induced by \mathfrak{G}_r . Furthermore, for states q of $\mathfrak{A}_{\varphi',\beta}$, we define $\delta(q, A) = \delta'(q, A)$ and for states q of $\mathfrak{A}_{\theta_j,\beta}$ we define $\delta(q, A) = \delta^j(q, A)$. The resulting automaton accepts w if and only if w has at least one r -match of degree β that is φ' -satisfying of degree β (cf. [6], where the correctness is proven for plain LDL).

820 Finally, we consider the box operator, which requires the most involved construction due to the case distinction defining the b'_i and the subsequent maximization to obtain the b_i . First, recall that the semantics of the box operator is not dual to the semantics of the diamond operator. Nevertheless, the dual construction of the one for the diamond operator is useful as a building block.

825 We first present this construction before tackling the construction for the box operator.

In the dual construction, one interprets \mathfrak{G}_r as a universal automaton whose transitions are ignored if the test on the source of the transition fails. Furthermore, each visit to a final state spawns a copy of the automaton $\mathfrak{A}_{\varphi',\beta}$, as every

830 r -match has to be φ' -satisfying. Thus, the states of \mathfrak{G}_r are now accepting, as all r -matches have to be considered, and the automata for the tests are dualized in order to check for the failure of the test.

Formally, this approach yields the weak alternating Büchi automaton $(Q, 2^P, q_I, \delta, F)$ where Q and q_I are as above, where

$$\delta(q, A) = \begin{cases} \bigwedge_{q' \in Q^r} \bigwedge_{\pi \in \Pi(q, q')} \bigwedge_{p \in \delta^r(q', A)} (p \vee \bigvee_{\theta_j \in t(\pi)} \overline{\delta^j}(q_I^j, A)) \\ \bigwedge_{q' \in F^r} \bigwedge_{\pi \in \Pi(q, q')} (\delta'(q_I^j, A) \vee \bigvee_{\theta_j \in t(\pi)} \overline{\delta^j}(q_I^j, A)) \end{cases}$$

835 for states q of \mathfrak{G}_r , where q_I^j and q'_I are the initial states of $\mathfrak{A}_{\theta_j, \beta}$ and $\mathfrak{A}_{\varphi', \beta}$, respectively. Additionally, we define $\delta(q, A) = \delta'(q, A)$ for states q of $\mathfrak{A}_{\varphi', \beta}$, and $\delta(q, A) = \overline{\delta^j}(q, A)$ for states q of $\mathfrak{A}_{\theta_j, \beta}$. Finally, the set of accepting states is the union of the set of all states of \mathfrak{G}_r , the rejecting states of the $\mathfrak{A}_{\theta_j, \beta}$, and the accepting states of $\mathfrak{A}_{\varphi', \beta}$. Recall that dualizing the transition relation
840 and swapping accepting and rejecting states of the automata $\mathfrak{A}_{\theta_j, \beta}$ amounts to complementation. This allows terminating runs of \mathfrak{G}_r if a test does not hold true. The resulting automaton accepts a trace if and only if every r -match of degree β is φ -satisfying of degree β (cf. [6], where the correctness is proven for plain LDL)).

845 Now, we fix $\varphi = [\cdot r \cdot] \varphi'$. Recall that we have $V^{\text{RD}}(w, \varphi) = b_1 b_2 b_3 b_4$ with $b_i = \max\{b'_1, \dots, b'_i\}$ for some bits b'_i . The maximization is easily implemented using the Boolean closure properties of alternating automata provided we have automata checking that some bit b'_i is equal to one. Two cases are trivial: Indeed, we have $b'_1 = 1$ if and only if every r -match of degree 1111 is φ -satisfying
850 of degree 1111. This property is checked by the dual automaton constructed above. Furthermore, $b'_4 = 1$ if and only if $V^{\text{RD}}(w, \langle \cdot r \cdot \rangle \varphi') \succeq 0001$ or if there is no r -match of degree 0001. The former language is recognized by $\mathfrak{A}_{\langle \cdot r \cdot \rangle \varphi', 0001}$, the latter one by an automaton we construct below. We then combine these two automata to obtain $\mathfrak{A}_{\varphi, 0001}$.

855 Hence, it remains to consider b'_2 and b'_3 , which are both defined by a case distinction over the number of r -matches of the trace. These case distinctions are implemented using alternation. To this end, we first show how to test for the three cases, i.e., we argue that the following languages are recognizable by weak alternating Büchi automata, where $i \in \{1, 2, 3, 4\}$:

- 860
1. $L_i^\emptyset(r) = \{w \in (2^P)^\omega \mid |\mathcal{R}_i^{\text{RD}}(w, r)| = 0\}$.
 2. $L_i^f(r) = \{w \in (2^P)^\omega \mid 0 < |\mathcal{R}_i^{\text{RD}}(w, r)| < \infty\}$.
 3. $L_i^\infty(r) = \{w \in (2^P)^\omega \mid |\mathcal{R}_i^{\text{RD}}(w, r)| = \infty\}$.

Let $\theta_1, \dots, \theta_n$ be the tests in r . By induction hypothesis, we have weak alternating Büchi automata $\mathfrak{A}_{\theta_j, \beta}$ for every θ_j and every truth value β .

865 The first case is already solved, as we have, for each $i \in \{1, 2, 3, 4\}$, $\mathcal{R}_i^{\text{RD}}(w, r) = \emptyset$ if and only if $V_i^{\text{RD}}(w, \langle \cdot r \cdot \rangle \mathbf{tt}) = 0$, which is in turn equivalent to $V^{\text{RD}}(w, \langle \cdot r \cdot \rangle \mathbf{tt}) \prec 0^{i-1} 1^{5-i}$, i.e., the complement of the automaton $\mathfrak{A}_{\langle \cdot r \cdot \rangle \mathbf{tt}, 0^{i-1} 1^{5-i}}$ recognizes $L_i^\emptyset(r)$.

Next, we construct an automaton for the language $L_i^\infty(r)$. Then, the au-
870 tomaton for $L_i^f(r)$ is obtained as the intersection of the complement automata

for the other two languages (for the given r and i). Thus, we need to construct an automaton that accepts w if there are infinitely many r -matches of degree $0^{i-1}1^{5-i}$.

The construction of an automaton for $L_i^\infty(r)$ is more involved than the previous one, as the automaton \mathfrak{G}_r checking for matches with r is non-deterministic. Nevertheless, we show that standard arguments about non-deterministic automata still yield the desired result. Intuitively, the automaton recognizing $L_i^\infty(r)$ has to determine whether infinitely many prefixes of w are accepted by \mathfrak{G}_r (while dealing with tests appropriately). This is implemented as follows: we start with \mathfrak{G}_r , eliminate ε -transitions as in the case of diamond formulas on Page 24 (i.e., resulting in a non-deterministic choice ranging over all ε -paths, each universally spawning a copy of $\mathfrak{A}_{\theta_j, 0^{i-1}1^{5-i}}$ for each test θ_j encountered along the corresponding ε -path). Furthermore, on each transition, the run branches universally into a disjoint copy of this structure where all states are rejecting. Also, all states that have a transition leading to a state that was accepting in \mathfrak{G}_r are equipped with a new transition leading to a fresh accepting sink state (with the same transition label and the same test automata being spawned). Finally, the states of the original copy are all accepting. Thus, the resulting automaton is weak.

Using König's Lemma, one can show that the resulting automaton accepts w if and only if $w \in L_i^\infty(r)$. We leave the details to the industrious reader and just note that we have now constructed all automata we need to capture the cases in the case distinction.

Extending the construction just presented also allows us to construct an automaton that accepts a trace w if and only if it has infinitely many φ' -satisfying r -matches (both of degree $0^{i-1}1^{5-i}$). To this end, the copies spawned to check for matches are not equipped with transitions leading to an accepting sink, but with transitions leading to the initial state of $\mathfrak{A}_{\varphi', 0^{i-1}1^{5-i}}$ to check for satisfaction of φ' . Similarly, we can construct an automaton that accepts a trace w if and only if it has infinitely many r -matches of degree $0^{i-1}1^{5-i}$ that are *not* φ' -satisfying of degree $0^{i-1}1^{5-i}$. Again, we leave the details to the reader.

These automata also allow us to construct an automaton that accepts a trace w if and only if $\mathcal{R}_i^{\text{RD}}(w, r)$ is infinite and almost all r -matches in $\mathcal{R}_i^{\text{RD}}(w, r)$ are φ' -satisfying (both of degree $0^{i-1}1^{5-i}$). This automaton is obtained by taking the automaton checking for infinitely many φ' -satisfying r -matches (both of degree β) and intersecting it with the complement of the one checking for infinitely many r -matches that are not φ' -satisfying of degree $0^{i-1}1^{5-i}$.

Combining the automata checking the cases of the case distinction with the automata checking for φ' -satisfiability yields the desired automata for b'_2 and b'_3 : A case distinction is easily implemented using the Boolean closure properties and all necessary auxiliary automata have been constructed above.

It is straightforward to verify that each automaton we construct is weak. Hence, it remains to argue that $\mathfrak{A}_{\varphi, \beta}$ is of linear size in $|\varphi|$. To this end, we say that a weak alternating Büchi automaton $(Q', \Sigma, q'_I, \delta', F')$ is a subautomaton of $(Q, \Sigma, q_I, \delta, F)$ if $Q' \subseteq Q$, $\delta'(q, A) = \delta(q, A)$ for every $q \in Q'$ and every $A \in \Sigma$, and $F' = Q' \cap F$.

Inspecting the construction above shows that an automaton $\mathfrak{A}_{\varphi,\beta}$ is built from automata for immediate subformulas (w.r.t. all truth values if necessary), a test automaton (if applicable), and a constant number of fresh states. Furthermore, if formulas share subformulas, then the construction can also share these subautomata. Hence, we obtain the desired linear upper bound on the size of $\mathfrak{A}_{\varphi,\beta}$. \square

It is not straightforward that the equivalent non-deterministic Büchi automata as in Theorem 3 can be constructed efficiently, as the definition of the alternating automaton involves ε -paths of arbitrary length. However, these can be restricted to paths of bounded lengths, as for every ε -path there is one that has the same tests, but no cycles between them. Then, as it is done for the similar construction for PLDL [6], one can show that the Büchi automata can be constructed on-the-fly in polynomial space, relying on the breakpoint construction [33] turning an alternating Büchi automaton into a non-deterministic one. This is sufficient for our applications later on.

Furthermore, as non-deterministic Büchi automata can be translated into deterministic parity automata (see, e.g., [30] for definitions), we obtain the following corollary of Theorem 3.

Corollary 1. *Let φ be an rLDL formula, $n = |\varphi|$, and $\beta \in \mathbb{B}_4$. There is a deterministic parity automaton $\mathfrak{P}_{\varphi,\beta}$ with $2^{2^{\mathcal{O}(n)}}$ states and with $2^{\mathcal{O}(n)}$ colors recognizing the language $\{w \in (2^P)^\omega \mid V^{\text{RD}}(w, \varphi) \succeq \beta\}$.*

The translations from logic to automata just proven allow us to study the expressiveness of rLDL and solve its model checking and synthesis problem.

4.3. Expressiveness

In this section, we compare the expressiveness of rLDL to that of rLTL(\square, \diamond) and LDL. Following Tabuada and Neider [22] we focus on the fragment rLTL(\square, \diamond) without next, until and release operators. While the next and until operator could be added easily, the robust semantics of the release operator is incompatible with our definition of the robust box operator. It turns out as expected, that rLDL subsumes rLTL(\square, \diamond) and rLDL. Conversely, every rLDL formula φ can be translated into four LDL formulas $\varphi_1, \dots, \varphi_4$ that encode φ in the following sense: We have $V_i^{\text{RD}}(w, \varphi) = V^{\text{D}}(w, \varphi_i)$ for every w .

Theorem 4. *Both rLTL(\square, \diamond) and LDL can be embedded into rLDL.*

Proof. Let us first embed rLTL(\square, \diamond) into rLDL by showing that the syntactic embedding of LTL into LDL extends to robust semantics. Recall that rLTL(\square, \diamond) only has temporal operators \diamond and \square , which we replace by $\langle \cdot \text{tt}^* \cdot \rangle$ and $[\cdot \text{tt}^* \cdot]$. Note that $\mathcal{R}_i^{\text{RD}}(w, \text{tt}^*) = \mathbb{N}$ holds true for every w and every i . Hence, a straightforward induction shows that the resulting rLDL formula is equivalent to the original rLTL(\square, \diamond) formula. In particular, only the first case in the case distinctions defining the semantics of the robust box operator of rLDL is used, which mimics the definition of the semantics of the always operator in rLTL(\square, \diamond).

Conversely, using a straightforward induction over the structure of LDL
960 formulas, we can show that $V^D(w, \varphi) = V_1^{\text{RD}}(w, \varphi')$ for every w and every LDL
formula φ , where φ' is the rLDL formula obtained from φ by replacing each $\langle r \rangle$
with $\langle \cdot r \cdot \rangle$, each $[r]$ with $[\cdot r \cdot]$, and each implication $\psi_1 \rightarrow \psi_2$ with $\neg\psi_1 \vee \psi_2$.
This shows that LDL can be embedded into rLDL. Note, however, that we need
to replace each implication with a negation and a disjunction. This is necessary
965 to account for the more complex definition of implications in $\text{rLTL}(\Box, \Diamond)/\text{rLDL}$. \square

As LTL is a semantic fragment of LDL, we immediately obtain that LTL
can be embedded into rLDL and, thus, rLDL inherits the lower bounds of LTL.

Our next theorem states that LDL and rLDL are of equal expressiveness.
970 The direction from LDL to rLDL was shown in Theorem 4, hence we focus on
the other one. Following Tabuada and Neider [22], we construct for every rLDL
formula φ four LDL formulas $\varphi_1, \dots, \varphi_4$ encoding φ as explained above. The
construction relies on Theorem 3, unlike the analogous result translating robust
LTL directly into LTL [22].

Theorem 5. *LDL and rLDL are equally expressive and the translations are
975 effective.*

Proof. As argued above, we only have to consider the direction from rLDL to
LDL. Hence, fix an rLDL formula φ and $i \in \{1, 2, 3, 4\}$. Due to Theorem 3, $\{w \in$
 $(2^P)^\omega \mid V_i^{\text{RD}}(w, \varphi) = 1\}$ is ω -regular. Hence, due to LDL being equi-expressive to
980 the ω -regular languages [8], there is also an LDL formula φ_i with $V^D(w, \varphi_i) = 1$
if and only if $V_i^{\text{RD}}(w, \varphi) = 1$. Hence, φ_i has the desired properties. \square

Let us analyze the complexity of the translation in more detail. A non-
deterministic Büchi automaton for an rLDL formula is in general of exponen-
tial size and has to be determinized before it can be translated into LDL (say
985 with max-parity acceptance), which incurs a second exponential blowup. The
resulting deterministic automaton can then be translated into LDL with an
exponential blowup. The resulting formula expresses that the unique run on
the input satisfies the following properties: there is an even color c and a po-
sition j such that after j no larger color appears, c appears at least once, and
990 every time color c appears, it is not the last occurrence of c (see [34] for a
similar construction). All three properties are easily expressed in LDL by con-
structing guards $r_{q,q'}$ that match infixes such that processing the infix starting
in q leads to the state q' . The number of subformulas is polynomial, but the
guards $r_{q,q'}$ have in general exponential size (both measured in the size of the
995 doubly-exponential deterministic automaton). Hence, the full construction in-
curs a triply-exponential blowup. It is open whether this is unavoidable. On
a more positive note, the resulting LDL formula is test-free, i.e., it does not
contain tests in its guards.

We leave the question of whether there are non-trivial lower bounds on the
1000 translation for future work. For the special case of translating $\text{rLTL}(\Box, \Diamond)$
into LTL mentioned above, there is only a linear blowup. This translation was

presented by Tabuada and Neider [22], but they only claimed an exponential upper bound. However, closer inspection shows that it is linear if the size of formulas is measured in the number of distinct subformulas, not the length of the formula.

4.4. Model Checking and Synthesis

Theorem 5 immediately provides solutions for typical applications of rLDL, such as model checking and synthesis, by reducing the problem from the domain of rLDL to that of LDL. However, the price to pay for this approach is a triply-exponential blow-up in the size of the resulting LDL formula, which is clearly prohibitive for any real-world application. For this reason, we now develop more efficient model checking and synthesis techniques that are based on our direct translation of rLDL into automata (Theorem 3).

We begin with the rLDL model checking problem, which is defined as follows.

Problem 3. *Let φ be an rLDL formula, \mathcal{S} a transition system, and let $\beta \in \mathbb{B}_4$. Does $V^{\text{RD}}(\lambda(\rho), \varphi) \succeq \beta$ hold true for all paths $\rho \in \Pi_{\mathcal{S}}$?*

Using the translation of rLDL formulas to weak alternating Büchi automata and subsequently to non-deterministic Büchi automata, Problem 3 can be solved as follows:

1. Translate the transition system \mathcal{S} into a non-deterministic Büchi automaton $\mathfrak{B}_{\mathcal{S}}$ with $L(\mathfrak{B}_{\mathcal{S}}) = \{\lambda(\rho) \in (2^P)^\omega \mid \rho \in \Pi_{\mathcal{S}}\}$ in the usual way: $\mathfrak{B}_{\mathcal{S}}$ has the same states as \mathcal{S} , the transitions are $\{(s, \lambda(s), s') \mid (s, s') \in E\}$, and all states are accepting.
2. Construct the weak alternating Büchi automaton $\mathfrak{A}_{\varphi, \beta}$ accepting the language $\{w \in (2^P)^\omega \mid V^{\text{RD}}(w, \varphi) \succeq \beta\}$.
3. Complement $\mathfrak{A}_{\varphi, \beta}$ to obtain a weak alternating Büchi automaton $\overline{\mathfrak{A}_{\varphi, \beta}}$ accepting the language $\{w \in (2^P)^\omega \mid V^{\text{RD}}(w, \varphi) \prec \beta\}$.
4. Convert $\overline{\mathfrak{A}_{\varphi, \beta}}$ into an equivalent non-deterministic Büchi automaton $\overline{\mathfrak{B}_{\varphi, \beta}}$ and compute the product automaton \mathfrak{B} with $L(\mathfrak{B}) = L(\mathfrak{B}_{\mathcal{S}}) \cap L(\overline{\mathfrak{B}_{\varphi, \beta}})$ in the usual way.
5. Check whether $L(\mathfrak{B}) = \emptyset$ using a standard algorithm such as a nested depth-first search [31]. The answer to Problem 3 is “yes” if and only if $L(\mathfrak{B}) = \emptyset$.

The number of states of the weak alternating Büchi automata in Step 2 and 3 is both in $\mathcal{O}(|\varphi|)$. Thus, the number of states of the non-deterministic Büchi automaton $\overline{\mathfrak{B}_{\varphi, \beta}}$ constructed in Step 4 is in $2^{\mathcal{O}(|\varphi|)}$, and that of \mathfrak{B} is in $|\mathcal{S}| \cdot 2^{\mathcal{O}(|\varphi|)}$, where $|\mathcal{S}|$ denotes the number of states of the transition system \mathcal{S} (cf. Theorem 3). Finally, the time required for the emptiness check in Step 5 is quadratic in the number of states of \mathfrak{B} (linear in the number of \mathfrak{B} 's transitions). Consequently, the rLDL model checking problem can be solved in time $|\mathcal{S}|^2 \cdot 2^{\mathcal{O}(|\varphi|)}$ and, hence, is in EXPTIME.

We now show that the problem is not only in EXPTIME, but that it is, in fact, PSPACE-complete. To this end, we leverage the exponential compilation property (see Theorem 3) and standard on-the-fly techniques for checking emptiness of exponentially-sized Büchi automata [5], which yield a PSPACE upper bound on the complexity of Problem 3. The matching lower bound follows from the subsumption of LDL shown above, as model checking LDL is PSPACE-complete.

Theorem 6. *rLDL model checking is PSPACE-complete.*

Proof. As shown above, the rLDL model checking problem is in EXPTIME. To show membership in PSPACE, we use the observation that given two states of \mathfrak{B} , one can decide in polynomial space whether the second state is a successor of the first one (cf. Vardi and Wolper [5]). Moreover, one can represent states of \mathfrak{B} in polynomial space. This allows running the classical model checking algorithm, which searches for a counterexample, in polynomial space by guessing an appropriate run.

PSPACE-hardness, on the other hand, follows immediately from the facts that (a) the LTL semantics is embedded in rLDL, via the embedding of LDL in rLDL (see Theorem 5), and (b) LTL model checking is PSPACE-hard [35]. Thus, rLDL model checking is PSPACE-hard as well. \square

Before we move on to reactive synthesis, let us briefly remark that the model checking problem for rLTL(\square, \diamond) is defined slightly differently. Instead of asking whether $V^R(\lambda(\rho), \varphi) \succeq \beta$, Tabuada and Neider [22] fix a set $B \subseteq \mathbb{B}_4$ and ask whether $V^R(\lambda(\rho), \varphi) \in B$ for all paths $\rho \in \Pi_S$. However, this slightly more general problem can easily be answered by a simple adaptation of Step 2 of the procedure above: given a (finite) set $B \subseteq \mathbb{B}_4$, we construct a weak alternating Büchi automaton accepting the language $\{w \in (2^P)^\omega \mid V^{RD}(w, \varphi) \in B\}$ using Boolean combinations of the automata $\mathfrak{A}_{\varphi, \beta}$. Then, it is not hard to verify that this variant of the rLDL model checking problem is also PSPACE-complete.

Similar to model checking, the translation from rLDL formulas to automata provides us with an effective means to synthesize reactive controllers from rLDL specifications, i.e., for the following problem, where an rLDL game has the form (G, φ, β) and Player 0 wins a play if and only if its trace w satisfies $V^{RD}(w, \varphi) \succeq \beta$.

Problem 4. *Let \mathcal{G} be an rLDL game and v a vertex. Determine whether Player 0 has a winning strategy for \mathcal{G} from v and compute a finite-state winning strategy if so.*

Corollary 1 provides a straightforward way to solve Problem 4 by reducing it to solving classical parity games (again, see [30, Chapter 2] for an introduction to parity games) while the lower bound follows from the subsumption of LDL.

Theorem 7. *Solving rLDL games is 2EXPTIME-complete.*

Proof. We proceed in several steps constructing a reduction to a parity game.

1. Construct the deterministic parity automaton $\mathfrak{P}_{\varphi, \beta}$ recognizing the language $\{w \in (2^P)^\omega \mid V^{RD}(w, \varphi) \succeq \beta\}$ according to Corollary 1.

- 1085 2. Construct the product of $\mathfrak{P}_{\varphi,\beta} = (Q, 2^P, q_I, \delta, \Omega)$ and the labeled game graph $G = (V_0, V_1, E, \lambda)$. This product is a classical (non-labeled) parity game $\mathcal{G}' = (G', \Omega')$ consisting of a game graph $G' = (V'_0, V'_1, E')$ with $V'_0 = V_0 \times Q$, $V'_1 = V_1 \times Q$, and $E' = \{((v, q), (v', \delta(q, \lambda(v)))) \mid (v, v') \in E\}$ as well as a parity winning condition Ω' with $\Omega'((v, q)) = \Omega(q)$ for each $(v, q) \in V'_0 \cup V'_1$. One obtains a play ρ in the original game \mathcal{G} from a play ρ' in the extended game \mathcal{G}' by projecting the vertices of ρ' onto the first component. Thus, Player 0 wins a play ρ' in \mathcal{G}' from a vertex (v, q_I) if and only if the trace $\lambda(\rho)$ obtained from the corresponding play ρ in \mathcal{G} satisfies $V^{\text{RP}}(\lambda(\rho), \varphi) \succeq \beta$.
- 1090
- 1095 3. Solve the game \mathcal{G}' with standard algorithms for parity games, e.g., the recent quasi-polynomial time algorithm [36]. Finally, check and return whether Player 0 has a winning strategy from vertex (v, q_I) and return a finite-state winning strategy if so.

The above reduction is a standard game reduction, whose correctness can be shown using standard techniques. In fact, it provides a 2EXPTIME algorithm to solve Problem 4: due to Corollary 1, the deterministic parity automaton $\mathfrak{P}_{\varphi,\beta}$ constructed in Step 1 has $2^{2^{\mathcal{O}(|\varphi|)}}$ states and $2^{\mathcal{O}(|\varphi|)}$ colors; consequently, the parity game \mathcal{G}' of Step 2 has $|V| \cdot 2^{2^{\mathcal{O}(|\varphi|)}}$ vertices and $2^{\mathcal{O}(|\varphi|)}$ colors; thus, using the quasi-polynomial algorithm for solving parity games [36] results in a doubly-exponential algorithm for Problem 4.

1100

1105

On the other hand, the fact that rLDL subsumes LDL and, hence, LTL immediately implies that solving rLDL games is 2EXPTIME-hard since solving LTL games is already 2EXPTIME-hard [37]. \square

5. Towards Robust and Prompt Linear Dynamic Logic

1110 In the previous sections, we studied robust LDL, i.e., we combined robustness and increased expressiveness, and robust Prompt-LTL, i.e., we combined robustness and quantitative operators. The third combination of two aspects, i.e., quantitative operators and increased expressiveness, has been studied before [6]. For all three resulting logics, model checking and synthesis have the same complexity as for plain LTL.

1115

Here, we consider the combination of all three extensions, obtaining the logic rPrompt-LDL, robust Prompt-LDL. As this logic is an extension of Prompt-LTL, it features negations only at the level of atomic propositions and does not allow implications. The formulas of rPrompt-LDL are given by the grammar

$$\begin{aligned} \varphi &::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle \cdot r \cdot \rangle \varphi \mid [\cdot r \cdot] \varphi \mid \langle \cdot r \cdot \rangle_{\mathbf{P}} \varphi \\ r &::= \phi \mid \varphi? \mid r + r \mid r ; r \mid r^* \end{aligned}$$

where p again ranges over the atomic propositions in P and ϕ over propositional formulas over P . Furthermore, the size of a formula is defined as for LDL and rLDL.

The semantics are defined as expected: We add a bound k to the semantics of rLDL, which bounds the scope of the prompt diamond operator.

- $V^{\text{RPD}}(w, k, p) = \begin{cases} 1111 & \text{if } p \in w(0), \\ 0000 & \text{if } p \notin w(0), \end{cases}$
- $V^{\text{RPD}}(w, k, \neg p) = \begin{cases} 0000 & \text{if } p \in w(0), \\ 1111 & \text{if } p \notin w(0), \end{cases}$
- $V^{\text{RPD}}(w, k, \varphi_0 \wedge \varphi_1) = \min\{V^{\text{RPD}}(w, k, \varphi_0), V^{\text{RPD}}(w, k, \varphi_1)\},$
- $V^{\text{RPD}}(w, k, \varphi_0 \vee \varphi_1) = \max\{V^{\text{RPD}}(w, k, \varphi_0), V^{\text{RPD}}(w, k, \varphi_1)\},$
- $V^{\text{RPD}}(w, k, \langle \cdot r \cdot \rangle \varphi) = b_1 b_2 b_3 b_4$ where $b_i = \max_{j \in \mathcal{R}_i^{\text{RPD}}(w, k, r)} V_i^{\text{RPD}}(w[j, \infty), k, \varphi),$
- 1125 • $V^{\text{RPD}}(w, k, [\cdot r \cdot] \varphi) = b_1 b_2 b_3 b_4$ with $b_i = \max\{b'_1, \dots, b'_i\}$ for every $i \in \{1, 2, 3, 4\}$, where
 - $b'_1 = \min_{j \in \mathcal{R}_1^{\text{RPD}}(w, k, r)} V_1^{\text{RPD}}(w[j, \infty), k, \varphi),$
 - $b'_2 = \begin{cases} \max_{j' \in \mathbb{N}} \min_{j \in \mathcal{R}_2^{\text{RPD}}(w, k, r) \cap \{j', j'+1, j'+2, \dots\}} V_2^{\text{RPD}}(w[j, \infty), k, \varphi) & \text{if } |\mathcal{R}_2^{\text{RPD}}(w, k, r)| = \infty, \\ \min_{j \in \mathcal{R}_2^{\text{RPD}}(w, k, r)} V_2^{\text{RPD}}(w[j, \infty), k, \varphi) & \text{if } 0 < |\mathcal{R}_2^{\text{RPD}}(w, k, r)| < \infty, \\ 1 & \text{if } |\mathcal{R}_2^{\text{RPD}}(w, k, r)| = 0, \end{cases}$
 - $b'_3 = \begin{cases} \min_{j' \in \mathbb{N}} \max_{j \in \mathcal{R}_3^{\text{RPD}}(w, k, r) \cap \{j', j'+1, j'+2, \dots\}} V_3^{\text{RPD}}(w[j, \infty), k, \varphi) & \text{if } |\mathcal{R}_3^{\text{RPD}}(w, k, r)| = \infty, \\ \max_{j \in \mathcal{R}_3^{\text{RPD}}(w, k, r)} V_3^{\text{RPD}}(w[j, \infty), k, \varphi) & \text{if } 0 < |\mathcal{R}_3^{\text{RPD}}(w, k, r)| < \infty, \\ 1 & \text{if } |\mathcal{R}_3^{\text{RPD}}(w, k, r)| = 0, \end{cases}$
 - 1130 – $b'_4 = \begin{cases} \max_{j \in \mathcal{R}_4^{\text{RPD}}(w, k, r)} V_4^{\text{RPD}}(w[j, \infty), k, \varphi) & \text{if } |\mathcal{R}_4^{\text{RPD}}(w, k, r)| > 0, \\ 1 & \text{if } |\mathcal{R}_4^{\text{RPD}}(w, k, r)| = 0, \end{cases}$ and
- $V^{\text{RPD}}(w, k, \langle \cdot r \cdot \rangle_{\mathbf{p}} \varphi) = b_1 b_2 b_3 b_4$ where $b_i = \max_{j \in \mathcal{R}_i^{\text{RPD}}(w, k, r) \cap \{0, \dots, k\}} V_i^{\text{RPD}}(w[j, \infty), k, \varphi).$

Here, we adapt the definition of $\mathcal{R}_i^{\text{RD}}$ to account for the parameter k : $\mathcal{R}_i^{\text{RPD}}(w, k, \varphi?) = \{0\}$ if $V_i^{\text{RPD}}(w, k, \varphi) = 1$ and $\mathcal{R}_i^{\text{RPD}}(w, k, \varphi?) = \emptyset$ otherwise. All other cases are defined as before, but propagate the parameter k .

1135 rLDL without negation and implication (and LDL, for which negation and implication can be eliminated) and Prompt-LDL formulas can easily be translated into equivalent rPrompt-LDL formulas. Prompt-LTL, however cannot necessarily be translated into equivalent rPrompt-LDL formulas, as the semantics of the release operator is not compatible with the semantics of rPrompt-LDL.

1140 **Example 4.** Consider the formula $[\cdot((\neg t)^* ; t ; (\neg t)^* ; t)^* \cdot] \langle \cdot \mathbf{tt}^* \cdot \rangle_{\mathbf{p}} s$ and interpret t as the tick of a clock and s as a synchronization. Then, the formula intuitively expresses that every other tick of the clock is followed after a bounded number of steps (not ticks!) by a synchronization.

1145 More formally, the different degrees of satisfaction of φ express the following possibilities, with respect to a given bound k : (i) every even clock tick is followed by a synchronization within k steps; (ii) almost every even clock tick is followed by a synchronization within k steps; (iii) infinitely many even clock ticks are followed by a synchronization within k steps; (iv) there is at least one even clock tick that is followed by a synchronization within k steps.

1150 *This property can neither be expressed in (robust) LDL nor in (robust) Prompt-LTL. Also note that unlike for the similar formula from Example 1, the last two possibilities are not trivial, as we now only consider positions with an even clock tick and not all positions.*

1155 In the previous sections, we have seen two approaches to translating robust logics into Büchi automata, the direct and the reduction-based one. Both are extensions of translations originally introduced by Tabuada and Neider for robust LTL. The former one translates a formula of a robust logic directly into an equivalent Büchi automaton while the latter one first translates a formula
1160 of a robust logic into an equivalent classical (non-robust) logic, for which a translation into equivalent Büchi automata is already known. For robust LTL, both approaches are applicable [22] and yield Büchi automata of exponential size. Here, out of necessity, we apply both approaches: for robust LDL, we present a direct translation while we present a reduction-based approach for robust Prompt-LTL. Let us quickly elaborate the reasons for this.

1165 First, consider the reduction-based approach for robust LTL, which translates a formula φ of robust LTL and a truth value $\beta \succ 0000$ into an LTL formula φ_β that captures φ with respect to β . To this end, the formula φ_β implements the intuitive meaning of the robust semantics for the always operator, e.g., we have $(\Box p)_{1111} = \Box p$, $(\Box p)_{0111} = \Diamond \Box p$, $(\Box p)_{0011} = \Box \Diamond p$, and
1170 $(\Box p)_{0001} = \Diamond p$.

Trying to apply this approach to the rLDL formula $\varphi = [\cdot r \cdot] p$, say for $\beta = 0111$, would imply using a formula of the form $\langle r_0 \cdot \rangle [\cdot r_1 \cdot] p$ where r_0 and r_1 are obtained by “splitting” up r . It captures the robust semantics of φ with respect to β on some trace w by expressing that there is an r_0 -match j such that
1175 every r_1 -match in $w[j, \infty)$ is p -satisfying with degree β . Thus, r_0 and r_1 have to be picked such that the r_1 -matches in $w[j, \infty)$ as above correspond exactly to the r -matches in w . Further, to obtain a translation of optimal complexity, r_0 and r_1 have to be of polynomial size in $|r|$. It is an open problem whether such a splitting is always possible, in particular in the presence of tests in r and
1180 guards with only finitely many r -matches.

Secondly, recall that the direct approach to robust LTL translates a formula φ of rLTL(\Box, \Diamond) into a Büchi automaton that captures φ with respect to all $\beta \in \mathbb{B}_4$ (by considering five initial states, one for each β). Trying to apply
1185 this approach to robust Prompt-LTL requires using a more general automaton model that is able to capture the quantitative nature of the prompt diamond operator while still yielding a model checking and a synthesis algorithm with the desired complexity. To the best of our knowledge, no such translation from Prompt-LTL to automata has been presented in the literature, which would be a special case of our construction here.

1190 Thus, according to the state-of-the-art, the direct approach is the only viable one for robust extensions of LDL while the reduction-based approach is the only viable one for robust extensions of Prompt-LTL. This leaves us with no viable approach for rPrompt-LDL.

Nevertheless, we identify a fragment of rPrompt-LDL for which both the

1195 model checking and the synthesis problem are decidable. We obtain this frag-
 ment by disallowing tests in guards and by requiring them to always have in-
 finitely many matches. For such formulas, one can translate the guard into a
 deterministic finite automaton (without tests) and then use this automaton to
 “split” r . However, this involves multiple exponential blowups and hence does
 1200 not prove that the fragment has the exponential compilation property. Nonethe-
 less, this translation shows that both model checking and synthesis are decidable
 for this fragment. The decidability of these problems for full rPrompt-LDL is
 left for further research and seemingly requires new approaches.

5.1. Restricting Guards in rPrompt-LDL

1205 We say that a guard r is test-free if it does not contain tests as atoms, but
 only propositional formulas over the atomic propositions. A formula is test-free
 if each of its guards is test-free. In the remainder, we only consider test-free
 formulas. As the adaptations made to define $\mathcal{R}_i^{\text{RPD}}$ are only concerned with tests,
 they can be ignored when reasoning about test-free formulas.

1210 **Remark 1.** *Let r be a test-free guard. Then, $\mathcal{R}_i^{\text{RPD}}(w, k, r)$ is independent of i
 and k for every trace w .*

Hence, in the following, we use $\mathcal{R}(w, r)$ (as defined for LDL) instead of
 $\mathcal{R}_i^{\text{RPD}}(w, k, r)$, since the definitions coincide for test-free guards.

1215 We say that a test-free guard r is limit-matching if we have $|\mathcal{R}(w, r)| = \infty$
 for every trace w . This is well-defined due to the previous remark. Again, a
 test-free formula is limit-matching if each of its guards is limit-matching.

Lemma 5. *The problem “Given a test-free formula φ , is φ limit-matching?”
 is in PSPACE.*

Proof. The problem is in PSPACE if one can decide in polynomial space whether
 1220 a single test-free guard is limit-matching. Hence, let r be such a guard, which
 is limit-matching if and only if infinitely many prefixes of each trace w match r .
 An application of König’s Lemma yields that the latter condition is equivalent
 to each w being Büchi-accepted by \mathfrak{G}_r . Due to test-freeness, \mathfrak{G}_r can indeed be
 seen as a Büchi automaton with ε -transitions. Hence, r is limit-matching if and
 1225 only if \mathfrak{G}_r is universal, which can be decided in polynomial space [38] (after
 eliminating ε -transitions). The automaton being of the same size as the guard
 and being efficiently constructible concludes the proof. \square

Example 5. *Recall the formula $\varphi = [((\neg t)^* ; t ; (\neg t)^* ; t)^* \cdot \langle \text{tt}^* \cdot \rangle_{\mathbf{P}} s$ from Ex-
 ample 4. It is test-free, but not limit-matching as traces with finitely many t
 1230 only have finitely many $((\neg t)^* ; t ; (\neg t)^* ; t)^*$ -matches.*

*Nevertheless, test-free and limit-matching rPrompt-LDL formulas can make
 use of arbitrary modulo counting, a significant advance in expressiveness over
 classical LTL, thus witnessing the usefulness of the fragment.*

1235 *For example, the formula $[r \cdot \langle r \cdot \rangle_{\mathbf{P}} s$ with $r = (\text{tt} ; \text{tt})^*$ expresses, when
 evaluated with respect to a bound k , that the distance between synchronizations*

at even positions is bounded by k , i.e., we use the test-free limit-matching guards to “filter out” the odd positions.

Let us note that for LDL, the test-free fragment is of equal expressive power as full LDL, albeit potentially less succinct. This claim follows easily from translating Büchi automata into LDL formulas, which results in test-free formulas.

In the following, we consider the model checking and the synthesis problem for test-free limit-matching formulas. To this end, we proceed as in the case of rPrompt-LTL: We reduce these problems to those for Prompt-LDL, i.e., we present a reduction-based translation to Büchi automata.

Due to only considering limit-matching formulas, we do not have to deal with the cases of having only finitely many matches of a guard. On the other hand, we have to “split” guards to capture the semantics of the robust diamond operator (recall the discussion in Section 5). Here, we exploit the formula under consideration being test-free.

The main technical result on this fragment states that the logic can be derobustified, i.e., translated into Prompt-LDL.

Theorem 8. *For every test-free limit-matching rPrompt-LDL formula φ and every $\beta \in \mathbb{B}_A$, there is a Prompt-LDL formula φ_β such that $V^{\text{RPD}}(w, k, \varphi) \succeq \beta$ if and only if $V^{\text{PD}}(w, k, \varphi_\beta) = 1$.*

Proof. Before we present the translation, we need to explain how to “split” guards, which is necessary to implement the semantics of the robust box operator (recall the discussion in Section 5). For example, we have to check that almost all r -matches are ψ -satisfying for some guard r and some subformula ψ . In LTL, “almost all” is expressed by $\diamond\Box$. We will use the analogous LDL operators, i.e., a formula of the form $\langle \cdot \rangle [\cdot]$. But now we need guards r_0 and r_1 for the diamond and the box operator so that the concatenation r_0r_1 is equivalent to r . To this end, we transform r into a deterministic automaton. Then, for each state q of that automaton there exists a guard $r_{q_I, q}$ capturing the words leading from the initial state to q , and a guard $r_{q, F}$ capturing all words leading from q to an accepting state. Ultimately, we end up with a formula of the form $\bigvee_q \langle r_{q_I, q} \rangle [r_{q, F}]$.

Let r be a test-free guard. Applying Lemma 3 to r yields an ε -NFA \mathfrak{G}_r without tests. Hence, eliminating ε -transitions and determinizing the resulting automaton yields a deterministic finite automaton \mathfrak{D}_r such that $w[0, j]$ is accepted by \mathfrak{D}_r if and only if $j \in \mathcal{R}(w, r)$. Furthermore, due to test-freeness, acceptance of $w[0, j]$ by \mathfrak{D}_r only depends on the prefix $w[0, j]$ of w , but not on the corresponding suffix $w[j, \infty)$. This property, which underlies the following construction, does not hold true for guards with tests.

Now, let Q be the set of states of \mathfrak{D}_r , q_I the initial state, and F the set of final states. Then, for every $q \in Q$, one can efficiently construct regular expressions (i.e., guards) $r_{q_I, q}$ and $r_{q, F}$ such that $w \in (2^P)^*$ is in the language of $r_{q_I, q}$ (of $r_{q, F}$) if and only if the unique run of \mathfrak{D}_r starting in q_I (in q) ends in q (in F).

Now, we are ready to construct φ_β . Again, the case $\beta = 0000$ is trivial.
 1280 Hence, we assume $\beta \succ 0000$ in the following. We proceed by induction over the construction of the formula:

- $p_\beta = p$ and $(\neg p)_\beta = \neg p$ for all atomic propositions $p \in P$ and all $\beta \succ 0000$.
- $(\varphi_0 \wedge \varphi_1)_\beta = (\varphi_0)_\beta \wedge (\varphi_1)_\beta$ for all $\beta \succ 0000$.
- $(\varphi_0 \vee \varphi_1)_\beta = (\varphi_0)_\beta \vee (\varphi_1)_\beta$ for all $\beta \succ 0000$.
- 1285 • $(\langle \cdot r \cdot \rangle \varphi)_\beta = \langle r \rangle (\varphi_\beta)$ for all $\beta \succ 0000$
- $([\cdot r \cdot] \varphi)_{1111} = [r] (\varphi_{1111})$,
- $([\cdot r \cdot] \varphi)_{0111} = \bigvee_{q \in Q} \langle r_{qI, q} \rangle [r_{q, F}] (\varphi_{0111})$, where $\mathfrak{D}_r = (Q, 2^P, q_I, \delta, F)$,
- $([\cdot r \cdot] \varphi)_{0011} = \bigwedge_{q \in Q} [r_{qI, q}] \langle r_{q, F} \rangle (\varphi_{0011})$, where $\mathfrak{D}_r = (Q, 2^P, q_I, \delta, F)$,
- $([\cdot r \cdot] \varphi)_{0001} = \langle r \rangle (\varphi_{0001})$, and
- 1290 • $(\langle \cdot r \cdot \rangle_{\mathbf{P}} \varphi)_\beta = \langle r \rangle_{\mathbf{P}} (\varphi_\beta)$ for all $\beta \succ 0000$.

In the construction of $([\cdot r \cdot] \varphi)_{0011}$, we rely on \mathfrak{D}_r being deterministic, since we quantify over all states reached by a prefix. In a non-deterministic automaton, there could be rejecting runs on accepted words, which would still be required to be completable to an accepting run by $([\cdot r \cdot] \varphi)_{0011}$.

1295 A straightforward induction over the construction of φ , relying on the fact that φ is limit-matching, yields the correctness of the translation. The fact that φ is limit-matching explains the construction of $([\cdot r \cdot] \varphi)_\beta$, which only has to implement the first case (“ $|\mathcal{R}(w, r)| = \infty$ ”) of the definition of the semantics. \square

Now, the model checking and the synthesis problem for rPrompt-LDL, which
 1300 are defined as expected, can be solved by reducing them to their analogues for Prompt-LDL (cf. Section 2.4). We obtain the following results.

Corollary 2. *The rPrompt-LDL model checking and synthesis problem are decidable for the test-free limit-matching fragment.*

We refrain from specifying the exact complexity of the algorithms, as we
 1305 conjecture them to be several exponents away from optimal algorithms: The guards $r_{qI, q}$ and $r_{q, F}$ are already of doubly-exponential size and we still have to translate the formula φ_β containing these guards into (deterministic) automata to solve the problems.

Note that our approach for the fragment, which relies on a translation to
 1310 Prompt-LDL, cannot easily be extended to formulas with tests and to formulas with non-limit-matching guards. The existence of tests complicates the construction of the deterministic automaton required to “split” the guards. Consider, for example, the guard $(\varphi_0?; a; \varphi'_0) + (\varphi_1?; a; \varphi'_1)$: after processing an a , depending on which tests hold true before the a , the automaton still has to distinguish whether φ'_0 or φ'_1 has to hold after processing the a . Implementing this
 1315

requires non-determinism that cannot be resolved while only reading a prefix of a trace.

1320 Complicating the situation even further, the lack of negations in prompt logics does not allow to “disambiguate” the guard. Similarly, allowing non-limit-matching guards requires us to implement the full case distinction in the definition of the semantics of the robust box operator. However, implementing a case distinction in Prompt-LDL is again complicated by the lack of negations.

6. Conclusion

1325 We addressed the problems of verification and synthesis with robust, expressive, and quantitative linear temporal specifications. Inspired by robust LTL, we have first developed robust extensions of the logics LDL and Prompt-LTL, named rLDL and rPrompt-LTL, respectively. Then, we combined rLDL and rPrompt-LTL into a third logic, named rPrompt-LDL, which has the expressiveness of ω -regular languages and allows robust reasoning about timing bounds.

1330 For rLDL and rPrompt-LTL, we have shown how to solve the model checking and synthesis problem relying on the exponential compilation property. Hence, all these problems are not harder than those for plain LTL. The situation for the combination of all three basic logics, i.e., for rPrompt-LDL, is less encouraging. We show the problems to be decidable for an important fragment, but due to a blowup of the formulas during the reduction, we (most likely) do not obtain 1335 optimal algorithms. Decidability for the full logic remains open.

In future work, we aim to determine the exact complexity of the model checking and synthesis problem for (full) rPrompt-LDL. One promising approach is to generalize the translation of rLDL into weak alternating Büchi 1340 automata. However, this requires a suitable quantitative alternating automata model with strong closure properties that can be transformed into equivalent non-deterministic and deterministic automata.

Another promising direction for further research is to study the semantics for the robust box operator proposed in Footnote 5 on Page 17. In particular, it 1345 is open whether the translation into alternating automata can be generalized to this setting without a blowup. Also, we leave open whether full robust LTL, i.e., with until and release, can be embedded into rLDL. As is, the robust semantics of the release operator (see [22]) is not compatible with our robust semantics for rLDL. In future work, we plan to study generalizations of full robust LTL.

1350 Another natural question is whether the techniques developed for rLDL can be applied to a robust version of the Property Specification Language [3].

Acknowledgements We would like to thank the reviewers for their detailed feedback, which improved the paper considerably.

References

- 1355 [1] D. Neider, A. Weinert, M. Zimmermann, Robust, expressive, and quantitative linear temporal logics: Pick any two for free, in: J. Leroux,

- J. Raskin (Eds.), GandALF 2019, Vol. 305 of EPTCS, 2019, pp. 1–16. doi:10.4204/EPTCS.305.1.
- 1360 [2] A. Pnueli, The temporal logic of programs, in: FOCS 1977, IEEE, 1977, pp. 46–57. doi:10.1109/SFCS.1977.32.
- [3] C. Eisner, D. Fisman, A Practical Introduction to PSL, Integrated Circuits and Systems, Springer, 2006. doi:10.1007/978-0-387-36123-9.
- 1365 [4] L. Fix, Fifteen years of formal property verification in intel, in: O. Grumberg, H. Veith (Eds.), 25 Years of Model Checking - History, Achievements, Perspectives, Vol. 5000 of LNCS, Springer, 2008, pp. 139–144. doi:10.1007/978-3-540-69850-0_8.
- [5] M. Y. Vardi, P. Wolper, Reasoning about infinite computations, Inf. Comput. 115 (1) (1994) 1–37. doi:10.1006/inco.1994.1092.
- 1370 [6] P. Faymonville, M. Zimmermann, Parametric linear dynamic logic, Inf. Comput. 253 (2017) 237–256. doi:10.1016/j.ic.2016.07.009.
- [7] M. Leucker, C. Sánchez, Regular linear temporal logic, in: C. B. Jones, Z. Liu, J. Woodcock (Eds.), ICTAC 2007, Vol. 4711 of LNCS, Springer, 2007, pp. 291–305. doi:10.1007/978-3-540-75292-9_20.
- 1375 [8] M. Y. Vardi, The rise and fall of LTL, in: G. D’Agostino, S. L. Torre (Eds.), GandALF 2011, Vol. 54 of EPTCS, 2011.
- [9] P. Wolper, Temporal logic can be more expressive, Information and Control 56 (1/2) (1983) 72–99. doi:10.1016/S0019-9958(83)80051-5.
- 1380 [10] R. Alur, K. Etessami, S. L. Torre, D. Peled, Parametric temporal logic for “model measuring”, ACM Trans. Comput. Log. 2 (3) (2001) 388–407. doi:10.1145/377978.377990.
- [11] R. Koymans, Specifying real-time properties with metric temporal logic, Real-Time Systems 2 (1990) 255–299. doi:10.1007/BF01995674.
- 1385 [12] O. Kupferman, N. Piterman, M. Y. Vardi, From liveness to promptness, Formal Methods in System Design 34 (2) (2009) 83–103. doi:10.1007/s10703-009-0067-z.
- [13] M. Zimmermann, Parameterized linear temporal logics meet costs: still not costlier than LTL, Acta Inf. 55 (2) (2018) 129–152. doi:10.1007/s00236-016-0279-9.
- 1390 [14] R. Bloem, K. Chatterjee, K. Greimel, T. A. Henzinger, B. Jobstmann, Robustness in the presence of liveness, in: CAV 2010, Vol. 6174 of LNCS, Springer, 2010, pp. 410–424. doi:10.1007/978-3-642-14295-6_36.

- [15] E. Dallal, D. Neider, P. Tabuada, Synthesis of safety controllers robust to unmodeled intermittent disturbances, in: CDC 2016, 2016, pp. 7425–7430. doi:10.1109/CDC.2016.7799416.
- 1395 [16] A. Donzé, O. Maler, Robust satisfaction of temporal logic over real-valued signals, in: K. Chatterjee, T. A. Henzinger (Eds.), FORMATS 2010, Vol. 6246 of LNCS, Springer, 2010, pp. 92–106. doi:10.1007/978-3-642-15297-9_9.
- 1400 [17] L. Doyen, T. A. Henzinger, A. Legay, D. Nickovic, Robustness of sequential circuits, in: L. Gomes, V. Khomenko, J. M. Fernandes (Eds.), ACSD 2010, IEEE Computer Society, 2010, pp. 77–84. doi:10.1109/ACSD.2010.26.
- [18] G. E. Fainekos, G. J. Pappas, Robustness of temporal logic specifications for continuous-time signals, *Theor. Comput. Sci.* 410 (42) (2009) 4262–4291. doi:10.1016/j.tcs.2009.06.021.
- 1405 [19] R. Majumdar, I. Saha, Symbolic robustness analysis, in: T. P. Baker (Ed.), RTSS 2009, IEEE Computer Society, 2009, pp. 355–363. doi:10.1109/RTSS.2009.17.
- [20] D. Neider, A. Weinert, M. Zimmermann, Synthesizing optimally resilient controllers, in: D. R. Ghica, A. Jung (Eds.), CSL 2018, Vol. 119 of LIPIcs, Schloss Dagstuhl - LZI, 2018, pp. 34:1–34:17. doi:10.4230/LIPIcs.CSL.2018.34.
- 1410 [21] P. Tabuada, S. Y. Caliskan, M. Rungger, R. Majumdar, Towards robustness for cyber-physical systems, *IEEE Trans. Automat. Contr.* 59 (12) (2014) 3151–3163. doi:10.1109/TAC.2014.2351632.
- 1415 [22] P. Tabuada, D. Neider, Robust linear temporal logic, in: J. Talbot, L. Regnier (Eds.), CSL 2016, Vol. 62 of LIPIcs, Schloss Dagstuhl - LZI, 2016, pp. 10:1–10:21. doi:10.4230/LIPIcs.CSL.2016.10.
- [23] T. Anevlavis, M. Philippe, D. Neider, P. Tabuada, Verifying rLTL formulas: now faster than ever before!, in: CDC 2018, IEEE, 2018, pp. 1556–1561. doi:10.1109/CDC.2018.8619014.
- 1420 [24] T. Anevlavis, D. Neider, M. Phillippe, P. Tabuada, Evrostos: the rLTL verifier, in: N. Ozay, P. Prabhakar (Eds.), HSCC 2019, ACM, 2019, pp. 218–223. doi:10.1145/3302504.3311812.
- [25] C. Mascle, D. Neider, M. Schwenger, P. Tabuada, A. Weinert, M. Zimmermann, From LTL to rLTL monitoring: Improved monitorability through robust semantics, in: HSCC 2020, SCM, New York, NY, USA, 2020. doi:10.1145/3365365.3382197.
- 1425 [26] R. Alur, S. La Torre, P. Madhusudan, Playing games with boxes and diamonds, in: R. M. Amadio, D. Lugiez (Eds.), CONCUR 2003, Vol. 2761 of LNCS, Springer, 2003, pp. 127–141. doi:10.1007/978-3-540-45187-7_8.
- 1430

- [27] R. Alur, S. L. Torre, Deterministic generators and games for LTL fragments, *ACM Trans. Comput. Log.* 5 (1) (2004) 1–25. doi:10.1145/963927.963928.
- 1435 [28] G. De Giacomo, M. Y. Vardi, Linear temporal logic and linear dynamic logic on finite traces, in: F. Rossi (Ed.), *IJCAI, IJCAI/AAAI*, 2013.
- [29] M. Zimmermann, Optimal bounds in parametric LTL games, *Theor. Comput. Sci.* 493 (2013) 30–45. doi:10.1016/j.tcs.2012.07.039.
- 1440 [30] E. Grädel, W. Thomas, T. Wilke (Eds.), *Automata, Logics, and Infinite Games: A Guide to Current Research*, Vol. 2500 of LNCS, Springer, 2002. doi:10.1007/3-540-36387-4.
- [31] C. Baier, J.-P. Katoen, *Principles of Model Checking*, The MIT Press, 2008.
- [32] P. Faymonville, M. Zimmermann, Parametric linear dynamic logic, in: A. Peron, C. Piazza (Eds.), *GandALF 2014*, Vol. 161 of EPTCS, 2014, pp. 60–73. doi:10.4204/EPTCS.161.8.
- 1445 [33] S. Miyano, T. Hayashi, Alternating finite automata on omega-words, *Theor. Comput. Sci.* 32 (1984) 321–330. doi:10.1016/0304-3975(84)90049-5.
- 1450 [34] A. Weinert, M. Zimmermann, Visibly linear dynamic logic, in: A. Lal, S. Akshay, S. Saurabh, S. Sen (Eds.), *FSTTCS 2016*, Vol. 65 of LIPIcs, Schloss Dagstuhl - LZI, 2016, pp. 28:1–28:14. doi:10.4230/LIPIcs.FSTTCS.2016.28.
- [35] A. P. Sistla, E. M. Clarke, The complexity of propositional linear temporal logics, *J. ACM* 32 (3) (1985) 733–749. doi:10.1145/3828.3837.
- 1455 [36] C. S. Calude, S. Jain, B. Khoussainov, W. Li, F. Stephan, Deciding parity games in quasipolynomial time, in: *STOC 2017*, ACM, 2017, pp. 252–263. doi:10.1145/3055399.3055409.
- 1460 [37] A. Pnueli, R. Rosner, On the synthesis of an asynchronous reactive module, in: G. Ausiello, M. Dezani-Ciancaglini, S. R. D. Rocca (Eds.), *ICALP 1989*, Vol. 372 of LNCS, Springer, 1989, pp. 652–671. doi:10.1007/BFb0035790.
- [38] A. P. Sistla, M. Y. Vardi, P. Wolper, The complementation problem for Büchi automata with applications to temporal logic (extended abstract), in: W. Brauer (Ed.), *ICALP 1985*, Vol. 194 of LNCS, Springer, 1985, pp. 465–474. doi:10.1007/BFb0015772.