On Relaxing Metric Information in Linear Temporal Logic*

Carlo A. Furia and Paola Spoletini

Abstract

This paper studies the equi-satisfiability of metric linear temporal logic (LTL) and its qualitative subset. Metric LTL formulas rely on the *next* operator to encode distances, whereas qualitative LTL formulas use only the *until* modality. The paper shows how to transform any metric LTL formula M into a qualitative one Q, such that Q and M are equi-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, these results may 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.

*First version June 2009; revised version April 2010. A preliminary version of this work has been presented at the 11th Italian Conference on Theoretical Computer Science, Cremona, Italy, 29 September 2009.

Contents

...
Learn well your grammar,
And never stammer,
Write well and neatly,
And sing most sweetly,
...
Drink tea, not coffee;
Never eat toffy.
Eat bread with butter.
Once more don't stutter.
...

Lewis Carrol, *Rules and regulations*

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 U — are instead purely *qualitative*: they exclusively 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 expressiveness and complexity relations between LTL and its qualitative subset are well known from the classic work by Lamport et al. [Lam83, PW97, Ete00, DS02, KS05]. The present paper studies instead the *equi-satisfiability* of quantitative (metric) and qualitative LTL formulas and determines under what conditions the satisfiability problem for LTL is reducible from metric to qualitative. 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; 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 first result of the paper shows that LTL and its qualitative subset are *equi-satisfiable*, hence every generic LTL formula ϕ admits a qualitative formula ϕ' that is satisfiable iff ϕ is. Given the aforementioned expressiveness gap between quantitative and qualitative LTL, ϕ' has to introduce new propositional letters. This result is a consequence of Etessami's [Ete00] and is not practically useful to reduce the complexity of formulas encoding large distances: the sizes of ϕ' and ϕ are polynomially correlated, so ϕ' essentially encodes, in a different form, the same information as ϕ.

This first construction is, however, the basis of a sharpening that introduces the notion of *bounded variability*. 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 second construction of the paper shows how to transform any LTL formula ϕ into a qualitative formula ϕ' that is equi-satisfiable to ϕ *over models with bounded variability.* The size of ϕ' does not depend on the distances (i.e., the number of occurrences of X) in ϕ but only on the maximum number of non-stuttering steps v. Intuitively, ϕ' drops some information encoded in ϕ; this information is not needed to decide satisfiability over models with bounded variability. On the technical level, this second construction introduces a normal form for quantitative LTL formulas and discrete-time generalized versions of the dense-time *Pnueli operators* [HR04].

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* [FMMR10] 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.

Related work. To our knowledge, this is the first work comparing LTL and its qualitative subset with respect to *satisfiability* (as opposed to expressiveness and complexity), and introducing and studying the notion of bounded variability and Pnueli operators *over discrete time* (as opposed to dense time).

The comparison between quantitative and qualitative LTL with respect to complexity and expressiveness is well understood within the framework of "classic" temporal logic (e.g., [GHR94, Eme90, DS02]). Lamport introduced the notion of *stuttering* to characterize qualitative LTL [Lam83]; the characterization was completed by Peled and Wilke [PW97] and generalized in [KS05]. The proof of the equi-satisfiability of qualitative and quantitative LTL — given in Section 5 — is linked to a result of Etessami's [Ete00].

Different notions of bounded variability for behaviors over *dense* (or continuous) time have been studied by several authors [Wil94, Frä96, FR08]. The proofs of the main results of the present paper use a technique that removes and adds stuttering steps in words to match some metric requirements; the notion of *stretching* — introduced in [HR05] and further used in [BMOW08] — exploits a related technique but over dense time models.

Hirshfeld and Rabinovich studied the expressiveness and decidability of Pnueli operators over dense time [HR04]; the operators themselves were first mentioned in a conjecture attributed to Pnueli [AH92, Wil94].

The problem of formalizing systems with heterogeneous time granularities using temporal logic has been studied by a few authors [CCM+91, BH10]. See also the survey [FMMR10] for a discussion on the problem of time granularity heterogeneity.

Outline. The rest of the paper is organized as follows. Section 2 introduces basic definitions, including LTL and bounded variability. Section 3 outlines a few examples that illustrate the results of the paper. Section 4 presents extensions of LTL with the Pnueli operators. Section 5 contains the first main result of the paper: the equi-satisfiability of LTL and qualitative LTL. Section 6 presents the second major result of the paper and details how the metric information can be relaxed while preserving satisfiability, for models with bounded variability. Finally, Section 7 outlines ideas for future work.

2 Definitions

This section introduces the syntax and semantics of LTL. $\mathbb{N}$ denotes the set of natural numbers $\{0, 1, 2, \ldots\}$.

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

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

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

Assume the standard abbreviations for $\top, \bot, \vee, \Rightarrow, \Leftrightarrow$ and for the derived temporal operators *eventually*: $\mathsf{F}\phi \triangleq \top \,\mathsf{U}\, \phi$; *always*: $\mathsf{G}\phi \triangleq \neg\mathsf{F}\neg\phi$; *release*: $\phi_1 \,\mathsf{R}\, \phi_2 \triangleq \neg(\neg\phi_1 \,\mathsf{U}\, \neg\phi_2)$; and *distance* $\mathsf{X}^k\phi = \underbrace{\mathsf{X}\,\mathsf{X}\cdots\mathsf{X}}_{k}\,\phi$ for $k \geq 0$.

Size and height. Let ϕ be an LTL formula. $|\phi|$ denotes the size of ϕ, encoded as a string. $\mathcal{P}(\phi) \subseteq \mathcal{P}$ denotes the (finite) set of propositional letters occurring in ϕ. For a temporal operator $\mathsf{H} \in \{\mathsf{U}, \mathsf{X}\}$, the *temporal height* (or *nesting depth*) $\mathcal{H}(\phi, \mathsf{H})$ of H in ϕ is the maximum number of nested occurrences of H in ϕ. For example, $\mathcal{H}(\phi, \mathsf{X}) = 0$ iff X is not used in ϕ. $\mathtt{L}\big(\mathsf{U}^{h_1}, \mathsf{X}^{h_2}\big)$ denotes the fragment of LTL whose formulas ψ are such that $\mathcal{H}(\psi, \mathsf{U}) \leq h_1$ and $\mathcal{H}(\psi, \mathsf{X}) \leq h_2$. Omit the superscript to mean that there is no bound on the temporal height of an operator. Hence, $\mathtt{L}(\mathsf{U}, \mathsf{X})$ is the same as all LTL; $\mathtt{L}\big(\mathsf{U}, \mathsf{X}^0\big) = \mathtt{L}(\mathsf{U})$ denotes *qualitative* LTL, where no *next* operator is used; and $\mathtt{L}\big(\mathsf{U}^0, \mathsf{X}^0\big) = \mathtt{P}(\mathcal{P})$ denotes purely propositional formulas where no temporal operators are used.

ω-words. An *ω-word* (or simply *word*) over S is a mapping $w : \mathbb{N} \to 2^S$ or, equivalently, a denumerable word $w(0)w(1)\cdots$ of elements $w(i) \subseteq S$. For $T \subseteq S$, $w|_T$ is the projection of w over T, defined as $w(0)|_T w(1)|_T \cdots$, where $w(i)|_T = w(i) \setminus (S \setminus T)$ 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 $w, i \models \phi_2$
$w, i \models \phi_1 \,\mathsf{U}\, \phi_2$	iff	there exists $j \geq i$ such that $w, j \models \phi_2$ and for all $i \leq k < j$ it is $w, k \models \phi_1$
$w, i \models \mathsf{X}\,\phi$	iff	$w, i+1 \models \phi$
$w \models \phi$	iff	$w, 0 \models \phi$

Strict until. Note that a non-strict *until* is assumed, as a strict *until* would make the *next* operator redundant. Indeed, let $w, i \models p \, \mathsf{U}^+ \, q$ be defined as there exists $j > i$ such that $w, j \models \phi_2$ and for all $i \leq k < j$ it is $w, k \models \phi_1$. Then, $\mathsf{X}\, p$ is equivalent to $(p \Rightarrow p \, \mathsf{U}^+ \, p) \wedge (\neg p \, \mathsf{U}^+ \, p \;\wedge\; p \, \mathsf{R}^+ \, p)$.

Satisfiability and validity. $[\![\phi]\!]$ denotes the set $\{w \in (2^{\mathcal{P}})^\omega \mid w \models \phi\}$ of all models of ϕ. ϕ is *satisfiable* iff $[\![\phi]\!] \neq \emptyset$ and is *valid* iff $[\![\phi]\!] = (2^{\mathcal{P}})^\omega$. Two formulas ϕ_1, ϕ_2 are *equivalent* iff $[\![\phi_1]\!] = [\![\phi_2]\!]$; they are *equi-satisfiable* iff they are either both satisfiable or both unsatisfiable.

Stuttering. A position $i \in \mathbb{N}$ is *redundant* in an ω-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* is any position i such that $w(i+1) \neq w(i)$ or $w(i+j) = w(i)$ for all $j \in \mathbb{N}$. Two words w_1, w_2 are *stutter-equivalent* (or equivalent under stuttering) iff one can be obtained from the other 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. A *stutter-free* word is one without stuttering steps.

Proposition 1. *Closure under stutter equivalence is a necessary and sufficient condition for qualitative* `LTL` *languages; that is:*

- *[Lam83]* $\phi \in \mathsf{L}(\mathsf{U})$ *implies that* $[\![\phi]\!]$ *is closed under stutter equivalence;*
- *[PW97]* W *closed under stutter equivalence and expressible in* `LTL` *implies there exists* $\phi \in \mathsf{L}(\mathsf{U})$ *such that* $[\![\phi]\!] = W$.

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 non-stuttering steps. $\mathbf{var}(P, v/k)$ 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 $v < k$.

3 Examples

This section presents two example `LTL` formalizations of systems with time granularity heterogeneity, where events occur sparsely over large time scales.

3.1 The school exam

The first example, in Table 1, formalizes a simple description of an exam session. Assume a one-to-one correspondence between hours and positions in a word; that is, the hour is the time unit. Propositions d and m hold at the beginning of every day (1) and month (2), respectively. Any exam finishes (e_{stop}) in exactly three hours, and always by the end of a day (3). There is at most one exam every day (4–5). Every student must turn in (t) her exam before the exam finishes (6). Once the exam is turned in, it is sent (n) to the TAs by the end of

$$d \Rightarrow \mathsf{X}(\neg d \,\mathsf{U}\, d) \wedge \mathsf{X}^{24} d \tag{1}$$

$$m \Rightarrow \mathsf{X}(\neg m \,\mathsf{U}\, m) \wedge \mathsf{X}^{24\cdot 30} m \tag{2}$$

$$e_{\text{start}} \Rightarrow \neg e_{\text{stop}} \wedge \neg d \wedge \mathsf{X}((\neg e_{\text{stop}} \wedge \neg d) \,\mathsf{U}\, e_{\text{stop}}) \wedge \mathsf{X}^{3} e_{\text{stop}} \tag{3}$$

$$e_{\text{start}} \Rightarrow \neg e_{\text{start}} \,\mathsf{U}\, d \tag{4}$$

$$e_{\text{stop}} \Rightarrow \neg(e_{\text{stop}} \vee e_{\text{start}}) \,\mathsf{U}\, d \tag{5}$$

$$e_{\text{start}} \Rightarrow t \,\mathsf{R}\, \neg e_{\text{stop}} \tag{6}$$

$$t \Rightarrow n \,\mathsf{R}\, \neg d \tag{7}$$

$$n \Rightarrow g \,\mathsf{R}\, \neg m \tag{8}$$

Table 1: The school exam example.

the day for grading (7). After being received by the TAs, it will be graded (g) by the end of the current month (8).

The complete system specification is given by:

$$\mathsf{G}((1) \wedge (2) \wedge (3) \wedge (4) \wedge (5) \wedge (6) \wedge (7) \wedge (8)) \tag{9}$$

The following formula is the separated-next form of (9), where the first conjunct corresponds to $\kappa \in \mathsf{L}(\mathsf{U})$.

$$\mathsf{G}\left(\begin{array}{l} (v_1 \Leftrightarrow \neg d \,\mathsf{U}\, d) \wedge (v_2 \Leftrightarrow \neg m \,\mathsf{U}\, m) \\ \wedge (v_3 \Leftrightarrow (\neg e_{\text{stop}} \wedge \neg d) \,\mathsf{U}\, e_{\text{stop}}) \\ \wedge (e_{\text{start}} \Rightarrow \neg e_{\text{start}} \,\mathsf{U}\, d) \\ \wedge (e_{\text{stop}} \Rightarrow \neg(e_{\text{stop}} \vee e_{\text{start}}) \,\mathsf{U}\, d) \\ \wedge (e_{\text{start}} \Rightarrow t \,\mathsf{R}\, \neg e_{\text{stop}}) \\ \wedge (t \Rightarrow n \,\mathsf{R}\, \neg d) \wedge (n \Rightarrow g \,\mathsf{R}\, \neg m) \\ \wedge (d \Rightarrow x_1 \wedge x_5) \wedge (m \Rightarrow x_2 \wedge x_6) \\ \wedge (e_{\text{start}} \Rightarrow \neg e_{\text{stop}} \wedge \neg d \wedge x_3 \wedge x_4) \end{array}\right) \wedge \mathsf{G}\left(\begin{array}{l} (x_1 \Leftrightarrow \mathsf{X}\, v_1) \\ \wedge (x_2 \Leftrightarrow \mathsf{X}\, v_2) \\ \wedge (x_3 \Leftrightarrow \mathsf{X}\, v_3) \\ \wedge (x_4 \Leftrightarrow \mathsf{X}^{3} e_{\text{stop}}) \\ \wedge (x_5 \Leftrightarrow \mathsf{X}^{24} d) \\ \wedge (x_6 \Leftrightarrow \mathsf{X}^{24\cdot 30} m) \end{array}\right) \tag{10}$$

Notice that: $|\mathcal{P}((10))| = 16$, $|(10)|_{\mathsf{G}} = M_{(10)} = 6$, $|(10)|_{\mathsf{X}} = 24 \cdot 30 = 720$; the latter dominates over the other size parameters, including $|(10)|_{\mathsf{U}} \simeq 100$.

3.2 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. The behavior is completely described by the following formula.

$$\begin{array}{l} q \wedge \mathsf{G}\left(q \Rightarrow \mathsf{X}\,(\neg q \wedge \neg q \,\mathsf{U}\, q) \wedge \mathsf{X}^{d_4} q\right) \wedge \mathsf{GF}e \\ \mathsf{G}\,(e \Rightarrow \neg q \wedge \mathsf{X}\,(\neg e \,\mathsf{U}\, q)) \wedge \mathsf{G}\left(e \Rightarrow \mathsf{X}^{d_2} q \vee \mathsf{X}^{d_3} q\right) \end{array} \tag{11}$$

It should be clear that the variability of every model of (11) is bounded: over every quadrennial (1460 days), at most one event p and one event e can happen — which corresponds to a variability of 4/1460 (each event is counted twice to account for its becoming true and then false again).

4 Extensions and normal forms for LTL

This section introduces some extensions of LTL that do not affect its expressiveness or its complexity; all constants are encoded in unary. It also shows how to transform any LTL formula in a normal form where nesting of temporal operators is limited; the results in the following sections will use this normal form.

4.1 Extending LTL with counting

$\mathtt{LTL}^{\mathsf{C}}$ extends LTL with the counting modalities C_k^n for $k, n \in \mathbb{N}$. $w, i \models \mathsf{C}_k^n\, \phi$ holds iff there exist (at least) k positions over the interval $[i+1..i+n]$ where ϕ holds. Counting modalities do not affect the expressiveness of LTL over discrete time[1] as they can be encoded as follows:

$$\mathsf{C}_k^n\, \phi \;\triangleq\; \bigvee_{\substack{S \subseteq [1..n] \\ |S| = k}} \bigwedge_{t \in S} \mathsf{X}^{\,t} \phi \tag{12}$$

The size of (12) is $\binom{n}{k}$, worst-case exponential in n.[2]

The complexity of the satisfiability problem for $\mathtt{LTL}^{\mathsf{C}}$ is, however, the same as for LTL. To see this, it is sufficient to show an encoding of the counting modality $\mathsf{C}_k^n\, \phi$ into an equi-satisfiable LTL formula which avoids the exponential blow-up. To this end, add n^2 fresh propositional letters p_j^i for $1 \leq i, j \leq n$. Intuitively p_j^i holds iff $\mathsf{C}_j^i\, \phi$ does; formally, a formula of size $\mathrm{O}(n \cdot |\phi|)$ defines p_j^i as:

$$p_j^i \Leftrightarrow \begin{cases} \left(p_{j-1}^{i-1} \wedge \mathsf{X}^{\,i}\phi\right) \vee p_j^{i-1} & 1 < j \leq i \leq n \\ \mathsf{X}\,\mathsf{F}_{\leq i-1}\phi & 1 = j < i \leq n \\ \mathsf{X}\,\phi & 1 = j = i \\ \bot & j > i \end{cases}$$

To check the satisfiability of an $\mathtt{LTL}^{\mathsf{C}}$ formula ψ with counting modalities, change ψ into $\overline{\psi}$ by replacing every occurrence of $\mathsf{C}_k^n\, \phi$ with p_k^n and then check the satisfiability of the new formula $\psi' \triangleq \overline{\psi} \wedge \bigwedge_{i,j} \mathsf{G}(p_j^i \Leftrightarrow \cdots)$ including the definitions of the p_j^i. The size of ψ' is polynomial in the size of ψ, and ψ' is satisfiable iff ψ is; this proves that LTL and $\mathtt{LTL}^{\mathsf{C}}$ have the same PSPACE worst-case complexity.

Pnueli modalities. $\mathtt{LTL}^{\mathsf{Pn}}$ extends LTL with the Pnueli modalities Pn_k^n for $k, n \in \mathbb{N}$. $w, i \models \mathsf{Pn}_k^n(\phi_1, \ldots, \phi_k)$ holds iff there exist k positions $i + 1 \leq k_1 < k_2 < \cdots < k_k \leq i + n$ such that $w, k_j \models \phi_j$ for all $1 \leq j \leq k$. $\mathtt{LTL}^{\mathsf{Pn}}$ subsumes $\mathtt{LTL}^{\mathsf{C}}$, because $\mathsf{C}_k^n\, \phi \equiv \mathsf{Pn}_k^n(\phi, \ldots, \phi)$; $\mathtt{LTL}^{\mathsf{Pn}}$ also has the same expressiveness and complexity as LTL. The proof follows an argument similar to the case of the counting modalities; in particular, to encode the satisfiability for $\mathtt{LTL}^{\mathsf{Pn}}$in LTL, introduce n^2 propositions p_j^i for $1 \leq i, j \leq n$, such that p_j^i holds iff $\mathsf{Pn}_j^i(\phi_1, \ldots, \phi_j)$ does.

[1] Unlike over dense time [HR04].

[2] Recall that $\frac{n^k}{k^k} \leq \binom{n}{k}$ [GKP94], hence in particular $\binom{n}{n/2} \geq 2^{n/2}$.

4.2 Flat-next form

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

$$\kappa \ \wedge \bigwedge_{i=1,\dots,N} \mathsf{G}(x_i \Leftrightarrow \mathsf{X}\,\pi_i) \tag{13}$$

where $\kappa \in \mathtt{L}(\mathsf{U})$, $x_i \in \mathcal{P}$, $\pi_i \in \mathtt{P}(\mathcal{P})$. The nesting depth of the X operator is one in (13); that is (13) $\in \mathtt{L}\big(\mathsf{U}, \mathsf{X}^1\big)$.

Lemma 2. *For any $\phi \in$ LTL it is possible to build, in polynomial time, an equi-satisfiable formula η in FNF.*

Proof. Initially, let $\mathcal{Q} = \mathcal{P}$ and $\phi' = \phi$. Repeat the following two steps, working recursively starting from the top-level of ϕ', until $\phi' \in \mathtt{L}(\mathsf{U})$:

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

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

Steps 1–2 are repeated at most a number of times proportional to the size $|\phi|$ of ϕ. Define κ as $\phi' \wedge \bigwedge_{\substack{p_\psi \in \mathcal{Q} \\ \psi \in \mathtt{L}(\mathsf{U})}} \mathsf{G}(p_\psi \Leftrightarrow \psi)$. Finally, η is $\kappa \wedge \bigwedge_{p_{\mathsf{X}\,\pi} \in \mathcal{Q}} \mathsf{G}(p_{\mathsf{X}\,\pi} \Leftrightarrow \mathsf{X}\,\pi)$. η is in FNF and it is equi-satisfiable to ϕ. □

Example 3. The following is formula (11) of the elections example in flat-next form; the first conjunct is the qualitative part κ, and x_k encodes $\mathsf{X}^{\,k} q$ for $k \geq 1$.

$$q \wedge \mathsf{G}\begin{pmatrix} (u \Leftrightarrow \neg e \,\mathsf{U}\, q) \wedge (v \Leftrightarrow \neg q \wedge \neg q \,\mathsf{U}\, q) \\ \wedge\, (q \Rightarrow x_v \wedge x_{d_4}) \wedge \mathsf{F} e \\ \wedge\, (e \Rightarrow \neg q \ \wedge \ x_u) \\ \wedge\, (e \Rightarrow x_{d_2} \vee x_{d_3}) \end{pmatrix} \wedge \mathsf{G}\begin{pmatrix} (x_u \Leftrightarrow \mathsf{X}\,u) \wedge (x_v \Leftrightarrow \mathsf{X}\,v) \\ \wedge (x_1 \Leftrightarrow \mathsf{X}\,q) \\ \bigwedge_{2 \leq k \leq d_4} (x_k \Leftrightarrow \mathsf{X}\,x_{k-1}) \end{pmatrix} \tag{14}$$

4.3 Separated-next form

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

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

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

Given that the FNF is a special case of the SNF, it is obvious that any LTL formula can transformed into an equi-satisfiable SNF one in polynomial time. However, the SNF becomes interesting when it isolates subformulas with a nesting depth of X as high as possible, as the rest of the paper demonstrates.

Given a formula ϕ in SNF (15), $|\phi|_{\mathsf{U}} \triangleq |\kappa|$, $|\phi|_{\mathsf{G}} \triangleq M$, $|\phi|_{\mathsf{P}} \triangleq \max_i \pi_i$, and $|\phi|_{\mathsf{X}} \triangleq \max_i \mathbb{D}(i) = \mathbb{D}(M)$ denote respectively the until, globally, propositional, and next size of the components of ϕ. Then, $|\phi|$ is in $\mathrm{O}(|\phi|_{\mathsf{U}} + |\phi|_{\mathsf{G}}(|\phi|_{\mathsf{P}} + |\phi|_{\mathsf{X}}))$.

Example 4. The following formula Ω is formula (11) in separated-next form, with $d_1 = d_2 = 1$, $d_3 = 40$, $d_4 = 41$, $d_5 = 1460$.

$$\Omega \triangleq \left(q \wedge \mathsf{G} \left(\begin{array}{l} (u \Leftrightarrow \neg e \,\mathsf{U}\, q) \\ \wedge\, (v \Leftrightarrow \neg q \wedge \neg q \,\mathsf{U}\, q) \\ \wedge\, (q \Rightarrow x_2 \wedge x_5) \wedge \mathsf{F} e \\ \wedge\, (e \Rightarrow \neg q \,\wedge\, x_1) \\ \wedge\, (e \Rightarrow x_3 \vee x_4) \end{array} \right) \right) \wedge \mathsf{G} \left(\begin{array}{l} (x_1 \Leftrightarrow \mathsf{X}^{\,d_1} u) \\ \wedge\, (x_2 \Leftrightarrow \mathsf{X}^{\,d_2} v) \\ \wedge\, (x_3 \Leftrightarrow \mathsf{X}^{\,d_3} q) \\ \wedge\, (x_4 \Leftrightarrow \mathsf{X}^{\,d_4} q) \\ \wedge\, (x_5 \Leftrightarrow \mathsf{X}^{\,d_5} q) \end{array} \right) \tag{16}$$

Notice that $\kappa_\Omega \in \mathsf{L}(\mathsf{U})$ is the first conjunct, $|\mathcal{P}(\Omega)| = 9$, $|\Omega|_{\mathsf{G}} = M_\Omega = 5$, $|\Omega|_{\mathsf{X}} = d_5$; the latter dominates over the other size parameters.

The following is a (partial) model of Ω.

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

In any such model there are at most 6 non-stuttering steps over $[1..1460]$: positions $1, 1419, 1420, 1421, 1459, 1460$; hence the variability is $6/1460$ and let $\mathbb{V}_\Omega = 6$ and $D_\Omega = 1460$.

5 LTL and qualitative LTL are equi-satisfiable

This section shows how to transform any `LTL` formula into an equi-satisfiable $\mathsf{L}(\mathsf{U})$ formula. The construction is the foundation for the novel results of Section 6.

Informal presentation

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 non-stuttering step coincides with a non-stuttering step of s. Then, replace any occurrence of a subformula $\mathsf{X}\, p$ with a suitable *until* formula that defines the value of p at the nest non-stuttering step of s. This changes the quantitative $\mathsf{X}\, p$ formula into a qualitative formula where the precise metric information is relaxed. Very informally, this basic idea can be regarded as a discrete-time analogue of the stretching introduced in [HR05] and further exploited in [BMOW08].

Formal presentation

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

$$\curlywedge(\pi) \quad \triangleq \quad \left(\pi \wedge \mathsf{F}\neg\pi \;\Rightarrow \left(\begin{array}{c} s\,\mathsf{U}\,(\neg\pi \wedge \neg s) \\ \vee \\ \neg s\,\mathsf{U}\,(\neg\pi \wedge s) \\ \vee \\ \bigvee_{q \in \mathcal{P}\setminus\{\pi\}} (q \wedge \pi)\,\mathsf{U}\,(\neg q \wedge \pi) \\ \vee \\ \bigvee_{q \in \mathcal{P}\setminus\{\pi\}} (\neg q \wedge \pi)\,\mathsf{U}\,(q \wedge \pi) \end{array}\right)\right)$$

$$\curlyvee(\mathcal{P}) \quad \triangleq \quad \left(\left(\bigwedge_{p \in \mathcal{P}} \mathsf{G}p \vee \mathsf{G}\neg p\right) \Rightarrow \mathsf{G}s\right)$$

$$\mathcal{X}^1(\pi) \quad \triangleq \quad s\,\mathsf{U}\,\pi \;\vee\; \neg s\,\mathsf{U}\,\pi$$

$$\mathcal{X}^2(\pi) \quad \triangleq \quad (\pi \wedge s \;\Rightarrow\; \neg s\,\mathsf{R}\,\pi) \;\wedge\; (\pi \wedge \neg s \;\Rightarrow\; s\,\mathsf{R}\,\pi)$$

$$\mathcal{X}(\pi) \quad \triangleq \quad \mathcal{X}^1(\pi) \wedge \mathcal{X}^2(\pi)$$

$\curlywedge(\pi)$ links any transition of the truth value of π to occur simultaneously with a transition of s. $\mathcal{X}(\pi)$ is instead essentially a qualitative relaxations of the *next* operator: $w, i \models \mathcal{X}(p)$ holds iff the next non-stuttering step of s is $j \geq i$ and $w, j+1 \models p$ holds.

Finally, build a qualitative formula ξ from η as:

$$\xi \quad \triangleq \quad \curlyvee(\mathcal{P}) \wedge \bigwedge_{p \in \mathcal{P}} \mathsf{G}\,(\curlywedge(p) \wedge \curlywedge(\neg p)) \;\wedge\; \kappa \;\wedge \bigwedge_{i=1,\ldots,N} \mathsf{G}\left(\begin{array}{c} x_i \Rightarrow \mathcal{X}(\pi_i) \\ \wedge \\ \neg x_i \Rightarrow \mathcal{X}(\neg\pi_i) \end{array}\right)$$

Lemma 5. *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$(\eta) \Rightarrow$ **SAT**(ξ)**.** In the first part we show that ξ is satisfiable if η is satisfiable. Hence, assume $w \models \eta$ for some $w \in (2^{\mathcal{P}(\eta)})^\omega$. Let us build $x \in (2^{\mathcal{P}(\xi)})^\omega$ such that $x \models \xi$.

Indeed, 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 non-stuttering steps — except possibly for an infinite tail of constant states.

For any $p \in \mathcal{P}(\eta)$ let us show that $x \models \curlyvee(\mathcal{P}) \wedge \mathsf{G}\,(\curlywedge(p) \wedge \curlywedge(\neg p))$. The proof of $x \models \curlyvee(\mathcal{P})$ is routine. Then, let $i \in \mathbb{N}$ be such that $x, i \models p$. If $x, i \models \mathsf{G}p$ then trivially $x, i \models \curlywedge(p)$, because the $\pi \wedge \mathsf{F}\neg\pi$ 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 $x, i \models s\,\mathsf{U}\,(\neg p \wedge \neg s)$ and $x, i \models \neg s\,\mathsf{U}\,(\neg p \wedge s)$ 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 $(q \wedge p)\,\mathsf{U}\,(\neg q \wedge p)$ holds at i, whereas in the latter case $(\neg q \wedge p)\,\mathsf{U}\,(q \wedge p)$ holds at i. Hence, if $x, i \models p$ then $x, i \models \curlywedge(p) \wedge \curlywedge(\neg p)$ is established. If

$x, i \models \neg p$ instead, a similar reasoning also proves that $x, i \models \curlywedge(p) \wedge \curlywedge(\neg p)$. In all, $x \models \mathsf{G}(\curlywedge(p) \wedge \curlywedge(\neg p))$ 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 $x, i \models s \,\mathsf{U}\, \pi_h$. In addition, assume that $x, i \models \pi_h$: we have to show that $\neg s \,\mathsf{R}\, \pi_h$. 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 $x, i \models \neg s \,\mathsf{U}\, \pi_h$. Similarly as we did in the previous case, we can establish also that if $x, i \models \pi_h$ then $s \,\mathsf{R}\, \pi_h$. 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 $x, i \models \mathcal{X}(\neg \pi_h)$. In all, $x \models \mathsf{G}(x_h \Rightarrow \mathcal{X}(\pi_h))$ and $x \models \mathsf{G}(\neg x_h \Rightarrow \mathcal{X}(\neg \pi_h))$ is established.

SAT(ξ) $\Rightarrow$ SAT(η). Let us now show that η is satisfiable if ξ is satisfiable. Hence, assume that $w \models \xi$ for some $w \in (2^{\mathcal{P}(\eta) \cup \{s\}})^\omega$. Let us build $x \in (2^{\mathcal{P}(\eta)})^\omega$ such that $x \models \eta$.

Build y from w by removing all its stuttering steps. 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 $y, i \models x_h \Leftrightarrow \mathsf{X}\, \pi_h$.

1. Assume $y, i \models x_h \wedge \mathcal{X}(\pi_h)$. Adversarially, 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, $\curlywedge(r) \wedge \curlywedge(\neg r)$ forces s to also change its truth value at i. In all we have the following situation:

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

   But then $\mathcal{X}^2(\pi_h)$ requires in particular $\neg s \,\mathsf{R}\, \pi_h$ 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 $\mathcal{X}^1(\pi_h)$ 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, $\curlywedge(r) \wedge \curlywedge(\neg r)$ forces s to also change its truth value at i. In all we have the following situation:

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

But then $\mathcal{X}^1(\pi_h)$ cannot hold at i, because neither $s \,\mathsf{U}\, \pi_h$ nor $\neg s \,\mathsf{U}\, \pi_h$ 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 $y \models \bigwedge_{i=1,\ldots,N} \mathsf{G}(x_i \Leftrightarrow \mathsf{X}\,\pi_i)$. In addition $w \models \kappa$ implies $y \models \kappa$ as well, because y is obtained from w only by removing stuttering steps and $\kappa \in \mathsf{L}(\mathsf{U})$ is closed under stuttering. Hence $x = y|_{\mathcal{P}(\eta)}$ is a model that satisfies η. □

Example 6. Let $\xi(14)$ be formula (14) modified according to the construction of the present section. The proof of Lemma 5 shows that the qualitative formula $\xi(14)$ preserves the stutter-free models of the equi-satisfiable LTL formula (14). On the other hand, consider a model of (14) with a sequence of $d_5 - 2$ stuttering steps $\underbrace{\neg q \cdots \neg q}_{d_5-2}\, q$, such as the one in Example 3; it corresponds to the following stutter-free model of $\xi(14)$:

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

This shows that the transformation of (14) into $\xi(14)$ — 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.

6 Relaxing distance formulas

This section sharpens the results of Section 5 by showing how to encode the redundancy of stuttering steps more succinctly in words with bounded variability. The following results require $\mathtt{LTL}^{\mathsf{qPn}}$ and $\mathtt{LTL}^{\mathsf{exqPn}}$, two extensions of $\mathsf{L}(\mathsf{U})$ with qualitative counterparts of the Pnueli operator discussed in Section 4.1. Section 6.1 introduces these extensions and shows that they have the same complexity as LTL. Then, Section 6.2 shows how to transform any LTL formula

ϕ (given in SNF) and a positive integer parameter $\mathbb{V}$ into a $\mathtt{LTL}^{\mathsf{exqPn}}$ formula ϕ' which is equi-satisfiable to ϕ over words with variability bounded by $\mathbb{V}/D$, where D is the largest distance used in ϕ. The size of ϕ' is polynomial in $\mathbb{V}$, the *number* of distance sub-formulas, and the size of propositional sub-formulas appearing ϕ; however, the size of ϕ' does not depend on the *values* of distances in ϕ. As a consequence, checking the satisfiability of ϕ' is much less complex than checking the original ϕ whenever the distances used in ϕ are very large and dominate over the other size parameters.

6.1 Qualitative Pnueli operators

$\mathtt{LTL}^{\mathsf{qPn}}$ extends $\mathtt{LTL}$ with the *qualitative Pnueli modalities* qPn_k^n for $k, n \in \mathbb{N}$. $w, i \models \mathsf{qPn}_k^n(\phi_1, \ldots, \phi_k)$ holds iff there exist k positions $i \leq k_1 < \cdots < k_k$ such that, for all $1 \leq j \leq k$: (1) k_j is a non-stuttering step; (2) $w, k_j + 1 \models \phi_j$; and (3) there are no more than n non-stuttering steps between i and k_k (both included).

Example 7. Consider the following word w; non-stuttering steps are in bold and underlined.

1	2	3	4	5	**6**	**7**	8	9	**10**	11	**12**	**13**	14
s	$\neg s$	$\neg s$	$\neg s$	$\neg s$	$\neg s$	s	$\neg s$	$\neg s$	$\neg s$	s	s	$\neg s$	s
$\neg v$	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	$\neg e$	$\neg e$	$\neg e$	$\neg e$

Then, $w, 1 \models \mathsf{qPn}_4^6(v, \neg q, e, q)$ holds for the positions $1, 6, 7, 13$; but $w, 8 \not\models \mathsf{qPn}_2^2(s, q)$ because $w, 13 \not\models q$.

Any $\mathtt{LTL}^{\mathsf{qPn}}$ formulas admits an equi-satisfiable $\mathtt{LTL}$ formula, and their sizes are polynomially correlated; hence $\mathtt{LTL}^{\mathsf{qPn}}$ and $\mathtt{LTL}$ have the same worst-case complexity. The proof introduces a fresh letter s, which is constrained to change truth value with the other letters in $\mathcal{P}$ (as detailed in Section 5). It then uses an encoding similar to that of "standard" Pnueli operators (see Section 4.1), where the qualitative $\mathcal{X}$ operator — introduced in Section 5 — replaces the quantitative X; namely, to encode a formula $\mathsf{qPn}_k^n(\phi_1, \ldots, \phi_n)$ introduce n^2 propositional letters $\{q_i^j \mid 1 \leq i, j \leq n\}$. Every q_i^j holds iff $\mathsf{qPn}_i^j(\phi_1, \ldots, \phi_i)$ does; formally, a formula of size $\mathrm{O}(n \cdot |\phi_i|)$ defines p_i^j as:

$$q_i^j \Leftrightarrow \begin{cases} \left(q_{i-1}^{j-1} \wedge \mathcal{X}^j(\phi_i)\right) \vee q_i^{j-1} & 1 < i \leq j \leq n \\ \underbrace{\mathcal{X}(\phi_1 \vee \mathcal{X}(\phi_1 \vee \cdots))}_{j \text{ nested } \mathcal{X}} & 1 = i < j \leq n \\ \mathcal{X}(\phi_1) & 1 = i = j \\ \bot & i > j \end{cases} \tag{17}$$

where $\mathcal{X}^j$ abbreviates $\underbrace{\mathcal{X}\mathcal{X}\cdots\mathcal{X}}_{j}$. In the following, $\tau(\underline{q}, \underline{i}, \underline{j}, \underline{n}, \underline{\phi_1}, \ldots)$ denotes (17) with $\underline{q}, \underline{i}, \underline{j}, \underline{n}, \underline{\phi_1}, \ldots$ respectively replacing $q, i, j, n, \phi_1, \ldots$.

The construction outlined is general, but the remainder of the paper only considers $\mathtt{LTL}^{\mathsf{qPn}}$ formulas in the form:

$$\Xi \triangleq \curlyvee(\mathcal{P}) \wedge \bigwedge_{p\in\mathcal{P}} \mathsf{G}(\curlywedge(p) \wedge \curlywedge(\neg p)) \wedge \kappa \wedge \bigwedge_{i=1,\ldots,M} \mathsf{G}\left((\xi_i \Rightarrow \mathsf{qPn}^{\mathbb{J}(i)}_{\mathbb{I}(i)}\left(\psi^i_1,\ldots,\psi^i_{\mathbb{I}(i)}\right)\right) \tag{18}$$

for a set of propositions $\mathcal{Q} = \mathcal{P}\cup\{s\}$, where $\mathbb{J}, \mathbb{I}$ are two mappings $[1..M] \to \mathbb{N}_{>0}$, $\kappa, \psi_i \in \mathsf{L}(\mathsf{U})$ for all $1 \leq i \leq M$, $\xi_i \in \mathsf{P}(\mathcal{P})$, and s does not occur in κ or in any ψ^j_i. Correspondingly, introduce fresh letters $q(i)^j_k$ and construct Ξ' from Ξ:

$$\Xi' \triangleq \begin{pmatrix} \curlyvee(\mathcal{P}) \wedge \bigwedge_{p\in\mathcal{P}} \mathsf{G}(\curlywedge(p) \wedge \curlywedge(\neg p)) \\ \wedge \quad \kappa \quad \wedge \\ \bigwedge_{i=1,\ldots,M} \mathsf{G}(\xi_i \Rightarrow q(i)^{\mathbb{J}(i)}_{\mathbb{I}(i)}) \end{pmatrix} \wedge \bigwedge_{\substack{1\leq k\leq M\\ 1\leq i,j\leq \mathbb{J}(k)}} \mathsf{G}\tau\left(q(k), i, j, \mathbb{J}(k), \psi^k_1, \ldots, \psi^k_{\mathbb{I}(k)}\right) \tag{19}$$

Lemma 8. *The following hold for* Ξ, Ξ'*:*

1. $\Xi' \in \mathsf{L}(\mathsf{U})$*;*
2. $|\Xi'|$ *is polynomially bounded by* $|\Xi|$ *(and can be built in polynomial time);*
3. $[\![\Xi]\!]$ *is closed under stuttering;*
4. Ξ' *and* Ξ *are equi-satisfiable;*
5. $[\![\Xi]\!]$ *equals the projection over* $\mathcal{Q}$ *of* $[\![\Xi']\!]$*.*

Proof sketch. 1: is clear by construction.

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

$$\curlyvee(\mathcal{P}) \wedge \bigwedge_{p\in\mathcal{P}} \mathsf{G}(\curlywedge(p) \wedge \curlywedge(\neg p)) \wedge \quad \kappa \quad \wedge \bigwedge_{i=1,\ldots,M} \mathsf{G}(\xi_i \Rightarrow q(i)^{\mathbb{J}(i)}_{\mathbb{I}(i)}) \tag{20}$$

has size bounded by $|\Xi|+M\cdot I\cdot J$, hence also $\mathrm{O}(|\Xi|)$. The second subformula of Ξ':

$$\bigwedge_{\substack{1\leq k\leq M\\ 1\leq i,j\leq \mathbb{J}(k)}} \mathsf{G}\tau(q(k), i, j, \mathbb{J}(k), \psi^k_1, \ldots, \psi^k_{\mathbb{I}(k)}) \tag{21}$$

has size bounded by $M \cdot J^2 \cdot \mathrm{O}(J)$, hence also polynomial in $|\Xi|$.

3: follows from 1 and 5.

4: follows from 5.

5: follows from the definition in (17) and the semantics of the $\mathcal{X}$ operator, along the lines of Lemma 5. □

Qualitative extended Pnueli operators

$\mathtt{LTL}^{\mathsf{exqPn}}$ is the last extension of $\mathtt{LTL}$ we consider. $\mathtt{LTL}^{\mathsf{exqPn}}$ introduces the *qualitative extended Pnueli modalities* $\mathsf{exqPn}_k^{n;\langle n_1,\ldots,n_k\rangle}$ for $k, n \in \mathbb{N}$ and $n_1, \ldots, n_k \in \mathbb{N} \cup \{*\}$. $w, i \models \mathsf{exqPn}_k^{n;\langle n_1,\ldots,n_k\rangle}(\phi_1, \ldots, \phi_k)$ holds iff there exist k positions $i \leq k_1 < \cdots < k_k$ such that, for all $1 \leq j \leq k$: (1) k_j is a non-stuttering step; (2) $w, k_j + 1 \models \phi_j$; (3) for $j > 1$, if $n_j \neq *$ then there are no more than n_j non-stuttering steps between k_{j-1} and $k_j - 1$ (both included); (4) if $n_1 \neq *$ then there are no more than n_1 non-stuttering steps between i and k_1 (both included); and (5) there are no more than n non-stuttering steps between i and k_k (both included). Intuitively, the qualitative extended Pnueli modalities generalize the qualitative Pnueli modalities by imposing an additional requirement on the relative distance of k non-stuttering steps; for example, if $n_1 = 1$, ϕ_1 must hold right after the first non-stuttering step that follows or is at i, independently of the other following $k - 1$ non-stuttering steps. If $\langle n_1, \ldots, n_k\rangle$ is $\langle *^k\rangle$ then the qualitative extended Pnueli modality reduces to the corresponding qualitative Pnueli modality.

Example 9. Consider again word w from Example 7. For the positions $1, 6, 7, 13$, $w, 1 \models \mathsf{exqPn}_4^{6;\langle 3,2,*,3\rangle}(v, \neg q, e, q)$ holds. However, $w, 1 \not\models \mathsf{exqPn}_4^{6;\langle 3,2,*,1\rangle}(v, \neg q, e, q)$; in fact, let $k_1, \ldots, k_4$ 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 non-stuttering step immediately before 13; but $w, 12 + 1 \not\models e$.

Any $\mathtt{LTL}^{\mathsf{exqPn}}$ formula has an equi-satisfiable $\mathtt{LTL}$ formula of polynomially correlated size. Again, the proof follows the same lines as for $\mathtt{LTL}^{\mathsf{qPn}}$: to encode a formula $\mathsf{exqPn}_k^{n;\langle n_1,\ldots,n_k\rangle}(\phi_1, \ldots, \phi_k)$ introduce n^2 letters $Q = \{q_i^j \mid 1 \leq i, j \leq n\}$. Every q_i^j holds iff $\mathsf{exqPn}_i^{j;\langle n_1,\ldots,n_i\rangle}(\phi_1, \ldots, \phi_i)$ does; formally, a formula of size $\mathrm{O}(n \cdot |\phi_i|)$ defines p_i^j as:

$$q_i^j \Leftrightarrow \begin{cases} \left(\begin{array}{l} q_{i-1}^{j-1} \wedge \mathcal{X}^j(\phi_i) \wedge \\ \underbrace{\mathcal{X}^{j-n_i}(\phi_{i-1} \vee \mathcal{X}(\phi_{i-1} \vee \cdots))}_{n_i \text{ nested } \mathcal{X}} \end{array} \right) \vee q_i^{j-1} & \begin{array}{l} 1 < i \leq j \leq n \\ \wedge\, 0 < n_i < j \end{array} \\ \left(q_{i-1}^{j-1} \wedge \mathcal{X}^j(\phi_i)\right) \vee q_i^{j-1} & \begin{array}{l} 1 < i \leq j \leq n \,\wedge \\ (n_i = * \vee n_i \geq j) \end{array} \\ \underbrace{\mathcal{X}(\phi_1 \vee \mathcal{X}(\phi_1 \vee \cdots))}_{\min(j,n_1) \text{ nested } \mathcal{X}} & \begin{array}{l} 1 = i < j \leq n \\ \wedge\, 0 < n_i \neq * \end{array} \\ \underbrace{\mathcal{X}(\phi_1 \vee \mathcal{X}(\phi_1 \vee \cdots))}_{j \text{ nested } \mathcal{X}} & 1 = i < j \leq n \wedge n_i = * \\ \mathcal{X}(\phi_1) & 1 = i = j \wedge n_i \neq 0 \\ \bot & i > j \vee n_i = 0 \end{cases} \tag{22}$$

The reminder of the paper deals with $\mathtt{LTL}^{\mathsf{exqPn}}$ formulas in the form:

$$\Lambda \triangleq \curlyvee(\mathcal{P}) \wedge \bigwedge_{p \in \mathcal{P}} \mathsf{G}\begin{pmatrix} \curlywedge(p) \wedge \\ \curlywedge(\neg p) \end{pmatrix} \wedge \kappa \wedge \bigwedge_{i=1,\ldots,M} \mathsf{G}\left(\xi_i \Rightarrow \mathsf{exqPn}_{\mathbb{I}(i)}^{\mathbb{J}(i);\langle\mathbb{K}(i)\rangle}\left(\psi_1^i, \ldots, \psi_{\mathbb{I}(i)}^i\right)\right) \tag{23}$$

for a set of propositions $\mathcal{Q} = \mathcal{P} \cup \{s\}$, where $\mathbb{J}, \mathbb{I}$ are two mappings $[1..M] \to \mathbb{N}_{>0}$, $\mathbb{K}$ is a mapping $[1..M] \to (\mathbb{N} \cup \{*\})^{\mathbb{I}(i)}$, $\kappa, \psi_i \in \mathsf{L}(\mathsf{U})$ for all $1 \le i \le M$, $\xi_i \in \mathsf{P}(\mathcal{P})$, and s does not occur in κ or in any ψ_i^j. It is not difficult to adapt Lemma 8 to the $\mathtt{LTL}^{\mathsf{exqPn}}$ formula Λ and prove the following.

Lemma 10. *Given Λ, it is possible to build an* LTL *formula Λ' such that:*

1. $\Lambda' \in \mathsf{L}(\mathsf{U})$;
2. $|\Lambda'|$ *is polynomially bounded by* $|\Lambda|$ *(and can be built in polynomial time);*
3. $[\![\Lambda]\!]$ *is closed under stuttering;*
4. Λ' *and* Λ *are equi-satisfiable;*
5. $[\![\Lambda]\!]$ *equals the projection over* $\mathcal{Q}$ *of* $[\![\Lambda']\!]$.

6.2 Relaxing distance formulas

This section contains the main result of the paper, summarized in the following.

Theorem 11. *Given an* LTL *formula ϕ in SNF over propositions in $\mathcal{Q}$, and an integer parameter $\mathbb{V} > 0$, it is possible to build, in polynomial time, another* LTL *formula ϕ'' such that:*

- $|\phi''|$ *is polynomial in* $|\mathcal{Q}|, \mathbb{V}, |\phi|_{\mathsf{U}}, |\phi|_{\mathsf{G}}, |\phi|_{\mathsf{P}}$ *but is independent of* $|\phi|_{\mathsf{X}}$;
- ϕ *and* ϕ'' *are equi-satisfiable over words in* $\mathbf{var}(\mathcal{Q}, \mathbb{V}/|\phi|_{\mathsf{X}})$.

The following construction shows how to build a $\phi' \in \mathtt{LTL}^{\mathsf{exqPn}}$ from ϕ such that Lemmas 13 and 14 hold. Theorem 11 follows after transforming ϕ' into ϕ'' by eliminating the qualitative extended Pnueli operators as in Lemma 10.

Consider a generic LTL formula ϕ in SNF:

$$\phi \quad \triangleq \quad \kappa \wedge \bigwedge_{i=1,\ldots,M} \mathsf{G}(x_i \Leftrightarrow \mathsf{X}^{\mathbb{D}(i)} \pi_i) \tag{24}$$

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

Introduce $\mathbb{V}$ letters $Y = \{y_i \mid 1 \le i \le \mathbb{V}\}$. Formula $\varkappa\langle y\rangle$ constrains y_i to occur synchronously with every i-th non-stuttering step:

$$\varkappa\langle y\rangle \quad \triangleq \quad y_1 \wedge \bigwedge_{1 \le i \le \mathbb{V}} \mathsf{G}\left(y_i \Rightarrow \left(\begin{array}{l} \mathcal{X}(y_{1+(i \bmod \mathbb{V})}) \\ \bigwedge_{j \neq i} \neg y_j \wedge \\ y_i \,\mathsf{U}\, y_{1+(i \bmod \mathbb{V})} \end{array} \right)\right) \tag{25}$$

Let $D_1, D_2, \ldots, D_m$ be the sequence of sets that partition $[1..M]$ and such that: $i, j \in D_k$ 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 $\{z_i^j \mid 1 \le i \le m, 1 \le j \le \mathbb{V}\}$. At every i-th non-stuttering step, marked by y_i, the sequence $z_1^j, \ldots, z_m^j$ must hold over m of the following $\mathbb{V}$ non-stuttering steps; moreover, between each z_j^i and its preceding z_{j-1}^i there must be no more than $d_i - d_{i-1}$ non-stuttering steps, unless $d_i - d_{i-1} > \mathbb{V} - i + 1$. After defining $\delta_i \triangleq d_i - d_{i-1}$ if $d_i - d_{i-1} \le \mathbb{V} - i + 1$ and

$\delta_i \triangleq *$ otherwise, the qualitative extended Pnueli modalities capture precisely this behavior.

$$\varkappa\langle z, \mathsf{Pn}\rangle \quad \triangleq \quad \bigwedge_{1\leq i\leq \mathbb{V}} \mathsf{G}\left(y_i \Rightarrow \mathsf{exqPn}_m^{\mathbb{V};\langle\delta_1,\ldots,\delta_m\rangle}\left(z_1^i,\ldots,z_m^i\right)\right) \tag{26}$$

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

$$\varkappa\langle z, \mathsf{U}\rangle \triangleq \left(\bigwedge_{h,j\neq 1} \neg z_h^j\right) \mathsf{U}\, z_1^1 \wedge \bigwedge_{\substack{1\leq j\leq \mathbb{V}\\1\leq i\leq m}} \mathsf{G}\left(z_i^j \Rightarrow \begin{pmatrix} \neg z^{1+(j \bmod \mathbb{V})}_{1+(i \bmod m)} \\ \wedge \bigwedge_{h\neq i} \neg z_h^j \end{pmatrix} \mathsf{U}\, z^j_{1+(i \bmod m)}\right) \tag{27}$$

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 non-stuttering step and before the $j+1$-th, then π_i has to hold at the k-th position in the sequence $z_1^j,\ldots,z_m^j$, with $k = D(i)$.

$$\varkappa\langle x, \pi\rangle \quad \triangleq \quad \bigwedge_{\substack{1\leq i\leq M\\1\leq j\leq \mathbb{V}}} \mathsf{G}\begin{pmatrix} x_i \wedge y_j \Rightarrow \neg z^j_{D(i)}\, \mathsf{U}\, z^j_{D(i)} \wedge \pi_i \\ \wedge \\ \neg x_i \wedge y_j \Rightarrow \neg z^j_{D(i)}\, \mathsf{U}\, z^j_{D(i)} \wedge \neg\pi_i \end{pmatrix} \tag{28}$$

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

$$\phi' \quad \triangleq \quad \kappa \wedge \Upsilon(\mathcal{Q}) \wedge \bigwedge_{p\in\mathcal{Q}} \mathsf{G}\begin{pmatrix} \curlywedge(p)\wedge \\ \curlywedge(\neg p)\end{pmatrix} \wedge \varkappa\langle y\rangle \wedge \varkappa\langle z,\mathsf{Pn}\rangle \wedge \varkappa\langle z,\mathsf{U}\rangle \wedge \varkappa\langle x,\pi\rangle \tag{29}$$

Example 12. In the elections example, formulas $\varkappa\langle y\rangle$, $\varkappa\langle z, \mathsf{U}\rangle$, and $\varkappa\langle x, \pi\rangle$ are instantiated with $\mathbb{V} = \mathbb{V}_\Omega = 6$ and $m = m_\Omega = 4$. To instantiate $\varkappa\langle z, \mathsf{Pn}\rangle$, consider the following table.

i	$\mathbb{V}_\Omega - i + 1$	d_i	$d_i - d_{i-1}$	δ_i
1	6	1	1	1
2	5	40	39	$*$
3	4	41	1	1
4	3	1460	1419	$*$

Lemma 13. *ϕ' and ϕ are equi-satisfiable over words in* $\mathbf{var}(\mathcal{Q}, \mathbb{V}/D)$.

Proof. The proof consists of two parts.

$\mathbf{SAT}(\phi) \Rightarrow \mathbf{SAT}(\phi')$. Let $w \in \mathbf{var}(\mathcal{Q}, \mathbb{V}/D)$ such that $w \models \phi$. w' adds propositions s, y_i, z_i^j, constrained as follows. s switches its truth value at every non-stuttering step, 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 non-stuttering step signaled by s. Whenever a given y_j holds, a sequence of z_i^j's hold over the following $\mathbb{V}$ non-stuttering step, 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 non-stuttering step 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^j_{i+i}. Once w' is built, the rest of the proof follows the lines of Lemma 5. It is clear that $w' \models \bigwedge_{p \in \mathcal{Q}} \curlyvee(\mathcal{P}) \wedge \mathsf{G}(\curlywedge(p) \wedge \curlywedge(\neg p))$ and $w' \models \kappa$. In addition, $w' \models \varkappa\langle y\rangle \wedge \varkappa\langle z, \mathsf{U}\rangle$ is a consequence of the set up of the y_j's and the z^j_i'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, $w, i \models \mathsf{X}^{\mathbb{D}(j)} \pi_j$ holds for all $j \in b$ and $w, i \models \mathsf{X}^{\mathbb{D}(k)} \neg \pi_k$ 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 $\mathbb{V}$ non-stuttering steps of item s over positions i to $i+D$. Let $i \leq t_1 < \cdots < t_{\mathbb{V}} \leq i+D$ be these transition instants. There are only stuttering steps between any such two consecutive t_i's, hence there exists a subset $u_1 < \cdots < u_m$ of the t_i's such that z^j_i 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 non-stuttering step. Because of how each z^j_i'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^j_i first holds; because i is generic, $w' \models \varkappa\langle x, \pi\rangle$ holds. Also, if $d_i - d_{i-1} \leq \mathbb{V} - i + 1$, there are no more than $d_i - d_{i-1}$ non-stuttering steps between u_{i-1} and u_i, for all $1 \leq i \leq m$ (and assuming $u_0 = d_0 = 0$); this establishes $w', i \models \mathsf{exqPn}_m^{\mathbb{V};\langle\delta_1,\ldots,\delta_m\rangle}\left(z^j_1, \ldots, z^j_m\right)$. In all, $w' \models \phi'$ holds.

SAT$(\phi') \Rightarrow$ **SAT**(ϕ). Let w' be a word in

$$\mathbf{var}(\mathcal{Q} \cup \{s\} \cup \{y_i \mid 1 \leq i \leq \mathbb{V}\} \cup \{z^j_i \mid 1 \leq i \leq m, 1 \leq j \leq \mathbb{V}\}, \}, \mathbb{V}/D)$$

such that $w' \models \phi'$. Let w be w' with all stuttering steps removed; $w \models \phi'$ as well from Lemma 10. 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 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 $i \leq t_1 < \cdots < t_{\mathbb{V}}$ be the following $\mathbb{V}$ non-stuttering steps of s — and hence of any proposition in $\mathcal{Q}$ as well, according to $\curlyvee(\mathcal{Q}) \wedge \bigwedge_{p\in\mathcal{Q}} \mathsf{G}(\curlywedge(p) \wedge \curlywedge(\neg p))$. $\varkappa\langle y\rangle$ implies that a unique y_j holds at i; correspondingly, $\varkappa\langle z, \mathsf{Pn}\rangle$ entails that there exists a subset of the $u_1 < \cdots < u_m$ of the sequence $t_1 < \cdots < t_{\mathbb{V}}$ such that z^j_k 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^j_g at $u_g + 1$. The inductive hypothesis implies that $u_{g-1} + 1 \leq i + d_{g-1} \leq u_g$, and $\varkappa\langle z, \mathsf{Pn}\rangle$ 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. Overall, induction guarantees that the finally modified w is such that $w \models \phi$. □

Lemma 14. *$|\phi'|$ is polynomial in* $|\mathcal{Q}|, \mathbb{V}, |\phi|_{\mathsf{U}}, |\phi|_{\mathsf{G}}, |\phi|_{\mathsf{P}}$.

Proof. The size of ϕ' is $|\kappa| + |\mathcal{Q}|^2 + |\varkappa\langle y\rangle| + |\varkappa\langle z, \mathsf{Pn}\rangle| + |\varkappa\langle z, \mathsf{U}\rangle| + |\varkappa\langle x, \pi\rangle|$, up to constant factors. Then, $|\kappa|$ is $|\phi|_{\mathsf{U}}$; $|\varkappa\langle y\rangle|$ is $\mathrm{O}(\mathbb{V}^3)$; $|\varkappa\langle z, \mathsf{Pn}\rangle|$ is $\mathrm{O}(m^2 \cdot \mathbb{V}^2)$, which is $\mathrm{O}(|\phi|^2_{\mathsf{G}} \cdot \mathbb{V}^2)$, then the qualitative extended Pnueli operators can be eliminated as shown in Section 6.1, getting an LTL formula of size still polynomial in $|\phi|_{\mathsf{G}} \cdot \mathbb{V} \cdot |\phi|_{\mathsf{P}}$; $|\varkappa\langle z, \mathsf{U}\rangle|$ is $\mathrm{O}(m^3 \cdot \mathbb{V}^2)$ which is $\mathrm{O}(|\phi|^3_{\mathsf{G}} \cdot \mathbb{V}^2)$; and $|\varkappa\langle x, \pi\rangle|$ is $\mathrm{O}(M \cdot \mathbb{V} \cdot (M + \mathbb{V} + |\phi|_{\mathsf{P}}))$, which is $\mathrm{O}(|\phi|^2_{\mathsf{G}} \cdot \mathbb{V}^2 \cdot |\phi|_{\mathsf{P}})$. □

Example 15. Consider the running elections example and transform Ω 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, non-stuttering steps are in bold and underlined, while a hat marks successors of non-stuttering steps.

$\underline{\mathbf{1}}$	$\widehat{2}$	3	4	5	$\underline{\mathbf{6}}$	$\widehat{\underline{\mathbf{7}}}$	$\widehat{8}$	9	$\underline{\mathbf{10}}$	$\widehat{11}$	$\underline{\mathbf{12}}$	$\widehat{\underline{\mathbf{13}}}$	$\widehat{14}$
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, for example the one in Example 4. For instance, the metric requirement of e occurring once at $1460 + 1 - 40 = 1421$ can be accommodate by removing all the stuttering steps at position 8 and by adding $1421 - 8 = 1413$ additional stuttering steps at position 2.

7 Future work

Future work will encompass implementation and generalizations. In particular, we plan to generalize the results to the case of subword stuttering, i.e., where a subword such as *abc* is repeated n times such as in $(abc)^n$ (cf. [KS05]). Also, adaptation of the technique to MTL will be considered. Finally, we will implement the reduction technique and assess its practicality experimentally.

References

[AH92] 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.

[BH10] Alan Burns and Ian J. Hayes. A timeband framework for modelling real-time systems. *Real-Time Systems*, Accepted for publication, 2010.

[BMOW08] 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.

[CCM+91] 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.

[DS02] Stéphane Demri and Philippe Schnoebelen. The complexity of propositional linear temporal logics in simple cases. *Information and Computation*, 174(1):84–103, 2002.

[Eme90] 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.

[Ete00] Kousha Etessami. A note on a question of Peled and Wilke regarding stutter-invariant LTL. *Information Processing Letters*, 75(6):261–263, 2000.

[FMMR10] 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.

[FR08] 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.

[Frä96] 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.

[GHR94] 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.

[GKP94] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A foundation for computer science*. Addison-Wesley, 1994.

[HR04] Yoram Hirshfeld and Alexander Moshe Rabinovich. Logics for real time: Decidability and complexity. *Fundamenta Informaticae*, 62(1):1–28, 2004.

[HR05] Yoram Hirshfeld and Alexander Moshe Rabinovich. Timer formulas and decidable metric temporal logic. *Information and Computation*, 198(2):148–178, 2005.

[KS05] Antonín Kučera and Jan Strejček. The stuttering principle revisited. *Acta Informatica*, 41(7/8):415–434, 2005.

[Lam83] 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.

[PW97] Doron Peled and Thomas Wilke. Stutter-invariant temporal properties are expressible without the next-time operator. *Information Processing Letters*, 63(5):243–246, 1997.

[Wil94] Thomas Wilke. Specifying timed state sequences in powerful decidable logics and timed automata. In Hans Langmaack, Willem P. de Roever, and Jan Vytopil, 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.