



Basic Research in Computer Science

BRICS RS-97-5 Etessami et al.: First-Order Logic with Two Variables and Unary Temporal Logic

First-Order Logic with Two Variables and Unary Temporal Logic

Kousha Etessami
Moshe Y. Vardi
Thomas Wilke

BRICS Report Series

ISSN 0909-0878

RS-97-5

March 1997

**Copyright © 1997, BRICS, Department of Computer Science
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**See back inner page for a list of recent BRICS Report Series publications.
Copies may be obtained by contacting:**

**BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK-8000 Aarhus C
Denmark
Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through the World Wide
Web and anonymous FTP through these URLs:**

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory RS/97/5/

First-Order Logic with Two Variables and Unary Temporal Logic*

Kousha Etessami[†] Moshe Y. Vardi[‡] Thomas Wilke[§]

Abstract

We investigate the power of first-order logic with only two variables over ω -words and finite words, a logic denoted by FO^2 . We prove that FO^2 can express precisely the same properties as linear temporal logic with only the unary temporal operators: “next”, “previously”, “sometime in the future”, and “sometime in the past”, a logic we denote by unary-TL. Moreover, our translation from FO^2 to unary-TL converts every FO^2 formula to an equivalent unary-TL formula that is at most exponentially larger, and whose operator depth is at most twice the quantifier depth of the first-order formula. We show that this translation is optimal.

While satisfiability for full linear temporal logic, as well as for unary-TL, is known to be PSPACE-complete, we prove that satisfiability for FO^2 is NEXP-complete, in sharp contrast to the fact that satisfiability for FO^3 has non-elementary computational complexity. Our NEXP time upper bound for FO^2 satisfiability has the advantage of being in terms of the *quantifier depth* of the input formula. It is obtained using a small model property for FO^2 of independent interest, namely: a satisfiable FO^2 formula has a model whose “size” is at most exponential in the quantifier depth of the formula. Using our translation from FO^2 to unary-TL we derive this small model property from a corresponding small model property for unary-TL. Our proof of the small model property for unary-TL is based on an analysis of unary-TL types.

*Part of the research reported here was conducted while the authors were visiting DIMACS as part of the Special Year on Logic and Algorithms.

[†]Basic Research in Computer Science (BRICS), Centre of the Danish National Research Foundation. This research supported by the ESPRIT Long Term Research Programme of the EU under project number 20244 (ALCOM-IT). email: kousha@brics.dk.

[‡]Department of Computer Science, Rice University. email: vardi@cs.rice.edu. During the preparation of this paper this author was partially supported by an NSF Grant.

[§]Institut für Informatik und Prakt. Math., Christian-Albrechts-Universität zu Kiel, Germany. email: tw@informatik.uni-kiel.de

1 Introduction

Over the past three decades a considerable amount of knowledge has accumulated regarding the relationship between first-order and temporal logic over both finite words and ω -words: the first-order expressible properties are exactly those expressible in temporal logic [Kam68, GPSS80, GHR94]; three variables suffice for expressing all the first-order expressible properties [Kam68, IK89]; while satisfiability for first-order logic with three variables has non-elementary computational complexity [Sto74], the satisfiability problem for temporal logic is PSPACE-complete [SC85]; moreover, there are classes of first-order formulas with three variables whose smallest equivalent temporal formulas require non-elementarily larger size (a consequence derivable from [Sto74]). In computer science the importance of this work stems from the practical relevance of temporal logic, which is used extensively today to specify and verify properties of reactive systems (see, e.g., [Pnu77] and [MP92]).

In this paper we provide a scaled down study of the relationship between first-order and temporal logic. Looking at first-order logic with only two variables, we show that the tight correspondence to temporal logic persists. We prove that first-order logic with two variables, denoted by FO^2 , has precisely the same expressive power as temporal logic with the usual future and past unary temporal operators: “next”, “previously”, “sometime in the future”, and “sometime in the past”, but without the binary operators “until” and “since”, a logic we denote by unary-TL. In other words, FO^2 coincides with the lowest level of the combined until/since hierarchy (which is known to be infinite [EW96]).

By contrast to the quite difficult proofs available for the correspondence between full first-order logic and temporal logic (cf., e.g., [Kam68, GPSS80, GHR94]), our proof that $\text{FO}^2 = \text{unary-TL}$ is an easily understood inductive translation. In fact, our proof yields the following much stronger assertions: (1) FO^2 formulas can be translated to equivalent unary-TL formulas that are at most exponentially larger and whose operator depth is at most twice the quantifier depth of the first-order formula, and (2) the translation can be carried out in time polynomial in the size of the output formula.

We show that our translation is essentially optimal by exhibiting a sequence of FO^2 formulas that require exponentially larger unary-TL formulas. Thus, while with just three variables there is already a non-elementary gap between the succinctness of first-order logic and full temporal logic, FO^2 remains more succinct than unary-TL but not nearly as much: an exponential blowup is exactly what is necessary in the worst-case.

The same result that shows that satisfiability for temporal logic is PSPACE-complete ([SC85]) also shows that satisfiability remains PSPACE-complete for unary-TL. We prove on the other hand that satisfiability for FO^2 is NEXP-complete. This again contrasts sharply with the non-elementary complexity of satisfiability for FO^3 . Moreover, this is surprising given that FO^2 is exponentially more succinct than unary-TL, and that satisfiability for unary-TL is PSPACE-complete, leading one to expect that FO^2 satisfiability will be EXPSpace-complete. Indeed, as a consequence of our NEXP bound it follows that FO^2 formulas that require “large” (exponentially bigger) unary-TL expressions necessarily have models that are “very small” (subexponential) with respect to the size of their unary-TL expression. Such “very small” models do not exist in general for unary-TL, as we can easily express with an $n^{O(1)}$ size unary-TL formula a “counter” whose smallest model has size 2^n .

An interesting and related aspect of our NEXP upper bound is that the time bound is only in terms of the *quantifier depth* of the FO^2 formula. This is because we prove our upper bound using an unusually strong small model property for FO^2 , one which states that every satisfiable FO^2 formula has a model whose “size” is at most exponential in the quantifier depth of the given formula, rather than the size of the entire formula. For large but shallow formulas the gap between these quantities can make a significant difference.

It should be noted here that in a recent result Grädel, Kolaitis, and Vardi [GKV97] have shown that satisfiability for two-variable first-order formulas over *arbitrary* relational structures is computable in NEXP time. Their results also rely on a small model property. They prove that every satisfiable two-variable formula over arbitrary structures has a finite model of size at most exponential in the size of the formula, improving on a previous doubly-exponential bound obtained by Mortimer [Mor74]. Despite the similarity between the statement of their result and ours, the two are essentially incompatible and neither result implies the other. The reasons for this are two-fold. First, our results hold over words, i.e., over a unary vocabulary with built-in ordering. In particular, unlike arbitrary structures, over words we do not have a *genuine* finite model property: with two variables one can say that for every position in the word there is a greater position. Secondly, our “small” model property (Theorem 4) shows that every satisfiable formula has a model whose “size” is bounded exponentially by the quantifier depth of the formula, whereas the small model property of [GKV97] depends on the size of the entire formula. Moreover, the proof techniques used in the two results are completely different.

Our proof of the “small” model property for FO^2 is facilitated by our transla-

tion. It is enough to prove the same small model property for unary-TL (in terms of operator depth instead of quantifier depth) because our translation from FO^2 to unary-TL at most doubles the quantifier/operator depth. The existence of small models for unary-TL is established by an analysis of unary-TL types; these types behave quite differently than types for temporal logic in general.

FO^2 provides built-in binary predicates for a total order and a successor relation (besides free unary predicates). As further evidence of the robust correspondence between first-order and temporal logic we show that even when FO^2 is further restricted by removing the successor predicate, the relationship to temporal logic still persists: the resulting logic has exactly the same power as temporal logic with temporal operators “sometime in the future” and “sometime in the past” only. Moreover, we determine the complexity of satisfiability for this further restricted first-order logic, and the corresponding temporal logic, as well as their difference in succinctness.

All our results hold both for finite words and ω -words with only minor technical changes. In this conference paper we only deal with the more interesting case of ω -words.

The paper is organized as follows. Section 2 introduces our notation and terminology. Section 3 presents the translation from FO^2 to unary-TL and shows it is optimal. Section 4 establishes NEXP-completeness of satisfiability for FO^2 . In Section 5, we establish the small model property. Section 6 is concerned with FO^2 without “successor” and unary-TL without “next” and “previously”. We conclude in Section 7.

2 Terminology and Notation

We work with first-order logic over ω -words, where our vocabulary contains unary predicates from $\rho_m = \{P_0, P_1, P_2, \dots, P_{m-1}\}$ for some m , and in addition contains the built-in predicates “suc” for “successor” and “<” for “less than”. We use FO^2 to denote the class of properties definable by first-order formulas $\varphi(x)$, where at most x occurs free, and where at most 2 variables occur in all of φ . We will also informally use FO^2 to refer to the set of such formulas. An FO^2 formula $\varphi(x)$ naturally defines a property of ω -words over the alphabet whose symbols are subsets of $\sigma_m = \{p_0, \dots, p_{m-1}\}$, namely the property $\{\alpha \in (2^{\sigma_m})^\omega \mid \alpha \models \varphi[0]\}$. Here $\alpha \models P_j(i)$ iff $p_j \in \alpha_i$ for $i \geq 0, j < m$.

We use unary-TL to denote the class of properties of ω -words definable by linear temporal logic formulas built from atomic propositions from σ_m , using

the boolean connectives and the unary temporal operators \oplus (“next”), \ominus (“previously”), \diamond (“eventually” or “sometime in the future”), and \diamondleftarrow (“sometime in the past”). We also use unary-TL to denote the set of such formulas. By convention, $(\alpha, i) \models \diamond\varphi$ if there exists a position j strictly greater than i such that $(\alpha, j) \models \varphi$. The same applies to \diamondleftarrow .

Formulas φ and ψ from FO^2 over ρ_m or unary-TL over σ_m are said to be *equivalent* if $(\alpha, i) \models \varphi$ iff $(\alpha, i) \models \psi$ for all $\alpha \in (2^{\sigma_m})^\omega$ and $i \geq 0$. A formula φ from FO^2 over ρ_m or unary-TL over σ_m is said to be *satisfiable* if the property over 2^{σ_m} it defines is non-empty. Let the *atomic order formulas* be $x = y$, $\text{suc}(x, y)$, $\text{suc}(y, x)$, $x < y$, and $y < x$.

The length of a formula φ is denoted by $|\varphi|$, its quantifier (operator) depth by $\text{qdp}(\varphi)$ (respectively $\text{odp}(\varphi)$).

3 Unary-TL = FO^2

We prove that the logics unary-TL and FO^2 are equally expressive. The non-trivial direction of this fact follows from the following much stronger statement:

Theorem 1 *Every FO^2 formula $\varphi(x)$ can be converted to an equivalent unary-TL formula φ' with $|\varphi'| \in 2^{\mathcal{O}(|\varphi|(\text{qdp}(\varphi)+1))}$ and $\text{odp}(\varphi') \leq 2 \text{qdp}(\varphi)$. Moreover, the translation is computable in time polynomially in $|\varphi'|$.*

The reverse translation is trivial, and linear in both size and operator/quantifier depth. Note the contrast between the theorem and what follows from the work in [Sto74]: there is a non-elementary lower bound in terms of blow-up in size for any translation of first-order formulas with three variables into temporal formulas.

Proof of Theorem 1. Given an FO^2 formula $\varphi(x)$ the translation procedure works as follows. When $\varphi(x)$ is atomic, i. e., of the form $P_i x$, it outputs p_i . When $\varphi(x)$ is of the form $\psi_1 \vee \psi_2$ or $\neg\psi$ —we say that $\varphi(x)$ is *composite*—it recursively computes ψ'_1 and ψ'_2 , or ψ' and outputs $\psi'_1 \vee \psi'_2$ or $\neg\psi'$. The two cases that remain are when $\varphi(x)$ is of the form $\exists x\varphi^*(x)$ or $\exists y\varphi^*(x, y)$. In both cases, we say that $\varphi(x)$ is *existential*. In the first case, $\varphi(x)$ is equivalent to $\exists y\varphi^*(y)$ and, viewing x as a dummy free variable in $\varphi^*(y)$, this reduces to the second case.

In the second case, we can rewrite $\varphi^*(x, y)$ in the form

$$\begin{aligned} \varphi^*(x, y) = & \beta(\chi_0(x, y), \dots, \chi_{r-1}(x, y), \\ & \xi_0(x), \dots, \xi_{s-1}(x), \zeta_0(y), \dots, \zeta_{t-1}(y)) \end{aligned}$$

where β is a propositional formula, each formula χ_i is an atomic order formula, each formula ξ_i is an atomic or existential FO² formula with $\text{qdp}(\xi_i) < \text{qdp}(\varphi)$, and each formula ζ_i is an atomic or existential FO² formula with $\text{qdp}(\zeta_i) < \text{qdp}(\varphi)$.

In order to be able to recurse on subformulas of φ we have to separate the ξ_i 's from the ζ_i 's. We first introduce a case distinction on which of the subformulas ξ_i hold or not. Let T and F denote true and false. We obtain the following equivalent formulation for φ :

$$\bigvee_{\vec{\gamma} \in \{\text{T}, \text{F}\}^s} \left(\bigwedge_{i < s} (\xi_i \leftrightarrow \gamma_i) \wedge \exists y \beta(\chi_0, \dots, \chi_{r-1}, \gamma_0, \dots, \gamma_{s-1}, \zeta_0, \dots, \zeta_{t-1}) \right)$$

We proceed by a case distinction on which order relation holds between x and y . We consider five mutually exclusive cases, determined by the following formulas, which we call *order types*: $x = y$, $\text{suc}(x, y)$, $\text{suc}(y, x)$, $x < y \wedge \neg \text{suc}(x, y)$, $y < x \wedge \neg \text{suc}(y, x)$. When we assume that one of these order types is true, each atomic order formula evaluates to either T or F, in particular, each of the χ_i 's evaluates to either T or F; we will denote this truth value by χ_i^τ . We can finally rewrite φ as follows, where Υ stands for the set of all order types:

$$\bigvee_{\vec{\gamma} \in \{\text{T}, \text{F}\}^s} \left(\bigwedge_{i < s} (\xi_i \leftrightarrow \gamma_i) \wedge \bigvee_{\tau \in \Upsilon} \exists y (\tau \wedge \beta(\chi_0^\tau, \dots, \chi_{r-1}^\tau, \vec{\gamma}, \vec{\zeta})) \right)$$

Notice now the following. If τ is an order type, $\psi(x)$ an FO² formula, and ψ' an equivalent unary-TL formula, there is an obvious way to obtain a unary-TL formula $\tau[\psi]$ equivalent to $\tau \wedge \psi(y)$, see the following table:

τ	$\tau[\psi]$
$x = y$	ψ'
$\text{suc}(x, y)$	$\oplus \psi'$
$\text{suc}(y, x)$	$\ominus \psi'$
$x < y \wedge \neg \text{suc}(x, y)$	$\oplus \diamond \psi'$
$y < x \wedge \neg \text{suc}(y, x)$	$\ominus \diamond \psi'$

Our procedure will therefore recursively compute ξ'_i for $i < s$ and ζ'_i for $i < t$

and output

$$\bigvee_{\bar{\gamma} \in \{\mathbb{T}, \mathbb{F}\}^s} \left(\bigwedge_{i < s} (\xi'_i \leftrightarrow \gamma_i) \wedge \right. \quad (1)$$

$$\left. \bigvee_{\tau \in \mathcal{T}} \tau [\beta(\chi_0^\tau, \dots, \chi_{r-1}^\tau, \bar{\gamma}, \bar{\zeta})] \right) \quad (2)$$

Now we verify that $|\varphi'|$ and $\text{odp}(\varphi')$ are bounded as stated in the theorem. The proof is inductive on the quantifier depth of φ . The basic observation is that there are $2^s \leq 2^{|\varphi|}$ possibilities for $\bar{\gamma}$ in the disjunction of line 1 above, and by the inductive hypothesis each disjunct has length at most $c \cdot |\varphi| 2^{\mathcal{O}(|\varphi| \cdot \text{qdp}(\varphi))}$, for a constant c . The stated bound for $|\varphi'|$ follows by induction. That $\text{odp}(\varphi') \leq 2 \text{qdp}(\varphi)$ follows from the conversions in the table above.

It is straightforward to verify that our translation to φ' can be computed in time polynomial in $|\varphi'|$. ■

An exponential blow-up, as incurred in the translation of Theorem 1, is necessary:

Theorem 2 *There is a sequence $(\varphi_n)_{n \geq 1}$ of FO^2 sentences over ρ_1 of size $\mathcal{O}(n^2)$ such that the shortest temporal logic formulas equivalent to φ_n have size $2^{\Omega(n)}$.*

Proof. We give only a proof for an unbounded vocabulary; in this case the formulas φ_n can be chosen to be of size $\mathcal{O}(n)$.

The formula φ_n is a formula over ρ_{n+1} that defines the following property, denoted L_n : “any two positions that agree on p_0, \dots, p_{n-1} also agree on p_n ”. This is easily defined in FO^2 within size linear in n :

$$\varphi_n = \forall x \forall y \left(\bigwedge_{i < n} (P_i x \leftrightarrow P_i y) \rightarrow (P_n x \leftrightarrow P_n y) \right).$$

As every property (language) defined by any temporal logic formula φ (even with “until” or “since”) is recognized by a non-deterministic Büchi automaton with $2^{\mathcal{O}(|\varphi|)}$ states, see [VW94], it is enough to show that every Büchi automaton for L_n requires at least 2^{2^n} states.

Suppose \mathcal{A} recognizes L_n . Let a_0, \dots, a_{2^n-1} be any sequence of the 2^n symbols of the alphabet 2^{σ_n} . For every subset K of $\{0, \dots, 2^n - 1\}$ let w_K be the word $b_0 \dots b_{2^n-1}$ with $b_i = a_i$ if $i \in K$ and else $b_i = a_i \cup \{p_n\}$. Notice that there are 2^{2^n} such words. Also, $w_K \models \varphi$ and $w_K w_{K'} \not\models \varphi_n$ for $K \neq K'$. Therefore, if $K \neq K'$ and q_K and $q_{K'}$ are the states assumed by \mathcal{A} in accepting runs for w_K and $w_{K'}$ after 2^n steps, then q_K and $q_{K'}$ have to be distinct, i. e., \mathcal{A} needs at least 2^{2^n} states. ■

4 The Complexity of Satisfiability for FO^2

We now show that the satisfiability problem for FO^2 over ω -words is NEXP-complete. This is in sharp contrast to the non-elementary lower bound for satisfiability of first-order logic with three variables over words which follows from [Sto74]. Satisfiability for unary-TL remains, as with full TL, PSPACE-complete [SC85].

Theorem 3 *Satisfiability for an FO^2 formula φ over ρ_m is decidable in non-deterministic time $2^{\mathcal{O}(\text{qdp}(\varphi)^2 m)}$, and thus satisfiability for FO^2 is in NEXP.*

As a main tool for our NEXP upper bound we prove a strong small model property for FO^2 which is of interest in its own right:

Theorem 4 *Every satisfiable FO^2 formula φ over ρ_m has a model of the form uv^ω , where both u and v have length bounded by $2^{\mathcal{O}(\text{qdp}(\varphi)^2 m)}$.*

We will prove Theorem 4 in Section 5.

The other ingredient in our NEXP upper bound is the following lemma, which allows us to find out, given strings u and v , whether uv^ω satisfies an FO^2 formula φ by just checking φ on the string uv^{2d+1} , where d is the quantifier depth of φ .

Lemma 1 *Let $\varphi(x)$ be an FO^2 formula, and let u and v be words with $|v| > 2$, and let $d = \text{qdp}(\varphi)$.*

1. *For $r \geq 0$ and $0 \leq s < |v|$,*

$$uv^\omega \models \varphi[|u| + 2d|v| + s] \text{ iff } uv^\omega \models \varphi[|u| + (2d + r)|v| + s] \quad (3)$$

2. *In particular, if $\varphi(x) = \exists y \varphi^*(x, y)$ and $uv^\omega \models \varphi[i]$ with $i < |uv^{2d+1}|$, then there exists $j \leq |uv^{2d+3}|$ such that $uv^\omega \models \varphi^*[i, j]$.*

Proof. Part 2 follows from the proof of part 1. The proof for part 1 is by induction on the quantifier depth d .

Base case: When $\varphi(x)$ is quantifier free, the only thing we can say about the only variable x is which predicates hold at x , and clearly the predicates that hold at a position $j = |uv| + q|v| + r$ are exactly those that hold at $|u| + r$ (simply because we are at the same position in the word v).

Inductive case: Assume true for d , we prove the assertion for $d + 1$. Our formula $\varphi(x)$ of depth $d + 1$ is a boolean combination of formulas $\varphi'(x)$ of the form:

$$\exists y \beta(\chi_1, \dots, \chi_l, \psi_1(x), \dots, \psi_m(x), \gamma_1(y), \dots, \gamma_c(y))$$

where β denotes a boolean combination of the given formulas and each $\chi_i(x, y)$ is an atomic order relation (i.e., one of $x < y$, $\text{succ}(y, x)$, etc.). We will argue that part 1 holds for formulas of the form φ' and it will follow that it holds for φ as well because the “iff” in part 1 is preserved under boolean combination.

(\Leftarrow) Suppose $\varphi[j]$ holds for $j = |u| + (2(d + 1) + r)|v| + s$, where $r \geq 0$ and $0 \leq s \leq |v|$. Then there is a witness for y , namely a position k at which $\beta(\chi_1[j, k], \dots, \chi_l[j, k], \psi_1[j], \dots, \psi_m[j], \gamma_1[k], \dots, \gamma_c[k])$ holds. We consider several cases based on the location of k in uv^ω . Let $j'_{d+1} = |u| + 2(d + 1)|v| + s$. We want to show that $\varphi[j'_{d+1}]$ also holds.

1. $j \leq k$: In this case we know by the inductive hypothesis that j'_{d+1} satisfies the same ψ_i 's as j , and that $j'_{d+1} + (k - j)$ satisfies the same γ_i 's as k , and thus is a witness for j'_{d+1} just as k is for j , because their juxtaposition is exactly the same.
2. $|u| + (2d + 1)|v| \leq k < j$: In this case the exact same argument as in case 1 works, with the roles of k and j reversed.
3. $k < |u| + (2d + 1)|v|$: In this case, we can fix k as a witness for both j and j'_{d+1} because, given that $|v| > 2$, the order type of (k, j'_{d+1}) and (k, j) is the same.

(\Rightarrow) Suppose that $\varphi[j]$ holds for j where $|uv^{2d+2}| \leq j < |uv^{2d+3}|$. Then the claim is that $\varphi[j']$ holds for $j' = j + r|v|$ and for all r . This is again split into cases based on the location of the witness k .

1. $j \leq k$: But then $j + r|v|$ has a witness at $k + r|v|$.
2. $|u| + (2d + 1)|v| \leq k < j$: In this case again $j + r|v|$ has $k + r|v|$ as a witness.
3. $k < |u| + (2d + 1)|v|$: Now again as in the second case above k is a witness for both j and $j + r|v|$ because, given that $|v| > 2$, the order types of (k, j) and $(k, j + r|v|)$ are the same. ■

Proof of Theorem 3. The non-deterministic algorithm determines the satisfiability of an FO^2 formula $\varphi(x)$ over ρ_m as follows. It first guesses u and v of length bounded by $2^{\mathcal{O}(\text{qdp}(\varphi)^2 m)}$. It then builds up a table that contains for every $i < |uv^{2d+1}|$ and for every subformula $\psi(z)$ of $\varphi(x)$ a bit saying whether $uv^i \models \psi[z]$. This is done inductively. The entry for an atomic or composite (see proof of Theorem 1) ψ is easily determined. From Lemma 1, part 2, it follows that in order to determine whether or not an existential formula (see proof of Theorem 1) of the form $\exists y \beta(\bar{\chi}(x, y), \bar{\xi}(x), \bar{\zeta}(y))$ holds at a position $i < |uv^{2d+1}|$ it suffices to consider only positions $< |uv^{2d+3}|$ for y . Whether or not a formula $\zeta(y)$ holds at such a position can be determined by a lookup in the table according to (3). The algorithm outputs the entry for position 0 and $\varphi(x)$. ■

Now to conclude that FO^2 satisfiability is NEXP-complete, we observe that it is NEXP-hard. This follows from the work of [Le80, F84]. We sketch the reduction for completeness:

Theorem 5 ([Le80, F84]) *FO^2 satisfiability is NEXP-hard. In fact, even satisfiability for FO^2 over ρ_1 , as well as satisfiability for FO^2 formulas that do not use “suc” and “<” are NEXP-hard.*

Proof. We only sketch the proof for showing that satisfiability for FO^2 without either “suc” or “<” is NEXP-hard. We give a reduction from the problem of determining whether for a given tiling system $T \subseteq \{0, 1, \dots, c-1\}^4$ with c colors and a given initial row $x \in T^+$ of length n there exists a tiling of a $2^n \times 2^n$ square consistent with T and with x occurring in the lower left corner. (Recall that an element $\langle c_1, c_2, c_3, c_4 \rangle \in T$ is considered a square tile with left edge colored by c_1 , right edge colored by c_2 , etc. A tiling is consistent if adjacent edges carry the same color.) This problem is known to be NEXP-complete, see, e. g., [F84]. We can, with a short FO^2 formula, name the adjacent positions in a tiling (and check their consistency) by exploiting the fact that addition has poly-sized propositional formulae. The predicates are used to specify the address coordinates, as well as tile content, of positions in the tiling. ■

5 A small model property for FO^2

Theorem 1 tells us that every FO^2 formula of depth k can be translated into an equivalent unary-TL formula of depth $2k$. Thus in order to prove Theorem 4 it suffices to prove the same small model property for unary-TL, namely:

Theorem 6 *Every satisfiable unary-TL formula φ over σ_m has a model uv^ω where the sizes of u and v are bounded by $2^{\mathcal{O}(\text{odp}(\varphi)^2 m)}$.*

We mention without proof that such a small model property does not hold for temporal logic in general. In fact, one can prove a non-elementary lower bound on the size of small models of temporal logic formulas in terms of operator depth. We also mention that there is a family of satisfiable unary-TL formulas φ_n over σ_m of depth $\mathcal{O}(n)$ where the smallest models have size $\Omega(2^{mn})$.

We first sketch a proof of Theorem 6 and then go into details. Let $k, k' \geq 0$. We say that a unary-TL formula φ is of *depth (at most) (k, k')* if it is of depth (at most) k in \diamond and depth (at most) k' in \oplus . Given an ω -word w and a position $i \geq 0$, the (k, k') *type of i in w* , denoted $\tau_{k,k'}^w(i)$, is the set of all unary-TL formulas of depth at most (k, k') that hold in w at i . This means that $w \models \varphi$ iff $\varphi \in \tau_{k,k'}^w(0)$ for every formula φ of operator depth (k, k') . It is thus enough to show that for every ω -word w there exist u and v of size bounded by $2^{\mathcal{O}((k+k'+1)^2 m)}$ such that $\tau_{k,k'}^w(0) = \tau_{k,k'}^{w'}(0)$ for $w' = uv^\omega$. In order to establish this, we first show that for every ω -word w one can find u and v such that w and uv^ω agree on the types of position 0 and such that u and v are bounded polynomially in the number of types that occur in w . We then show that the number of types occurring in a given ω -word is bounded by $2^{\mathcal{O}((k+k'+1)^2 m)}$. (Notice that the number of $(0, 0)$ types occurring in any ω -word over σ_m is already in $\Omega(2^m)$.)

The following lemma establishes that the $(k+1, k')$ type of a position i in a given word w is determined uniquely by (1) i 's local neighborhood, (2) the (k, k') types that occur to its right, and (3) the (k, k') types that occur to its left.

Lemma 2 *Let w and w' be ω -words and $i, i' \geq 0$.*

1. $\tau_{0,k'}^w(i) = \tau_{0,k'}^{w'}(i')$ iff $w_{i-k'} \dots w_i \dots w_{i+k'} = w'_{i'-k'} \dots w'_{i'} \dots w'_{i'+k'}$ where, by convention, $w_j = \$$ and $w'_j = \$$ for $j < 0$ ($\$$ being a special symbol).
2. $\tau_{k+1,k'}^w(i) = \tau_{k+1,k'}^{w'}(i')$ if and only if $\tau_{0,k'}^w(i) = \tau_{0,k'}^{w'}(i')$, $\{\tau_{k,k'}^w(j) \mid j < i\} = \{\tau_{k,k'}^{w'}(j) \mid j < i'\}$, and $\{\tau_{k,k'}^w(j) \mid j > i\} = \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$.

Proof. Part 1 is clear: A depth k' formula that uses no \diamond operator can describe completely the content of the k' -neighborhood of the current position, and nothing more.

To prove part 2 we proceed by induction on k . The base case, $k = 0$, is immediate. Assume true for k , we prove the claim for $k+1$.

(\Rightarrow) If $\tau_{k+1,k'}^w(i) = \tau_{k+1,k'}^{w'}(i')$ then in particular (w, i) and (w', i') agree on all depth $(0, k)$ formulas, and thus $\tau_{0,k'}^w(i) = \tau_{0,k'}^{w'}(i')$.

To show $\{\tau_{k,k'}^w(j) \mid j > i\} \subseteq \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$, let $\tau' \in \{\tau_{k,k'}^w(j) \mid j > i\}$. There are only a bounded number of inequivalent formulas of depth (k, k') (this can be proved by a straightforward induction). Let $\Lambda_{k,k'}$ denote the set of such formulas. There is thus a formula

$$\gamma = \bigwedge_{\varphi \in \tau'} \varphi \wedge \bigwedge_{\psi \in \Lambda_{k,k'} \setminus \tau'} \neg \psi$$

which holds in a word precisely at those positions having (k, k') -type τ' . But then $\diamond \gamma$ is a depth $(k+1, k')$ formula that holds at (w', i') precisely when $\tau' \in \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$. A symmetric proof shows that $\{\tau_{k,k'}^w(j) \mid j > i\} \supseteq \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$, and thus $\{\tau_{k,k'}^w(j) \mid j > i\} = \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$. A similar proof shows that $\{\tau_{k,k'}^w(j) \mid j < i\} = \{\tau_{k,k'}^{w'}(j) \mid j < i'\}$.

(\Leftarrow) Assume that $\tau_{0,k'}^w(i) = \tau_{0,k'}^{w'}(i')$, $\{\tau_{k,k'}^w(j) \mid j < i\} = \{\tau_{k,k'}^{w'}(j) \mid j < i'\}$, and $\{\tau_{k,k'}^w(j) \mid j > i\} = \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$.

First observe that every unary-TL formula ψ of depth (k, k') is equivalent to a formula ψ' also of depth (k, k') , where ψ' is in a normal form where all \oplus have been “moved in”, i.e., appear without any \diamond operators in their scope. In other words, every unary-TL formula of depth (k, k') is equivalent to a boolean combination of formulas of depth (k, k') starting with \diamond , and formulas of depth $(0, k')$. Thus we can restrict our attention to normal form formulas.

Now, let ψ be a depth $(k+1, k')$ formula in normal form.

If the outermost connective of ψ is \oplus , then it is a depth $(0, k')$ formula. Thus, since by assumption $\tau_{0,k'}^w(i) = \tau_{0,k'}^{w'}(i')$, $\psi \in \tau_{k+1,k'}^w(i) \Leftrightarrow \psi \in \tau_{k+1,k'}^{w'}(i')$.

If the outermost connective of ψ is \diamond then $\psi = \diamond \gamma$. Now $(w, i) \models \psi$ iff there exists a $j > i$ such that $\gamma \in \tau_{k,k'}^w(j)$. Hence, since by assumption $\{\tau_{k,k'}^w(j) \mid j > i\} = \{\tau_{k,k'}^{w'}(j) \mid j > i'\}$, we have $\psi \in \tau_{k+1,k'}^w(i) \Leftrightarrow \psi \in \tau_{k+1,k'}^{w'}(i')$. The case when $\psi = \diamond \gamma$ is symmetric.

To conclude the proof, note that the “ \Leftrightarrow ” in the previous two paragraphs is preserved under boolean combination. \blacksquare

Using Lemma 2, we can now establish the following lemma which shows how to collapse ω -words in order to get “smaller” ω -words without changing the type structure of the ω -word in an essential way. In the following lemma k' will be fixed, and we adopt the shorthand notation τ_k^w for $\tau_{(k,k')}^w$.

Lemma 3 *Let $w = u_0 u_1 u_2 \dots$ be an ω -word.*

1. Assume i and j are positions such that $i < j$ and $\tau_k^w(i) = \tau_k^w(j)$. Let $w' = u_1 \dots u_i u_{j+1} u_{j+2} \dots$.
Then $\tau_k^{w'}(l) = \tau_k^w(l)$ for $l \leq i$ and $\tau_k^{w'}(l) = \tau_k^w(l + (j - i))$ for $l > i$.
2. Assume i and j are positions such that
 - (a) $i < j$,
 - (b) $\tau_k^w(i) = \tau_k^w(j)$,
 - (c) $\{\tau_k^w(l) \mid l \geq 0\} = \{\tau_k^w(l) \mid l < i\}$, and
 - (d) $\{\tau_k^w(l) \mid i \leq l < j\} = \{\tau_k^w(l) \mid \exists \omega l' (\tau_k^w(l) = \tau_k^w(l'))\}$.

Let $w' = u_0 \dots u_i (u_{i+1} \dots u_j)^\omega$.

Then $\tau_k^{w'}(l) = \tau_k^w(l)$ for $l \leq i$ and $\tau_k^{w'}(i + r(j - i) + s) = \tau_k^w(i + s)$ for $r \geq 0$ and $0 \leq s < j - i$.

Proof. We prove part 1 by induction on k . Base case, $k = 0$. When we cut out a piece of a word, we don't change any of the characters we didn't cut out, and moreover the characters in the k' -neighborhoods of a point remain the same, thus we don't change $(0, k')$ -types of any point.

Assume true for k . Suppose $\tau_{k+1}^w(i) = \tau_{k+1}^w(j)$. By part 2 of Lemma 2 it follows that

$$\{\tau_k^w(i + 1), \dots, \tau_k^w(j - 1)\} \subseteq \{\tau_k^w(m) \mid m < i\} \quad (4)$$

$$\{\tau_k^w(i + 1), \dots, \tau_k^w(j - 1)\} \subseteq \{\tau_k^w(m) \mid m > j\} \quad (5)$$

Let $\pi(l)$ be the mapping defined by:

$$\pi(l) = \begin{cases} l & \text{if } l \leq i \\ l + (j - i) & \text{otherwise} \end{cases}$$

By the inductive hypothesis we know that for all l , $\tau_k^{w'}(l) = \tau_k^w(\pi(l))$.

But then $\{\tau_k^{w'}(m) \mid m > l\} = \{\tau_k^w(\pi(m)) \mid m > l\} = \{\tau_k^w(m) \mid m > \pi(l)\}$, the last equality following from containment 5. Similarly, using containment 4, we have $\{\tau_k^{w'}(m) \mid m \leq l\} = \{\tau_k^w(\pi(m)) \mid m \leq l\} = \{\tau_k^w(m) \mid m \leq \pi(l)\}$. But then by part 2 of Lemma 2 we have $\tau_{k+1}^{w'}(l) = \tau_{k+1}^w(\pi(l))$, which is what we wanted to prove.

The proof of part 2 is again by induction on k . Base case, $k = 0$. For $l \leq i$, given that i and j have the same k' -neighborhood, the k' -neighborhood of position l in w' is the same as the k' neighborhood of l in w . Also, for $l = i + r(j - i) + s$,

by the same fact, l has the same k' -neighborhood as $i + s$. The base case then follows from part 1 of Lemma 2.

Suppose true for k , we prove the claim for $k + 1$. First note that $(k + 1, k')$ -types constitute a refinement of (k, k') -types, meaning that two positions with the same $(k + 1, k')$ -type have the same (k, k') -type. Thus, given that (a) through (d) hold for $k + 1$, by the inductive hypothesis we know that $\tau_k^{w'}(l) = \tau_k^w(l)$ for $l \leq i$ and $\tau_k^{w'}(i + r(j - i) + s) = \tau_k^w(i + s)$ for $r \geq 0$ and $0 \leq s < j - i$.

Thus, in particular, we claim that for $l \leq i$, $\{\tau_k^{w'}(m) \mid m > l\} = \{\tau^w(m) \mid m > l\}$. This is so because by (d) the infinitely recurring $(k + 1, k')$ -types (and thus also infinitely recurring (k, k') -types) are precisely those that already occur at $i, i + 1, \dots, i + (j - i - 1) = j - 1$. In a similar way it follows that $\{\tau_k^{w'}(m) \mid m \leq l\} = \{\tau^w(m) \mid m \leq l\}$. Thus, by part 2 of Lemma 2, it follows that $\tau_{k+1}^w(l) = \tau_{k+1}^{w'}(l)$.

A similar proof shows that $\tau_{k+1}^{w'}(i + r(j - i) + s) = \tau_{k+1}^w(i + s)$, for $r \geq 0$ and $0 \leq s < (j - i)$. ■

From this lemma, we conclude:

Lemma 4 *Let w be an ω -word over 2^{σ_m} and t the number of (k, k') types occurring in w . There exists w' of the form uv^ω such that the length of u and v is less than $(t + 1)^2$ and such that $\tau_{k,k'}^w(0) = \tau_{k,k'}^{w'}(0)$.*

Proof. Part 2 of Lemma 3 immediately implies we can assume $w = uv^\omega$ for some u and v . We can also assume that u and v are chosen such that the assumptions of part 2 of Lemma 3 are given with $i = |u|$ and $j = |uv|$. Now, let u and v be such that for every other such pair u' and v' we have $|uv| \leq |u'v'|$. For contradiction, assume first $|v| \geq (t + 1)^2$. For every (k, k') type τ of a position s with $i \leq s < j$ pick a position i_τ such that $i \leq i_\tau < j$ and $\tau_{k,k'}^w(i_\tau) = \tau$. Since $|v| \geq (t + 1)^2$, we can find two positions l and l' carrying the same type such that $i \leq l < l' < j$ and either $i_\tau < l$ or $l' < i_\tau$ for each of the i_τ 's. Thus, by part 2 of Lemma 3, $u' = u$ and $v' = v_0v_1 \dots v_{l-|u|}v_{l'-|u|+1} \dots v_{|v|-|u|-1}$ would be a smaller pair. If $|u| \geq (t + 1)^2$ we obtain a contradiction in a similar way using part 1 of Lemma 3. ■

We now upper bound the number of types that can occur in a given ω -word:

Lemma 5 *The number of (k, k') types occurring in any ω -word over σ_m is at most $2^{3((2k'+1)(m+1)+1)}$, i. e., $|\{\tau_{k,k'}^w(i) \mid i \geq 0\}| \leq 2^{3((2k'+1)(m+1)+1)}$ for every $w \in (2^{\sigma_m})^\omega$.*

Proof. The proof is by induction on k . Let w be any ω -word over σ_m . Let $t_{(k,k')}$ be the number of (k, k') types occurring in w . For the base case, from Lemma 2, part 1, it is easy to see that $t_{(0,k')} \leq 2^{(2k'+1)(m+1)}$. Now observe that the sequence $(\{\tau_{k,k'}^w(i) \mid j < i\})_{i \geq 0}$ is an increasing sequence containing at most $t_{(k,k')}$ distinct elements. Similarly, the sequence $(\{\tau_{k,k'}^w(j) \mid j > i\})_{i \geq 0}$ is a decreasing sequence containing at most $t_{(k,k')} + 1$ distinct elements. Therefore, there are only $2t_{(k,k')} + 1$ many distinct pairs of the form $(\{\tau_{k,k'}^w(j) \mid j < i\}, \{\tau_{k,k'}^w(j) \mid j > i\})$, and thus, using Lemma 2, part 1, $t_{(k+1,k')} \leq (2t_{(k,k')} + 1)2^{(2k'+1)(m+1)}$, where, again, $2^{(2k'+1)(m+1)}$ accounts for the number of distinct $(0, k')$ types. The lemma follows by induction. ■

Theorem 6 now follows from Lemma 4 together with Lemma 5; this also proves Theorem 4. ■

We also prove a different small model property in terms of formula size (proof omitted):

- Theorem 7** 1. *Every satisfiable FO^2 formula φ has a model of size $2^{\mathcal{O}(|\varphi|)}$.*
 2. *Satisfiability for an FO^2 formula φ can also be decided in non-deterministic time $2^{\mathcal{O}(|\varphi|)}$.*

6 Further Restricting FO^2 and Unary-TL

It is only natural to further restrict FO^2 by allowing $<$ as the only built-in predicate; we denote the resulting logic by $\text{FO}^2[<]$. Correspondingly, one can consider the logical language that is obtained from unary-TL by disallowing the use of \oplus and \ominus ; we denote this language by $\text{TL}[\diamond]$.

Obviously, $\text{TL}[\diamond]$ can easily be translated into $\text{FO}^2[<]$. A slight modification of the translation from FO^2 to unary-TL described in the proof of Theorem 1 yields the reverse translation, i. e., $\text{TL}[\diamond] = \text{FO}^2[<]$. Corresponding to Theorem 1, we have:

Theorem 8 *Every FO^2 formula $\varphi(x)$ can be converted to an equivalent $\text{TL}[\diamond]$ formula φ' with $|\varphi'| \in 2^{\mathcal{O}(|\varphi| \text{ qdp}(\varphi))}$ and $\text{odp}(\varphi') \leq \text{qdp}(\varphi)$.*

As $\text{FO}^2[<]$ is a sublogic of FO^2 , the upper bounds for the complexity of the satisfiability for FO^2 carry over to $\text{FO}^2[<]$. Moreover, as we have seen in Theorem 5, even the hardness result carries over:

Theorem 9 *Satisfiability for $\text{FO}^2[<]$ is NEXP-complete. In fact, satisfiability for an $\text{FO}^2[<]$ formula φ over ρ_m is decidable in non-deterministic time $2^{\mathcal{O}(\text{qdp}(\varphi)m)}$ and also in non-deterministic time $2^{\mathcal{O}(|\varphi|)}$.*

Note the difference between $2^{\mathcal{O}(\text{qdp}(\varphi)^2m)}$ in Theorem 3 and $2^{\mathcal{O}(\text{qdp}(\varphi)m)}$ in the above theorem.

That satisfiability for $\text{FO}^2[<]$ is no less difficult than satisfiability for FO^2 (both are NEXP-complete) contrasts with what happens to satisfiability when passing from unary-TL to $\text{TL}[\diamond]$. In [SC85], it was shown that satisfiability for the temporal logic where the only temporal operator is “at present or sometime in the future” is in NP. We show that satisfiability for $\text{TL}[\diamond]$ (which now includes the past operator) remains in NP, and thus is NP-complete. This is obtained by proving a linear-size model property: for every satisfiable $\text{TL}[\diamond]$ formula φ there exist u and v with $|u|, |v| \leq |\varphi|$ such that $uv^\omega \models \varphi$ (proof omitted).

7 Conclusion

We have shown that the close correspondence between first-order and temporal logic over words persists when looking at first-order formulas with only two variables, and we have presented an easily understood translation of these formulas into temporal formulas. Our translation is optimal: the formulas incur at most an exponential blow-up in size and we have proved that this is necessary in the worst case.

The satisfiability problem for unary-TL is known to remain, as with full TL, PSPACE-complete, but we have shown that FO^2 satisfiability is drastically simpler than FO^3 satisfiability: the former is NEXP-complete, while the latter is known to require non-elementary complexity. Moreover, our NEXP upper bound for FO^2 satisfiability, and the corresponding small model properties for FO^2 and unary-TL, have the advantage of being only in terms of quantifier/operator depth and the number of propositions in the vocabulary, rather than the size of the entire formula, a fact that may be of potential use when dealing with large but shallow formulas.

Some remaining questions: (1) Given a regular language L (say, as a Büchi or finite automaton), can we decide whether L is FO^2 (and thus unary-TL) definable? [EW96, TW96b] obtained closely related results, but neither yields this particular fact. (2) Is the FO^2 quantifier alternation hierarchy strict? This question can also be phrased in terms of operator alternation in unary-TL. (3) Does satisfiability remain NEXP-hard for $\text{FO}^2[<]$ formulas (without successor) over a bounded

number of predicates? (4) Can the upper bound of the small model property for FO^2 be improved to $2^{\mathcal{O}(\text{qdp}(\varphi)+m)}$? This would make (the proof of) Theorem 7 obsolete.

References

- [EW96] K. Etessami and Th. Wilke. An Until hierarchy for temporal logic. In *11th Annual IEEE Symposium on Logic in Computer Science*, New Brunswick, New Jersey, pages 108–117, 1996.
- [Fü84] M. Fürer. The computational complexity of the unconstrained domino problem (with implications for logical decision problems). In *Logic and Machines: Decision Problems and Complexity*, pages 312–319. Volume 171 in Lect. Notes in Comput. Sci., Springer, 1984.
- [GHR94] D. M. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic*, volume 1. Clarendon Press, Oxford, 1994.
- [GPSS80] D. M. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. On the temporal analysis of fairness. In *Conference Record of the 7th Annual ACM Symposium on Principles of Programming Languages*, Las Vegas, Nev., pages 163–173, 1980.
- [GKV97] E. Grädel, Ph. G. Kolaitis, and M. Y. Vardi. On the Decision Problem for Two-Variable First-Order Logic. To appear in *Bulletin of the Assoc. for Symbolic Logic*.
- [IK89] N. Immerman and D. Kozen. Definability with bounded number of bound variables. *Information and Computation*, 83(2):121–139, 1989.
- [Kam68] J. A. W. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, University of California, Los Angeles, 1968.
- [Le80] H. R. Lewis. Complexity Results for Classes of Quantificational Formulas. *J. Comput. System Sci.*, 21: 317–353, 1980.
- [MP92] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*, Springer-Verlag, Berlin/New York, 1992.

- [Mor74] M. Mortimer. On languages with two variables. *Z. Math. Logik Grundlagen Math.*, 21:135-140, 1975.
- [Pnu77] A. Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, Providence, Rhode Island, pages 46–57, 1977.
- [SC85] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *J. Assoc. Comput. Mach.*, 32(3):733–749, 1985.
- [Sto74] L. J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, Department of Electrical Engineering, MIT, 1974.
- [TW96b] D. Thérien and Th. Wilke. Temporal logic and semidirect products: An effective characterization of the until hierarchy. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, Burlington, Vermont, pages 256–263, 1996.
- [VW94] M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation* 115: 1–37, 1994.

Recent BRICS Report Series Publications

- RS-97-5 Kousha Etessami, Moshe Y. Vardi, and Thomas Wilke. *First-Order Logic with Two Variables and Unary Temporal Logic*. March 1997. 18 pp. To appear in *Twelfth Annual IEEE Symposium on Logic in Computer Science, LICS '97 Proceedings*.
- RS-97-4 Richard Blute, Josée Desharnais, Abbas Edalat, and Prakash Panangaden. *Bisimulation for Labelled Markov Processes*. March 1997. 48 pp. To appear in *Twelfth Annual IEEE Symposium on Logic in Computer Science, LICS '97 Proceedings*.
- RS-97-3 Carsten Butz and Ieke Moerdijk. *A Definability Theorem for First Order Logic*. March 1997. 10 pp.
- RS-97-2 David A. Schmidt. *Abstract Interpretation in the Operational Semantics Hierarchy*. March 1997. 33 pp.
- RS-97-1 Olivier Danvy and Mayer Goldberg. *Partial Evaluation of the Euclidian Algorithm (Extended Version)*. January 1997. 16 pp. To appear in the journal *Lisp and Symbolic Computation*.
- RS-96-62 P. S. Thiagarajan and Igor Walukiewicz. *An Expressively Complete Linear Time Temporal Logic for Mazurkiewicz Traces*. December 1996. i+13 pp. To appear in *Twelfth Annual IEEE Symposium on Logic in Computer Science, LICS '97 Proceedings*.
- RS-96-61 Sergei Soloviev. *Proof of a Conjecture of S. Mac Lane*. December 1996. 53 pp. Extended abstract appears in Pitt, Rydeheard and Johnstone, editors, *Category Theory and Computer Science: 6th International Conference, CTCS '95 Proceedings*, LNCS 953, 1995, pages 59–80.
- RS-96-60 Johan Bengtsson, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, and Wang Yi. *UPPAAL in 1995*. December 1996. 5 pp. Appears in Margaria and Steffen, editors, *Tools and Algorithms for The Construction and Analysis of Systems: 2nd International Workshop, TACAS '96 Proceedings*, LNCS 1055, 1996, pages 431–434.
- RS-96-59 Kim G. Larsen, Paul Pettersson, and Wang Yi. *Compositional and Symbolic Model-Checking of Real-Time Systems*. December 1996. 12 pp. Appears in *16th IEEE Real-Time Systems Symposium, RTSS '95 Proceedings*, 1995.