

ON RELAXING METRIC INFORMATION IN LINEAR TEMPORAL LOGIC

Carlo A. Furia* and Paola Spoletini†

June 3, 2011

Abstract

Metric LTL formulas rely on the *next* operator to encode time distances, whereas qualitative LTL formulas use only the *until* operator. This paper shows how to transform any metric LTL formula M into a qualitative formula Q , such that Q is satisfiable if and only if M is satisfiable over words with variability bounded with respect to the largest distances used in M (i.e., occurrences of *next*), but the size of Q is independent of such distances. Besides the theoretical interest, this result can help simplify the verification of systems with time-granularity heterogeneity, where large distances are required to express the coarse-grain dynamics in terms of fine-grain time units.

1 Introduction and motivation

Linear temporal logic (LTL) supports a simple model of *metric time* through the *next* operator X . Under the assumption of a one-to-one correspondence between consecutive states and discrete instants of time, nested occurrences of X “count” instants to express time distances. LTL formulas without X — using only the *until* operator — are instead purely *qualitative*: they constrain the ordering of events, not their absolute distance. Therefore, qualitative LTL formulas express models that are insensitive to additions or removals of *stuttering steps*: consecutive repetitions of the same state. The fundamental properties of LTL with respect to its qualitative subset are well known from classic work: quantitative (metric) LTL is strictly more expressive [24, 27, 15, 23], but reasoning has the same worst-case complexity [32, 12].

The present paper investigates when the metric information, encoded by nested occurrences of X , is redundant and can be relaxed. The relaxation transforms a quantitative formula into an equi-satisfiable qualitative one that is independent of the number of X in the original formula; reasoning on the transformed formula is thus simpler by a factor proportional to the amount of metric information stripped.

The motivation behind this study refers to an informal notion of *redundancy*, which stuttering steps seem to encode. Consider a metric LTL formula ϕ describing models characterized by many stuttering steps distributed over large time distances; for example, the formalization of an event for elections that occur every four years in November, in a variable day of the month, with the day as time unit. Formula ϕ is large because it encodes large time distances in unary form with many occurrences of the X operator;

*Chair of Software Engineering, ETH Zürich, Switzerland. E-mail: caf@inf.ethz.ch.

†Università degli Studi dell’Insubria, Italy. E-mail: paola.spoletini@uninsubria.it.

for example, a four-year distance requires at least 1460 “next”, one for each day. However, the information carried by ϕ is prominently redundant as every stuttering step is a duplication that only pads uneventful time instants. Is it possible, under a rigorous assumption of “sparse events”, to simplify ϕ into an equi-satisfiable formula ϕ' which does not encode explicitly the redundant information?

The notion of *bounded variability*, adapted from dense-time models, provides a suitable formalization of the intuitive notion of “sparse events”: models with bounded variability have, over every interval of fixed length, only a limited number v of steps that are not stuttering (i.e., redundant repetitions). The main result of the paper (in Section 5) shows how to transform efficiently any LTL formula ϕ into a qualitative formula ϕ' such that ϕ is satisfiable *over models with bounded variability* iff ϕ' is satisfiable over models of any variability. The size of ϕ' does not depend on the distances (i.e., the number of nested occurrences of X) in ϕ but only on the maximum number of non-stuttering steps v . In other words, ϕ' drops some information encoded in ϕ ; this information is not needed to decide satisfiability over models with bounded variability.

On the technical level, the construction that eliminates metric information relies on a normal form for LTL formulas and on discrete-time generalized versions of the dense-time *Pnueli operators* [20]. The correctness proof follows the idea of adding and removing stuttering steps to re-introduce the metric information dropped in models satisfying only qualitative constraints; it is reminiscent of the notion of *stretching*, also originally introduced for dense-time models [21, 5]. Section 4 first demonstrates this technique by showing how to transform any metric LTL formula into a qualitative one which is equi-satisfiable (for generic models) and of polynomial size. The feasibility of such a construction is unsurprising in hindsight, given the complexity results about qualitative LTL [32] and Etessami’s construction [15]. However, it is the necessary basis of the techniques used to derive the main result for models with bounded variability.

Besides the theoretical interest, the results of the present paper may be practically useful to simplify the temporal-logic analysis of systems characterized by heterogeneous components evolving over wildly different time scales, such as minutes, weeks, and years. Assuming incommensurable distances are not a concern, such heterogeneity of *time granularities* [17] can, in principle, be modeled in terms of the finest-grain time units; but this solution comes with a significant price to pay to accommodate the largest time units in terms of the smallest, resulting in huge formulas. If, however, the dynamics of the components with faster time scales are “sparse” enough, there is a redundancy in the global behavior of the system that the notion of bounded variability captures. Hence, the analysis can be carried out more efficiently by leveraging the results of the present paper.

Outline. The paper is organized as follows. The rest of the present section recalls related work. Section 2 introduces notation and basic definitions. Section 3 presents normal forms for LTL formulas. Section 4 proves the equi-satisfiability of LTL and its qualitative subset. Section 5 shows how the metric information can be relaxed while preserving satisfiability, for models with bounded variability. Section 6 concludes and outlines future work.

1.1 Related work

The expressiveness and complexity of LTL and of its qualitative subset have been thoroughly investigated in the classic framework of temporal logic [19, 13, 22]. With respect to expressiveness, Lamport introduced the notion of *stuttering* to characterize qualitative LTL [24]; the characterization was completed by Peled and Wilke [27], perfected by Etessami and others [14, 15, 30, 11], and generalized by Kučera and Strejček [23]. With respect to complexity, the seminal work of Sistla and Clarke established the PSPACE-completeness of both LTL and qualitative LTL [32], and other authors have generalized or specialized the result [12, 26, 2].

To our knowledge, the present paper is the first investigating satisfiability-preserving relaxations of metric information in temporal logic formulas. More generally, the problem of formalizing systems with heterogeneous time granularities using temporal logic [17] has been studied by only a few authors [9, 10, 8, 6]; [8], in particular, presents an encoding of temporal granularities in LTL, but it does not discuss efficiency of the encoding.

Some of the techniques used in the the paper borrow from existing approaches in the literature. The normal forms for LTL introduced in Section 3 are related to a construction used in *temporal testers* [28]. The definition of *bounded variability* in Section 2 translates to discrete time a notion introduced for dense (or continuous) time models [33, 16, 18, 7].

Hirshfeld and Rabinovich studied the expressiveness and decidability of Pnueli operators over dense time [20]; the operators themselves were first mentioned in a conjecture attributed to Pnueli [1, 33]. Section 5 introduces discrete-time qualitative variants of such operators. *Counting operators* [25] are somehow similar to discrete-time Pnueli operators in that they both facilitate the expression of concise counting requirements; both extensions do not increase the expressive power of LTL, nor its complexity under a unary encoding. [4] introduce a much more expressive counting extension of LTL, which is decidable only in special cases.

The proofs of Lemmas 11 and 20 use a technique that removes and adds stuttering steps in words to match some metric requirements; the notion of *stretching* — introduced in [21] and further used in [5] — is similar but for dense-time models.

2 Definitions

This section introduces the syntax and semantics of LTL and other basic definitions.

\mathbb{N} denotes the set of natural numbers $\{0, 1, 2, \dots\}$ and $\mathbb{N}_{>0}$ denotes the positive naturals $\{n \in \mathbb{N} \mid n > 0\}$. For any two natural numbers $a \leq b$, $[a..b]$ denotes the interval of naturals $a, a + 1, \dots, b$.

2.1 LTL formulas

LTL syntax. The following grammar defines the set of LTL formulas:

$$\text{LTL } \ni \phi ::= x \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \text{ U } \phi_2 \mid \text{X } \phi$$

where x ranges over a set $\mathcal{P} = \{p, q, r, \dots\}$ of propositional letters.

Assume the standard abbreviations for $\top, \perp, \vee, \Rightarrow, \Leftrightarrow$ and for the derived temporal operators:

- *eventually*: $\text{F}\phi \triangleq \top \text{ U } \phi$;

- *always*: $G\phi \triangleq \neg F\neg\phi$;
- *release*: $\phi_1 R \phi_2 \triangleq \neg(\neg\phi_1 U \neg\phi_2)$;
- *distance* $X^k \phi = \underbrace{XX \cdots X}_k \phi$ for $k \geq 0$.

Size and height. Let ϕ be an LTL formula. $\mathcal{P}(\phi) \subseteq \mathcal{P}$ denotes the (finite) set of propositional letters occurring in ϕ . $|\phi|$ denotes the size of ϕ . Three features determine the size of ϕ : the size $|\phi|_p$ of its propositional structure; the size $|\phi|_U$ of its *until* subformulas; and the size $|\phi|_X$ of its *next* subformulas. They are defined inductively as follows.

$$(|\phi|_p, |\phi|_U, |\phi|_X) = \begin{cases} \langle 1, 0, 0 \rangle & \phi = x \\ \langle 1 + |\phi'|_p, |\phi'|_U, |\phi'|_X \rangle & \phi = \neg\phi' \\ \langle 1 + |\phi_1|_p + |\phi_2|_p, |\phi_1|_U + |\phi_2|_U, |\phi_1|_X + |\phi_2|_X \rangle & \phi = \phi_1 \wedge \phi_2 \\ \langle |\phi_1|_p + |\phi_2|_p, 1 + |\phi_1|_U + |\phi_2|_U, |\phi_1|_X + |\phi_2|_X \rangle & \phi = \phi_1 U \phi_2 \\ \langle |\phi'|_p, |\phi'|_U, 1 + |\phi'|_X \rangle & \phi = X\phi' \end{cases}$$

Correspondingly, $|\phi|$ is $|\phi|_p + |\phi|_U + |\phi|_X$.

For a temporal operator $H \in \{U, X\}$, the *temporal height* (or *nesting depth*) $\mathcal{H}(\phi, H)$ of H in ϕ is the maximum number of nested occurrences of H in ϕ . For example, $\mathcal{H}(\phi, X) = 0$ iff X is not used in ϕ . $d(\phi)$ denotes instead the maximum number of *consecutive* nested occurrences of the *next* operator, that is the largest n such that X^n occurs in ϕ ; clearly, $d(\phi) \leq \mathcal{H}(\phi, X)$. Finally, $s(\phi)$ is the number of *distinct subformulas* of the form $X^m \phi$ with $m \geq 1$. Notice that $|\phi|_X$ is bounded by $d(\phi) \cdot s(\phi)$, hence $|\phi|$ is in $O(|\phi|_p + |\phi|_U + d(\phi) \cdot s(\phi))$.

$L(U^{h_1}, X^{h_2})$ denotes the fragment of LTL whose formulas ψ are such that $\mathcal{H}(\psi, U) \leq h_1$ and $\mathcal{H}(\psi, X) \leq h_2$. Omit the superscript to mean that there is no bound on the temporal height of an operator. Hence, $L(U, X)$ is the same as all LTL; $L(U, X^0) = L(U)$ denotes *qualitative* LTL, where no *next* operator is used; and $L(U^0, X^0) = P(\mathcal{P})$ denotes *propositional formulas* without any temporal operator.

Example 1. Consider the two formulas:

$$\Gamma_1 \triangleq X(p \wedge X((p U q) \wedge Xq)) \quad \Gamma_2 \triangleq Xp \wedge X^2(p U q) \wedge X^3q.$$

Γ_1 and Γ_2 are semantically equivalent (see Example 3) but syntactically different; in fact, some size parameters differ in the two formulas: $|\Gamma_1|_p = |\Gamma_2|_p = 5$; $|\Gamma_1|_U = |\Gamma_2|_U = 1$; $|\Gamma_1|_X = 3$, $|\Gamma_2|_X = 6$; $|\Gamma_1| = 12$, $|\Gamma_2| = 9$; $\mathcal{H}(\Gamma_1, U) = \mathcal{H}(\Gamma_2, U) = 1$; $\mathcal{H}(\Gamma_1, X) = \mathcal{H}(\Gamma_2, X) = 3$; $d(\Gamma_1) = 1$, $d(\Gamma_2) = 3$; $s(\Gamma_1) = s(\Gamma_2) = 3$. \diamond

ω -words. An ω -word (or simply *word*) over a set S of propositional letters is a mapping $w : \mathbb{N} \rightarrow 2^S$ or, equivalently, a denumerable sequence $w(0)w(1)\cdots$ of elements $w(i) \subseteq S$. The set of all ω -words over S is denoted by $\mathcal{W}[S]$.

For $T \subseteq S$, $w|_T$ is the projection of w over T , defined as $w(0)|_T w(1)|_T \cdots$, where $w(i)|_T = w(i) \cap T$ for all $i \in \mathbb{N}$. The projection is extended to sets of words as expected.

For $i, j \in \mathbb{N}$, w_i denotes the suffix $w(i)w(i+1)\cdots$ of w ; $w(i, j)$ denotes the subword of w of length j starting at $w(i)$ (with $w(i, 0) = \epsilon$ for all i); and $w(i : j)$ denotes the subword $w(i)w(i+1)\cdots w(j)$ (with $w(i, j) = \epsilon$ for all $j < i$).

LTL semantics. The satisfaction relation \models is defined as usual, for an LTL formula ϕ , interpreted over an ω -word w over \mathcal{P} , at position $i \in \mathbb{N}$.

$w, i \models p$	iff	$p \in w(i)$
$w, i \models \neg\phi$	iff	$w, i \not\models \phi$
$w, i \models \phi_1 \wedge \phi_2$	iff	$w, i \models \phi_1$ and $w, i \models \phi_2$
$w, i \models \phi_1 \cup \phi_2$	iff	there exists $j \geq i$ such that $w, j \models \phi_2$ and for all $i \leq k < j$ it is $w, k \models \phi_1$
$w, i \models X\phi$	iff	$w, i + 1 \models \phi$
$w \models \phi$	iff	$w, 0 \models \phi$

Satisfiability and validity. $\llbracket \phi \rrbracket$ denotes the set $\{w \in \mathcal{W}[\mathcal{P}] \mid w \models \phi\}$ of all models of ϕ . ϕ is *satisfiable* iff $\llbracket \phi \rrbracket \neq \emptyset$ and is *valid* iff $\llbracket \phi \rrbracket = \mathcal{W}[\mathcal{P}]$. Two formulas ϕ_1, ϕ_2 are *equivalent* iff $\llbracket \phi_1 \rrbracket = \llbracket \phi_2 \rrbracket$; they are *equi-satisfiable* iff they are either both satisfiable or both unsatisfiable.

Proposition 2 ([32]). *Checking the satisfiability of an LTL or qualitative LTL formula is complete for PSPACE; it can be done in time exponential in the size of the formula.*

Example 3. Consider again Γ_1, Γ_2 in Example 1. If \mathcal{S} denotes the set of words w such that $p \in w(1), q \in w(2)$ or $p \in w(2),$ and $q \in w(3)$, then $\llbracket \Gamma_1 \rrbracket = \llbracket \Gamma_2 \rrbracket = \mathcal{S}$. \diamond

2.2 Stuttering

A position $i \in \mathbb{N}$ is *redundant* in a word w iff $w(i+1) = w(i)$ and there exists a $j > i$ such that $w(j) \neq w(i)$; a redundant position is also called *stuttering step*. Conversely, a *non-stuttering step (nss)* is any position i such that $w(i+1) \neq w(i)$ or $w(i+j) = w(i)$ for all $j \in \mathbb{N}$.

A *stutter-free* word is one without stuttering steps. Two words w_1, w_2 are *stutter-equivalent* (or equivalent under stuttering) iff they are reducible to the same stutter-free word by removing an arbitrary number of stuttering steps.

A set of words W is *closed under stuttering* (or stutter-invariant) iff for every word $w \in W$, for all words w' such that w and w' are stutter-equivalent, $w' \in W$ too.

Recall the following fundamental results about stuttering and LTL.

Proposition 4. *Closure under stutter equivalence is a necessary and sufficient condition for qualitative LTL languages; that is:*

- [24] $\phi \in \mathbf{L}(\mathbf{U})$ implies that $\llbracket \phi \rrbracket$ is closed under stutter equivalence;
- [27] W closed under stutter equivalence and expressible in LTL implies there exists $\phi \in \mathbf{L}(\mathbf{U})$ such that $\llbracket \phi \rrbracket = W$.

2.3 Variability

Let W be a set of words and v, k two positive integers. A set of propositional letters $P \subseteq \mathcal{P}$ has *variability bounded by v/k in W* iff: for every $w \in W$, the projection $w(i, k)|_P$ over P of every subword $w(i, k)$ of length k has at most v nss. $\mathbf{var}(P, v/k)$ denotes the set of all words where P has variability bounded by v/k . Note that $\mathbf{var}(P, v/k)$ is not closed under stuttering for any $v < k$.

Example 5 (The elections). Consider elections that occur every four years, in one of two consecutive days. The example is deliberately kept simple to be able to demonstrate it with the various constructions of the paper. Proposition q marks the first day of every quadrennial, hence it holds initially and then precisely every $d_4 = 365 \cdot 4 = 1460$ days. The elections e occur once within every quadrennial; precisely they occur $d_2 = 40$ or $d_3 = 41$ days before the end of the quadrennial. Assuming models with variability bounded by $5/1460$, the behavior is completely described by the following formula.

$$q \tag{1}$$

$$\wedge \mathbf{G} (q \Rightarrow \mathbf{X} (\neg q \wedge \neg q \mathbf{U} q) \wedge \mathbf{X}^{d_4} q) \tag{2}$$

$$\wedge \mathbf{G} (q \Rightarrow \mathbf{X} \neg (\neg e \mathbf{U} q)) \tag{3}$$

$$\wedge \mathbf{G} (e \Rightarrow \neg q \wedge \mathbf{X} (\neg e \mathbf{U} q)) \tag{4}$$

$$\wedge \mathbf{G} (e \Rightarrow \mathbf{X}^{d_2} q \vee \mathbf{X}^{d_3} q) \tag{5}$$

The proposition q marks the beginning of every quadrennial: q holds initially (1) and then always *at least* every d_4 steps (2). The elections, marked by proposition e , must occur once before the next quadrennial starts (3). They must also occur not at the beginning of a new quadrennial and at most once during the quadrennial (4); precisely, they occur d_2 or d_3 days before the end of the current quadrennial (5). A variability of $5/1460$ makes such model tight, as it allows *at most* 5 nss over a windows of length 1460: 2 of them accounts for q becoming true and then false again once, and the other 3 nss mark a similar double transition of e . \diamond

3 Normal forms for LTL

This section presents two normal forms for LTL where the nesting of temporal operators is limited; the results in the following sections will use these normal forms.

3.1 Flat-next form

An LTL formula is in *flat-next form* (FNF) when it is written as:

$$\kappa \wedge \mathbf{G} \left(\bigwedge_{i=1, \dots, N} (x_i \Leftrightarrow \mathbf{X} \pi_i) \right) \tag{6}$$

where $\kappa \in \mathbf{L}(\mathbf{U})$, $x_i \in \mathcal{P}$, $\pi_i \in \mathbf{P}(\mathcal{P})$. Clearly, (6) $\in \mathbf{L}(\mathbf{U}, \mathbf{X}^1)$.

The nesting depth of the \mathbf{X} operators can always be reduced to one without affecting satisfiability or complexity.

Lemma 6. *For any $\phi \in \mathbf{LTL}$ it is possible to build, in polynomial time, an equisatisfiable formula η in FNF such that $|\eta|$ and $|\mathcal{P}(\eta)|$ are polynomial in $|\phi|$.*

Proof. Initially, let $\mathcal{Q} = \mathcal{P}(\phi)$ and $\phi' = \phi$. Repeat the following two steps until $\phi' \in \mathbf{L}(\mathbf{U})$, with step 1 having higher precedence than step 2:

1. Replace a sub-formula of ϕ' in the form $\mathbf{X} \pi$, with $\pi \in \mathbf{P}(\mathcal{Q})$, by a fresh propositional letter $p_{\mathbf{X} \pi}$, and add $p_{\mathbf{X} \pi}$ to \mathcal{Q} .

2. Replace a maximal qualitative sub-formula $\psi \in \mathbf{L}(\mathbf{U})$ of ϕ' that is within the scope of some \mathbf{X} operator by a fresh propositional letter p_ψ , and add p_ψ to \mathcal{Q} .

Define κ as $\phi' \wedge \mathbf{G}(\bigwedge_{\substack{p_\psi \in \mathcal{Q} \\ \psi \in \mathbf{L}(\mathbf{U})}} (p_\psi \Leftrightarrow \psi))$; and η as $\kappa \wedge \mathbf{G}(\bigwedge_{p_{\mathbf{X}\pi} \in \mathcal{Q}} (p_{\mathbf{X}\pi} \Leftrightarrow \mathbf{X}\pi))$. η is in FNF and equi-satisfiable to ϕ . Moreover, steps 1–2 are repeated at most a number of times proportional to $|\phi|$, hence $|\eta|$ and $|\mathcal{Q}(\eta)|$ are polynomial in $|\phi|$. \square

Example 7. The following is the formula in the elections Example 5 in flat-next form.

$$\underbrace{\left(q \wedge \mathbf{G} \left(\begin{array}{l} (u \Leftrightarrow \neg e \mathbf{U} q) \\ \wedge (v \Leftrightarrow \neg q \wedge \neg q \mathbf{U} q) \\ \wedge (q \Rightarrow x_v \wedge x_{d_4}) \\ \wedge (q \Rightarrow \neg x_u) \\ \wedge (e \Rightarrow \neg q \wedge x_u) \\ \wedge (e \Rightarrow x_{d_2} \vee x_{d_3}) \end{array} \right) \right)}_{\kappa} \wedge \mathbf{G} \left(\begin{array}{l} (x_u \Leftrightarrow \mathbf{X} u) \wedge (x_v \Leftrightarrow \mathbf{X} v) \\ \wedge (x_1 \Leftrightarrow \mathbf{X} q) \\ \wedge_{2 \leq k \leq d_4} (x_k \Leftrightarrow \mathbf{X} x_{k-1}) \end{array} \right) \quad (7)$$

The first conjunct is the qualitative part κ , and x_k encodes $\mathbf{X}^k q$ for $k \geq 1$. \diamond

3.2 Separated-next form

An LTL formula is in *separated-next form* (SNF) when it is written as:

$$\kappa \wedge \mathbf{G} \left(\bigwedge_{i=1, \dots, M} (x_i \Leftrightarrow \mathbf{X}^{\mathbb{D}(i)} \pi_i) \right) \quad (8)$$

where $\kappa \in \mathbf{L}(\mathbf{U})$, $x_i \in \mathcal{P}$, $\pi_i \in \mathbf{P}(\mathcal{P})$, and \mathbb{D} is a monotonically non-decreasing mapping $[1..M] \rightarrow \mathbb{N}_{>0}$.

Given that the FNF is a special case of the SNF, it is obvious that any LTL formula can be transformed into an equi-satisfiable SNF one in polynomial time. The SNF, however, becomes interesting when it isolates subformulas with a nesting depth of \mathbf{X} as high as possible, as stated in the following.

Lemma 8. *For any $\phi \in \mathbf{LTL}$ it is possible to build, in polynomial time, an equi-satisfiable formula η in SNF (8) such that $|\kappa|$, $\max_i |\pi_i|$, and $|\mathcal{P}(\eta)|$ are in $O(|\phi|_{\mathbf{p}} + |\phi|_{\mathbf{U}} + s(\phi))$, $M = s(\phi)$, and $d(\eta) = \max_i \mathbb{D}(i) = \mathbb{D}(M) = d(\phi)$.*

Proof. The construction mirrors the proof of Lemma 6, with step 1 replaced by:

- 1'. Replace a sub-formula of ϕ' in the form $\mathbf{X}^n \pi$ for a maximal $n \geq 1 \dots$

κ introduces at most a proposition for each of the $s(\phi)$ maximal *next*-subformulas of ϕ and does not otherwise increase the propositional or *until* structure of ϕ up to constant factors. A similar reasoning applies to the maximum size of the π_i 's, which is independent of $d(\phi)$. Finally, notice that $|\phi|_{\mathbf{p}}$ bounds $|\mathcal{P}(\phi)|$, and $|\mathcal{P}(\eta)|$ is no larger than $2s(\phi) + |\mathcal{P}(\phi)|$. \square

Example 9. The following formula Ω is the formula of Example 5 in separated-next form, with $d_1 = d_2 = 1$, $d_3 = 40$, $d_4 = 41$, $d_5 = 1460$.

$$\Omega \triangleq \underbrace{\left(q \wedge \mathbf{G} \left(\begin{array}{l} (u \Leftrightarrow \neg e \mathbf{U} q) \\ \wedge (v \Leftrightarrow \neg q \wedge \neg q \mathbf{U} q) \\ \wedge (q \Rightarrow x_2 \wedge x_5) \wedge (q \Rightarrow \neg x_1) \\ \wedge (e \Rightarrow \neg q \wedge x_1) \wedge (e \Rightarrow x_3 \vee x_4) \end{array} \right) \right)}_{\kappa_\Omega} \wedge \mathbf{G} \left(\begin{array}{l} (x_1 \Leftrightarrow \mathbf{X}^{d_1} u) \\ \wedge (x_2 \Leftrightarrow \mathbf{X}^{d_2} v) \\ \wedge (x_3 \Leftrightarrow \mathbf{X}^{d_3} q) \\ \wedge (x_4 \Leftrightarrow \mathbf{X}^{d_4} q) \\ \wedge (x_5 \Leftrightarrow \mathbf{X}^{d_5} q) \end{array} \right) \quad (9)$$

Notice that $\kappa_\Omega \in \mathbf{L}(\mathbf{U})$ is the first conjunct, $|\mathcal{P}(\Omega)| = 9$, $M_\Omega = 5$, $d(\Omega) = d_5$; the last one dominates over the other size parameters. The following is a model of Ω .

1	2	3	4	...	1420	1421	1422	1423	...	1460	1461
q	$\neg q$	$\neg q$	$\neg q$...	$\neg q$	$\neg q$	$\neg q$	$\neg q$...	$\neg q$	q
$\neg e$	$\neg e$	$\neg e$	$\neg e$...	$\neg e$	e	$\neg e$	$\neg e$...	$\neg e$	$\neg e$
u	$\neg u$	$\neg u$	$\neg u$...	$\neg u$	$\neg u$	u	u	...	u	u
$\neg v$	v	v	v	...	v	v	v	v	...	v	$\neg v$
$\neg x_1$	$\neg x_1$	$\neg x_1$	$\neg x_1$...	$\neg x_1$	x_1	x_1	x_1	...	x_1	$\neg x_1$
x_2	x_2	x_2	x_2	...	x_2	x_2	x_2	x_2	...	$\neg x_2$	x_2
$\neg x_3$	$\neg x_3$	$\neg x_3$	$\neg x_3$...	$\neg x_3$	x_3	$\neg x_3$	$\neg x_3$...	$\neg x_3$	$\neg x_3$
$\neg x_4$	$\neg x_4$	$\neg x_4$	$\neg x_4$...	x_4	$\neg x_4$	$\neg x_4$	$\neg x_4$...	$\neg x_4$	$\neg x_4$
x_5	$\neg x_5$	$\neg x_5$	$\neg x_5$...	$\neg x_5$	$\neg x_5$	$\neg x_5$	$\neg x_5$...	$\neg x_5$	x_5

In any model of Ω corresponding to a model of (1–5) with variability 5/1460 there are at most 6 nss over [1..1460]: 1, 1419, 1420, 1421, 1459, 1460, corresponding to a variability of 6/1460. \diamond

4 Reducing LTL to qualitative LTL

This section shows how to transform any LTL formula into an equi-satisfiable $\mathbf{L}(\mathbf{U})$ formula of polynomially correlated size. The following theorem summarizes the result.

Theorem 10. *Given an LTL formula ϕ , it is possible to build, in polynomial time, a qualitative LTL formula $\xi \in \mathbf{L}(\mathbf{U})$ such that ϕ and ξ are equi-satisfiable and have polynomially correlated size.*

Let η be ϕ in FNF; the remainder of this section shows the construction of ξ from η and proves its correctness in Lemma 11. Theorem 10 then follows from Lemma 6.

Informal presentation. Informally, the construction to turn an LTL formula into an equi-satisfiable qualitative one works as follows. Introduce a fresh propositional letter s . Constrain s to change truth value with any propositional letter in \mathcal{P} ; in other words, any nss coincides with a nss of s . Then, replace any occurrence of a subformula $\mathbf{X} p$ with a suitable *until* formula that defines the value of p at the next nss of s . In practice, this means that a formula such as $\mathbf{X} p$ forces p to hold in the next state (with a new state of s) only if this is necessary, i.e., if this requires a nss. This changes the quantitative $\mathbf{X} p$ formula into a qualitative formula where the precise metric information is relaxed.

Formal presentation. Formally, for an LTL formula in FNF $\eta \in \mathbf{L}(\mathbf{U}, \mathbf{X}^1)$ over $\mathcal{P} = \mathcal{P}(\eta)$, we build another formula $\xi \in \mathbf{L}(\mathbf{U})$ that is equi-satisfiable to η . To this end, let $s \notin \mathcal{P}$ be a fresh propositional letter. For every propositional formula $\pi \in \mathcal{P}(\mathcal{P})$, and every $\phi \in \mathbf{LTL}$ define:

$$\begin{aligned} \lambda(\pi) &\triangleq \pi \wedge \mathbf{F}\neg\pi \Rightarrow \left(\begin{array}{c} s \mathbf{U} (\neg\pi \wedge \neg s) \\ \vee \\ \neg s \mathbf{U} (\neg\pi \wedge s) \\ \vee \\ \bigvee_{q \in \mathcal{P} \setminus \{\pi\}} (q \wedge \pi) \mathbf{U} (\neg q \wedge \pi) \\ \vee \\ \bigvee_{q \in \mathcal{P} \setminus \{\pi\}} (\neg q \wedge \pi) \mathbf{U} (q \wedge \pi) \end{array} \right) \\ \Upsilon(\mathcal{P}) &\triangleq \left(\bigwedge_{p \in \mathcal{P}} \mathbf{G}p \vee \mathbf{G}\neg p \right) \Rightarrow \mathbf{G}s \\ \mathcal{U}(\phi) &\triangleq s \mathbf{U} \phi \vee \neg s \mathbf{U} \phi \\ \mathcal{R}(\phi) &\triangleq (\phi \wedge s \Rightarrow \neg s \mathbf{R} \phi) \wedge (\phi \wedge \neg s \Rightarrow s \mathbf{R} \phi) \\ \mathcal{X}(\phi) &\triangleq \mathcal{U}(\phi) \wedge \mathcal{R}(\phi) \end{aligned}$$

$\lambda(\pi)$ links any transition of the truth value of π to occur simultaneously with a transition of s . $\Upsilon(\mathcal{P})$ deals with the special case where no proposition ever changes truth value. $\mathcal{X}(\pi)$ is instead essentially a qualitative relaxations of the *next* operator: $w, i \models \mathcal{X}(p)$ holds iff the next nss of s is $j \geq i$ and $w, j+1 \models p$ holds. In particular, if s never changes truth value from position i on, $w, i \models \mathcal{X}(\phi)$ iff $w, i \models \mathbf{F}\phi \wedge (\phi \Rightarrow \mathbf{G}\phi)$, for every ϕ .

Finally, build a qualitative formula ξ from η as:

$$\xi \triangleq \Upsilon(\mathcal{P}) \wedge \bigwedge_{p \in \mathcal{P}} \mathbf{G}(\lambda(p) \wedge \lambda(\neg p)) \wedge \kappa \wedge \bigwedge_{i=1, \dots, N} \mathbf{G} \left(\begin{array}{c} x_i \Rightarrow \mathcal{X}(\pi_i) \\ \wedge \\ \neg x_i \Rightarrow \mathcal{X}(\neg\pi_i) \end{array} \right)$$

It should be clear that $|\xi|$ is in $O(|\eta|^2)$. Then, the following lemma justifies the correctness of the construction given.

Lemma 11. *η and ξ are equi-satisfiable formulas.*

Proof. Remind that $\mathcal{P}(\eta) = \mathcal{P}$ and $\mathcal{P}(\xi) = \mathcal{P} \cup \{s\}$. The proof is in two parts.

SAT(η) \Rightarrow **SAT**(ξ). In the first part show that ξ is satisfiable if η is satisfiable. Hence, assume $w \models \eta$ for some $w \in \mathcal{W}[\mathcal{P}(\eta)]$. Build an $x \in \mathcal{W}[\mathcal{P}(\xi)]$ such that $x \models \xi$ as follows. x coincides with w over $\mathcal{P}(\eta)$, hence $x \models \kappa$ because $s \notin \mathcal{P}(\kappa)$. In addition, s is added to x according to the following recursive definition: $s \in x(0)$ and, for $i > 0$, if $w(i-1) = w(i)$ then $s \in x(i) \Leftrightarrow s \in x(i-1)$, whereas if $w(i-1) \neq w(i)$ then $s \in x(i) \Leftrightarrow s \notin x(i-1)$. In other words, s switches its truth value at nss — except possibly for an infinite tail of constant states.

For any $p \in \mathcal{P}(\eta)$ let us show that $x \models \Upsilon(\mathcal{P}) \wedge \mathbf{G}(\lambda(p) \wedge \lambda(\neg p))$. The proof of $x \models \Upsilon(\mathcal{P})$ is routine. Then, let $i \in \mathbb{N}$ be such that $x, i \models p$. If $x, i \models \mathbf{G}p$ then trivially $x, i \models \lambda(p)$, because the $\pi \wedge \mathbf{F}\neg\pi$ is false at i . Otherwise, let $j > i$ be the least integer such that $x, j \models \neg p$. If no other proposition changes truth value over $x(i:j)$, that is if $x(i:j-1)|_{\mathcal{P}(\eta)}$ is a sequence of stuttering steps, then s switches its truth value precisely

at j . Hence, one of $x, i \models s \cup (\neg p \wedge \neg s)$ and $x, i \models \neg s \cup (\neg p \wedge s)$ holds. Otherwise, there exist $q \neq p$ and $k < j$ such that either $x, i \models q$ and $x, k \models \neg q$ or $x, i \models \neg q$ and $x, k \models q$. In the former case $(q \wedge p) \cup (\neg q \wedge p)$ holds at i , whereas in the latter case $(\neg q \wedge p) \cup (q \wedge p)$ holds at i . Hence, if $x, i \models p$ then $x, i \models \lambda(p) \wedge \lambda(\neg p)$ is established. If $x, i \models \neg p$ instead, a similar reasoning also proves that $x, i \models \lambda(p) \wedge \lambda(\neg p)$. In all, $x \models G(\lambda(p) \wedge \lambda(\neg p))$ holds.

Finally, let us prove the last conjunct of ξ , for a generic $h \in [1..N]$. Let $i \in \mathbb{N}$ such that $x, i \models x_h$: we prove that $x, i \models \mathcal{X}(\pi_h)$.

Since we are assuming η , $x, i+1 \models \pi_h$ holds. If $x, i \models s$ then clearly $x, i \models s \cup \pi_h$.

Now, assume that $x, i \models \pi_h$: we have to show that $\neg s \text{ R } \pi_h$. That is, for a generic $j \geq i$, either $x, j \models \pi_h$ or there exists $i \leq k < j$ such that $x, k \models \neg s$. The goal is trivial for $j = i$, as $x, i \models \pi_h$ by assumption. It is also trivial for $j = i + 1$, as $x, i + 1 \models \pi_h$ also holds. For $j > i + 1$, assume adversarially that $x, j \models \neg \pi_h$. Notice that this implies that $x, j - 1 \models \neg x_h$, hence x_h changes its truth value from true to false at some $i \leq m < j - 1$. Then, $x(m) \neq x(m + 1)$ is not a stuttering step, which implies that s also changes its truth value at m . Since s is true at i , s must be false at some $i < k \leq m + 1 \leq j - 1 < j$. So $x, k \models \neg s$ which closes the current branch of the proof.

Let us now consider the case $x, i \models \neg s$ hence $x, i \models \neg s \cup \pi_h$. Similarly as we did in the previous case, we can establish also that if $x, i \models \pi_h$ then $s \text{ R } \pi_h$. In all, we have shown that $x, i \models \mathcal{X}(\pi_h)$.

For $i \in \mathbb{N}$ such that $x, i \models \neg x_h$, a very similar reasoning shows that $x, i \models \mathcal{X}(\neg \pi_h)$.

i is generic, which entails the last conjuncts of ξ : $x \models G(x_h \Rightarrow \mathcal{X}(\pi_h))$ and $x \models G(\neg x_h \Rightarrow \mathcal{X}(\neg \pi_h))$.

SAT(ξ) \Rightarrow SAT(η). In the second part, show that η is satisfiable if ξ is satisfiable. Hence, assume that $w \models \xi$ for some $w \in \mathcal{W}[\mathcal{P}(\eta) \cup \{s\}]$. Build an $x \in \mathcal{W}[\mathcal{P}(\eta)]$ such that $x \models \eta$ as follows. First, let y be w with all stuttering steps removed. Then, let $i \in \mathbb{N}$ be a generic position and $h \in [1..N]$; since $y \models \xi$ then in particular $y, i \models x_h \Rightarrow \mathcal{X}(\pi_h)$ and $y, i \models \neg x_h \Rightarrow \mathcal{X}(\neg \pi_h)$. Let us show that $y, i \models x_h \Leftrightarrow \mathcal{X}(\pi_h)$.

1. Assume $y, i \models x_h \wedge \mathcal{X}(\pi_h)$. Ad absurdum, let $y, i + 1 \models \neg \pi_h$. We now discuss two cases, whether $y, i \models \pi_h$ or $y, i \models \neg \pi_h$, and we show that in both cases we reach a contradiction, hence $y, i + 1 \models \pi_h$.

- (a) Assume $y, i \models \pi_h$. Also, assume that $y, i \models s$; this is without loss of generality because $\mathcal{X}(\pi_h)$ is symmetric with respect to the truth value of s . Since π_h switches from true to false at i , some proposition $r \neq s$ changes its truth value at i . Hence, $\lambda(r) \wedge \lambda(\neg r)$ forces s to also change its truth value at i . In all we have the following situation:

$$\begin{array}{ccccccc} & & x_h & & & & \\ & \cdots & \pi_h & \neg \pi_h & \cdots & & \\ & & s & \neg s & & & \\ \hline & & i & i + 1 & & & \end{array}$$

But then $\mathcal{R}(\pi_h)$ requires in particular $\neg s \text{ R } \pi_h$ to hold at i ; this is however false because neither $y, i + 1 \models \pi_h$ nor $y, i \models \neg s$. Hence, the contradiction.

- (b) Assume $y, i \models \neg \pi_h$. Also, assume that $y, i \models s$; this is without loss of generality because $\mathcal{X}(\pi_h)$ is symmetric with respect to the truth value of

s . Note that $\mathcal{U}(\pi_h)$ implies that π_h must eventually hold; let $j > i + 1$ be the least instant such that $y, j \models \pi_h$. So, π_h does not hold at all positions in $[i..j - 1]$ and becomes true at j . From the assumption that y has no stuttering steps it must be $y(i + 1) \neq y(i)$. Hence there exists some atomic proposition r that changes its truth value at i . Correspondingly, $\wedge(r) \wedge \wedge(\neg r)$ forces s to also change its truth value at i . In all we have the following situation:

$$\begin{array}{ccccccc}
 & & x_h & & & & \\
 \cdots & \neg\pi_h & \neg\pi_h & \cdots & \pi_h & \cdots & \\
 & r & \neg r & & & & \\
 & s & \neg s & & & & \\
 \hline
 & i & i + 1 & & j & &
 \end{array}$$

But then $\mathcal{U}(\pi_h)$ cannot hold at i , because neither $s \cup \pi_h$ nor $\neg s \cup \pi_h$ holds at i . Hence, the contradiction.

2. The proof of the other case $y, i \models \neg x_h \wedge \mathcal{X}(\neg\pi_h)$ can be obtained by symmetry from the previous case.

Since i and h are generic, we have established $y \models \bigwedge_{i=1, \dots, N} \mathbf{G}(x_i \Leftrightarrow \mathbf{X}\pi_i)$. In addition $w \models \kappa$ implies $y \models \kappa$ as well, because y is obtained from w only by removing stuttering steps and $\kappa \in \mathbf{L}(\mathbf{U})$ is closed under stuttering. Hence $x = y|_{\mathcal{P}(\eta)}$ is a model that satisfies η . \square

Example 12. Let $\xi(7)$ be formula (7) modified according to the construction of the current section. The proof of Lemma 11 shows that the qualitative formula $\xi(7)$ preserves the stutter-free models of the equi-satisfiable LTL formula (7). On the other hand, consider a model of (7) with a sequence of $d_5 - 2$ stuttering steps $\neg q \cdots \neg q q$, such as the one in Example 7; it corresponds to the following stutter-free model of $\xi(7)$:

$$\begin{array}{ccccccc}
 \neg q & \neg q & \cdots & \neg q & \neg q & \neg q & q \\
 \neg x_1 & \neg x_1 & \cdots & \neg x_1 & \neg x_1 & x_1 & \cdots \\
 \neg x_2 & \neg x_2 & \cdots & \neg x_2 & x_2 & \cdots & \cdots \\
 \neg x_3 & \neg x_3 & \cdots & x_3 & \cdots & \cdots & \cdots \\
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 \neg x_{d_5-3} & x_{d_5-3} & \cdots & \cdots & \cdots & \cdots & \cdots
 \end{array}$$

This shows that the transformation of (7) into $\xi(7)$ — and more generally of η into ξ — does not represent the redundancy of words with bounded variability more succinctly, but merely encodes it in a different form. \diamond

5 LTL with bounded variability

This section specializes the results of Section 4 by showing how to more succinctly encode the redundancy of stuttering steps in words with bounded variability. The following results require $\mathbf{LTL}(\mathbf{U}, \text{exqPn})$: an extension of $\mathbf{L}(\mathbf{U})$ with a qualitative variant of the Pnueli operators.

Section 5.1 recalls “standard” Pnueli operators, introduces $\mathbf{LTL}(\mathbf{U}, \text{exqPn})$, and shows that $\mathbf{LTL}(\mathbf{U}, \text{exqPn})$ has the same complexity as LTL. Then, Section 5.2 shows

how to transform any LTL formula ϕ and a positive integer parameter \mathbb{V} into an LTL(U, exqPn) formula ϕ' which is satisfiable (over unconstrained words) iff ϕ is satisfiable over words with variability bounded by $\mathbb{V}/d(\phi)$ — recall that $d(\phi)$ is the largest distance used in ϕ . The size of ϕ' is polynomial in \mathbb{V} , the *number* $s(\phi)$ of distance sub-formulas, and the size of qualitative sub-formulas appearing ϕ ; however, the size of ϕ' does not depend on $d(\phi)$ — the *values* of distances in ϕ . As a consequence, checking the satisfiability of ϕ' — which can be done with standard LTL algorithms — is significantly less complex than checking the original ϕ whenever the distances used in ϕ are very large and dominate over the other size parameters.

5.1 Pnueli operators

Pnueli operators have been introduced for dense-time models [20]; this section considers their discrete-time counterparts and variations thereof.

5.1.1 Pnueli operators

For $k, n \in \mathbb{N}$, the *Pnueli operator* Pn_k^n is a k -ary temporal operator with the following semantics:

$$w, i \models \text{Pn}_k^n(\phi_1, \dots, \phi_k)$$

holds iff there exist k positions $i + 1 \leq k_1 < k_2 < \dots < k_k \leq i + n$ such that $w, k_j \models \phi_j$ for all $1 \leq j \leq k$. It is not difficult to show that extending LTL with Pnueli operators does not affect its expressiveness or complexity (this is not the case over dense time [20]), under the assumption of a unary encoding of the integer constants.

Example 13. Consider the following word w (nss are in bold and underlined).

<u>1</u>	2	3	4	5	<u>6</u>	<u>7</u>	8	9	<u>10</u>	11	<u>12</u>	<u>13</u>	14
<i>s</i>	$\neg s$	$\neg s$	$\neg s$	$\neg s$	$\neg s$	<i>s</i>	$\neg s$	$\neg s$	$\neg s$	<i>s</i>	<i>s</i>	$\neg s$	<i>s</i>
$\neg v$	<i>v</i>	<i>v</i>	<i>v</i>	<i>v</i>	<i>v</i>	$\neg v$	<i>v</i>	<i>v</i>	<i>v</i>	<i>v</i>	<i>v</i>	$\neg v$	<i>v</i>
<i>q</i>	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	<i>q</i>
$\neg e$	$\neg e$	$\neg e$	$\neg e$	$\neg e$	$\neg e$	$\neg e$	<i>e</i>	<i>e</i>	<i>e</i>	$\neg e$	$\neg e$	$\neg e$	$\neg e$

Then, $w, 1 \models \text{Pn}_4^{13}(v, s, e, s)$ and $w, 8 \models \text{Pn}_2^3(\neg s, \mathbf{X} s)$ hold. On the contrary, $w, 8 \not\models \text{Pn}_3^3(\neg s, \neg s, \neg s)$ because in particular $w, 8 + 3 \not\models \neg s$. \diamond

The results of the present paper are based on a qualitative version of the Pnueli operators: the *qualitative extended Pnueli operators* $\text{exqPn}_k^{n; \langle n_1, \dots, n_k \rangle}$ for $k, n \in \mathbb{N}$ and $n_1, \dots, n_k \in \mathbb{N} \cup \{*\}$. Their semantics is defined as follows:

$$w, i \models \text{exqPn}_k^{n; \langle n_1, \dots, n_k \rangle}(\phi_1, \dots, \phi_k)$$

holds iff there exist k positions $i \leq k_1 < \dots < k_k$ such that all the following hold, for all $1 \leq j \leq k$:

1. k_j is a nss;
2. $w, k_j + 1 \models \phi_j$;
3. for $j > 1$, if $n_j \neq *$ then there are no more than n_j nss between k_{j-1} and $k_j - 1$ (both included);

4. if $n_1 \neq *$ then there are no more than n_1 nss between i and k_1 (both included);
5. there are no more than n nss between i and k_k (both included).

Intuitively, the qualitative extended Pnueli operators are qualitative counterparts to the standard Pnueli operators, which are further generalized by imposing an additional requirement on the relative distance of k nss. For example, if $n_1 = 1$, ϕ_1 must hold right after the first nss that follows or is at i , independently of the other following $k - 1$ nss.

Example 14. Consider again word w from Example 13, where nss are in bold and underlined. For the positions 1, 6, 7, 13, $w, 1 \models \text{exqPn}_4^{6;\langle 3,2,*,3 \rangle}(v, \neg q, e, q)$ holds. On the contrary, $w, 1 \not\models \text{exqPn}_4^{6;\langle 3,2,*,1 \rangle}(v, \neg q, e, q)$; in fact, let k_1, \dots, k_4 be the positions that match the semantics of the operator. Then, $k_4 = 13$ as q only holds at 14, so that the last component of the constraint $\langle 3, 2, *, 1 \rangle$ forces k_3 to be 12, the nss immediately before 13; but $w, 12 + 1 \not\models e$. \diamond

5.1.2 LTL with Pnueli operators

LTL(U, exqPn) is the extension of qualitative LTL with qualitative extended Pnueli operators. Any LTL(U, exqPn) formula has an equi-satisfiable L(U) formula of polynomially correlated size that can be built in polynomial time, when the integer constants used in the Pnueli operators use a unary encoding. Precisely, to encode a formula $\text{exqPn}_k^{n;\langle n_1, \dots, n_k \rangle}(\phi_1, \dots, \phi_k)$ introduce n^2 letters $\{q_i^j \mid 1 \leq i, j \leq n\}$. Every q_i^j holds iff $\text{exqPn}_i^{j;\langle n_1, \dots, n_i \rangle}(\phi_1, \dots, \phi_i)$ does; formally, a formula of size $O(n \cdot \max_i |\phi_i|)$ defines q_i^j as follows, where \mathcal{X}^j abbreviates $\underbrace{\mathcal{X}\mathcal{X}\dots\mathcal{X}}_j$.

$$q_i^j \Leftrightarrow \begin{cases} \perp & i > j \vee n_i = 0 \\ \mathcal{X}(\phi_1) & 1 = i = j \wedge n_i \neq 0 \\ \underbrace{\mathcal{X}(\phi_1 \vee \mathcal{X}(\phi_1 \vee \dots))}_{j \text{ nested } \mathcal{X}} & 1 = i < j \leq n \wedge n_i = * \\ \underbrace{\mathcal{X}(\phi_1 \vee \mathcal{X}(\phi_1 \vee \dots))}_{\min(j, n_1) \text{ nested } \mathcal{X}} & 1 = i < j \leq n \\ (q_{i-1}^{j-1} \wedge \mathcal{X}^j(\phi_i)) \vee q_i^{j-1} & \wedge 0 < n_i \neq * \\ \left(\underbrace{\left(q_{i-1}^{j-1} \wedge \mathcal{X}^j(\phi_i) \wedge \mathcal{X}^{j-n_i}(\phi_{i-1} \vee \mathcal{X}(\phi_{i-1} \vee \dots)) \right)}_{n_i \text{ nested } \mathcal{X}} \right) \vee q_i^{j-1} & \begin{array}{l} 1 < i \leq j \leq n \wedge \\ (n_i = * \vee n_i \geq j) \\ 1 < i \leq j \leq n \\ \wedge 0 < n_i < j \end{array} \end{cases} \quad (10)$$

Informally, the recursive definition (10) works as follows (from top to bottom):

- If $n_i = 0$ or $i > j$, q_i^j is unsatisfiable.
- If there is a positive numeric constraint $n_1 \neq *$ on when ϕ_1 must occur, then q_1^1 is equivalent to $\mathcal{X}(\phi_1)$: ϕ_1 holds exactly after the next nss.
- If $n_1 = *$ and any $j > 1$, q_1^j requires that ϕ_1 holds for some of the following j nss.

- If $n_1 \neq *$ but $n_1 \neq 0$, and $j > 1$, q_1^j also requires that ϕ_1 holds correspondingly to some of the following n_1 nss; if, however, $j < n_1$, then the requirement on the following next j nss prevails over the constraint n_1 .
- If $i > 1$ and $n_i = *$ or $n_i \geq j$, the constraint n_i is subsumed by the constraint j ; then, q_i^j holds iff: (a) the more constraining q_i^{j-1} holds (that is, ϕ_1, \dots, ϕ_i hold over the following $j-1$ nss), or (b) ϕ_i holds exactly after the j -th nss, and q_{i-1}^{j-1} holds as well (that is, $\phi_1, \dots, \phi_{i-1}$ hold over the following $j-1$ nss).
- If $i > 1$ and $n_i < j$ then q_i^j is reducible to two cases; either the more constraining q_i^{j-1} holds, or all of the following hold: (a) q_{i-1}^{j-1} holds (which takes care of the first $i-1$ arguments $\phi_1, \dots, \phi_{i-1}$), (b) ϕ_i holds exactly after the j -th nss, and (c) there are no more than n_i nss between an occurrence of ϕ_{i-1} and the occurrence of ϕ_i .

Example 15. Continuing Example 14, $w, 1 \models \text{exqPn}_4^{6;(3,2,*,3)}(v, \neg q, e, q)$ is rewritten as $w, 1 \models q_4^6$. This in turns reduces to checking the following sequence of formulas at position 1: $q_3^5 \wedge \mathcal{X}^6(q) \wedge \mathcal{X}^3(\mathcal{X}^3(q)); q_3^4; q_3^3; q_2^2 \wedge \mathcal{X}^3(e); q_1^1 \wedge \mathcal{X}^2(\neg q); \mathcal{X}(v)$. \diamond

We can transform LTL(U, exqPn) formulas into equi-satisfiable qualitative formulas using definition (10). The construction is general, but the remainder will use an LTL(U, exqPn) formula Λ over $\mathcal{Q} = \mathcal{P} \cup \{s\}$ in the form:

$$\Lambda \triangleq \Upsilon(\mathcal{P}) \wedge \bigwedge_{p \in \mathcal{P}} \mathbf{G} \left(\begin{array}{c} \lambda(p) \wedge \\ \lambda(\neg p) \end{array} \right) \wedge \kappa \wedge \bigwedge_{i=1, \dots, M} \mathbf{G} \left(\xi_i \Rightarrow \text{exqPn}_{\mathbb{I}(i)}^{\mathbb{J}(i); \langle \mathbb{K}(i) \rangle} (\psi_1^i, \dots, \psi_{\mathbb{I}(i)}^i) \right) \quad (11)$$

where \mathbb{J}, \mathbb{I} are two mappings $[1..M] \rightarrow \mathbb{N}_{>0}$; \mathbb{K} is a mapping $[1..M] \rightarrow (\mathbb{N} \cup \{*\})^{\mathbb{I}(i)}$; $\kappa, \psi_i \in \mathbf{L}(\mathbf{U})$ for all $1 \leq i \leq M$; $\xi_i \in \mathbf{P}(\mathcal{P})$; and s does not occur in κ or in any ψ_i^j .

Lemma 16. $\llbracket \Lambda \rrbracket$ is closed under stuttering.

Proof. A consequence of the semantics of the extended qualitative Pnueli operators — which introduce no metric constraint — and of the particular form of Λ — where only qualitative subformulas appear. \square

Lemma 17. It is possible to build, in polynomial time, a formula Λ' such that:

1. $\Lambda' \in \mathbf{L}(\mathbf{U})$;
2. $|\Lambda'|$ is polynomially bounded by $|\Lambda|$;
3. Λ' and Λ are equi-satisfiable.

Proof. Let $\tau \langle q, \underline{i}, \underline{j}, \underline{n}, \underline{n}_1, \dots, \underline{\phi}_1, \dots \rangle$ denote (10) with $q, \underline{i}, \underline{j}, \underline{n}, \underline{n}_1, \dots, \underline{\phi}_1, \dots$ respectively replacing $q, i, j, n, n_1, \dots, \phi_1, \dots$. Introduce fresh propositional letters $q[i]_k^j$ for $1 \leq i \leq M, 1 \leq j, k \leq \mathbb{J}(i)$ and construct Λ' from Λ as:

$$\Lambda' \triangleq \left(\begin{array}{c} \Upsilon(\mathcal{P}) \wedge \bigwedge_{p \in \mathcal{P}} \mathbf{G}(\lambda(p) \wedge \lambda(\neg p)) \\ \wedge \quad \kappa \quad \wedge \\ \bigwedge_{i=1, \dots, M} \mathbf{G}(\xi_i \Rightarrow q[i]_{\mathbb{I}(i)}^{\mathbb{J}(i)}) \end{array} \right) \wedge \bigwedge_{\substack{1 \leq k \leq M \\ 1 \leq i, j \leq \mathbb{J}(k)}} \mathbf{G} \tau \left\langle q[k], i, j, \mathbb{J}(k), \mathbb{K}(k), \psi_1^k, \dots, \psi_{\mathbb{I}(k)}^k \right\rangle \quad (12)$$

Then, facts 1–3 are straightforward to prove:

1: is clear by construction.

2: The size $|\Lambda|$ of Λ is polynomial in $|\mathcal{Q}|, I, J, K, M$, where J, I, K denote $\max_i \mathbb{J}(i)$, $\max_i \mathbb{I}(i)$, and $\max_i \max \mathbb{K}(i)$, respectively. The first subformula of Λ' :

$$\Upsilon(\mathcal{P}) \wedge \bigwedge_{p \in \mathcal{P}} \mathbf{G}(\lambda(p) \wedge \lambda(\neg p)) \wedge \kappa \wedge \bigwedge_{i=1, \dots, M} \mathbf{G}(\xi_i \Rightarrow q[i]_{\mathbb{I}(i)}^{\mathbb{J}(i)}) \quad (13)$$

clearly has size bounded by $|\Lambda|$. The second subformula of Λ' :

$$\bigwedge_{\substack{1 \leq k \leq M \\ 1 \leq i, j \leq \mathbb{J}(k)}} \mathbf{G}\tau \left\langle q[k], i, j, \mathbb{J}(k), \mathbb{K}(k), \psi_1^k, \dots, \psi_{\mathbb{I}(k)}^k \right\rangle \quad (14)$$

has size bounded by $M \cdot J^2 \cdot O(J \max_{i,j} |\psi_i^j|)$, hence also polynomial in $|\Lambda|$.

3: follows from the definition in (10) and the semantics of the \mathcal{X} operator, along the lines of Lemma 11. \square

5.2 Relaxing distance formulas

This section proves the following.

Theorem 18. *Given an LTL formula ϕ and an integer parameter $\mathbb{V} > 0$, it is possible to build, in polynomial time, a qualitative LTL formula ϕ'' such that:*

- $|\phi''|$ is polynomial in $\mathbb{V}, |\phi|_{\mathcal{P}}, |\phi|_{\mathcal{U}}, s(\phi)$ but is independent of $d(\phi)$;
- ϕ is satisfiable over words in $\text{var}(\mathcal{Q}, \mathbb{V}/d(\phi))$ iff ϕ'' is satisfiable over unconstrained words.

Let η be ϕ in SNF; the following construction builds a $\phi' \in \text{LTL}(\mathcal{U}, \text{exqPn})$ from η such that Lemmas 20 and 21 hold. Theorem 18 follows after transforming ϕ' into ϕ'' by eliminating the qualitative extended Pnueli operators according to Lemma 17.

Informal presentation. Let us first informally sketch the ideas behind the transformation from η to ϕ' , with the aid of a few examples referring to word w in Example 13 and formula Ω in Example 9.

The basic idea consists of relaxing every distance formula $\mathbf{X}^d a$ into a qualitative formula $\mathcal{X}^{d'}(a)$ with $d' \leq \mathbb{V}$, so that consecutive nss take the role of consecutive positions. The elimination or addition of stuttering steps reconciles words in the quantitative and qualitative transformed formulas. For example, $w, 1 \models \mathcal{X}^6(q)$ holds because q holds at position 14; adding $41 - 14 = 27$ repetitions of position 2 transforms w into a word w' where the quantitative requirement $w', 1 \models \mathbf{X}^{41}q$ holds as well.

The transformation must also preserve the ordering among events: if $\mathbf{X}^d a$ and $\mathbf{X}^e b$ both hold for some $d < e$, then $\mathcal{X}^{d'}(a)$ and $\mathcal{X}^{e'}(b)$ should hold for suitable $d' < e'$. Another constraint requires that $e' - d' \leq e - d$; otherwise, the transformed formula admits words with $e' - d' > e - d$ non-stuttering steps between consecutive occurrences of a and b , which may not be removable to put a and b at an absolute distance of $e - d$. For example, $\mathcal{X}(s) \wedge \mathcal{X}^2(q)$ is a suitable relaxation of $\mathbf{X}^{30}s \wedge \mathbf{X}^{31}q$, whereas $\mathcal{X}(s) \wedge \mathcal{X}^3(q)$ is not: $w, 10 \models \mathcal{X}(s) \wedge \mathcal{X}^3(q)$ but the nss 13 makes it impossible to pad w with stuttering steps such that s and q hold at positions $10 + 30$ and $10 + 31$.

Using these ideas, a formula η in SNF (8) is transformed by replacing the distance formulas with qualitative “snapshots” using the qualitative extended Pnueli operators: the predicates $\pi_1, \pi_2, \dots, \pi_M$ hold orderly over some of the the following \mathbb{V} nss, with the additional constraint that, between any two consecutive π_i, π_{i+1} , no more nss than the difference of the corresponding distances occur. For example, if $\neg x_1 \wedge x_2 \wedge \neg x_3 \wedge \neg x_4 \wedge x_5$ (corresponding to predicates $\neg u, v, \neg q, \neg q, q$) then formula $\text{exqPn}_4^{6; \langle 1, *, 1, * \rangle}(\neg u \wedge v, \neg q, \neg q, q)$ must hold, where $\neg u \wedge v$ occurs after the next nss and $\neg q, \neg q$ occupy consecutive nss.

This approach can be made rigorous, but introduces an exponential blow-up because it considers each of the 2^M subset of propositions x_1, \dots, x_M . The following construction avoids this blow-up by introducing auxiliary propositions y_i 's and z_i^j 's that mark nss and decouple them from the propositions that must hold therein.

Each y_i holds precisely from the i -th nss until the next $1 + (i \bmod \mathbb{V})$ nss. Then, for each given h , the propositions $z_1^h, z_2^h, \dots, z_m^h$ (where m is the number of different distances used in η) hold sequentially and cyclically from when y^h holds. Each z_k^h marks a position in the sequence that satisfies the qualitative extended Pnueli operator under consideration; correspondingly, for each index k' corresponding to a distance with index k , $\pi_{k'}$ holds with z_k^h iff $x_{k'}$ holds with y^h . For example, if $\neg x_1 \wedge x_2 \wedge \neg x_3 \wedge \neg x_4 \wedge x_5$ holds when some y_k holds, then the corresponding predicates $\neg u \wedge v, \neg q, \neg q, q$ hold orderly with the next occurrences of $z_1^k, z_2^k, z_3^k, z_4^k$.

The following construction formalizes these ideas.

Detailed construction. Consider a generic LTL formula η in SNF:

$$\eta \triangleq \kappa \wedge \bigwedge_{i=1, \dots, M} G(x_i \Leftrightarrow X^{\mathbb{D}(i)} \pi_i) \quad (15)$$

where $\kappa \in L(U)$, $x_i \in \mathcal{M} = \{x_i \mid 1 \leq i \leq M\}$, $\mathcal{Q} = \mathcal{P} \cup \mathcal{M}$, $\pi_i \in P(\mathcal{Q})$, and \mathbb{D} is a monotonically non-decreasing mapping $[1..M] \rightarrow \mathbb{N}_{>0}$.

Introduce \mathbb{V} letters $\{y_i \mid 1 \leq i \leq \mathbb{V}\}$. Formula $\varkappa(y)$ constrains y_i to occur synchronously with every i -th nss:

$$\varkappa(y) \triangleq y_1 \wedge \bigwedge_{1 \leq i \leq \mathbb{V}} G \left(y_i \Rightarrow \left(\begin{array}{l} \mathcal{X}(y_{1+(i \bmod \mathbb{V})}) \\ \bigwedge_{j \neq i} \neg y_j \wedge \\ y_i \cup y_{1+(i \bmod \mathbb{V})} \end{array} \right) \right) \quad (16)$$

Let D_1, D_2, \dots, D_m be the sequence of sets that partition $[1..M]$ in such a way that indices involving the same number of consecutive nested X 's are in the same set, and the sets appear in the sequence in increasing order of nested X 's; formally: $i, j \in D_k$ with $k \triangleq D(i) = D(j)$ for some k iff $\mathbb{D}(i) = \mathbb{D}(j) \triangleq d_k$ (and d_0 is defined as 0); and $i \in D_{k_1}$ and $j \in D_{k_2}$ with $k_1 < k_2$ implies $\mathbb{D}(i) < \mathbb{D}(j)$.

Then, introduce another $m \cdot \mathbb{V}$ letters $\{z_i^j \mid 1 \leq i \leq m, 1 \leq j \leq \mathbb{V}\}$. At every i -th nss, marked by y_i , the sequence z_1^j, \dots, z_m^j must hold over m of the following \mathbb{V} nss; moreover, between each z_j^i and its preceding z_{j-1}^i there must be no more than $d_i - d_{i-1}$ nss, unless $d_i - d_{i-1} > \mathbb{V} - i + 1$. After defining, for $1 \leq i \leq m$:

$$\delta_i \triangleq \begin{cases} d_i - d_{i-1} & \text{if } d_i - d_{i-1} \leq \mathbb{V} - i + 1 \\ * & \text{otherwise} \end{cases}$$

the qualitative extended Pnueli operators capture this behavior of the z_i^j 's.

$$\varkappa\langle z, \text{Pn} \rangle \triangleq \bigwedge_{1 \leq i \leq \mathbb{V}} \text{G} \left(y_i \Rightarrow \text{exqPn}_m^{\mathbb{V}; \langle \delta_1, \dots, \delta_m \rangle} (z_1^i, \dots, z_m^i) \right) \quad (17)$$

Additionally, constrain the z_i^j 's to hold sequentially, according to the following.

$$\varkappa\langle z, \text{U} \rangle \triangleq \left(\bigwedge_{h, j \neq 1} \neg z_h^j \right) \text{U} z_1^1 \wedge \bigwedge_{\substack{1 \leq j \leq \mathbb{V} \\ 1 \leq i \leq m}} \text{G} \left(z_i^j \Rightarrow \left(\neg z_{1+(i \bmod m)}^{1+(j \bmod \mathbb{V})} \right) \text{U} z_{1+(i \bmod m)}^j \right) \quad (18)$$

Once the z_i^j 's and the y_i 's are constrained, link the x_i 's to the values of the π_i 's in the distance formulas. If some x_i holds, after or at the j -th nss and before the $j+1$ -th, then π_i has to hold at the k -th position in the sequence z_1^j, \dots, z_m^j , with $k = D(i)$.

$$\varkappa\langle x, \pi \rangle \triangleq \bigwedge_{\substack{1 \leq i \leq M \\ 1 \leq j \leq \mathbb{V}}} \text{G} \left(\begin{array}{c} x_i \wedge y_j \Rightarrow \neg z_{D(i)}^j \text{U} z_{D(i)}^j \wedge \pi_i \\ \wedge \\ \neg x_i \wedge y_j \Rightarrow \neg z_{D(i)}^j \text{U} z_{D(i)}^j \wedge \neg \pi_i \end{array} \right) \quad (19)$$

Finally, combine the various \varkappa formulas to transform η into ϕ' :

$$\phi' \triangleq \kappa \wedge \Upsilon(\mathcal{Q}) \wedge \bigwedge_{p \in \mathcal{Q}} \text{G} \left(\begin{array}{c} \lambda(p) \wedge \\ \lambda(\neg p) \end{array} \right) \wedge \varkappa\langle y \rangle \wedge \varkappa\langle z, \text{Pn} \rangle \wedge \varkappa\langle z, \text{U} \rangle \wedge \varkappa\langle x, \pi \rangle \quad (20)$$

Example 19. In the elections example, $\mathbb{V} = 6$, $m = 4$ instantiate $\varkappa\langle y \rangle$, $\varkappa\langle z, \text{U} \rangle$, and $\varkappa\langle x, \pi \rangle$. Then, $\delta_1 = \delta_3 = 1$ and $\delta_2 = \delta_4 = *$ instantiate $\varkappa\langle z, \text{Pn} \rangle$. \diamond

The correctness of the above construction and the proof of Theorem 18 rely on the following two lemmas.

Lemma 20. η is satisfiable over words in $\text{var}(\mathcal{Q}, \mathbb{V}/d(\phi))$ iff ϕ' is satisfiable over unconstrained words.

Proof. Let D be $d(\eta)$, which equals $d(\phi)$ by Lemma 8. The proof consists of two parts.

SAT(η) \Rightarrow **SAT**(ϕ'). Let $w \in \text{var}(\mathcal{Q}, \mathbb{V}/D)$ such that $w \models \eta$. w' adds propositions s, y_i, z_i^j , constrained as follows. s switches its truth value at every nss, except for possibly an infinite tail of constant values over w . Exactly one of the y_i 's holds at every instant, and they rotate at every nss signaled by s . Whenever a given y_j holds, a sequence of z_i^j 's hold over the following \mathbb{V} nss, in a sequential fashion. Namely, let k be the first step where a certain y_j holds, let h_i be the last non-stuttering before position $k + d_i$, and let l_i be the δ_i -th nss after h_{i-1} (included, with $h_0 = k$); then, z_i^j starts to hold at $\min(h_i, l_i, \mathbb{V} - k + 1) + 1$, and holds until the next z_{i+1}^j .

Once w' is built, the rest of the proof follows the lines of Lemma 11. It is clear that $w' \models \bigwedge_{p \in \mathcal{Q}} \Upsilon(P) \wedge \text{G}(\lambda(p) \wedge \lambda(\neg p))$ and $w' \models \kappa$. In addition, $w' \models \varkappa\langle y \rangle \wedge \varkappa\langle z, \text{U} \rangle$ is a consequence of the set up of the y_j 's and the z_i^j 's. Then, let i be the current generic instant and $b \subseteq [1..M]$ be a generic subset such that $\bigwedge_{i \in b} x_i \wedge \bigwedge_{i \notin b} \neg x_i$ holds at i . Hence, $w, i \models \text{X}^{\mathbb{D}(j)} \pi_j$ holds for all $j \in b$ and $w, i \models \text{X}^{\mathbb{D}(k)} \neg \pi_k$ holds for all $k \notin b$. The variability of w — and that of w' — is bounded by \mathbb{V}/D ; hence, there are at most \mathbb{V}

nss of item s over positions i to $i+D$. Let $i \leq t_1 < \dots < t_{\mathbb{V}} \leq i+D$ be these transition instants. There are only stuttering steps between any such two consecutive t_i 's, hence there exists a subset $u_1 < \dots < u_m$ of the t_i 's such that z_i^j holds at u_i for all i 's and some unique j . Now, for all g such that $D(g) = i$, π_g holds at $k + d_i$ and (at least) since the previous and until the next nss. Because of how each z_i^j 's mark the stuttering positions before $k+d_i$, for every g such that $D(g) = i$, π_i must in particular hold where z_i^j first holds; because i is generic, $w' \models \varkappa(x, \pi)$ holds. Also, if $d_i - d_{i-1} \leq \mathbb{V} - i + 1$, there are no more than $d_i - d_{i-1}$ nss between u_{i-1} and u_i , for all $1 \leq i \leq m$ (and assuming $u_0 = d_0 = 0$); this establishes $w', i \models \text{exqPn}_m^{\mathbb{V}; \langle \delta_1, \dots, \delta_m \rangle} (z_1^j, \dots, z_m^j)$. In all, $w' \models \phi'$ holds.

SAT(ϕ') \Rightarrow **SAT**(η). Let w' be an unconstrained word in

$$\mathcal{W}[\mathcal{Q} \cup \{s\} \cup \{y_i \mid 1 \leq i \leq \mathbb{V}\} \cup \{z_i^j \mid 1 \leq i \leq m, 1 \leq j \leq \mathbb{V}\}]$$

such that $w' \models \phi'$. Initially, let w be w' with all stuttering steps removed; $w \models \phi'$ as well from Lemma 16. Modify w as follows, until $w \models \eta$ is the case.

Let i be the current generic instant and $b \subseteq [1..M]$ be a generic subset such that $\bigwedge_{i \in b} x_i \wedge \bigwedge_{i \notin b} \neg x_i$ holds at i on w . The rest of the proof works inductively on $1 \leq h \leq M$; let us focus on the more interesting inductive step.

Let $i \leq t_1 < \dots < t_{\mathbb{V}}$ be the following \mathbb{V} nss of s — and hence of any proposition in \mathcal{Q} as well, according to $\Upsilon(\mathcal{Q}) \wedge \bigwedge_{p \in \mathcal{Q}} \mathsf{G}(\lambda(p) \wedge \lambda(\neg p))$. $\varkappa\langle y \rangle$ implies that a unique y_j holds at i ; correspondingly, $\varkappa\langle z, \text{Pn} \rangle$ entails that there exists a subset of the $u_1 < \dots < u_m$ of the sequence $t_1 < \dots < t_{\mathbb{V}}$ such that z_k^j holds at $u_k + 1$ for all $1 \leq k \leq m$. Assume x_h holds at i (the case of $\neg x_h$ is clearly symmetrical and is omitted), with $g = D(h)$; then, $\varkappa\langle x, \pi \rangle$ requires that π_h holds with z_g^j at $u_g + 1$. The inductive hypothesis implies that $u_{g-1} + 1 \leq i + d_{g-1} \leq u_g$, and $\varkappa\langle z, \text{Pn} \rangle$ and the definition of δ_g guarantee that $u_g < i + d_g$. Correspondingly, add $\theta \triangleq i + d_g - u_g - 1$ stuttering steps at position u_g in w . This “shifts” the previous position $u_g + 1$ to the new position $i + d_g$; hence $w, i + d_g \models \pi_h$ and $i + d_g \leq d_{g+1}$ because we added only stuttering steps. Also, $w \models \phi'$ is still the case, because Lemma 16 guarantees that the removal or addition of stuttering steps to w do not affect the satisfiability of ϕ' . Finally, observe that we introduced no more than m nss over every subword of w of length \mathbb{V} , and $m \leq M \leq D$ because of the pigeonhole principle, hence the variability of propositions \mathcal{Q} in w is bounded by \mathbb{V}/D .

In all, induction proves that the finally modified w is such that $w \models \eta$ and $w \in \text{var}(\mathcal{Q}, \mathbb{V}/D)$. \square

Lemma 21. $|\phi'|$ is polynomial in $\mathbb{V}, |\phi|_{\text{p}}, |\phi|_{\text{u}}, s(\phi)$.

Proof. The size of ϕ' is $|\kappa| + |\mathcal{Q}|^2 + |\varkappa\langle y \rangle| + |\varkappa\langle z, \text{Pn} \rangle| + |\varkappa\langle z, \text{U} \rangle| + |\varkappa\langle x, \pi \rangle|$, up to constant multiplicative factors. Then, κ is unchanged from η ; $|\varkappa\langle y \rangle|$ is $\mathcal{O}(\mathbb{V}^3)$; $|\varkappa\langle z, \text{Pn} \rangle|$ is $\mathcal{O}(m^2 \cdot \mathbb{V}^2)$, which is $\mathcal{O}(M^2 \cdot \mathbb{V}^2)$; $|\varkappa\langle z, \text{U} \rangle|$ is $\mathcal{O}(m^3 \cdot \mathbb{V}^2)$ which is $\mathcal{O}(M^3 \cdot \mathbb{V}^2)$; and $|\varkappa\langle x, \pi \rangle|$ is $\mathcal{O}(M \cdot \mathbb{V} \cdot (M + \mathbb{V} + \max_i \pi_i))$. The statement follows by Lemma 8. \square

Example 22. Consider the running elections example and transform Ω (Example 9) into Ω' according to the above construction. The following is a partial model for Ω' , where all propositions not appearing at some position are assumed to be false there, nss

are in bold and underlined, while a hat marks successors of nss.

<u>1</u>	<u>2</u>	3	4	5	<u>6</u>	<u>7</u>	<u>8</u>	9	<u>10</u>	<u>11</u>	<u>12</u>	<u>13</u>	<u>14</u>
s	$\neg s$	$\neg s$	$\neg s$	$\neg s$	$\neg s$	s	$\neg s$	$\neg s$	$\neg s$	s	s	$\neg s$	s
y_1	y_2	y_2	y_2	y_2	y_2	y_3	y_4	y_4	y_4	y_5	y_5	y_6	y_1
	z_1^1	z_1^1	z_1^1	z_1^1	z_2^1	z_3^1	z_3^1	z_3^1	z_3^1	z_3^1	z_3^1	z_3^1	z_4^1
						z_1^2	z_2^2	z_2^2	z_2^2	z_2^2	z_2^2	z_2^2	z_3^2
							z_1^3	z_1^3	z_1^3	z_1^3	z_1^3	z_2^3	z_2^3
										z_1^4	z_1^4	z_1^4	z_2^4
												z_1^5	z_1^5
													z_1^6
$x_2 \wedge x_5$						x_4	x_3	x_3	$x_1 \wedge x_3$	x_1	x_1	x_1	x_1
$\neg v$	v	v	v	v	v	v	v	v	v	v	v	v	$\neg v$
q	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	$\neg q$	q
$\neg e$	$\neg e$	$\neg e$	$\neg e$	$\neg e$	$\neg e$	$\neg e$	e	e	e	$\neg e$	$\neg e$	$\neg e$	$\neg e$
u	$\neg u$	$\neg u$	$\neg u$	$\neg u$	$\neg u$	$\neg u$	$\neg u$	$\neg u$	$\neg u$	u	u	u	u

It should be clear that the model can be transformed into one satisfying Ω , such as the one in Example 9. For example, the metric requirement that e occur once at $1460 + 1 - 40 = 1421$ can be accommodated by removing all the stuttering steps at position 8 and by adding $1421 - 8 = 1413$ additional stuttering steps at position 2. \diamond

6 Conclusion and future work

This paper investigated satisfiability-preserving transformations of LTL into its qualitative subset. For generic models, every LTL formula admits an equi-satisfiable qualitative formula of polynomially correlated size. For models with bounded variability, every LTL formula admits an equi-satisfiable qualitative formula whose size does not depend on the *distances* used in the original formula, where distances are defined by nested occurrences of *next* operators. Models with bounded variability can describe the behavior of systems with time granularity heterogeneity, where components evolving with wildly different time scales coexist. Under this assumption, the result of the present paper can be leveraged to simplify the automated reasoning of temporal logic specifications.

Future work will investigate possible generalizations and consider implementations. Concerning theoretical aspects, we will consider extensions of the results of the present paper to:

- *subword* stuttering [23], where a subword is repeated multiple times, such as in the word $abc\ abc\ abc\ \dots$;
- Büchi automata and the classical linear-time model-checking problem.

On the practical side, we will assess the practical usefulness of the results of the present paper. To this end, we plan to:

- implement a translator from LTL to formulas equi-satisfiable over words with bounded variability and combine it with off-the-shelf LTL satisfiability checking tools [34, 3, 29, 31];
- formalize systems characterized by time granularity heterogeneity, in order to determine how often the assumption of “sparse” events is compatible with accurate models thereof.

Acknowledgements. A preliminary version of this work has been presented at the 11th Italian Conference on Theoretical Computer Science (Cremona, Italy, 28–30 September 2009). The authors thank the conference attendees for their comments and for interesting discussions.

References

- [1] Rajeev Alur and Thomas A. Henzinger. Logics and models of real time: A survey. In J. W. de Bakker, Cornelis Huizing, and Willem P. de Roever, editors, *Proceedings of the Real-Time: Theory in Practice, REX Workshop*, volume 600 of *Lecture Notes in Computer Science*, pages 74–106. Springer-Verlag, 1992.
- [2] Michael Bauland, Thomas Schneider, Henning Schnoor, Ilka Schnoor, and Heribert Vollmer. The complexity of generalized satisfiability for linear temporal logic. *Logical Methods in Computer Science*, 5(1), 2008.
- [3] Domenico Bianculli, Angelo Morzenti, Matteo Pradella, Pierluigi San Pietro, and Paola Spoletini. Trio2Promela: A model checker for temporal metric specifications. In *ICSE Companion*, pages 61–62. IEEE Computer Society, 2007.
- [4] Ahmed Bouajjani, Rachid Echahed, and Peter Habermehl. On the verification problem of nonregular properties for nonregular processes. In *Proceedings, Tenth Annual IEEE Symposium on Logic in Computer Science, San Diego, California, USA, 26-29 June 1995*, pages 123–133. IEEE Computer Society, 1995.
- [5] Patricia Bouyer, Nicolas Markey, Joël Ouaknine, and James Worrell. On expressiveness and complexity in real-time model checking. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP'08)*, volume 5126 of *Lecture Notes in Computer Science*, pages 124–135. Springer, 2008.
- [6] Alan Burns and Ian J. Hayes. A timeband framework for modelling real-time systems. *Real-Time Systems*, Accepted for publication, 2010.
- [7] Christian Colombo, Gordon J. Pace, and Gerardo Schneider. Safe runtime verification of real-time properties. In Joël Ouaknine and Frits W. Vaandrager, editors, *Formal Modeling and Analysis of Timed Systems, 7th International Conference*, volume 5813 of *Lecture Notes in Computer Science*, pages 103–117, 2009.
- [8] Carlo Combi, Massimo Franceschet, and Adriano Peron. Representing and reasoning about temporal granularities. *Journal of Logic and Computation*, 14(1):51–77, 2004.
- [9] Edoardo Corsetti, Ernani Crivelli, Dino Mandrioli, Angelo Morzenti, Angelo Montanari, Pierluigi San Pietro, and Elena Ratto. Dealing with different time scales in formal specifications. In *In Proceedings of the 6th International Workshop on Software Specification and Design*, pages 92–101, 1991.
- [10] Ugo Dal Lago and Angelo Montanari. Calendars, time granularities, and automata. In *Advances in Spatial and Temporal Databases, 7th International Symposium, SSTD 2001, Redondo Beach, CA, USA, July 12-15, 2001, Proceedings*, volume 2121 of *Lecture Notes in Computer Science*, pages 279–298, 2001.

- [11] Christian Dax, Felix Klaedtke, and Stefan Leue. Specification languages for stutter-invariant regular properties. In Zhiming Liu and Anders P. Ravn, editors, *Automated Technology for Verification and Analysis, 7th International Symposium, ATVA 2009, Macao, China, October 14-16, 2009. Proceedings*, volume 5799 of *Lecture Notes in Computer Science*, pages 244–254. Springer, 2009.
- [12] Stéphane Demri and Philippe Schnoebelen. The complexity of propositional linear temporal logics in simple cases. *Information and Computation*, 174(1):84–103, 2002.
- [13] E. Allen Emerson. Temporal and modal logic. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 996–1072. Elsevier Science, 1990.
- [14] Kousha Etessami. Stutter-invariant languages, omega-automata, and temporal logic. In Nicolas Halbwachs and Doron Peled, editors, *Computer Aided Verification, 11th International Conference, CAV '99, Trento, Italy, July 6-10, 1999, Proceedings*, volume 1633 of *Lecture Notes in Computer Science*, pages 236–248. Springer, 1999.
- [15] Kousha Etessami. A note on a question of Peled and Wilke regarding stutter-invariant LTL. *Information Processing Letters*, 75(6):261–263, 2000.
- [16] Martin Fränzle. Decidability of duration calculi on restricted model classes. Technical Report Kiel MF 21/1, Institut für Informatik und Praktische Mathematik, Christian-Albrechts Universität Kiel, July 1996.
- [17] Carlo A. Furia, Dino Mandrioli, Angelo Morzenti, and Matteo Rossi. Modeling time in computing: a taxonomy and a comparative survey. *ACM Computing Surveys*, 42(2):1–59, February 2010. Article 6. Also available as <http://arxiv.org/abs/0807.4132>.
- [18] Carlo A. Furia and Matteo Rossi. MTL with bounded variability: Decidability and complexity. In Franck Cassez and Claude Jard, editors, *Proceedings of the 6th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'08)*, volume 5215 of *Lecture Notes in Computer Science*, pages 109–123. Springer-Verlag, September 2008.
- [19] Dov M. Gabbay, Ian Hodkinson, and Mark Reynolds. *Temporal Logic (vol. 1): mathematical foundations and computational aspects*, volume 28 of *Oxford Logic Guides*. Oxford University Press, 1994.
- [20] Yoram Hirshfeld and Alexander Moshe Rabinovich. Logics for real time: Decidability and complexity. *Fundamenta Informaticae*, 62(1):1–28, 2004.
- [21] Yoram Hirshfeld and Alexander Moshe Rabinovich. Timer formulas and decidable metric temporal logic. *Information and Computation*, 198(2):148–178, 2005.
- [22] Yonit Kesten, Zohar Manna, and Amir Pnueli. Temporal verification of simulation and refinement. In J. W. de Bakker, Willem P. de Roever, and Grzegorz Rozenberg, editors, *A Decade of Concurrency, Reflections and Perspectives, REX School/Symposium*, volume 803 of *Lecture Notes in Computer Science*. Springer, 1994.

- [23] Antonín Kučera and Jan Strejček. The stuttering principle revisited. *Acta Informatica*, 41(7/8):415–434, 2005.
- [24] Leslie Lamport. What good is temporal logic? In R. E. A. Mason, editor, *Proceedings of the IFIP 9th World Computer Congress on Information Processing (IFIP'83)*, pages 657–668. North Holland/IFIP, 1983.
- [25] François Laroussinie, Antoine Meyer, and Eudes Petonnet. Counting LTL. In *Proceedings of the 17th International Symposium on Temporal Representation and Reasoning (TIME'10)*. IEEE Computer Society Press, 2010.
- [26] Nicolas Markey. Past is for free: on the complexity of verifying linear temporal properties with past. *Acta Informatica*, 40(6-7):431–458, 2004.
- [27] Doron Peled and Thomas Wilke. Stutter-invariant temporal properties are expressible without the next-time operator. *Information Processing Letters*, 63(5):243–246, 1997.
- [28] Amir Pnueli and Aleksandr Zaks. On the merits of temporal testers. In *25 Years of Model Checking*, volume 5000 of *Lecture Notes in Computer Science*, pages 172–195. Springer, 2008.
- [29] Matteo Pradella, Angelo Morzenti, and Pierluigi San Pietro. A metric encoding for bounded model checking. In *FM 2009: Formal Methods*, volume 5850 of *Lecture Notes in Computer Science*, pages 741–756. Springer, 2009.
- [30] Alexander Moshe Rabinovich. Expressive completeness of temporal logic of action. In Lubos Brim, Jozef Gruska, and Jiri Zlatuska, editors, *Mathematical Foundations of Computer Science 1998, 23rd International Symposium, MFCS'98, Brno, Czech Republic, August 24-28, 1998, Proceedings*, volume 1450 of *Lecture Notes in Computer Science*, pages 229–238. Springer, 1998.
- [31] Kristin Y. Rozier and Moshe Y. Vardi. LTL satisfiability checking. In Dragan Bosnacki and Stefan Edelkamp, editors, *Proceedings of the 14th International SPIN Workshop on Model Checking Software (SPIN'07)*, volume 4595 of *Lecture Notes in Computer Science*, pages 149–167. Springer, 2007.
- [32] A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM*, 32(3):733–749, 1985.
- [33] Thomas Wilke. Specifying timed state sequences in powerful decidable logics and timed automata. In Hans Langmaack, Willem P. de Roever, and Jan Vytupil, editors, *Proceedings of the 3rd International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'94)*, volume 863 of *Lecture Notes in Computer Science*, pages 694–715. Springer-Verlag, 1994.
- [34] Martin De Wulf, Laurent Doyen, Nicolas Maquet, and Jean-François Raskin. Alaska. In Sung Deok Cha, Jin-Young Choi, Moonzoo Kim, Insup Lee, and Mahesh Viswanathan, editors, *6th International Symposium on Automated Technology for Verification and Analysis (ATVA'08)*, volume 5311 of *Lecture Notes in Computer Science*, pages 240–245. Springer, 2008.