

First-Order and Temporal Logics for Nested Words

Rajeev Alur

University of Pennsylvania

Marcelo Arenas

PUC Chile

Pablo Barceló

Universidad de Chile

Kousha Etessami

University of Edinburgh

Neil Immerman

University of Massachusetts

Leonid Libkin

University of Edinburgh

Abstract

Nested words are a structured model of execution paths in procedural programs, reflecting their call and return nesting structure. Finite nested words also capture the structure of parse trees and other tree-structured data, such as XML.

We provide new temporal logics for finite and infinite nested words, which are natural extensions of LTL, and prove that these logics are first-order expressively-complete. One of them is based on adding a “within” modality, evaluating a formula on a subword, to a logic CaRet previously studied in the context of verifying properties of recursive state machines. The other logic is based on the notion of a summary path that combines the linear and nesting structures. For that logic, both model-checking and satisfiability are shown to be EXPTIME-complete.

Finally, we prove that first-order logic over nested words has the three-variable property, and we present a temporal logic for nested words which is complete for the two-variable fragment of first-order.

1 Introduction

An execution of a procedural program can reveal not just a linear sequence of program states encountered during the execution, but also the correspondence between each point during the execution at which a procedure is called and the point when we return from that procedure call. This leads naturally to the notion of a finite or infinite nested word ([4, 3, 2]). A nested word is simply a finite or ω -word supplied with an additional binary matching relation which relates corresponding call and return points (and of course satisfies “well-bracketing” properties). Finite nested words offer an alternative way to view any data which has both a sequential string structure as well as a tree-like hierarchical structure. Examples of such data are XML documents and parse trees.

Pushdown systems (PDSs), Boolean Programs, and Recursive State Machines (RSMs), are equivalent abstract models of procedural programs, with finite data abstraction but unbounded call stack. Software model checking technology is by now thoroughly developed for checking ω -regular properties of runs for these models, when the runs are viewed as ordinary words (see [5, 8, 1]). Unfortunately, temporal logic and ω -regular properties over ordinary words are inadequate for expressing a variety of properties of program executions that are useful in interprocedural program analysis and software verification. These include Hoare-like pre/post conditions on procedures, stack inspection properties, and other useful program analysis properties that go well beyond ω -regular (see [2] for some examples). On the other hand, many such program analysis properties can easily be expressed when runs are viewed as nested words. Runs of Boolean Programs and RSMs can naturally be viewed as nested words once we add “summary edges” between matching calls and returns, and we can thus hope to extend model checking technology for procedural programs using richer temporal logics over nested words which remain tractable for analysis.

These considerations motivated the definition of Visibly Pushdown Languages (VPLs) [3] and the call-return temporal logic CaRet [2]. CaRet is a temporal logic over nested words which extends LTL with new temporal operators that allow for navigation through a nested word both via its ordinary sequential structure, as well as its matching call-return summary structure. The standard LTL model checking algorithms for RSMs and PDSs can be extended to allow model checking of CaRet, with essentially the same complexity [2]. VPLs [3] are a richer class of languages that capture MSO-definable properties of nested words. Recently, results about VPLs have been recast in light of nested words, and in particular in terms of Nested Word Automata [4] which offer a machine acceptor for (ω -regular nested words, with all the expected closure properties.

Over ordinary words, LTL has long been considered the

temporal logic of choice for program verification, not only because its temporal operators offer the right abstraction for reasoning about events over time, but because it provides a good balance between expressiveness (first-order complete), conciseness (can be exponentially more succinct compared to automata), and the complexity of model-checking (linear time in the size of the finite transition system, and PSPACE in the size of the temporal formula).

This raises the question: *What is the right temporal logic for nested words?*

The question obviously need not have a unique answer, particularly since nested words can arise in various application domains: for example, program verification, as we already discussed, or navigation and querying XML documents under “sequential” representation (see, e.g., [27]). However, it is reasonable to hope that any good temporal logic for nested words should possess the same basic qualities that make LTL a good logic for ordinary words, namely: (1) *first-order expressive completeness*: LTL has the same expressive power as first-order logic over words, and we would want the same over nested words; (2) *reasonable complexity for model checking and satisfiability*; and (3) *nice closure properties*: LTL is closed under boolean combinations including negation without any blow-up, and we would want the same for a logic over nested words. Finally (and perhaps least easy to quantify), we want (4) *natural temporal operators with simple and intuitive semantics*.

Unfortunately, the logic CaRet appears to be deficient with respect to some of these criteria: although it is easily first-order expressible, proving incompleteness – a widely believed conjecture – appears to be quite difficult. Also, some temporal operators in CaRet (such as the past-time call modalities), motivated by program analysis, may not be viewed as particularly natural in other applications. There is much related work in the XML community on logics for trees (see, e.g., surveys [14, 15, 28]), but they tend to have different kinds of deficiency for our purposes: they concentrate on the hierarchical structure of the data and largely ignore its linear structure; also, they are designed for finite trees.

We introduce and study new temporal logics over nested words. The main logic we consider, *Nested Word Temporal Logic* (NWTL) extends LTL with both a future and past variant of the standard Until operator, which is interpreted over *summary paths* rather than the ordinary linear sequence of positions. A summary path is the unique shortest directed path one can take between a position in a run and some future position, if one is allowed to use both successor edges and matching call-return summary edges. We show that NWTL possesses all the desirable properties we want from a temporal logic on nested words. In particular, it is both first-order expressively complete and has good model checking complexity. Indeed we provide

a tableaux construction which translates an NWTL formula into a Nested Word Automaton, enabling the standard automata theoretic approach to model checking of Boolean Programs and RSMs with complexity that is polynomial in the size the model and EXPTIME in the size of the formula.

We then explore some alternative temporal logics, which extend variants of CaRet with variants of unary “Within” operators proposed in [2], and we show that these extensions are also FO-complete. However, we observe that the model checking and satisfiability problems for these logics are 2EXPTIME-complete. These logics are – provably – more concise than NWTL, but we pay for conciseness with added complexity.

It follows from our proof of FO-completeness for NWTL that over nested words, every first-order formula with one free variable can be expressed using only 3 variables. More generally, we show, using EF games, that 3 variables suffice for expressing any first order formula with two or fewer free variables, similarly to the case of words [13] or finite trees [19]. Finally, we show that a natural unary temporal logic over nested words is expressively complete for first-order logic with 2 variables, echoing a similar result known for unary temporal logic over ordinary words [9].

Related Work. VPLs and nested words were introduced in [3, 4]. The logic CaRet was defined in [2] with the goal of expressing and checking some natural non-regular program specifications. The theory of VPLs and CaRet has been recast in light of nested words in [4]. Other aspects of nested words (automata characterizations, games, model-checking) were further studied in [1, 4, 2, 16]. It was also observed that nested words are closely related to a sequential, or “event-based” API for XML known as SAX [24] (as opposed to a tree-based DOM API [7]). SAX representation is very important in streaming applications, and questions related to recognizing classes of nested words by the usual word automata have been addressed in [27, 6].

While finite nested words can indeed be seen as XML documents under the SAX representation, and while much effort has been spent over the past decade on languages for tree-structured data (see, e.g. [14, 15, 28] for surveys), adapting the logics developed for tree-structured data is not as straightforward as it might seem, even though from the complexity point of view, translations between the DOM and the SAX representations are easy [26]. The main problem is that most such logics rely on the tree-based representation and ignore the linear structure, making the natural navigation through nested words rather unnatural under the tree representation. Translations between DOM and SAX are easy for first-order properties, but verifying navigational properties expressed in first-order is necessarily non-elementary even for words if one wants to keep the data complexity linear [10]. On the other hand, logics for XML tend to have good model-checking properties (at least

in the finite case), typically matching the complexity of LTL [11, 21]. We do employ such logics (e.g., those in [18, 19, 25]) in the proof of the expressive completeness of NWTL, first by using syntactic translations that reconcile both types of navigation, and then by combining them with a composition game argument that extends the result to the infinite case, which is not considered in the XML setting. This, however, involves a nontrivial amount of work. Furthermore, “within” operators do not have any natural analog on trees, and the proof for them is done by a direct composition argument on nested words.

Organization. Basic notations are given in Section 2. Section 3 defines temporal logics on nested words, and Section 4 presents expressive completeness results. We study model-checking in Section 5, and in Section 6 we prove the 3-variable property and present a logic for the 2-variable fragment. Due to space limitations, proofs are only sketched here.

2 Nested Words

A *matching* on \mathbb{N} or an interval $[1, n]$ of \mathbb{N} consists of a binary relation μ and two unary relations call and ret , satisfying the following: (1) if $\mu(i, j)$ holds then $\text{call}(i)$ and $\text{ret}(j)$ and $i < j$; (2) if $\mu(i, j)$ and $\mu(i, j')$ hold then $j = j'$ and if $\mu(i, j)$ and $\mu(i', j)$ hold then $i = i'$; (3) if $i \leq j$ and $\text{call}(i)$ and $\text{ret}(j)$ then there exists $i \leq k \leq j$ such that either $\mu(i, k)$ or $\mu(k, j)$.

Let Σ be a finite alphabet. A *finite nested word* of length n over Σ is a tuple $\bar{w} = (w, \mu, \text{call}, \text{ret})$, where $w = a_1 \dots a_n \in \Sigma^*$, and $(\mu, \text{call}, \text{ret})$ is a matching on $[1, n]$. A *nested ω -word* is a tuple $\bar{w} = (w, \mu, \text{call}, \text{ret})$, where $w = a_1 \dots \in \Sigma^\omega$, and $(\mu, \text{call}, \text{ret})$ is a matching on \mathbb{N} .

We say that a position i in a nested word \bar{w} is a *call* position if $\text{call}(i)$ holds; a *return* position if $\text{ret}(i)$ holds; and an *internal* position if it is neither a call nor a return. If $\mu(i, j)$ holds, we say that i is the matching call of j , and j is the matching return of i , and write $c(j) = i$ and $r(i) = j$. Calls without matching returns are *pending* calls, and returns without matching calls are *pending* returns. A nested word is said to be *well-matched* if no calls or returns are pending. Note that for well-matched nested words, the unary predicates call and ret are uniquely specified by the relation μ .

A nested word $\bar{w} = (w, \mu, \text{call}, \text{ret})$ is represented as a first-order structure: $\langle U, (P_a)_{a \in \Sigma}, <, \mu, \text{call}, \text{ret} \rangle$, where U is $\{1, \dots, n\}$ if w is a finite word of length n and \mathbb{N} if \bar{w} is a nested ω -word; $<$ is the usual ordering, P_a is the set of positions labeled a , and $(\mu, \text{call}, \text{ret})$ is the matching relation. When we talk about first-order logic (FO) over nested words, we assume FO over such structures.

For a nested word \bar{w} , and two elements i, j of \bar{w} , we denote by $\bar{w}[i, j]$ the substructure of \bar{w} (i.e. a finite nested

word) induced by elements ℓ such that $i \leq \ell \leq j$. If $j < i$ we assume that $\bar{w}[i, j]$ is the empty nested word. For nested ω -words \bar{w} , $\bar{w}[i, \infty]$ denotes the substructure induced by elements $\ell \geq i$. When this is clear from the context, we do not distinguish references to positions in subwords $\bar{w}[i, j]$ and \bar{w} itself, e.g. we shall often write $(\bar{w}[i, j], i) \models \varphi$ to mean that φ is true at the first position of $\bar{w}[i, j]$.

3 Temporal Logics over Nested Words

We now describe our approach to temporal logics for nested words. It is similar to the approach taken by the logic CaRet [2]. Namely, we shall consider LTL-like logics that define the next/previous and until/since operators for various types of paths in nested words.

All the logics will be able to refer to propositional letters, including the base unary relations call and ret , and will be closed under all Boolean combinations. We shall write \top for true and \perp for false. For all the logics we shall define the notion of satisfaction with respect to a position in a nested word, writing $(\bar{w}, i) \models \varphi$ when the formula φ is true in the position i of the word \bar{w} .

Since nested words are naturally represented as transition systems with two binary relations – the successor and the matching relation – in all our logics we introduce *next operators* \bigcirc and \bigcirc_μ . The semantics of those is standard: $(\bar{w}, i) \models \bigcirc\varphi$ iff $(\bar{w}, i+1) \models \varphi$, $(\bar{w}, i) \models \bigcirc_\mu\varphi$ iff i is a call with a matching return j (i.e. $\mu(i, j)$ holds) and $(\bar{w}, j) \models \varphi$. Likewise, we shall have *past operators* \ominus and \ominus_μ : that is, $\ominus\varphi$ is true in position $i > 1$ iff φ is true in position $i - 1$, and $\ominus_\mu\varphi$ is true in position j if j is a return position with matching call i and φ is true at i .

The *until/since operators* depend on what a path is. In general, there are various notions of paths through a nested word. We shall consider until/since operators for paths that are unambiguous: that is, for every pair of positions i and j with $i < j$, there could be at most one path between them. Then, with respect to any such given notion of a path, we have the until and since operators with the usual semantics:

- $(\bar{w}, i) \models \varphi\mathbf{U}\psi$ iff there is a position $j \geq i$ and a path $i = i_0 < i_1 < \dots < i_k = j$ between them such that $(\bar{w}, j) \models \psi$ and $(\bar{w}, i_p) \models \varphi$ for every $0 \leq p < k$.
- $(\bar{w}, i) \models \varphi\mathbf{S}\psi$ iff there is a position $j \leq i$ and a path $j = i_0 < i_1 < \dots < i_k = i$ between them such that $(\bar{w}, j) \models \psi$ and $(\bar{w}, i_p) \models \varphi$ for every $0 < p \leq k$.

The approach of CaRet was to introduce three types of paths, based on the linear successor (called *linear paths*), the call-return relation (called *abstract paths*), and the innermost call relation (called *call paths*).

To define those, we need the notions $\mathcal{C}(i)$ and $\mathcal{R}(i)$ for each position i – these are the innermost call within which

the current action i is executed, and its corresponding return. Formally, $\mathcal{C}(i)$ is the greatest matched call position $j < i$ whose matching return is after i (if such a call position exists), and $\mathcal{R}(i)$ is the least matched return position $\ell > i$ whose matching call is before i .

Definition 3.1 (Linear, call and abstract paths) Given positions $i < j$, a sequence $i = i_0 < i_1 < \dots < i_k = j$ is

- a linear path if $i_{p+1} = i_p + 1$ for all $p < k$;
- a call path if $i_p = \mathcal{C}(i_{p+1})$ for all $p < k$;
- an abstract path if

$$i_{p+1} = \begin{cases} r(i_p) & \text{if } i_p \text{ is a matched call} \\ i_p + 1 & \text{otherwise.} \end{cases}$$

We shall denote until/since operators corresponding to these paths by \mathbf{U}/\mathbf{S} for linear paths, $\mathbf{U}^c/\mathbf{S}^c$ for call paths, and $\mathbf{U}^a/\mathbf{S}^a$ for abstract paths.¹

Our logics will have some of the next/previous and until/since operators. Some examples are:

- When we restrict ourselves to the purely linear fragment, our operators are \bigcirc and \ominus , and \mathbf{U} and \mathbf{S} , i.e. precisely LTL (with past operators).
- The logic CaRet [2] has the following operators: the next operators \bigcirc and \bigcirc_μ ; the linear and abstract untils (i.e., \mathbf{U} and \mathbf{U}^a), the call since (i.e., \mathbf{S}^c) and a previous operator \ominus_c that will be defined in Section 4.2.

Another notion of a path combines both the linear and the nesting structure. It is the shortest path between two positions i and j . Unlike an abstract path, it decides when to skip a call based on position j . Basically, a summary path from i to j moves along successor edges until it finds a call position k . If k has a matching return ℓ such that j appears after ℓ , then the summary path skips the entire call from k to ℓ and continues from ℓ ; otherwise the path continues as a successor path. Note that every abstract path is a summary path, but there are summary paths that are not abstract paths.


Definition 3.2 A summary path between $i < j$ in a nested word \bar{w} is a sequence $i = i_0 < i_1 < \dots < i_k = j$ such that for all $p < k$,

$$i_{p+1} = \begin{cases} r(i_p) & \text{if } i_p \text{ is a matched call and } j \geq r(i_p) \\ i_p + 1 & \text{otherwise} \end{cases}$$

The corresponding until/since operators are denoted by \mathbf{U}^σ and \mathbf{S}^σ .

¹Our definition of abstract path differs very slightly from that in [2]: there if i_p is not a call and $i_p + 1$ is a return, the path stops. This does not affect the results in any significant way: in fact for summary paths, to be defined shortly, adding the same stopping condition results in an equivalent logic that is used heavily in the proof of expressive completeness.

For example, in the figure below, $\langle 2, 4, 5 \rangle$ is a call path, $\langle 3, 4, 6, 7, 8, 10 \rangle$ is both an abstract and a summary path; and $\langle 3, 4, 6, 7, 8, 9 \rangle$ is a summary path but not an abstract path (as 9 occurs inside a call $\mu(8, 10)$, there is actually no abstract path from 3 to 9).



4 Expressive Completeness

In this section we study logics that are expressively complete for FO, i.e. temporal logics that have exactly the same power as FO formulas in one free variable over finite and infinite nested words. In other words, for every formula φ of an expressively complete temporal logic there is an FO formula $\varphi'(x)$ such that $(\bar{w}, i) \models \varphi$ iff $\bar{w} \models \varphi'(i)$ for every nested word \bar{w} and position i in it, and conversely, for every FO formula $\psi(x)$ there is a temporal formula ψ' such that $\bar{w} \models \psi(i)$ iff $(\bar{w}, i) \models \psi'$.

Our starting point is a logic NWTL (nested-word temporal logic) based on summary paths introduced in the previous section. We show that this logic is expressively complete (and of course remains expressively complete under the addition of operators present in logics inspired by verification of properties of execution paths in programs). This latter remark will be of importance later, when we study the complexity of model checking.

We then look at logics close to those in the verification literature, i.e. with operators such as call and abstract until and since, and ask what needs to be added to them to get expressive completeness. We confirm a conjecture of [2] that a *within* operator is what's needed: such an operator evaluates a formula on a nested subword.

4.1 Expressive completeness and NWTL

The logic NWTL (*nested words temporal logic*) has next and previous operators, as well as until and since with respect to summary paths. That is, its formulas are given by:

$$\varphi, \varphi' := \top \mid a \mid \text{call} \mid \text{ret} \mid \neg\varphi \mid \varphi \vee \varphi' \mid \bigcirc\varphi \mid \bigcirc_\mu\varphi \mid \ominus\varphi \mid \ominus_\mu\varphi \mid \varphi\mathbf{U}^\sigma\varphi' \mid \varphi\mathbf{S}^\sigma\varphi'$$

where a ranges over Σ . We use abbreviations int for $\neg\text{call} \wedge \neg\text{ret}$ (true in an internal position). Note that in the absence of pending calls and returns, call and ret are definable as $\bigcirc_\mu\top$ and $\ominus_\mu\top$, respectively.

Theorem 4.1 NWTL = FO over both finite and infinite nested words.

Proof sketch. Translation of NWTL into FO is quite straightforward, but we show how to do it carefully, to get the 3-variable property. For the converse, we define yet another notion of path, called a strict summary path, that is different from summary paths in two ways. First, if it skips a call, it jumps from i not to $r(i)$ but $r(i) + 1$. Second, if it reaches a matched return position, it stops. We then look at the logic NWTL^s in which the semantics of until and since is modified so that they refer to strict summary paths. We then show that $\text{NWTL}^s \subseteq \text{NWTL}$ and $\text{FO} \subseteq \text{NWTL}^s$.

The former is by a direct translation. The proof of $\text{FO} \subseteq \text{NWTL}^s$ is in two parts. First we deal with the finite case. We look at the standard translation from nested words into binary trees. If a matched call position i is translated into a node s of a tree, then the first position inside the call is translated into the right successor of s , and the linear successor of $r(i)$ is translated into the left successor of s . If i is an internal position, or an unmatched call or return position, its linear successor is translated into the left successor of s . With this translation, strict summary paths become paths in a tree.

We next use until/since-based logics for trees from [18, 25]. By a slight adaptation of techniques from these papers (in particular using the separation property from [18]), we prove expressive completeness of a translation of NWTL^s into a tree logic, and then derive expressive completeness of NWTL^s for finite nested words.

In the infinite case, we combine the finite case and the separation property of [18] with Kamp's theorem and the separation property of LTL. Note that a nested ω -word is translated into an infinite tree with exactly one infinite branch. A composition argument that labels positions of that branch with types of subtrees reduces each FO formula to an LTL formula over that branch in which propositions are types of subtrees, expressible in NWTL^s by the proof in the finite case. Using the separation properties, we then show how to translate such a description into NWTL^s . \square

Recall that FO^k stands for a fragment of FO that consists of formulas which use at most k variables in total. First, from our translation from NWTL to FO we get:

Corollary 4.2 *Over nested words, every FO formula with at most one free variable is equivalent to an FO^3 formula.*

Furthermore, for FO sentences, we can eliminate the since operator.

Corollary 4.3 *For every FO sentence Φ over finite or infinite nested words, there is a formula φ of NWTL that does not use the since operator \mathbf{S}^σ such that $\bar{w} \models \Phi$ iff $(\bar{w}, 1) \models \varphi$.*

The previous operators \ominus and \ominus_μ , however, are needed even for FO sentences over nested words. This situation is quite different thus from LTL, for which the separation

property says that FO sentences over the usual, unnested, words can be evaluated without using the previous \ominus and since \mathbf{S} operators. Let $\text{NWTL}^{\text{future}}$ be the fragment of NWTL that does not use \mathbf{S}^σ and the operators \ominus and \ominus_μ .

Proposition 4.4 *There are FO sentences over nested words that cannot be expressed in $\text{NWTL}^{\text{future}}$.*

Proof sketch. Let \bar{w}_1 and \bar{w}_2 be two well-matched nested words, of length n_1 and n_2 respectively. We first show that, for every $\text{NWTL}^{\text{future}}$ formula, there is an integer k such that $\bar{w}_1[i_1, n_1] \equiv_k \bar{w}_2[i_2, n_2]$ implies $(\bar{w}_1, i_1) \models \varphi$ iff $(\bar{w}_2, i_2) \models \varphi$. Here \equiv_k means that Player II has a win in the k -round Ehrenfeucht-Fraïssé game. This follows from expressive-completeness and properties of future formulas. Using this, we show that there is no $\text{NWTL}^{\text{future}}$ formula equivalent to $\bigcirc_\mu \top \wedge \bigcirc_\mu \ominus a$ (checking whether the first position is a call, and the position preceding its matching return is labeled a). \square

Note also that adding all other until/since pairs to NWTL does not change its expressiveness. That is, if we let NWTL^+ be $\text{NWTL} + \{\mathbf{U}, \mathbf{S}, \mathbf{U}^c, \mathbf{S}^c, \mathbf{U}^a, \mathbf{S}^a\}$, then:

Corollary 4.5 $\text{NWTL}^+ = \text{FO}$.

Later, when we deal with model-checking, we shall prove upper bound results for NWTL^+ that, while expressively complete for FO, allows more operators.

4.2 The *within* operator

We now go back to the three until/since operators originally proposed for temporal logics on nested words, based on the the linear, call, and abstract paths. In other words, our basic logic, denoted by LTL^μ , is

$$\begin{aligned} \varphi, \varphi' := & \top \mid a \mid \text{call} \mid \text{ret} \mid \neg\varphi \mid \varphi \vee \varphi' \mid \\ & \bigcirc\varphi \mid \bigcirc_\mu\varphi \mid \ominus\varphi \mid \ominus_\mu\varphi \mid \\ & \varphi\mathbf{U}\varphi' \mid \varphi\mathbf{S}\varphi' \mid \varphi\mathbf{U}^c\varphi' \mid \varphi\mathbf{S}^c\varphi' \mid \varphi\mathbf{U}^a\varphi' \mid \varphi\mathbf{S}^a\varphi' \end{aligned}$$

We now extend this logic with the following *within* operator proposed in [2]. If φ is a formula, then $\mathcal{W}\varphi$ is a formula, and $(\bar{w}, i) \models \mathcal{W}\varphi$ iff i is a call, and $(\bar{w}[i, j], i) \models \varphi$, where $j = r(i)$ if i is a matched call and $j = |\bar{w}|$ if i is an unmatched call. In other words, $\mathcal{W}\varphi$ evaluates φ on a subword restricted to a single procedure. We denote such an extended logic by $\text{LTL}^\mu + \mathcal{W}$.

Theorem 4.6 $\text{LTL}^\mu + \mathcal{W} = \text{FO}$ over both finite and infinite nested words.

The inclusion of $\text{LTL}^\mu + \mathcal{W}$ into FO is routine. The converse is done by encoding NWTL into $\text{LTL}^\mu + \mathcal{W}$.

CaRet and other within operators The logic CaRet, as defined in [2], did not have all the operators of LTL^μ. In fact it did not have the previous operators \ominus and \ominus_μ , and it only had linear and abstract until operators, and the call since operator. That is, CaRet was defined as

$$\begin{aligned} \varphi, \varphi' := & \top \mid a \mid \text{call} \mid \text{ret} \mid \neg\varphi \mid \varphi \vee \varphi' \mid \\ & \bigcirc\varphi \mid \bigcirc_\mu\varphi \mid \ominus_c\varphi \mid \\ & \varphi\mathbf{U}\varphi' \mid \varphi\mathbf{U}^a\varphi' \mid \varphi\mathbf{S}^c\varphi', \end{aligned}$$

and we assume that a ranges over $\Sigma \cup \{\text{pret}\}$, where pret is true in pending returns (which is not definable with the remaining operators). Here \ominus_c is the previous operator corresponding to call paths. Formally, $(\bar{w}, i) \models \ominus_c\varphi$ if $\mathcal{C}(i)$ is defined and $(\bar{w}, \mathcal{C}(i)) \models \varphi$.

A natural question is whether there is an expressively-complete extension of this logic. It turns out that two *within* operators based on \mathcal{C} and \mathcal{R} (the innermost call and its return) functions provide such an extension. We define two new formulas $\mathcal{C}\varphi$ and $\mathcal{R}\varphi$ with the semantics as follows:

- $(\bar{w}, i) \models \mathcal{C}\varphi$ iff $\bar{w}[j, i] \models \varphi$, where $j = \mathcal{C}(i)$ if $\mathcal{C}(i)$ is defined, and $j = 1$ otherwise.
- $(\bar{w}, i) \models \mathcal{R}\varphi$ if $\bar{w}[i, j] \models \varphi$, where $j = \mathcal{R}(i)$ if $\mathcal{R}(i)$ is defined, and $j = |\bar{w}|$ (if \bar{w} is finite) or ∞ (if \bar{w} is infinite) otherwise.

Theorem 4.7 $\text{CaRet} + \{\mathcal{C}, \mathcal{R}\} = \text{FO}$ over both finite and infinite nested words.

The proof of this result is somewhat involved, and relies on different techniques. The operators used in CaRet do not correspond naturally to tree translations of nested words, and the lack of all until/since pairs makes a translation from NWTL hard. We thus use a composition argument *directly* on nested words.

5 Model-Checking and Satisfiability

In this section we show that both model-checking and satisfiability are single-exponential-time for NWTL. In fact we prove this bound for NWTL⁺, an FO-complete extension of NWTL with all of $\mathbf{U}, \mathbf{S}, \mathbf{U}^c, \mathbf{S}^c, \mathbf{U}^a, \mathbf{S}^a$. We use automata-theoretic techniques, by translating formula into equivalent automata on nested words. We then show that a different expressively complete logic based on adding the *within* operator to CaRet requires doubly-exponential time for model-checking, but is exponentially more succinct.

5.1 Nested word automata

A *nondeterministic nested word automaton* (NWA) A over an alphabet Σ is a structure $(Q, Q_0, F, F_c, \delta_c, \delta_i, \delta_r, \delta_{pr})$ consisting of a finite set of

states Q , a set of initial states $Q_0 \subseteq Q$, a set of (linear) accepting states $F \subseteq Q$, a set of pending call accepting states $F_c \subseteq Q$, a call-transition relation $\delta_c \subseteq Q \times \Sigma \times Q \times Q$, an internal-transition relation $\delta_i \subseteq Q \times \Sigma \times Q$, a return-transition relation $\delta_r \subseteq Q \times Q \times \Sigma \times Q$, and a pending-return-transition relation $\delta_{pr} \subseteq Q \times \Sigma \times Q$. The automaton A starts in the initial state and reads the nested word from left to right. The state is propagated along the linear edges as in case of a standard word automaton. However, at a call, the nested word automaton propagates states along the linear edges and also along the nesting edge (if there is no matching return, then the latter is required to be in F_c for acceptance). At a matched return, the new state is determined based on the states propagated along the linear as well as the nesting incoming edges.

Formally, a *run* r of the automaton A over a nested word $\bar{w} = (a_1 a_2 \dots, \mu, \text{call}, \text{ret})$ is a sequence q_0, q_1, \dots of states along the linear edges, and a sequence q'_i , for every call position i , of states along nesting edges, such that $q_0 \in Q_0$ and for each position i , if i is a call then $(q_{i-1}, a_i, q_i, q'_i) \in \delta_c$; if i is internal, then $(q_{i-1}, a_i, q_i) \in \delta_i$; if i is a return such that $\mu(j, i)$, then $(q_{i-1}, q'_j, a_i, q_i) \in \delta_r$; and if i is an unmatched return then $(q_{i-1}, a_i, q_i) \in \delta_{pr}$. The run r is accepting if (1) for all pending calls i , $q'_i \in F_c$, and (2) the final state $q_\ell \in F$ for finite word of length ℓ , and for infinitely many positions i , $q_i \in F$, for nested ω -words. The automaton A accepts the nested word \bar{w} if it has an accepting run over \bar{w} .

Nested word automata have the same expressiveness as the monadic second order logic over nested words, and the language emptiness can be checked in polynomial-time [4].

5.2 Tableau construction

We now show how to build an NWA accepting the satisfying models of a formula of NWTL⁺. This leads to decision procedures for satisfiability and model checking.

Let us first consider special kinds of summary paths: *summary-down* paths are allowed to use only call edges (from a call to the first position inside the call), nesting edges (from a call to its matching return), and internal edges (from an internal or return position to a call or internal position). *Summary-up* paths are allowed to use only return edges (from a position preceding a return to the return), nesting edges and internal edges. We will use $\mathbf{U}^{\sigma\downarrow}$ and $\mathbf{U}^{\sigma\uparrow}$ to denote the corresponding until operators. Observe that $\varphi\mathbf{U}^{\sigma\downarrow}\psi$ is equivalent to $\varphi\mathbf{U}^{\sigma\uparrow}(\varphi\mathbf{U}^{\sigma\downarrow}\psi)$.

Given a formula φ , we wish to construct a nested word automaton A_φ whose states correspond to sets of subformulas of φ . Intuitively, given a nested word \bar{w} , a run r , which is a linear sequence $q_0 q_1 \dots$ of states and states q'_i labeling nesting edges from call positions, should be such that each state q_i is precisely the set of formulas that hold at position

$i + 1$. The label q'_i is used to remember abstract-next formulas that hold at position i and the abstract-previous formulas that hold at matching return. For clarity of presentation, we focus on formulas with next operators \bigcirc and \bigcirc_μ , and until over summary-down paths. It is easy to modify the construction to allow other types of untils and past operators.

Given a formula φ , the closure of φ , denoted by $cl(\varphi)$, is the smallest set that satisfies the following properties: $cl(\varphi)$ contains φ , **call**, **ret**, **int**, and $\bigcirc\text{ret}$; if either $\neg\psi$, or $\bigcirc\psi$ or $\bigcirc_\mu\psi$ is in $cl(\varphi)$ then $\psi \in cl(\varphi)$; if $\psi \vee \psi' \in cl(\varphi)$, then $\psi, \psi' \in cl(\varphi)$; if $\psi \mathbf{U}^{\sigma\downarrow}\psi' \in cl(\varphi)$, then $\psi, \psi', \bigcirc(\psi \mathbf{U}^{\sigma\downarrow}\psi')$, and $\bigcirc_\mu(\psi \mathbf{U}^{\sigma\downarrow}\psi')$ are in $cl(\varphi)$; and if $\psi \in cl(\varphi)$ and ψ is not of the form $\neg\theta$ (for any θ), then $\neg\psi \in cl(\varphi)$. It is straightforward to see that the size of $cl(\varphi)$ is only linear in the size of φ . Henceforth, we identify $\neg\neg\psi$ with the formula ψ .

An *atom* of φ is a set $\Phi \subseteq cl(\varphi)$ that satisfies:

- For every $\psi \in cl(\varphi)$, $\psi \in \Phi$ iff $\neg\psi \notin \Phi$.
- For every formula $\psi \vee \psi' \in cl(\varphi)$, $\psi \vee \psi' \in \Phi$ iff ($\psi \in \Phi$ or $\psi' \in \Phi$).
- For every formula $\psi \mathbf{U}^{\sigma\downarrow}\psi' \in cl(\varphi)$, $\psi \mathbf{U}^{\sigma\downarrow}\psi' \in \Phi$ iff either $\psi' \in \Phi$ or ($\psi \in \Phi$ and $\bigcirc\text{ret} \notin \Phi$ and either $\bigcirc(\psi \mathbf{U}^{\sigma\downarrow}\psi') \in \Phi$ or $\bigcirc_\mu(\psi \mathbf{U}^{\sigma\downarrow}\psi') \in \Phi$).
- Φ contains exactly one of the elements in the set $\{\text{call}, \text{ret}, \text{int}\}$.

These clauses capture local consistency requirements.

Given a formula φ , we build a nested word automaton A_φ as follows. The alphabet Σ is 2^{AP} , where AP is the set of atomic propositions.

1. Atoms of φ are states of A_φ ;
2. An atom Φ is an initial state iff $\varphi \in \Phi$;
3. For atoms Φ, Ψ and a symbol $a \subseteq AP$, (Φ, a, Ψ) is an internal transition of A_φ iff (a) **int** $\in \Phi$; and (b) for $p \in AP$, $p \in a$ iff $p \in \Phi$; and (c) for each $\bigcirc\psi \in cl(\varphi)$, $\psi \in \Psi$ iff $\bigcirc\psi \in \Phi$; (d) for each $\bigcirc_\mu\psi \in cl(\varphi)$, $\bigcirc_\mu\psi \notin \Phi$.
4. For atoms Φ, Ψ_l, Ψ_h and a symbol $a \subseteq AP$, $(\Phi, a, \Psi_l, \Psi_h)$ is a call transition of A_φ iff (a) **call** $\in \Phi$; and (b) for $p \in AP$, $p \in a$ iff $p \in \Phi$; and (c) for each $\bigcirc\psi \in cl(\varphi)$, $\psi \in \Psi_l$ iff $\bigcirc\psi \in \Phi$; and (d) for each $\bigcirc_\mu\psi \in cl(\varphi)$, $\bigcirc_\mu\psi \in \Psi_h$ iff $\bigcirc_\mu\psi \in \Phi$.
5. For atoms Φ_l, Φ_h, Ψ and a symbol $a \subseteq AP$, $(\Phi_l, \Phi_h, a, \Psi)$ is a return transition of A_φ iff (a) **ret** $\in \Phi_l$; and (b) for $p \in AP$, $p \in a$ iff $p \in \Phi_l$; and (c) for each $\bigcirc\psi \in cl(\varphi)$, $\psi \in \Psi$ iff $\bigcirc\psi \in \Phi_l$; and (d) for each $\bigcirc_\mu\psi \in cl(\varphi)$, $\bigcirc_\mu\psi \in \Phi_h$ iff $\psi \in \Phi_l$.
6. For atoms Φ, Ψ and a symbol $a \subseteq AP$, (Φ, a, Ψ) is a pending-return transition of A_φ iff (a) **ret** $\in \Phi$; and (b) for $p \in AP$, $p \in a$ iff $p \in \Phi$; and (c) for each

$\bigcirc\psi \in cl(\varphi)$, $\psi \in \Psi$ iff $\bigcirc\psi \in \Phi$; (d) for each $\bigcirc_\mu\psi \in cl(\varphi)$, $\bigcirc_\mu\psi \notin \Phi$.

The transition relation ensures that the current symbol is consistent with the atomic propositions in the current state, and next operators requirements are correctly propagated.

An atom Φ belongs to the set F_c iff Φ does not contain any abstract-next formula, and this ensures that, in an accepting run, at a pending call, no requirements are propagated along the nesting edge. For each until-formula ψ in the closure, let F_ψ be the set of atoms that either do not contain ψ or contain the second argument of ψ . Then a nested word \bar{w} over the alphabet 2^{AP} satisfies φ iff there is a run r of A_φ over \bar{w} such that for each until-formula $\psi \in cl(\varphi)$, for infinitely many positions $i, q_i \in F_\psi$. Thus,

Theorem 5.1 *For a formula φ of NWTL^+ , one can effectively construct a nondeterministic Büchi nested word automaton A_φ of size $2^{O(|\varphi|)}$ accepting the models of φ .*

Since the automaton A_φ is exponential in the size of φ , we can check satisfiability of φ in exponential-time by testing emptiness of A_φ . EXPTIME-hardness follows from the corresponding hardness result for CaRet.

Corollary 5.2 *The satisfiability problem for NWTL^+ is EXPTIME-complete.*

When programs are modeled by nested word automata A (or equivalently, pushdown automata, or recursive state machines), and specifications are given by formulas φ of NWTL^+ , we can use the classical automata-theoretic approach: negate the specification, build the NWA $A_{\neg\varphi}$ accepting models that violate φ , take the product with the program A , and test for emptiness of $L(A) \cap L(A_{\neg\varphi})$. Note that the program typically will be given more compactly, say, as a Boolean program [5], and thus, the NWA A may itself be exponential in the size of the input.

Corollary 5.3 *Model checking NWTL^+ specifications with respect to Boolean programs is EXPTIME-complete.*

5.3 Checking within operator

We now show that adding *within* operators makes model-checking doubly exponential. Given a formula φ of NWTL or NWTL^+ , let p_φ be a special proposition that does not appear in φ . Let W_φ be the language of nested words \bar{w} such that for each position i , $(\bar{w}, i) \models p_\varphi$ iff $(\bar{w}, i) \models \mathcal{W}\varphi$. We construct a doubly-exponential automaton B that captures W_φ . First, using the tableau construction for NWTL^+ , we construct an exponential-size automaton A that captures nested words that satisfy φ . Intuitively, every time a proposition p_φ is encountered, we want to start a new copy of A , and a state of B keeps track of states of multiple copies of

A. At a call, B guesses whether the call has a matching return or not. In the latter case, we need to maintain pairs of states of A so that the join at return positions can be done correctly. A state of B , then, is either a set of states of A or a set of pairs of states of A . We explain the latter case. A pair (q, q') belongs to the state of B , while reading position i of a nested word \bar{w} , if the subword from i to the first unmatched return can take A from q to q' . When reading an internal symbol a , a summary (q, q') in the current state can be updated to (u, q') , provided A has an internal transition from q to u on a . Let B read a call symbol a . Consider a summary (q, q') in the current state, and a call-transition (q, a, q_l, q_h) of A . Then B guesses the return transition (u_l, q_h, b, u) that will be used by A at the matching return, and sends the summary (q_l, u_l) along the call edge and the triple (b, u, q') along the nesting edge. While processing a return symbol b , the current state of B must contain summaries only of the form (q, q) where the two states match, and for each summary (b, u, q') retrieved from the state along the nesting edge, the new state contains (u, q') . Finally, B must enforce that $\mathcal{W}\varphi$ holds when p_φ is read. Only a call symbol a can contain p_φ , and when reading such a symbol, B guesses a call transition (q_0, a, q_l, q_h) , where q_0 is the initial state of A , and a return transition (u_l, q_h, b, q_f) , where q_f is an accepting state of A , and sends the summary (q_l, u_l) along the call edge and the symbol b along the nesting edge.

Lemma 5.4 *For every formula φ of NWTL^+ , there is a nested word automaton that accepts the language W_φ and has size doubly-exponential in $|\varphi|$.*

Consider a formula φ of $\text{NWTL}^+ + \mathcal{W}$. For every within-subformula $\mathcal{W}\psi$ of φ , let ψ' be obtained from ψ by substituting each top-level subformula $\mathcal{W}\psi$ in ψ by the proposition p_ψ . Each of these primed formulas is a formula of NWTL^+ . Then, if we take the product of the nested word automata accepting $W_{\psi'}$ corresponding to all the within-subformulas ψ , together with the nested word automaton $A_{\varphi'}$, the resulting language captures the set of models of φ . Intuitively, the automaton for $W_{\varphi'}$ is ensuring that the truth of the proposition p_φ reflects the truth of the subformula $\mathcal{W}\varphi$. If φ itself has a within-subformula $\mathcal{W}\psi$, then the automaton for φ treats it as an atomic proposition p_ψ , and the automaton checking p_ψ , running in parallel, makes sure that the truth of p_ψ correctly reflects the truth of $\mathcal{W}\psi$.

For the lower bound, the decision problem for LTL games can be reduced to the satisfiability problem for formulas with linear untils and within operators [17], and this shows that for CaRet extended with the within operator, the satisfiability problem is 2EXPTIME-hard. We thus obtain:

Proposition 5.5 *For the logic NWTL^+ extended with the within operator \mathcal{W} the satisfiability problem and the model checking problem with respect to Boolean programs, are both 2EXPTIME-complete.*

Remark: checking $\bar{w} \models \varphi$ for finite nested words For finite nested words, one evaluates the complexity of checking whether the given word satisfies a formula, in terms of the length $|\bar{w}|$ of the word and the size of the formula. A straightforward recursion on subformulas shows that for NWTl formulas the complexity of this check is $O(|\bar{w}| \cdot |\varphi|)$, and for both logics with *within* operators, $\text{CaRet} + \{\mathcal{C}, \mathcal{R}\}$ and $\text{LTL}^\mu + \mathcal{W}$, it is $O(|\bar{w}|^2 \cdot |\varphi|)$.

5.4 On *within* and succinctness

We saw that adding within operators to NWTL^+ increases the complexity of model-checking by one exponent. In particular, there could be no polynomial-time translation from $\text{NWTL}^+ + \mathcal{W}$ to NWTL^+ . We now prove a stronger result that gives a space bound as well: while $\text{NWTL}^+ + \mathcal{W}$ has the same power as NWTL^+ , its formulae can be exponentially more succinct than formulas NWTL^+ . That is, there is a sequence $\varphi_n, n \in \mathbb{N}$, of $\text{NWTL}^+ + \mathcal{W}$ formulas such that φ_n is of size $O(n)$, and the smallest formula of NWTL^+ equivalent to φ_n is of size $2^{\Omega(n)}$. For this result, we require nested ω -words to be over the alphabet 2^{AP} .

Theorem 5.6 *$\text{NWTL}^+ + \mathcal{W}$ is exponentially more succinct than NWTL^+ .*

The proof is based upon succinctness results in [9, 22], by adapting their examples to nested words.

6 Finite-Variable Fragments

We have already seen that FO formulas in one free variable over nested words can be written using just three distinct variables, as in the case of the usual, unnested, words. For finite nested words this is a consequence of a tree representation of nested words and the three-variable property for FO over finite trees [19], and for infinite nested words this is a consequence Theorem 4.1.

In this section we prove two results. First, we give a model-theoretic proof that FO formulas with zero, one, or two free variables over nested words (finite or infinite) are equivalent to FO^3 formulas. Given the $\text{FO} = \text{FO}^3$ collapse, we ask whether there is a temporal logic expressively complete for FO^2 , the two-variable fragment. We adapt techniques from [9] to find a temporal logic that has the same expressiveness as FO^2 over nested words (in a vocabulary that has successor relations corresponding to the “next” temporal operators).

6.1 The three-variable property

We give a model-theoretic, rather than a syntactic, argument, that uses Ehrenfeucht-Fraïssé games and shows that

over nested words, formulas with at most two free variables are equivalent to FO^3 formulas. Note that for finite nested words, the translation into trees, already used in the proof of Theorem 4.1, can be done using at most three variables. This means that the result of [19] establishing the 3-variable property for finite ordered unranked trees gives us the 3-variable property for finite nested words. We prove that $\text{FO} = \text{FO}^3$ over arbitrary nested words.

Theorem 6.1 *Over finite or infinite nested words, every FO formula with at most 2 free variables is equivalent to an FO^3 formula.*

Proof: We look at infinite nested words since the finite case was settled in [19]. It is more convenient to prove the result for ordered unranked forests in which every subtree is finite. We translate a nested ω -word into such a forest as follows: when $\mu(i, j)$ holds, the subword, $\bar{w}[i, j]$ is mapped to a subtree with root i , $i + 1$ as the first child of i , and $j + 1$ as i 's next sibling (note that this is different from the translation into binary trees we used before). If i is an internal position, or a pending call or return position, then it has no descendants and its next sibling is $i + 1$.

It is routine to define, in FO, relations \preceq_{desc} and \preceq_{sib} for descendant and younger sibling in such a forest. Furthermore, from these relations, we can define the usual \leq and μ in nested words using at most 3 variables as follows. The formulas for $x \leq y$ and $\mu(x, y)$ are given by

$$(y \preceq_{\text{desc}} x) \vee \exists z(x \preceq_{\text{desc}} z \wedge \exists x(z \prec_{\text{sib}} x \wedge y \preceq_{\text{desc}} x)) \\ (y \preceq_{\text{desc}} x) \wedge \forall z((z \preceq_{\text{desc}} x) \rightarrow z \leq y).$$

Thus, it suffices to prove the three-variable property for such ordered forests, which will be referred to as \mathcal{A} , \mathcal{B} , etc. We shall use pebble games. Let $\mathbf{G}_m^v(\mathcal{A}, a_1, b_1, \mathcal{B}, b_1, b_2)$ be the m -move, v -pebble game on structures \mathcal{A} and \mathcal{B} where initially pebbles x_i are placed on a_i in \mathcal{A} and b_i in \mathcal{B} . Player II has a winning strategy for $\mathbf{G}_m^v(\mathcal{A}, a_1, b_1, \mathcal{B}, b_1, b_2)$ iff \mathcal{A}, a_1, a_2 and \mathcal{B}, b_1, b_2 agree on all formulas with at most v variables and quantifier-depth m . We know from [13] that to prove Theorem 6.1, it suffices to show that, for all k , if Player II has a winning strategy for the game $\mathbf{G}_{3k+2}^3(\mathcal{A}, a_1, a_2; \mathcal{B}, b_1, b_2)$, then she also has a winning strategy for the game $\mathbf{G}_k^k(\mathcal{A}, a_1, a_2; \mathcal{B}, b_1, b_2)$.

We show that Player II can win the k -pebble game by maintaining a set of 3-pebble subgames on which she copies Player I's moves and decides on responses using her winning strategy for these smaller 3-pebble games. The choice of these sub-games will partition the universe $|\mathcal{A}| \cup |\mathcal{B}|$ so that each play by Player I in the k -pebble game will be answered in one 3-pebble game. This is similar to the proof that linear orderings have the 3-variable property [13]. \square

6.2 The two-variable fragment

In this section, we construct a temporal logic that captures the two-variable fragment of FO. Note that for finite unranked trees, a navigational logic capturing FO^2 is known [20, 19]: it corresponds to a fragment of XPath. However, translating the basic predicates over trees into the vocabulary of nested words requires 3 variables, and thus we cannot apply existing results even in the finite case.

Since FO^2 over a linear ordering cannot define the successor relation but temporal logics have next operators, we explicitly introduce successors into the vocabulary of FO. These successor relations in effect partition the linear edges into three disjoint types; *interior* edges, *call* edges, and *return* edges, and the nesting edges (except those from a position to its linear successor) into two disjoint types; *call-return* summaries, and *call-interior-return* summaries.

- $S^i(i, j)$ holds iff $j = i + 1$ and either $\mu(i, j)$ or i is not a call and j is not a return.
- $S^c(i, j)$ holds iff i is a call and $j = i + 1$ is not a return;
- $S^r(i, j)$ holds iff i is not a call and $j = i + 1$ is a return.
- $S^{cr}(i, j)$ holds iff $\mu(i, j)$ and there is a path from i to j using only call and return edges.
- $S^{cir}(i, j)$ holds iff $\mu(i, j)$ and neither $j = i + 1$ nor $S^{cr}(i, j)$.

Let T denote the set $\{c, i, r, cr, cir\}$ of all edge types. In addition to the built-in predicates S^t for $t \in T$, we add the *transitive closure* of all unions of subsets of these relations. That is, for each non-empty set $\Gamma \subseteq T$ of edge types, let S^Γ stand for the union $\cup_{t \in \Gamma} S^t$, and let \leq^Γ be the reflexive-transitive closure of S^Γ . Now when we refer to FO^2 over nested words, we mean FO in the vocabulary of the unary predicates plus all the \leq^Γ 's, the five successor relations, and the built-in unary `call` and `ret` predicates.

We define a temporal logic unary-NWTL that has future and past versions of next operators parameterized by edge types, and eventually operators parameterized by a set of edge types. Its formulas are given by:

$$\varphi := \top \mid a \mid \text{call} \mid \text{ret} \mid \neg\varphi \mid \varphi \vee \varphi' \mid \\ \bigcirc^t \varphi \mid \ominus^t \varphi \mid \diamond^\Gamma \varphi \mid \heartsuit^\Gamma \varphi$$

where a ranges over Σ , t ranges over T , and Γ ranges over non-empty subsets of T . The semantics is defined in the obvious way; for example, $(\bar{w}, i) \models \heartsuit^\Gamma \varphi$ iff for some position $i \leq^\Gamma j$, $(\bar{w}, j) \models \varphi$.

For an FO^2 formula $\varphi(x)$ with one free variable x , let $\text{qdp}(\varphi)$ be its quantifier depth, and for a unary-NWTL formula φ' , let $\text{odp}(\varphi')$ be its operator depth.

Theorem 6.2 *1. unary-NWTL is expressively complete for FO^2 .*

2. If formulas are viewed as DAGs (i.e. identical sub-formulas are shared), then every FO^2 formula $\varphi(x)$ can be converted to an equivalent unary-NWTL formula φ' of size $2^{\mathcal{O}(|\varphi|(\text{qdp}(\varphi)+1))}$ and $\text{odp}(\varphi') \leq 10 \text{qdp}(\varphi)$. The translation is computable in time polynomial in the size of φ' .
3. Model checking of unary-NWTL can be carried out with the same worst case complexity as for NWTL.

Proof sketch. The translation from unary-NWTL into FO^2 is standard. For the other direction we adapt techniques of [9]. Given an FO^2 formula $\varphi(x)$, the translation works as follows. When $\varphi(x)$ is of the form $a(x)$, for a proposition a , it outputs a . The cases of Boolean connectives are straightforward. 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 $\varphi^*(x, y)$ as $\beta(\chi_0(x, y), \dots, \chi_{r-1}(x, y), \xi_0(x), \dots, \xi_{s-1}(x), \zeta_0(y), \dots, \zeta_{t-1}(y))$, where β is a propositional formula, each χ_i is an atomic order formula, and ξ_i 's and ζ_i 's are atomic or existential FO^2 formulas with quantifier depth $< \text{qdp}(\varphi)$. In order to be able to recurse on subformulas of φ we have to separate the ξ_i 's from the ζ_i 's. For that, we consider mutually exclusive and complete *order types* that enumerate possible order relations between x and y with respect to different S^t 's. Under each order type, each atomic order formula evaluates to either \top or \perp . Furthermore, if τ is an order type, $\psi(x)$ an FO^2 formula, and ψ' an equivalent unary-NWTL formula, one can obtain a unary-NWTL formula $\tau\langle\psi\rangle$ equivalent to $\exists y(\tau \wedge \psi(y))$. Using this and the hypothesis for ξ_i' for $i < s$ and $\zeta_i'(x)$ we can compute φ' .

Model checking for unary-NWTL can be carried out with the same complexity as NWTL, by adapting the tableaux construction in Section 5. \square

Acknowledgments The authors were supported by: Alur – NSF CPA award 0541149; Arenas – FONDECYT grants 1050701, 7060172 and 1070732; Arenas and Barceló – grant P04-067-F from the Millennium Nucleus Centre for Web Research; Immerman – NSF grants CCF-0541018 and CCF-0514621; Libkin – the EC grant MEXC-CT-2005-024502, EPSRC grant E005039, and a grant from NSERC (while on leave from U. of Toronto).

References

- [1] R. Alur, M. Benedikt, K. Etessami, P. Godefroid, T. Reps, M. Yannakakis. Analysis of recursive state machines. *ACM TOPLAS* 27(4): 786–818 (2005).
- [2] R. Alur, K. Etessami and P. Madhusudan. A temporal logic of nested calls and returns. In *TACAS'04*, pages 467–481.
- [3] R. Alur and P. Madhusudan. Visibly pushdown languages. In *STOC'04*, pages 202–211.
- [4] R. Alur and P. Madhusudan. Adding nesting structure to words. In *DTL'06*, pages 1–13.
- [5] T. Ball and S. Rajamani. Bebop: A symbolic model checker for boolean programs. In *SPIN'00*, pages 113–130.
- [6] V. Bárány, C. Lóding, O. Serre. Regularity problems for visibly pushdown languages. *STACS 2006*, pages 420–431.
- [7] Document Object Model DOM. <http://www.w3.org/DOM>.
- [8] J. Esparza and S. Schwoon. A BDD-based model checker for recursive programs. In *CAV'01*, pages 324–336.
- [9] K. Etessami, M. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logic. *Information and Computation* 179(2): 279–295, 2002.
- [10] M. Frick, M. Grohe. The complexity of first-order and monadic second-order logic revisited. *LICS 2002*, 215–224.
- [11] G. Gottlob, C. Koch. Monadic datalog and the expressive power of languages for web information extraction. *Journal of the ACM* 51 (2004), 74–113.
- [12] H.W. Kamp. Tense Logic and the Theory of Linear Order. Ph.D. Thesis, UCLA, 1968.
- [13] N. Immerman and D. Kozen. Definability with bounded number of bound variables. *Information and Computation*, 83 (1989), 121-139.
- [14] N. Klarlund, T. Schwentick and D. Suciu. XML: model, schemas, types, logics, and queries. In *Logics for Emerging Applications of Databases*, Springer, 2003, pages 1–41.
- [15] L. Libkin. Logics for unranked trees: an overview. In *ICALP 2005*, pages 35-50.
- [16] C. Lóding, P. Madhusudan, O. Serre. Visibly pushdown games. In *FSTTCS 2004*, pages 408–420.
- [17] P. Madhusudan, personal communication.
- [18] M. Marx. Conditional XPath, the first order complete XPath dialect. In *PODS'04*, pages 13–22.
- [19] M. Marx. Conditional XPath. *TODS* 30(4): 929–959, 2005.
- [20] M. Marx and M. de Rijke. Semantic characterizations of navigational XPath. In *TDM'04*, pages 67–73.
- [21] F. Neven, T. Schwentick. Expressive and efficient pattern languages for tree-structured data. *PODS'00*, pages 145-156.
- [22] F. Laroussinie, N. Markey, and P. Schnoebelen. Temporal logic with forgettable past. In *LICS'02*, pages 383–392.
- [23] F. Neven and T. Schwentick. Query automata over finite trees. *Theor. Comput. Sci.* 275(1-2): 633–674, 2002.
- [24] SAX: A Simple API for XML. <http://www.saxproject.org>.
- [25] B.-H. Schlingloff. Expressive completeness of temporal logic of trees. *J. Appl. Non-Classical Log.* 2: 157-180, 1992.
- [26] L. Segoufin. Typing and querying XML documents: some complexity bounds. In *PODS'03*, pages 167–178.
- [27] L. Segoufin, V. Vianu. Validating streaming XML documents. In *PODS'02*, pages 53–64.
- [28] V. Vianu. A web Odyssey: from Codd to XML. In *ACM PODS'01*, pages 1–15.

A Proofs and Intermediate Results

Some terminology

The *quantifier rank* (or *quantifier depth*) of an FO formula φ is the depth of quantifier nesting in φ . The *rank- k type* of a structure \mathfrak{M} over a relational vocabulary is the set $\{\varphi \mid \mathfrak{M} \models \varphi \text{ and the quantifier rank of } \varphi \text{ is } k\}$, where φ ranges over FO sentences over the vocabulary. It is well-known that there are finitely many rank- k types for all k , and for each rank- k type τ there is an FO sentence φ_τ such that $\mathfrak{M} \models \varphi_\tau$ iff the rank- k type of \mathfrak{M} is τ . Sometimes we associate types with formulas that define them.

Many proofs in this paper make use of *Ehrenfeucht-Fraïssé* (EF) games. This game is played in two structures, \mathfrak{M} and \mathfrak{M}' , over the same vocabulary, by two players, *Player I* and the *Player II*. In round i Player I selects a structure, say \mathfrak{M} , and an element c_i in the domain of \mathfrak{M} ; Player II responds by selecting an element e_i in the domain of \mathfrak{M}' . Player II *wins* in k rounds, for $k \geq 0$, if $\{(c_i, e_i) \mid i \leq k\}$ defines a partial isomorphism between \mathfrak{M} and \mathfrak{M}' . Also, if \bar{a} is an m -tuple in the domain of \mathfrak{M} and \bar{b} is an m -tuple in the domain of \mathfrak{M}' , where $m \geq 0$, we write $(\mathfrak{M}, \bar{a}) \equiv_k (\mathfrak{M}', \bar{b})$ whenever Player II wins in k rounds no matter how Player I plays, but starting from position (\bar{a}, \bar{b}) .

We write $\mathfrak{M} \equiv_k \mathfrak{M}'$ iff \mathfrak{M} and \mathfrak{M}' have the same rank- k type, that is for every FO sentence φ of quantifier rank- k , $\mathfrak{M} \models \varphi \Leftrightarrow \mathfrak{M}' \models \varphi$. It is well-known that $\mathfrak{M} \equiv_k \mathfrak{M}'$ iff Player II has a winning strategy in the k -round Ehrenfeucht-Fraïssé game on \mathfrak{M} and \mathfrak{M}' .

In the proof of Theorem 6.1, we shall also use *k -pebble games*. In such a game, Player I and Player II have access k matching pebbles each, and each round consists of Player I either removing, or placing, or replacing a pebble in one structure, and Player II replicating the move in the other structure. The correspondence given by the matching pebbles should be a partial isomorphism. If Player II can play while maintaining partial isomorphism for m rounds, then the structures agree on all FO ^{k} sentences of quantifier rank up to m ; if Player II can play while maintaining partial isomorphism forever, then the structures agree on all FO ^{k} sentences.

Proof of Theorem 4.1

We start with the easy direction $\text{NWTL} \subseteq \text{FO}$.

Lemma A.1 *For every NWTL formula φ , there exists an FO formula $\alpha_\varphi(x)$ that uses at most three variables x, y, z such that for every nested word \bar{w} (finite or infinite), we have $(\bar{w}, i) \models \varphi$ iff $\bar{w} \models \alpha_\varphi(i)$.*

Proof of Lemma A.1: The proof is by induction on the formulas and very simple for all the cases except \mathbf{U}^σ and \mathbf{S}^σ : for example,

$$\alpha_{\bigcirc_\mu \varphi}(x) = \exists y (\mu(x, y) \wedge \exists x (x = y \wedge \alpha_\varphi(x))).$$

For translating \mathbf{U}^σ , we need a few auxiliary formulas. Our first goal is to define a formula $\gamma_r(x, z)$ saying that x is $\mathcal{R}(z)$, i.e. the return of the innermost call within which z is executed. For that, we start with $\delta(y, z) = z < y \wedge \text{ret}(y) \wedge \forall x (\mu(x, y) \rightarrow x < z)$ saying that y is a return that is preceded by z and whose matching call, if exists, precedes z , that is, y is a candidate for $\mathcal{R}(z)$. Then the formula $\gamma_r(x, z)$ is given by

$$\exists y (y = x \wedge \delta(y, z)) \wedge \forall y (\delta(y, z) \rightarrow y \geq x).$$

Likewise, we define $\gamma_c(y, z)$ stating that that y equals $\mathcal{C}(z)$, that is, the innermost call within which z is executed. Now define

$$\chi_1(y, z) = \exists x (\gamma_r(x, z) \wedge x \leq y), \quad \chi_2(x, z) = \exists y (\gamma_c(y, z) \wedge y \geq x)$$

and $\chi(x, y, z)$ as $\chi_1(y, z) \wedge \chi_2(x, z)$. Then this formula says that the summary path from x to y does not pass through z , assuming $x < z < y$. With this, $\alpha_{\varphi \mathbf{U}^\sigma \psi}(x)$ is given by

$$\alpha_\psi(x) \vee \exists y \left(y > x \wedge \alpha_\varphi(x) \wedge \exists x (x = y \wedge \alpha_\psi(x)) \wedge \forall z ((x < z < y \wedge \neg \chi(x, y, z)) \rightarrow \exists x (x = z \wedge \alpha_\varphi(x))) \right)$$

The proof for $\varphi \mathbf{S}^\sigma \psi$ is similar. This concludes the proof of the lemma. \square

In the proof of the other direction, $\text{FO} \subseteq \text{NWTL}$, we shall need two variants of summary paths. A *Semi-Strict summary path* between positions i and j , with $i < j$, in a nested word \bar{w} , is a sequence $i = i_0 < i_1 < \dots < i_k = j$ such that

$$i_{p+1} = \begin{cases} r(i_p) + 1 & \text{if } i_p \text{ is a matched call and } j > r(i_p) \\ i_p + 1 & \text{otherwise.} \end{cases}$$

That is, when skipping a call, instead of jumping to the matching return position, a semi-strict path will jump to its successor.

A *strict summary path* is a semi-strict summary path $i = i_0 < i_1 < i_2 < \dots < i_k = j$ in which no i_p with $p < k$ is a matched return position. In other words, a strict summary path stops if it reaches a matched return position. In particular there may be positions $i < j$ in a nested word such that no strict summary path exists between them. The until/since operators for semi-strict summary paths and strict summary paths will be denoted by $\mathbf{U}_{ss}^\sigma/\mathbf{S}_{ss}^\sigma$ and $\mathbf{U}_s^\sigma/\mathbf{S}_s^\sigma$, respectively. Versions of NWTL in which $\mathbf{U}^\sigma/\mathbf{S}^\sigma$ are replaced by $\mathbf{U}_{ss}^\sigma/\mathbf{S}_{ss}^\sigma$ (or $\mathbf{U}_s^\sigma/\mathbf{S}_s^\sigma$) will be denoted by NWTL^{ss} as NWTL^s .

We will use mret for $\text{ret} \wedge \ominus_\mu \top$, and mcall for $\text{call} \wedge \bigcirc_\mu \top$, to capture matching return and call positions, respectively.

The proof is based on two lemmas.

Lemma A.2 $\text{NWTL}^s \subseteq \text{NWTL}^{ss} \subseteq \text{NWTL}$.

Lemma A.3 $\text{FO} \subseteq \text{NWTL}^s$.

This of course implies the theorem: $\text{NWTL} \subseteq \text{FO} \subseteq \text{NWTL}^s \subseteq \text{NWTL}^{ss} \subseteq \text{NWTL}$. Note that as a corollary we also obtain $\text{NWTL}^s = \text{NWTL}^{ss} = \text{FO}$.

Proof of Lemma A.2: For translating an NWTL^s formula φ into an equivalent formula α_φ of NWTL^{ss} we need to express $\psi \mathbf{U}_s^\sigma \theta$ with \mathbf{U}_{ss}^σ , which is simply $(\alpha_\psi \wedge \neg \text{ret}) \mathbf{U}_{ss}^\sigma \alpha_\theta$, and likewise for the since operators. For translating each NWTL^{ss} formula φ into an equivalent NWTL formula β_φ , again we need to consider only the case of until/since operators. The formula $\psi \mathbf{U}_{ss}^\sigma \theta$ is translated into

$$\beta_\theta \vee \left(\beta_\psi \wedge \left(\left((\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\psi) \right) \mathbf{U}^\sigma \left((\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\theta) \wedge (\text{call} \rightarrow (\bigcirc \beta_\theta \vee \bigcirc_\mu \bigcirc \beta_\theta \vee \bigcirc (\neg \text{ret} \wedge \gamma))) \right) \right) \right), \quad (1)$$

where γ is a formula defined as follows:

$$\left((\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\psi) \wedge (\bigcirc \text{ret} \rightarrow \text{call}) \right) \mathbf{U}^\sigma \left((\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\theta) \wedge (\text{call} \rightarrow (\bigcirc \beta_\theta \vee \bigcirc_\mu \bigcirc \beta_\theta)) \right)$$

The translation for \mathbf{S}_{ss}^σ is similar. The proof that the translation is correct is by induction on the structure of NWTL^{ss} formulas. Again we need to consider only the case of until/since operators. Assume that ψ, θ are equivalent to β_ψ and β_θ , respectively. We need to prove that $\psi \mathbf{U}_{ss}^\sigma \theta$ is equivalent to (1). We only show here that if (\bar{w}, i) satisfies (1), then $(\bar{w}, i) \models \psi \mathbf{U}_{ss}^\sigma \theta$, as the other direction is similar. Given that (\bar{w}, i) satisfies (1), either $(\bar{w}, i) \models \beta_\theta$ or (\bar{w}, i) satisfies the right-hand-side of the outer disjunction of (1). Given that β_θ and θ are equivalent, in the former case $(\bar{w}, i) \models \psi \mathbf{U}_{ss}^\sigma \theta$. Thus, assume that the latter case holds. Then $(\bar{w}, i) \models \psi$, since ψ and β_ψ are equivalent, and there exists a summary path $i = i_0 < i_1 < \dots < i_p$ such that:

$$\begin{aligned} (\bar{w}, i_k) &\models (\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\psi) & 0 \leq k < p \\ (\bar{w}, i_p) &\models (\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\theta) \wedge (\text{call} \rightarrow (\bigcirc \beta_\theta \vee \bigcirc_\mu \bigcirc \beta_\theta \vee \bigcirc (\neg \text{ret} \wedge \gamma))) \end{aligned}$$

We consider three cases.

1. First, assume that i_p is not a call position. Then given that $(\bar{w}, i_p) \models (\neg \text{call} \rightarrow \bigcirc \beta_\theta)$, we have that $(\bar{w}, i_p + 1) \models \beta_\theta$. Only one semi-strict summary path with endpoints i and $i_p + 1$ can be obtained from the sequence $i_0 < i_1 < \dots < i_p < i_p + 1$ by removing some positions; let $i = j_0 < j_1 < \dots < j_\ell = i_p + 1$ be that semi-strict summary path. Then we have that $(\bar{w}, j_\ell) \models \theta$ since $(\bar{w}, i_p + 1) \models \beta_\theta$. Next we show that $(\bar{w}, j_k) \models \psi$ for every $k \in [0, \ell - 1]$.

If $k = 0$, then the property holds since $(\bar{w}, i) \models \psi$. Assume that $k \in [1, \ell - 1]$. If j_k is not a return position, then $(\bar{w}, j_k) \models \psi$ since $(\bar{w}, j_k) \models (\beta_\psi \vee \text{ret})$. If j_k is a return position, then given that j_k is in the semi-strict summary path from j_0 to j_ℓ , we have that $j_k - 1$ is also in this semi-strict summary path and is not a call position. Thus, given that $(\bar{w}, j_k - 1) \models (\neg \text{call} \rightarrow \bigcirc \beta_\psi)$, we have that $(\bar{w}, j_k) \models \beta_\psi$. We conclude that $(\bar{w}, j_k) \models \psi$ and, hence, $(\bar{w}, i) \models \psi \mathbf{U}_{ss}^\sigma \theta$.

2. Second, assume that i_p is a call position and $(\bar{w}, i_p) \models \bigcirc \beta_\theta \vee \bigcirc_\mu \bigcirc \beta_\theta$. Then there exists a position $i_{p+1} > i_p$ such that $(\bar{w}, i_{p+1}) \models \beta_\theta$ and i_{p+1} is either $i_p + 1$ or the linear successor of the matching return of i_p . Only one semi-strict summary path with endpoints i and i_{p+1} can be obtained from the sequence $i_0 < i_1 < \dots < i_p < i_{p+1}$ by removing some positions; let $i = j_0 < j_1 < \dots < j_\ell = i_{p+1}$ be that semi-strict summary path. Then we have that $(\bar{w}, j_\ell) \models \theta$ since $(\bar{w}, i_{p+1}) \models \beta_\theta$. Next we show that $(\bar{w}, j_k) \models \psi$ for every $k \in [0, \ell - 1]$. If $k = 0$, then the property holds since $(\bar{w}, i) \models \psi$. Assume that $k \in [1, \ell - 1]$. If j_k is not a return position, then $(\bar{w}, j_k) \models \psi$ since $(\bar{w}, j_k) \models (\beta_\psi \vee \text{ret})$. If j_k is a return position, then given that j_k is in the semi-strict summary path from j_0 to j_ℓ , we have that $j_k - 1$ is also in this semi-strict summary path and is not a call position. Thus, given that $(\bar{w}, j_k - 1) \models (\neg \text{call} \rightarrow \bigcirc \beta_\psi)$, we have that $(\bar{w}, j_k) \models \beta_\psi$. We conclude that $(\bar{w}, j_k) \models \psi$ and, hence, $(\bar{w}, i) \models \psi \mathbf{U}_{ss}^\sigma \theta$.
3. Third, assume that i_p is a call position and $(\bar{w}, i_p) \models \bigcirc(\neg \text{ret} \wedge \gamma)$. Then there exists a path $i_p + 1 = i_{p+1} < i_{p+2} \dots < i_q$ such that:

$$\begin{aligned} (\bar{w}, i_k) &\models (\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\psi) \wedge (\bigcirc \text{ret} \rightarrow \text{call}) & p + 1 \leq k < q \\ (\bar{w}, i_q) &\models (\beta_\psi \vee \text{ret}) \wedge (\neg \text{call} \rightarrow \bigcirc \beta_\theta) \wedge (\text{call} \rightarrow (\bigcirc \beta_\theta \vee \bigcirc_\mu \bigcirc \beta_\theta)) \end{aligned}$$

Next we show that if i_p is a matched call with return j_p , then $i_q < j_p$. On the contrary, assume that $i_q \geq j_p$. On the contrary, assume that $i_q \geq j_p$. Then there exists $k \in [p + 1, q]$ such that $i_k = j_p$. Given that i_{p+1} is not a return position, we have that $q > p + 1$ and, therefore, $i_k - 1$ is also a position in the path $i_{p+1} < i_{p+2} \dots < i_q$. But given that $i_k = j_p$ is the matching return of i_p and $i_p + 1 \leq i_k - 1$, we have that $i_k - 1$ is not a call position. Thus, $(\bar{w}, i_k - 1) \not\models \bigcirc \text{ret} \rightarrow \text{call}$, which contradicts the fact that $i_{p+1} < i_{p+2} \dots < i_q$ witnesses formula γ .

To finish the proof of the lemma we need to consider two cases.

- (a) Assume that i_q is not a call position. Then given that $(\bar{w}, i_q) \models (\neg \text{call} \rightarrow \bigcirc \beta_\theta)$, we have that $(\bar{w}, i_q + 1) \models \beta_\theta$. Furthermore, given that if i_p is a matched call with return j_p , then we necessarily have that $i_q < j_p$, we conclude that only one semi-strict summary path with endpoints $i = i_0$ and $i_q + 1$ can be obtained from the sequence $i_0 < i_1 < \dots < i_q < i_q + 1$ by removing some positions; let $i = j_0 < j_1 < \dots < j_\ell = i_q + 1$ be that semi-strict summary path. Then we have that $(\bar{w}, j_\ell) \models \theta$ since $(\bar{w}, i_q + 1) \models \beta_\theta$. Next we show that $(\bar{w}, j_k) \models \psi$ for every $k \in [0, \ell - 1]$. If $k = 0$, then the property holds since $(\bar{w}, i) \models \psi$. Assume that $k \in [1, \ell - 1]$. If j_k is not a return position, then $(\bar{w}, j_k) \models \psi$ since $(\bar{w}, j_k) \models (\beta_\psi \vee \text{ret})$. If j_k is a return position, then given that j_k is in the semi-strict summary path from j_0 to j_ℓ , we have that $j_k - 1$ is also in this semi-strict summary path and is not a call position. Thus, given that $(\bar{w}, j_k - 1) \models (\neg \text{call} \rightarrow \bigcirc \beta_\psi)$, we have that $(\bar{w}, j_k) \models \beta_\psi$. We conclude that $(\bar{w}, j_k) \models \psi$ and, hence, $(\bar{w}, i) \models \psi \mathbf{U}_{ss}^\sigma \theta$.
- (b) Finally assume that i_q is a call position. Then given that $(\bar{w}, i_q) \models (\text{call} \rightarrow (\bigcirc \beta_\theta \vee \bigcirc_\mu \bigcirc \beta_\theta))$, there exists a position $i_{q+1} > i_q$ such that $(\bar{w}, i_{q+1}) \models \beta_\theta$ and i_{q+1} is either $i_q + 1$ or the linear successor of the matching return of i_q . Furthermore, given that if i_p is a matched call with return j_p , then we necessarily have that $i_q < j_p$, we conclude that only one semi-strict summary path with endpoints $i = i_0$ and i_{q+1} can be obtained from the sequence $i_0 < i_1 < \dots < i_q < i_{q+1}$ by removing some positions; let $i = j_0 < j_1 < \dots < j_\ell = i_{q+1}$ be that semi-strict summary path. Then we have that $(\bar{w}, j_\ell) \models \theta$ since $(\bar{w}, i_{q+1}) \models \beta_\theta$. Next we show that $(\bar{w}, j_k) \models \psi$ for every $k \in [0, \ell - 1]$. If $k = 0$, then the property holds since $(\bar{w}, i) \models \psi$. Assume that $k \in [1, \ell - 1]$. If j_k is not a return position, then $(\bar{w}, j_k) \models \psi$ since $(\bar{w}, j_k) \models (\beta_\psi \vee \text{ret})$. If j_k is a return position, then given that j_k is in the semi-strict summary path from j_0 to j_ℓ , we have that $j_k - 1$ is also in this semi-strict summary path and is not a call position. Thus, given that $(\bar{w}, j_k - 1) \models (\neg \text{call} \rightarrow \bigcirc \beta_\psi)$, we have that $(\bar{w}, j_k) \models \beta_\psi$. We conclude that $(\bar{w}, j_k) \models \psi$ and, hence, $(\bar{w}, i) \models \psi \mathbf{U}_{ss}^\sigma \theta$.

This concludes the proof of the lemma. \square

Proof of Lemma A.3: We start with the finite case, and then show how the inclusion extends to nested ω -words.

As a tool we shall need a slight modification of a result from [25, 18] providing an expressively complete temporal logic for trees with at most binary branching. We consider binary trees whose domain D is a prefix-closed subset of $\{0, 1\}^*$, and we impose a condition that if $s \cdot 1 \in D$ then $s \cdot 0 \in D$. When we refer to FO on trees, we assume they have two successor relations S_0, S_1 and the descendant relation \preceq (which is just the prefix relation on strings) plus the labeling predicates, which include two new labels `pcall` and `pret` (for pending calls and returns). Each node can be labeled by either a letter from Σ , or by a letter from Σ and `pcall`, or by a letter from Σ and `pret` (i.e. labels `pcall` and `pret` need not be disjoint from other labels).

We also consider the following logic TL^{tree} :

$$\begin{aligned} \varphi \quad := \quad & a \mid \varphi \vee \varphi \mid \neg \varphi \mid \\ & \text{O}_{\downarrow} \varphi \mid \text{O}_{\uparrow} \varphi \mid \text{O}_{\rightarrow} \varphi \mid \text{O}_{\leftarrow} \varphi \mid \\ & \varphi \text{U}_{\downarrow} \psi \mid \varphi \text{S}_{\downarrow} \psi \end{aligned}$$

where a ranges over $\Sigma \cup \{\text{pcall}, \text{pret}\}$, with the following semantics:

- $(T, s) \models \text{O}_{\downarrow} \varphi$ iff $(T, s \cdot i) \models \varphi$ either for $i = 0$ or $i = 1$;
- $(T, s \cdot i) \models \text{O}_{\uparrow} \varphi$ iff $(T, s) \models \varphi$ (where i is either 0 or 1);
- $(T, s \cdot 0) \models \text{O}_{\rightarrow} \varphi$ iff $(T, s \cdot 1) \models \varphi$;
- $(T, s \cdot 1) \models \text{O}_{\leftarrow} \varphi$ iff $(T, s \cdot 0) \models \varphi$;
- $(T, s) \models \varphi \text{U}_{\downarrow} \psi$ iff there exists s' such that $s \preceq s'$, $(T, s') \models \psi$, and $(T, s'') \models \varphi$ for all s'' such that $s \preceq s'' \prec s'$;
- $(T, s) \models \varphi \text{S}_{\downarrow} \psi$ iff there exists s' such that $s' \preceq s$, $(T, s') \models \psi$, and $(T, s'') \models \varphi$ for all s'' such that $s' \prec s'' \preceq s$.

Lemma A.4 (see [18]) *For unary queries over finite binary trees, $\text{TL}^{\text{tree}} = \text{FO}$.*

This lemma is an immediate corollary of expressive completeness of logic $\mathcal{X}_{\text{until}}$ from [18] on ordered unranked trees, as for a fixed number of siblings, the until and since operators can be expressed in terms of the next and previous operators. The result of [18] applies to arbitrary alphabets, and thus in particular to our labelings that may use `pcall` and `pret`.

Next we need a translation from nested words into binary trees, essentially the same as in [4]. For each nested word \bar{w} we have a tree $T_{\bar{w}}$ and a function $\iota_{w-t} : \bar{w} \rightarrow T_{\bar{w}}$ that maps each position of \bar{w} to a node of $T_{\bar{w}}$ as follows:

- the first position of \bar{w} is mapped into the root of $T_{\bar{w}}$;
- if $s = \iota_{w-t}(i)$ then:
 1. if i is an internal, or an unmatched call, or an unmatched return, and is not the last position of \bar{w} , then s has only child $s \cdot 0$ and $\iota_{w-t}(i+1) = s \cdot 0$;
 2. if i is a matched call, then s has both children $s \cdot 0$ and $s \cdot 1$ and $\iota_{w-t}(r(i)+1) = s \cdot 0$, and $\iota_{w-t}(i+1) = s \cdot 1$.
 3. if i is a matched return, then s has no children.

Of course the Σ -labels of i and $\iota_{w-t}(i)$ are the same. If i was a pending call, we label $\iota_{w-t}(i)$ with `pcall`, and if i was a pending return, we label $\iota_{w-t}(i)$ with `pret`.

Note that ι_{w-t} is a bijection, and that labels `pcall` and `pret` may only occur on the leftmost branch of $T_{\bar{w}}$. The following is immediate by straightforward translations.

Claim A.5 *For every FO formula $\varphi(x)$ over nested words there is an FO formula $\varphi'(x)$ over trees such that for every nested word \bar{w} and a position i in it, we have $\bar{w} \models \varphi(i)$ iff $T_{\bar{w}} \models \varphi'(\iota_{w-t}(i))$.*

Since $\text{FO} = \text{TL}^{\text{tree}}$ by Lemma A.4, all that remains to prove is the following claim.

Claim A.6 *For every TL^{tree} formula φ , there exists an NWTL^s formula φ° such that for every nested word \bar{w} and every position i in it, we have*

$$(\bar{w}, i) \models \varphi^\circ \Leftrightarrow (T_{\bar{w}}, \iota_{w-t}(i)) \models \varphi.$$

This is now done by induction, omitting the obvious cases of propositional letters and Boolean connectives. We note that the path in the tree between $\iota_{w-t}(i)$ and $\iota_{w-t}(j)$ corresponds precisely to the strict summary path from i to j (that is, if such a strict summary path is $i = i_0, i_1, \dots, i_k = j$, then $\iota_{w-t}(i_0), \iota_{w-t}(i_1), \dots, \iota_{w-t}(i_k)$ is the path from $\iota_{w-t}(i)$ to $\iota_{w-t}(j)$ in $T_{\bar{w}}$). Hence, the translations of until and since operators are:

$$(\varphi \mathbf{U}_\downarrow \psi)^\circ = \varphi^\circ \mathbf{U}_s^\sigma \psi^\circ, \quad (\varphi \mathbf{S}_\downarrow \psi)^\circ = \varphi^\circ \mathbf{S}_s^\sigma \psi^\circ.$$

For translating next and previous operators, and pending calls/returns, define:

$$\begin{aligned} \text{mcall} &\equiv \bigcirc_\mu \top \text{ (true in a matched call position);} \\ \text{mret} &\equiv \ominus_\mu \top \text{ (true in a matched return position).} \end{aligned}$$

Then the rest of the translation is as follows:

$$\begin{aligned} \text{pcall}^\circ &\equiv \text{call} \wedge \neg \text{mcall} \\ \text{pret}^\circ &\equiv \text{ret} \wedge \neg \text{mret} \\ (\bigcirc_\downarrow \varphi)^\circ &\equiv \neg \text{ret} \wedge \left(\bigcirc \varphi^\circ \vee (\text{call} \wedge \bigcirc_\mu \bigcirc \varphi^\circ) \right) \\ (\bigcirc_\uparrow \varphi)^\circ &\equiv \left(\ominus \text{ret} \wedge \ominus \ominus_\mu \varphi^\circ \right) \vee \left(\ominus \neg \text{ret} \wedge \ominus \varphi^\circ \right) \\ (\bigcirc_{\leftarrow} \varphi)^\circ &\equiv \ominus \text{ret} \wedge \ominus \ominus_\mu \bigcirc \varphi^\circ \\ (\bigcirc_{\rightarrow} \varphi)^\circ &\equiv \ominus \text{call} \wedge \ominus \bigcirc_\mu \bigcirc \varphi^\circ \end{aligned}$$

Now with the proof completed for finite nested words, we extend it to the case of nested ω -words. Note that Claim A.5 continues to hold, and Claim A.6 provides a syntactic translation that applies to both finite and infinite nested words, and thus we need to prove an analog of Lemma A.4 for trees of the form $T_{\bar{w}}$, where \bar{w} ranges over nested ω -words.

If \bar{w} is an nested ω -word, then $T_{\bar{w}}$ has exactly one infinite branch, which consists precisely of $\iota_{w-t}(i)$ where i is an *outer* position, i.e., not inside any (matched) call. We say that i is inside a call if there exists a call j with a matching return k such that $j < i \leq k$. If i is an outer position, then we shall call $\iota_{w-t}(i)$ an *outer* node in the tree $T_{\bar{w}}$ as well.

If i is an outer position which is not a matched call, then so is $i+1$ and $\iota_{w-t}(i+1)$ is the left successor of $\iota_{w-t}(i)$. If i is an outer position and a call with $j > i$ being its matching return, then the left successor of $\iota_{w-t}(i)$ on the infinite path is $\iota_{w-t}(j+1)$. Furthermore, the subtree $t^{\bar{w}}(i)$ which has i as the root, its right child, and all the descendants of the right child is finite and isomorphic to $T_{\bar{w}[i,j]}$ (note that $\bar{w}[i,j]$ has no pending calls/returns). If i an internal or pending call/return outer position, we let $t^{\bar{w}}(i)$ be a single node tree labeled as i in \bar{w} .

Let \bar{w} now be an nested ω -word. For each outer position i we let $\tau_m^{\bar{w}}(i)$ be the rank- m type of $t^{\bar{w}}(i)$. If i is not a matched call, such a type is completely described by i 's label (which consists of a label in Σ and potentially pcall or pret).

If j is not an outer position, and i is an outer position such that $i < j \leq k$, where k is the matching return of i , then $\tau_m^{\bar{w}}(j)$ is the rank- m type of $(T_{\bar{w}[i,k]}, \iota_{w-t}(j))$ (i.e., the rank of $T_{\bar{w}[i,k]}$ with a distinguished node corresponding to j).

Next, for an nested ω -word \bar{w} , let s be a node in $T_{\bar{w}}$ such that $s = \iota_{w-t}(i)$. Let i_1, i_2, \dots enumerate all the outer positions of \bar{w} , and assume that i_p is such that $i_p \leq i < i_{p+1}$ – that is, $\iota_{w-t}(i)$ is a node in the subtree $t^{\bar{w}}(i_p)$. We now define a finite word $s_m^{\leftarrow}(\bar{w}, s)$ of length $p-1$ such that its positions $1, \dots, p-1$ are labeled $\tau_m^{\bar{w}}(i_1), \dots, \tau_m^{\bar{w}}(i_{p-1})$, and an ω -word $s_m^{\rightarrow}(\bar{w}, s)$ such that its positions $1, 2, \dots$ are labeled by $\tau_m^{\bar{w}}(i_{p+1}), \tau_m^{\bar{w}}(i_{p+2}), \dots$. Next the standard composition argument shows the following.

Claim A.7 *Let \bar{w}, \bar{w}' be two nested ω -words, and $s = \iota_{w-t}(i), s' = \iota_{w-t}(i')$ two nodes in $T_{\bar{w}}$ and $T_{\bar{w}'}$ such that:*

- (a) $s_m^{\leftarrow}(\bar{w}, s) \equiv_m s_m^{\leftarrow}(\bar{w}', s')$;
- (b) $s_m^{\rightarrow}(\bar{w}, s) \equiv_m s_m^{\rightarrow}(\bar{w}', s')$;
- (c) $\tau_m^{\bar{w}}(i) = \tau_m^{\bar{w}'}(i')$.

Then $(T_{\bar{w}}, s) \equiv_m (T_{\bar{w}'}, s')$.

The win for Player II is straightforward. If i_1, i_2, \dots enumerate outer positions in \bar{w} and $i_p \leq i < i_{p+1}$, then a move by Player I, say, in $T_{\bar{w}}$, occurs either in $t^{\bar{w}}(j)$ with $j < i$, or in $t^{\bar{w}}(i)$, or in $t^{\bar{w}}(j)$ with $j > i$. Player II then selects j' so that

the response is in $t^{\bar{w}'}(j')$ according to his winning strategy in games either (a) or (b) (if j is in $t^{\bar{w}}(i)$, then j' is in $\tau_m^{\bar{w}'}(i')$), and then, since the rank- m types of $t^{\bar{w}}(j)$ and the chosen $t^{\bar{w}'}(j')$ are the same, selects the actual response according to the winning strategy $t^{\bar{w}}(j) \equiv_m t^{\bar{w}'}(j')$.

Next we show how Claim A.7 proves that FO is expressible in TL^{tree} over infinite trees $T_{\bar{w}}$. First note that being an outer node is expressible: since $\bigcirc_{\leftarrow} \top$ is true in right children of matched calls, then

$$\alpha_{\text{outer}} = \neg(\top \mathbf{S}_\downarrow (\bigcirc_{\leftarrow} \top))$$

is true if no node on the path to the root is inside a call, that is, precisely in outer nodes.

Next note that for each rank- m type τ of a tree there is a formula β_τ such that if $s = \iota_{w-t}(i)$ is an outer node of $T_{\bar{w}}$, then $(T_{\bar{w}}, s) \models \beta_\tau$ iff the rank- m type of $t^{\bar{w}}(i)$ is τ . If i is not a matched call, then such a type is uniquely determined by i 's label and perhaps `pcall` or `pret`, and thus is definable in TL^{tree} .

If i is a matched call, the existence of such a formula β_τ follows from the fact the rank- m type of $t^{\bar{w}}(i)$ is completely determined by the label of i and the rank- m type τ' of the subtree $t_0^{\bar{w}}(i)$ of $t^{\bar{w}}(i)$ rooted at the right child of s (recall that the root has only (right) child, by the definition of $t^{\bar{w}}(i)$). Type τ' is expressible in FO and, since $t^{\bar{w}}(i)$ is finite, is expressible in TL^{tree} as well by Lemma A.4. Furthermore, by the separation property of [18], it is expressible by a formula $\beta'_{\tau'}$ that does not use \bigcirc_\uparrow and \mathbf{S}_\downarrow . This means that $(T_{\bar{w}}, s \cdot 1) \models \beta'_{\tau'}$ iff the rank- m type of $t_0^{\bar{w}}(i)$ is τ' . Hence, β_τ is expressible in TL^{tree} as a Boolean combination of propositional letters from Σ and formulas $\bigcirc_\downarrow \beta'_{\tau'}$. Note that in this case, β_τ does not use `pcall` and `pret`.

By Claim A.7, we need to express, for each node $s = \iota_{w-t}(i)$, the rank- m types of $s_m^{\leftarrow}(\bar{w}, s)$ and $s_m^{\rightarrow}(\bar{w}, s)$ in TL^{tree} over $T_{\bar{w}}$, as well as the rank- m type of $\tau^{\bar{w}}(i)$, in order to express a quantifier-rank m formula, as it will be a Boolean combination of such formulas. Given s , we need to define $\iota_{w-t}(i_p)$ – the outer position in whose scope s occurs – and then from that point evaluate two FO formulas, defining rank- m types of words over the alphabet of rank- m types of finite trees. By Kamp's theorem, each such FO formula is equivalent to an LTL formula whose propositional letters are rank- m types of trees.

Assume we have an LTL formula γ expressing the rank- m type τ_0 of $s_m^{\leftarrow}(\bar{w}, s)$. By Kamp's theorem and the separation property for LTL, it is written using only propositional letters, Boolean connectives, \bigcirc and \mathbf{U} (that is, no \ominus and \mathbf{S}). We now inductively take conjunction of each subformula of γ with $\neg(\bigcirc_{\leftarrow} \top)$ (i.e., a TL^{tree} formula which is true in left successors), replace LTL connectives \bigcirc and \mathbf{U} by \bigcirc_\downarrow and \mathbf{U}_\downarrow , and replace each propositional letter τ by β_τ , to obtain a TL^{tree} formula γ' . Then $(T_{\bar{w}}, \iota_{w-t}(i_p)) \models \gamma'$ iff $s_m^{\leftarrow}(\bar{w}, s)$ has type τ_0 . Thus, for a formula

$$\gamma'' = (\alpha_{\text{outer}} \wedge \gamma') \vee \neg \alpha_{\text{outer}} \mathbf{S}_\downarrow (\alpha_{\text{outer}} \wedge \gamma')$$

is true in $(T_{\bar{w}}, \iota_{w-t}(i))$ iff the rank- m type of $s_m^{\leftarrow}(\bar{w}, s)$ is τ_0 .

The proof for $s_m^{\rightarrow}(\bar{w}, s)$ is similar. Since this word is finite, by Kamp's theorem and the separation property, there is an LTL formula γ that uses \ominus , \mathbf{S} , propositional letters and Boolean connectives such that γ evaluated in the last position of the word expresses its rank- m type. Since there is exactly one path from each node to the root, to translate γ into a TL^{tree} formula γ' we just need to replace propositional letters by the corresponding formulas β_τ , and \ominus by \bigcirc_\uparrow . Then, as for the case of $s_m^{\leftarrow}(\bar{w}, s)$, we have that γ' evaluated in $\iota_{w-t}(i_p)$ expresses the type of $s_m^{\rightarrow}(\bar{w}, s)$. Then finally the same formula as in the case of $s_m^{\leftarrow}(\bar{w}, s)$ evaluated in s expresses that type.

Finally we need a TL^{tree} formula that expresses $\tau_m^{\bar{w}}(i)$, the rank- m type of $t^{\bar{w}}(i)$, when evaluated in $(T_{\bar{w}}, \iota_{w-t}(i))$. We can split this into two cases. If α_{outer} is true in $\iota_{w-t}(i)$, then, as explained earlier, the rank- m type of $t^{\bar{w}}(i)$ is a Boolean combination of propositional letters, and thus definable.

So we now consider the case when α_{outer} is not true in $\iota_{w-t}(i)$. Then $\tau_m^{\bar{w}}(i)$ is given by a Boolean combination of formulas that specify (1) the label of i_p , and (2) the rank- m type of $(t_0^{\bar{w}}(i_p), s)$, the subtree of $t^{\bar{w}}(i_p)$ rooted at the child of i_p with s as a distinguished node. This type can be expressed by a formula γ in TL^{tree} over $t_0^{\bar{w}}(i_p)$ by [18]. Hence if in γ we recursively take conjunction of each subformula with $\neg \alpha_{\text{outer}}$, we obtain a formula γ' of TL^{tree} that expresses the type of $(t_0^{\bar{w}}(i_p), s)$ when evaluated in $(T_{\bar{w}}, s)$. Thus, $\tau_m^{\bar{w}}(i)$ is expressible by a Boolean combination of formulas γ' and $\neg \alpha_{\text{outer}} \mathbf{S}_\downarrow (\alpha_{\text{outer}} \wedge a)$ where a is a propositional letter.

This completes the proof of translation of FO into TL^{tree} over nested ω -words, and thus the proof of the lemma. \square

Proof of Corollary 4.3

In the proof of Theorem 4.1, we show that every FO sentence over a nested word \bar{w} can be translated into an FO sentence over tree $T_{\bar{w}}$, and then, by the separation property of TL^{tree} [18] is equivalent to a TL^{tree} formula that does not use \mathbf{S}_\downarrow and \bigcirc_\uparrow .

Then, given that in the translation of TL^{tree} into NWTL^s we only use \mathbf{S}_s^σ in the rule $(\varphi \mathbf{S}_l \psi)^\circ = \varphi^\circ \mathbf{S}_s^\sigma \psi^\circ$, we see that the equivalent NWTL^s formula does not use \mathbf{S}_s^σ . Thus, given that in the proof of Lemma A.2, no since operator is used in the translations of \mathbf{U}_s^σ into \mathbf{U}_{ss}^σ and \mathbf{U}_{ss}^σ into \mathbf{U}^σ , the corollary follows for the finite case.

For the infinite case, we note that in the proof of Theorem 4.1, for the case of FO sentences we only need to specify the type of $s_m^\rightarrow(\bar{w}, s)$ where s is the root. Thus, one can see that in this case the use of \mathbf{S}_l in TL^{tree} formulas is not required, and hence the resulting formulas are translated into NWTL^s formulas without \mathbf{S}_s^σ .

Proof of Proposition 4.4

We shall look at finite nested words; the proof for the infinite case applies verbatim. To evaluate a formula φ of $\text{NWTL}^{\text{future}}$ in position i of a nested word \bar{w} of length n one only needs to look at $\bar{w}[i, n]$. That is, if \bar{w} and \bar{w}' of length n and n' respectively are such that $\bar{w}[i, n] \cong \bar{w}'[i', n']$, then $(\bar{w}, i) \models \varphi$ iff $(\bar{w}', i') \models \varphi$ for every formula φ of $\text{NWTL}^{\text{future}}$.

Furthermore, for every collection of $\text{NWTL}^{\text{future}}$ formulas $\Psi = \{\psi_1, \dots, \psi_l\}$, one can find a number $k = k(\Psi)$ such that

$$\bar{w}[i, n] \equiv_k \bar{w}'[i', n'] \text{ implies } (\bar{w}, i) \models \psi_p \Leftrightarrow (\bar{w}', i') \models \psi_p, \text{ for all } p \leq l.$$

In particular, if b^r stands for the word of length r in which all positions are labeled b and the matching relation is empty, we derive that there are numbers $k_1 = k_1(\Psi)$ and $k_2 = k_2(\Psi)$ with $k_1 > k_2$ such that

$$b^{k_1} \models \psi_p \Leftrightarrow b^{k_2} \models \psi_p, \text{ for all } p \leq l.$$

Now consider the following NWTL formula:

$$\alpha = \bigcirc_\mu \top \wedge \bigcirc_\mu \ominus a,$$

saying that the first position is a call, and the predecessor of its matching return is labeled a . We claim that this is not expressible in $\text{NWTL}^{\text{future}}$.

Assume to the contrary that there is a formula β of $\text{NWTL}^{\text{future}}$ equivalent to α . Let Ψ be the collection of all subformulas of β , including β itself, and let k_1 and k_2 be constructed as above. We now consider two nested words \bar{w}_1 and \bar{w}_2 of length $k_1 + 2$ whose underlying words are bab^{k_1} of length $n = k_1 + 2$, such that the matching relation μ_1 of \bar{w}_1 has one edge $\mu_1(1, 3)$, and the matching relation μ_2 of \bar{w}_2 has one edge $\mu_2(1, n + 1 - k_2)$. In other words, the only return position of \bar{w}_1 is $r_1 = 3$, and the only return position of \bar{w}_2 is $r_2 = n + 1 - k_2$, and thus $\bar{w}_1[r_1, n] = b^{k_1}$ and $\bar{w}_2[r_2, n] = b^{k_2}$. Further notice that for every $i > 1$ we have $\bar{w}_1[i, n] \cong \bar{w}_2[i, n]$.

Observe that $(\bar{w}_1, 1) \models \alpha$ and $(\bar{w}_2, 1) \models \neg\alpha$.

We now prove by induction on formulas in Ψ that for each such formula γ we have $(\bar{w}_1, 1) \models \gamma$ iff $(\bar{w}_2, 1) \models \gamma$, thus proving that β and α cannot be equivalent.

- The base case of propositional letters is immediate.
- The Boolean combinations are straightforward too.
- Let $\gamma = \bigcirc \psi$. Then

$$\begin{aligned} & (\bar{w}_1, 1) \models \gamma \\ \Leftrightarrow & (\bar{w}_1, 2) \models \psi \\ \Leftrightarrow & (\bar{w}_2, 2) \models \psi \\ \Leftrightarrow & (\bar{w}_2, 1) \models \gamma, \end{aligned}$$

since $\bar{w}_1[2, n] \cong \bar{w}_2[2, n]$.

- Let $\gamma = \bigcirc_\mu \psi$. Then

$$\begin{aligned} & (\bar{w}_1, 1) \models \gamma \\ \Leftrightarrow & (\bar{w}_1, 3) \models \psi \\ \Leftrightarrow & b^{k_1} \models \psi \\ \Leftrightarrow & b^{k_2} \models \psi \\ \Leftrightarrow & (\bar{w}_2, n + 1 - k_2) \models \psi \\ \Leftrightarrow & (\bar{w}_2, 1) \models \gamma, \end{aligned}$$

since $\psi \in \Psi$.

- Let $\gamma = \varphi \mathbf{U}^\sigma \psi$. Assume $(\bar{w}_1, 1) \models \gamma$. Consider three cases.

Case 1: $(\bar{w}_1, 1) \models \psi$. By the hypothesis $(\bar{w}_2, 1) \models \psi$ and we are done.

Case 2: The witness for $\varphi \mathbf{U}^\sigma \psi$ occurs beyond the only return. Then $(\bar{w}_1, 1) \models \varphi$ and $(\bar{w}_1, r_1) \models \varphi \mathbf{U}^\sigma \psi$. Since $\varphi \mathbf{U}^\sigma \psi \in \Psi$ we have $(\bar{w}_2, r_2) \models \varphi \mathbf{U}^\sigma \psi$, and by the hypothesis, $(\bar{w}_2, 1) \models \varphi$, so $(\bar{w}_2, 1) \models \varphi \mathbf{U}^\sigma \psi$.

Case 3: The witness for $\varphi \mathbf{U}^\sigma \psi$ occurs inside the call. Since for every position $i > 1$ we have $(\bar{w}_1, i) \models \varphi$ iff $(\bar{w}_2, i) \models \varphi$ and likewise for ψ , the same summary path witnesses $\varphi \mathbf{U}^\sigma \psi$ in \bar{w}_2 .

Thus, $(\bar{w}_2, 1) \models \gamma$.

Now assume $(\bar{w}_2, 1) \models \gamma$. In the proof of $(\bar{w}_1, 1) \models \gamma$ is the same as above in Cases 1 and 2. For Case 3, assume that in the path which is a witness for $\varphi \mathbf{U}^\sigma \psi$ the position in which ψ is true is the 2nd or the 3rd position in the word. Then the same path witnesses $(\bar{w}_1, 1) \models \gamma$, as in the proof of Case 3 above. Next assume it is a position with index j higher than 3 (which is still labeled b) where ψ first occurs. Then φ must be true in all positions i with $3 \leq i \leq j$ in \bar{w}_2 . Hence φ is true in all such positions in \bar{w}_1 as well, and thus the summary path in \bar{w}_1 that skips the first call (i.e. jumps from 1 to 3) witnesses $\varphi \mathbf{U}^\sigma \psi$. Hence, in all the cases $(\bar{w}_2, 1) \models \gamma$ implies $(\bar{w}_1, 1) \models \gamma$, which completes the inductive proof, and thus shows the inexpressibility of α in $\text{NWTL}^{\text{future}}$. \square

Proof of Theorem 4.6

The translation from $\text{LTL}^\mu + \mathcal{W}$ into FO is similar to the translation used in the proof of Theorem 4.7. To prove the other direction, we show how to translate NWTL into $\text{LTL}^\mu + \mathcal{W}$. More precisely, for every formula φ in NWTL , we show how to construct a formula α_φ in $\text{LTL}^\mu + \mathcal{W}$ such that for every nested word \bar{w} (finite or infinite) and position i in it, we have that $(\bar{w}, i) \models \varphi$ if and only if $(\bar{w}, i) \models \alpha_\varphi$.

Since LTL^μ includes the same past modalities as NWTL , α_φ is trivial to define for the atomic formulas, Boolean combinations and next and previous modalities:

$$\begin{aligned}
\alpha_{\top} &:= \top, \\
\alpha_{\text{call}} &:= \text{call}, \\
\alpha_{\text{ret}} &:= \text{ret}, \\
\alpha_a &:= a, \\
\alpha_{\neg\varphi} &:= \neg\alpha_\varphi, \\
\alpha_{\varphi \vee \psi} &:= \alpha_\varphi \vee \alpha_\psi, \\
\alpha_{\bigcirc\varphi} &:= \bigcirc\alpha_\varphi, \\
\alpha_{\bigcirc_\mu\varphi} &:= \bigcirc_\mu\alpha_\varphi, \\
\alpha_{\ominus\varphi} &:= \ominus\alpha_\varphi, \\
\alpha_{\ominus_\mu\varphi} &:= \ominus_\mu\alpha_\varphi.
\end{aligned}$$

Thus, we only need to show how to define $\alpha_{\varphi \mathbf{U}^\sigma \psi}$ and $\alpha_{\varphi \mathbf{S}^\sigma \psi}$. Formula $\alpha_{\varphi \mathbf{U}^\sigma \psi}$ is defined as:

$$\alpha_{\varphi \mathbf{U}^\sigma \psi} := \alpha_\varphi \mathbf{U}^a (\alpha_\psi \vee (\alpha_\varphi \wedge \bigcirc\alpha_\psi) \vee \mathcal{W} \diamond (\alpha_\psi \wedge \bigcirc\top \wedge (\neg\text{ret} \rightarrow \ominus(\beta \mathbf{S}^c \gamma)) \wedge (\text{ret} \rightarrow \ominus_\mu(\beta \mathbf{S}^c \gamma)))),$$

where β and γ are formulas defined as:

$$\begin{aligned}
\beta &:= \alpha_\varphi \mathbf{S}^a (\alpha_\varphi \wedge \neg\text{ret} \wedge \ominus((\alpha_\varphi \wedge \neg\text{ret}) \mathbf{S} (\alpha_\varphi \wedge \text{call}))), \\
\gamma &:= \neg\ominus\top.
\end{aligned}$$

Moreover, formula $\alpha_{\varphi \mathbf{S}^\sigma \psi}$ is defined as:

$$\alpha_{\varphi \mathbf{S}^\sigma \psi} := \delta \mathbf{S}^c (\alpha_\varphi \mathbf{S}^a \alpha_\psi),$$

where δ is a formula defined as:

$$\delta := \alpha_\varphi \mathbf{S}^a (\alpha_\varphi \wedge \neg\text{ret} \wedge \ominus((\alpha_\varphi \wedge \neg\text{ret}) \mathbf{S} \text{call})).$$

This concludes the proof of the theorem.

Proof of Theorem 4.7

In order to prove Theorem 4.7 we use the composition argument presented below.

Let \bar{w} be a nested word and i an element in \bar{w} . Let c_1, \dots, c_m , where $m \geq 0$, be all elements in \bar{w} such that, for each $j \in [1, m]$, $c_j < i$ and there is an element r_j such that $\mu(c_j, r_j)$ and $i \leq r_j$. Assume without loss of generality that $c_1 < c_2 < \dots < c_m$.

Fix $k \geq 0$. Let Γ be the set of all rank- k types of nested words with one distinguished constant. We define the word $\Omega_k(\bar{w}, i) = a_0 a_1 \dots a_m$ over alphabet $\Gamma \times \Gamma$ as follows:

- The element a_0 is labeled with the tuple whose first component is the rank- k type of $(\bar{w}[1, c_1 - 1], 1)$ and whose second component is the rank- k type of $(\bar{w}[r_1, \infty], 1)$ if $m \neq 0$; otherwise, it is labeled with the tuple whose first component is the rank- k type of $(\bar{w}[1, i - 1], 1)$ and whose second component is the rank- k type of $(\bar{w}[i, \infty], 1)$
- for each $0 < j < m$, the element a_j is labeled with the tuple whose first component is the rank- k type of $(\bar{w}[c_j, c_{j+1} - 1], 1)$ and whose second component is the rank- k type of $(\bar{w}[r_{j+1}, r_j - 1], 1)$; and
- if $m \neq 0$ then the element a_m is labeled with the the tuple whose first component is the rank- k type of $(\bar{w}[c_m, i - 1], 1)$ and whose second component is the rank- k type of $(\bar{w}[i, r_m - 1], 1)$.

The following is our composition argument:

Lemma A.8 (Composition Method) *Let \bar{w}_1 and \bar{w}_2 be two nested ω -words, and let i and i' be two elements in \bar{w}_1 and \bar{w}_2 , respectively, that share the same label in Σ , and such that i is a call (resp., return) iff i' is a call (resp., return). Then $\Omega_k(\bar{w}_1, i) \equiv_{k+2} \Omega_k(\bar{w}_2, i')$ implies $(\bar{w}_1, i) \equiv_k (\bar{w}_2, i')$.*

Proof: First we need to introduce some terminology. Let \bar{w} be a nested ω -word and i a position in \bar{w} . Assume elements $c_1, \dots, c_m, r_1, \dots, r_m$ are defined as above. With each element s of \bar{w} we associate an element $[s]$ of $\Omega_k(\bar{w}, i)$ as follows:

- If $m \neq 0$ and s belongs to $\bar{w}[0, c_1 - 1]$ or $\bar{w}[r_1, \infty]$ then $[s]$ is the first element of $\Omega_k(\bar{w}, i)$. In such case we say that $\bar{w}[0, c_1 - 1]$ and $\bar{w}[r_1, \infty]$ are the left and right intervals represented by $[s]$, respectively. If $m = 0$ and s belongs to $\bar{w}[0, i - 1]$ or $\bar{w}[i, \infty]$ then $[s]$ is also the first element of $\Omega_k(\bar{w}, i)$. In such case we say that $\bar{w}[0, i - 1]$ and $\bar{w}[i, \infty]$ are the left and right intervals represented by $[s]$, respectively.
- If $m \neq 0$ and s belongs to $\bar{w}[c_m, i - 1]$ or $\bar{w}[i, r_m - 1]$ then $[s]$ is the last element of $\Omega_k(\bar{w}, i)$. In such case we say that $\bar{w}[c_m, i - 1]$ and $\bar{w}[i, r_m - 1]$ are the left and right intervals represented by $[s]$, respectively.
- If $m \neq 0$ and s belongs to $\bar{w}[c_\ell, c_{\ell+1} - 1]$ or $\bar{w}[r_{\ell+1}, r_\ell - 1]$, for some $1 \leq \ell < m$, then $[s]$ is the $(\ell + 1)$ -th element of $\Omega_k(\bar{w}, i)$. In such case we say that $\bar{w}[c_\ell, c_{\ell+1} - 1]$ and $\bar{w}[r_{\ell+1}, r_\ell - 1]$ are the left and right intervals represented by $[s]$, respectively.

We denote by $[s]^L$ and $[s]^R$ the left and right intervals represented by $[s]$, respectively.

We now prove the lemma. For each round j ($0 \leq j \leq k$) of the k -round game on (\bar{w}_1, i) and (\bar{w}_2, i') , Player II's response b_j in \bar{w}_2 to an element a_j in \bar{w}_1 , played by Player I is defined as follows (the strategy for the case when Player I picks a point in \bar{w}_2 is completely symmetric). Assume that Player I plays element $[a_j]$ in $\Omega_k(\bar{w}_1, i)$ in the round j of the $(k + 2)$ -round game on $\Omega_k(\bar{w}_1, i)$ and $\Omega_k(\bar{w}_2, i')$. Then given that $\Omega_k(\bar{w}_1, i) \equiv_{k+2} \Omega_k(\bar{w}_2, i')$, Player II uses her winning strategy to choose a response $[q_j]$ in $\Omega_k(\bar{w}_2, i')$ to $[a_j]$. Thus, by definition of Ω_k , we have that the right and left intervals represented by $[a_j]$ have the same rank- k type than the right and left intervals represented by $[q_j]$ (with the first element distinguished as a constant), respectively. Hence, if a_j belongs to the left interval represented by $[a_j]$, then the Player II can find response b_j to a_j according to the winning strategy for the k -round game on $[a_j]^L$ and $[q_j]^L$, and if a_j belongs to the right interval represented by $[a_j]$, then the Player II can find response b_j to a_j according to the winning strategy for the k -round game on $[a_j]^R$ and $[q_j]^R$.

Assume that for round $0 \leq j < k$ the elements played by following this strategy are (1) $([p_1], \dots, [p_j])$ in $\Omega_k(\bar{w}_1, i)$, (2) $([q_1], \dots, [q_j])$ in $\Omega_k(\bar{w}_2, i')$, (3) (a_1, \dots, a_j) in \bar{w}_1 , and (4) (b_1, \dots, b_j) in \bar{w}_2 . We note that by definition of the strategy, for every $i \in [1, j]$, we have that $a_i = p_i$ or $b_i = q_i$. Since we assume that the $[p_j]$'s and $[q_j]$'s are played according to a winning strategy for Player II in the $(k + 2)$ -round game on $\Omega_k(\bar{w}_1, i)$ and $\Omega_k(\bar{w}_2, i')$, it is the case that:

$$\begin{aligned} (\Omega_k(\bar{w}_1, i), [p_1], \dots, [p_j]) &\equiv_{k-j+2} \\ (\Omega_k(\bar{w}_2, i'), [q_1], \dots, [q_j]). \end{aligned}$$

By the way the strategy is defined, for each $\ell \in [1, j]$, if \bar{a}_ℓ^L and \bar{a}_ℓ^R are the subtuples of (a_1, \dots, a_j) containing the elements from (a_1, \dots, a_j) that belong to $[a_\ell]^L$ and $[a_\ell]^R$, respectively, then the corresponding subtuples \bar{b}_ℓ^L and \bar{b}_ℓ^R of (b_1, \dots, b_j) contain the elements from (b_1, \dots, b_j) that belong to $[b_\ell]^L$ and $[b_\ell]^R$, respectively. Further, by definition of the strategy, we also have that $([a_\ell]^L, \bar{a}_\ell^L, 1) \equiv_{k-j} ([b_\ell]^L, \bar{b}_\ell^L, 1)$ and $([a_\ell]^R, \bar{a}_\ell^R, 1) \equiv_{k-j} ([b_\ell]^R, \bar{b}_\ell^R, 1)$ (where 1 represents the first element of the interval).

We now show how to define Player II's response in the round $j+1$. Let us assume without loss of generality that for round $j+1$ of the game on (\bar{w}_1, i) and (\bar{w}_2, i') , Player I picks an element a_{j+1} in \bar{w}_1 that belongs to the left interval represented by $[a_{j+1}]$ (all the other cases can be treated in a similar way). Player II response b_{j+1} in \bar{w}_2 is defined as follows. First, there must be an element $[s]$ in $\Omega_k(\bar{w}_2, i')$ such that

$$\begin{aligned} (\Omega_k(\bar{w}_1, i), [p_1], \dots, [p_j], [p_{j+1}]) &\equiv_{k-j+1} \\ (\Omega_k(\bar{w}_2, i'), [q_1], \dots, [q_j], [s]), \end{aligned}$$

where $p_{j+1} = a_{j+1}$. The latter, together with the way that the strategy is defined, implies that there is an element b in $[s]^L$ such that $([a_{j+1}]^L, \bar{a}', a_{j+1}, 1) \equiv_{k-j-1} ([s]^L, \bar{b}', b, 1)$, where \bar{a}' is the subtuple of (a_1, \dots, a_j) containing all the elements in (a_1, \dots, a_j) that belong to $[a_{j+1}]^L$ and \bar{b}' is the corresponding subtuple of (b_1, \dots, b_j) . We then set $b_{j+1} = b$.

We show by induction that, for each $j \leq k$, if (a_1, \dots, a_j) and (b_1, \dots, b_j) are the first j moves played by Player I and Player II on (\bar{w}_1, i) and (\bar{w}_2, i') , respectively, according to the strategy defined above, then $((a_1, \dots, a_j), (b_1, \dots, b_j))$ defines a partial isomorphism between (\bar{w}_1, i) and (\bar{w}_2, i') . This is sufficient to show that $(\bar{w}_1, i) \equiv_k (\bar{w}_2, i')$.

Assume $j = 0$. Then it follows from the statement of the lemma that i and i' share the same label in Σ , and that i is a call (resp., return) iff i' is a call (resp., return).

Assume that the property holds for j . Also assume without loss of generality that for the round $j+1$ of the game on (\bar{w}_1, i) and (\bar{w}_2, i') , Player I picks an element a_{j+1} in \bar{w}_1 that belongs to the right interval represented by $[a_{j+1}]$ (all the other cases can be treated in a similar way). We prove that b_{j+1} as defined above preserves the partial isomorphism. It is not hard to see that $a_{j+1} = i$ iff $b_{j+1} = i'$. Indeed, assume first that $i \neq r_m$. Then $[a_{j+1}]$ is the last element of $\Omega_k(\bar{w}_1, i)$, and $\Omega_k(\bar{w}_1, i) \equiv_{k+2} \Omega_k(\bar{w}_2, i')$ implies that $[b_{j+1}]$ is the last element of $\Omega_k(\bar{w}_2, i')$. Since $a_{j+1} = i$ is the first element of $[a_{j+1}]^R$, b_{j+1} is the first element of $[b_{j+1}]^R$, which is i' . Assume now that $i = r_m$. Then both the last elements of $\Omega_k(\bar{w}_1, i)$ and $\Omega_k(\bar{w}_2, i')$ are of the form (τ, τ_ε) , where τ_ε is the rank- k type of the empty nested word. Thus, i' is also a non-pending return, and $[a_{j+1}]$ is the penultimate element of $\Omega_k(\bar{w}_1, i)$. Since, $\Omega_k(\bar{w}_1, i) \equiv_{k+2} \Omega_k(\bar{w}_2, i')$, $[b_{j+1}]$ is the penultimate element of $\Omega_k(\bar{w}_2, i')$. But since i' is a non-pending return, it implies that the first element of the right interval associated with the penultimate element of $\Omega_k(\bar{w}_2, i')$ is also $b_{j+1} = i'$.

Further, it is also clear that the label of a_{j+1} in \bar{w}_1 is a iff the label of b_{j+1} in \bar{w}_2 is a , for each $a \in \text{Sigma}$. Next we consider the remaining cases.

- Assume that $a_{j+1} \in \text{call}$. Then $([a_{j+1}]^L, \bar{a}', a_{j+1}, 1) \equiv_{k-j-1} ([b_{j+1}]^L, \bar{b}', b_{j+1}, 1)$, where \bar{a}' is the subtuple of (a_1, \dots, a_j) containing all the elements in (a_1, \dots, a_j) that belong to $[a_{j+1}]^L$ and \bar{b}' is the corresponding subtuple of (b_1, \dots, b_j) . This immediately implies that $b_{j+1} \in \text{call}$. The converse is proved analogously.
- Assume that $a_{j+1} \in \text{ret}$. This is similar to the previous case.

- First, assume that $a_{j+1} < a_\ell$ holds for some $\ell \in [1, j]$. Since a_{j+1} belongs to $[a_{j+1}]^R$, we have that a_ℓ belongs to $[a_\ell]^R$ and, thus, we only need to consider the cases $[a_\ell] = [a_{j+1}]$ and $[a_{j+1}] < [a_\ell]$. If $[a_\ell] = [a_{j+1}]$, then $([a_\ell]^L, a_\ell, a_{j+1}) \equiv_0 ([b_\ell]^L, b_\ell, b_{j+1})$ and, therefore, $b_{j+1} < b_\ell$ also holds. If $[a_{j+1}] < [a_\ell]$, then $[b_{j+1}] < [b_\ell]$ and, thus, $b_{j+1} < b_\ell$ holds since b_ℓ and b_{j+1} belong to $[b_\ell]^L$ and $[b_{j+1}]^L$, respectively.

Second, assume that $a_\ell < a_{j+1}$ holds for some $\ell \in [1, j]$. We need to consider three cases: $[a_\ell] = [a_{j+1}]$, $[a_\ell] < [a_{j+1}]$ and $[a_{j+1}] < [a_\ell]$. If $[a_\ell] = [a_{j+1}]$, then $([a_\ell]^L, a_\ell, a_{j+1}) \equiv_0 ([b_\ell]^L, b_\ell, b_{j+1})$ and, therefore, $b_\ell < b_{j+1}$ also holds. If $[a_{j+1}] < [a_\ell]$, then a_ℓ belongs to $[a_\ell]^L$ and $[b_{j+1}] < [b_\ell]$ and, thus, $b_\ell < b_{j+1}$ holds since b_ℓ belongs to $[b_\ell]^L$ while b_{j+1} belongs to $[b_{j+1}]^R$. Finally, if $[a_\ell] < [a_{j+1}]$, then $[b_\ell] < [b_{j+1}]$ and, thus, $b_\ell < b_{j+1}$ holds since b_{j+1} belongs to $[b_{j+1}]^R$ and every element in $[b_{j+1}]^R$ is bigger than every element in either $[b_\ell]^R$ or $[b_\ell]^L$.

- First, assume that $\mu(a_{j+1}, a_\ell)$ holds for some $\ell \in [1, j]$. Since a_{j+1} belongs to the right interval represented by $[a_{j+1}]$, we have that $[a_\ell] = [a_{j+1}]$. Thus, given that $([a_\ell]^L, a_\ell, a_{j+1}) \equiv_0 ([b_\ell]^L, b_\ell, b_{j+1})$, we conclude that $\mu(b_{j+1}, b_\ell)$ holds.

Second, assume that $\mu(a_\ell, a_{j+1})$ holds for some $\ell \in [1, j]$. We need to consider two cases. If both a_ℓ and a_{j+1} belong to the same interval, then $([a_\ell]^L, a_\ell, a_{j+1}) \equiv_0 ([b_\ell]^L, b_\ell, b_{j+1})$, and thus, $\mu(b_\ell, b_{j+1})$ holds. If a_ℓ and a_{j+1} belong to

$$\begin{aligned}
\alpha_a(x) &:= P_a(x), \\
\alpha_{\text{call}}(x) &:= \text{call}(x), \\
\alpha_{\text{ret}}(x) &:= \text{ret}(x), \\
\alpha_{\text{int}}(x) &:= \neg \text{call}(x) \wedge \neg \text{ret}(x), \\
\alpha_{\text{pret}}(x) &:= \text{ret}(x) \wedge \neg \exists y \mu(y, x), \\
\alpha_{\neg \varphi}(x) &:= \neg \alpha_\varphi(x), \\
\alpha_{\varphi \vee \psi}(x) &:= \alpha_\varphi(x) \vee \alpha_\psi(x), \\
\alpha_{\bigcirc_\varphi}(x) &:= \exists y (x < y \wedge \neg \exists z (x < z \wedge z < y) \wedge \alpha_\varphi(y)), \\
\alpha_{\bigcirc_\mu \varphi}(x) &:= \exists y (\mu(x, y) \wedge \alpha_\varphi(y)), \\
\alpha_{\bigcirc_c \varphi}(x) &:= \exists y \exists z (y < x \wedge x < z \wedge \mu(y, z) \wedge \alpha_\varphi(y) \wedge \\
&\quad \forall u \forall v (u < x \wedge x < v \wedge \mu(u, v) \rightarrow u = y \vee u < y)), \\
\alpha_{\varphi \mathbf{U} \psi}(x) &:= \exists y ((x < y \vee x = y) \wedge \alpha_\psi(y) \wedge \\
&\quad \forall z (z < y \wedge (z = x \vee x < z) \rightarrow \alpha_\varphi(z))), \\
\alpha_{\varphi \mathbf{U}^a \psi}(x) &:= \exists y ((x < y \vee x = y) \wedge \alpha_\psi(y) \wedge \forall u \forall v (u < y \wedge y < v \wedge \mu(u, v) \rightarrow u < x) \wedge \\
&\quad \forall z (z < y \wedge (z = x \vee x < z) \wedge \forall u \forall v (u < z \wedge z < v \wedge \mu(u, v) \rightarrow u < x) \rightarrow \alpha_\varphi(z))), \\
\alpha_{\varphi \mathbf{S}^c \psi}(x) &:= \alpha_\psi(x) \vee \exists y (y < x \wedge \alpha_{\text{call}}(y) \wedge \forall z (\mu(y, z) \rightarrow x < z) \wedge \alpha_\psi(y) \wedge \\
&\quad \forall z (((z = x) \vee (\alpha_{\text{call}}(z) \wedge z < x \wedge y < z \wedge \forall u (\mu(z, u) \rightarrow x < u))) \rightarrow \alpha_\varphi(z))), \\
\alpha_{\mathcal{C} \varphi}(x) &:= (\neg \exists y \exists z (\mu(y, z) \wedge y < x \wedge x < z) \wedge \forall z (\neg \exists u (u < z) \rightarrow \alpha_\varphi(z)^{[z, x]}) \vee \\
&\quad (\exists y \exists z (\mu(y, z) \wedge y < x \wedge x < z \wedge \\
&\quad \quad \forall u \forall v (u < x \wedge x < v \wedge \mu(u, v) \rightarrow u = y \vee u < y) \wedge \alpha_\varphi(y)^{[y, x]})), \\
\alpha_{\mathcal{R} \varphi}(x) &:= (\neg \exists y \exists z (\mu(y, z) \wedge y < x \wedge x < z) \wedge \forall z (\neg \exists u (z < u) \rightarrow \alpha_\varphi(x)^{[x, z]}) \vee \\
&\quad (\exists y \exists z (\mu(y, z) \wedge y < x \wedge x < z \wedge \\
&\quad \quad \forall u \forall v (u < x \wedge x < v \wedge \mu(u, v) \rightarrow u = y \vee u < y) \wedge \alpha_\varphi(x)^{[x, z]})).
\end{aligned}$$

Figure 1. Translating CaRet + $\{\mathcal{C}, \mathcal{R}\}$ into FO.

distinct intervals, then $[a_\ell] = [a_{j+1}] + 1$, a_ℓ is the first element of $[a_\ell]^L$ and a_{j+1} is the first element of $[a_{j+1}]^R$. Thus, given that $(\Omega_k(\bar{w}_1, i), [a_\ell], [a_{j+1}]) \equiv_1 (\Omega_k(\bar{w}_2, i'), [b_\ell], [b_{j+1}])$, it is the case that $[b_\ell] = [b_{j+1}] + 1$. Furthermore, given that the rank- k type of $[a_\ell]^L$ is the same as the rank- k type of $[b_\ell]^L$ with the first element distinguished as a constant, we conclude that b_ℓ is the first element of $[b_\ell]^L$. In the same way, we conclude that b_{j+1} is the first element of $[b_{j+1}]^R$ and, therefore, $\mu(b_\ell, b_{j+1})$ holds.

This concludes the proof of the lemma. \square

We now present the proof of Theorem 4.7.

Proof of Theorem 4.7: We first show that every CaRet + $\{\mathcal{R}, \mathcal{C}\}$ formula φ is equivalent to an FO formula $\alpha_\varphi(x)$ over nested words, that is, for every nested word \bar{w} (finite or infinite), we have $(\bar{w}, i) \models \varphi$ iff $\bar{w} \models \alpha_\varphi(i)$. The translation is standard and can be done by recursively defining $\alpha_\varphi(x)$ from φ as shown in Figure A. In the following we assume that the FO formula $\theta(x)^{[y, z]}$ is the *relativization* of $\theta(x)$ to elements in the interval $[y, z]$, that is, $\theta(x)^{[y, z]}$ is obtained from $\theta(x)$ by replacing each subformula of the form $\exists u \beta$ with $\exists u (y \leq u \wedge u \leq z \wedge \beta)$ and each subformula of the form $\forall u \beta$ with $\forall u (y \leq u \wedge u \leq z \rightarrow \beta)$.

We now show the other direction, that is, $\text{FO} \subseteq \text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$. We extend the vocabulary with an extra atomic predicate min interpreted as the first element of each nested word \bar{w} . We start by proving this result for FO sentences (that is, we prove that for every FO sentence φ there is an CaRet + $\{\mathcal{R}, \mathcal{C}\}$ formula ψ , such that $\bar{w} \models \varphi$ iff $(\bar{w}, 1) \models \psi$), and then extend it to the case of FO formulas with one free variable. Let φ be an FO sentence. We use induction on the quantifier rank to prove that φ is equivalent to an CaRet + $\{\mathcal{R}, \mathcal{C}\}$ formula.

For $k = 0$ the property trivially holds, as φ is a Boolean combination of formulas of the form $P_a(\text{min})$, $\text{min} < \text{min}$,

$\min = \min$, and $\mu(\min, \min)$.

We now prove for $k + 1$ assuming that the property holds for k . Since every FO sentence of quantifier rank $k + 1$ is a Boolean combination of FO sentences of the form $\exists x\varphi(x)$, where $\varphi(x)$ is a formula of quantifier rank k , we just have to show how to express in $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ a sentence of this form.

Let Γ be the set of all rank- k types of nested words over alphabet Σ . By induction hypothesis, for each $\tau \in \Gamma$ there is an $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ formula ξ_τ such that $(\bar{w}, 1) \models \xi_\tau$ iff the rank- k type of \bar{w} is τ . It is not hard to see then that there is an $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ formula ξ_τ^ℓ such that $(\bar{w}, 1) \models \xi_\tau^\ell$ iff the rank- k type of \bar{w}^ℓ is τ , where \bar{w}^ℓ is the nested word obtained from \bar{w} by removing its last element (note that in order to construct ξ_τ^ℓ we require the temporal operator \bigcirc).

Let Λ be the set of all rank- $(k+2)$ types of words over alphabet $\Gamma \times \Gamma$. We first construct, for each $\lambda \in \Lambda$, an $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ formula α_λ over alphabet Σ such that,

$$(\bar{w}, i) \models \alpha_\lambda \iff \text{the rank-}(k+2) \text{ type of } \Omega_k(\bar{w}, i) \text{ is } \lambda,$$

for each nested word \bar{w} and position i of \bar{w} .

Fix $\lambda \in \Lambda$. From Kamp's theorem [12] there is an LTL formula β_λ over alphabet $\Gamma \times \Gamma$ such that a word s satisfies β_λ evaluated on its last element iff the rank- $(k+2)$ type of s is λ . By the separation property, we can assume that β_λ only mentions past modalities \ominus and \mathbf{S} . Moreover, given that $\varphi \mathbf{S} \psi \equiv \psi \vee (\varphi \wedge \ominus(\varphi \mathbf{S} \psi))$, we can also assume that β_λ is a Boolean combination of formulas of the form either φ or $\ominus\psi$, where φ does not mention any temporal modality and ψ is an arbitrary past LTL formula. Thus, since $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ is closed under Boolean combinations, to show how to define α_λ from β_λ , we only need to consider two cases: (1) β_λ is an LTL formula over $\Gamma \times \Gamma$ without temporal modalities, and (2) β_λ is of the form $\ominus\psi$, where ψ is an arbitrary past LTL formula over $\Gamma \times \Gamma$. Next we consider these two cases.

- Assume that β_λ is an LTL formula without temporal modalities. Then α_λ is defined to be β_λ° , where $(\)^\circ$ is defined recursively as follows. Given $(\tau, \tau') \in \Gamma \times \Gamma$,

$$\begin{aligned} (\tau, \tau')^\circ &:= (\neg \text{ret} \wedge \mathcal{C}\xi_\tau^\ell \wedge \mathcal{R}\xi_{\tau'}^\ell) \vee \\ &\quad (\text{ret} \wedge \mathcal{C}\xi_\tau^\ell \wedge \xi_{\tau_\varepsilon}), \end{aligned}$$

with τ_ε the rank- k type of the empty word. Furthermore, if ψ and φ are LTL formulas without temporal modalities, then

$$\begin{aligned} (\neg\varphi)^\circ &:= \neg\varphi^\circ, \\ (\varphi \vee \psi)^\circ &:= \varphi^\circ \vee \psi^\circ. \end{aligned}$$

- Assume that β_λ is a formula of the form $\ominus\varphi$, where φ is an arbitrary past LTL formula. Then α_λ is defined to be β_λ^* , where $(\)^*$ is defined recursively as follows. Given $(\tau, \tau') \in \Gamma \times \Gamma$,

$$\begin{aligned} (\tau, \tau')^* &:= (\ominus_c \top \wedge \mathcal{C}\xi_\tau^\ell \wedge \bigcirc_\mu \mathcal{R}\xi_{\tau'}^\ell) \vee \\ &\quad (\neg \ominus_c \top \wedge \mathcal{C}\xi_\tau^\ell \wedge \bigcirc_\mu \mathcal{R}\xi_{\tau'}) \end{aligned}$$

Furthermore, if ψ and φ are past LTL formulas, then

$$\begin{aligned} (\neg\varphi)^* &:= \neg\varphi^*, \\ (\varphi \vee \psi)^* &:= \varphi^* \vee \psi^*, \\ (\ominus\varphi)^* &:= \ominus_c \varphi^*, \\ (\varphi \mathbf{S} \psi)^* &:= \varphi^* \mathbf{S}^c \psi^*. \end{aligned}$$

Now, let $\exists x\varphi(x)$ be an FO sentence such that the quantifier rank of $\varphi(x)$ is k . Then, from our composition method $\varphi(x)$ can be expressed in $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ as the formula $\bigvee_{\lambda \in \Lambda'} \alpha_\lambda$, where $\Lambda' \subseteq \Lambda$ is the set of all rank- $(k+2)$ types of words over alphabet $\Gamma \times \Gamma$ that belong to $\{\Omega_k(\bar{w}, i) \mid \bar{w} \models \varphi(i)\}$. Thus, $\exists x\varphi(x)$ can be expressed as the following $\text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ formula: $\top \mathbf{U} (\bigvee_{\lambda \in \Lambda'} \alpha_\lambda)$. This concludes the proof of the theorem.

Finally, from the composition method and the previous proof we see that the equivalence $\text{FO} = \text{CaRet} + \{\mathcal{R}, \mathcal{C}\}$ holds for unary queries over nested words. \square

Proof of Theorem 5.6

From the FO completeness of NWTL^+ , we have that $\text{NWTL}^+ + \mathcal{W}$ can be translated into NWTL^+ . We show that at least an exponential blow-up is necessary for such translation. More precisely, we construct a sequence $\{\varphi_n\}_{n \geq 1}$ of $\text{NWTL}^+ + \mathcal{W}$ formulas of size $O(n)$, such that the shortest NWTL^+ formula that is equivalent to φ_n is of size $2^{\Omega(n)}$. Our proof is a modification of similar proofs given in [9, 22]. Assume $\Sigma = \{a_0, \dots, a_n\}$, and let φ_n be the following $\text{NWTL}^+ + \mathcal{W}$ formula (here, $\Box^\sigma \theta$ and $\Diamond^\sigma \theta$ are abbreviations for $\neg(\top \mathbf{U}^\sigma \neg \theta)$ and $\top \mathbf{S}^\sigma \theta$, respectively):

$$\Box^\sigma \left(\text{call} \rightarrow \mathcal{W} \Box^\sigma \left(\left(\bigwedge_{i=1}^n (a_i \leftrightarrow \Diamond^\sigma (a_i \wedge \neg \ominus \top)) \right) \rightarrow (a_0 \leftrightarrow \Diamond^\sigma (a_0 \wedge \neg \ominus \top)) \right) \right).$$

It is not hard to see that $\bar{w} \models \varphi_n$ iff for every positions i, j in \bar{w} such that $\mu(i, j)$ holds, if position ℓ in $\bar{w}[i, j]$ coincides with i on a_1, \dots, a_n , then ℓ also coincides with i on a_0 .

It is shown in Theorem 5.1 that for each NWTL^+ formula α , the language

$$L_\alpha = \{ \bar{w} \mid \bar{w} \text{ is a nested } \omega\text{-word such that } \bar{w} \models \alpha \}$$

is recognized by a nondeterministic nested word automaton with $2^{O(|\alpha|)}$ states. Thus, it is enough to show that every such automaton for L_{φ_n} requires at least $2^{2^{\Omega(n)}}$ states. Let A be a nondeterministic nested word automaton for L_{φ_n} . Assume that b_0, \dots, b_{2^n-1} is an enumeration of the symbols in $2^\Sigma \setminus \{a_0\}$. For every $K \subseteq \{0, \dots, 2^n - 1\}$ let \bar{w}_K be the word $c_0 \cdots c_{2^n-1}$ over alphabet 2^Σ , where for each $i \leq 2^n - 1$:

$$c_i = \begin{cases} b_i & i \in K \\ b_i \cup \{a_0\} & \text{otherwise} \end{cases}$$

It is not hard to see that for each $K \subseteq \{0, \dots, 2^n - 1\}$, the nested ω -word (\bar{w}_K^ω, μ) , where $\mu = \{(c_j, c_{3 \cdot 2^n - 1 - j}) \mid 0 \leq j \leq 2^n - 1\}$, is such that $(\bar{w}_K^\omega, \mu) \models \varphi_n$. Let (q_K^1, q_K^2) and $(q_{K'}^1, q_{K'}^2)$ be pairs of states such that (1) A is in states q_K^1 and q_K^2 in an accepting run for (\bar{w}_K^ω, μ) after reading 2^n and $2 \cdot 2^n$ symbols from \bar{w}_K^ω , respectively, and (2) A is in states $q_{K'}^1$ and $q_{K'}^2$ in an accepting run for $(\bar{w}_{K'}^\omega, \mu)$ after reading 2^n and $2 \cdot 2^n$ symbols from $\bar{w}_{K'}^\omega$, respectively. Next we show that $(q_K^1, q_K^2) \neq (q_{K'}^1, q_{K'}^2)$ if $K \neq K'$. On the contrary, assume that $(q_K^1, q_K^2) = (q_{K'}^1, q_{K'}^2)$. Then A accepts $(\bar{w}_K \bar{w}_{K'} \bar{w}_K^\omega, \mu')$, where $\mu' = \{(c_j, c_{3 \cdot 2^n - 1 - j}) \mid 0 \leq j \leq 2^n - 1\}$, which is a contradiction since $(\bar{w}_K \bar{w}_{K'} \bar{w}_K^\omega, \mu') \not\models \varphi_n$. Given that the number of different K 's is 2^{2^n} , the latter implies that the number of different pairs of states of A is at least 2^{2^n} . Thus, if A has m states, then $m^2 \geq 2^{2^n}$ and, hence, $m \geq 2^{2^{n-1}}$. Therefore, the number of different states of A is at least $2^{2^{\Omega(n)}}$. This concludes the proof of the theorem.

Proof of Theorem 6.1

As we mentioned already, in the finite case this is a direct consequence of [19] so we concentrate on the infinite case. It is more convenient for us to prove the result for ordered unranked forests in which a subtree rooted at every node is finite. The way to translate a nested ω -word into such a forest is as follows: when a matched call i with $\mu(i, j)$ is encountered, it defines a subtree with i as its root, and $j + 1$ as the next sibling (note that this is different from the translation into binary trees we used before). If i is an internal position, or a pending call or a pending return position, then it has no descendants and its next sibling is $i + 1$. Matched returns do not have next sibling, nor do they have any descendants. The nodes in the forest are labeled with `call`, `ret`, and the propositions in Σ , as in the original nested word.

It is routine to define, in FO, relations \preceq_{desc} and \preceq_{sib} for descendant and younger sibling in such a forest. Furthermore, from these relations, we can define the usual \leq and μ in nested words using at most 3 variables as follows. For $x \leq y$, the definition is given by

$$(y \preceq_{\text{desc}} x) \vee \exists z \left(x \preceq_{\text{desc}} z \wedge \exists x (z \prec_{\text{sib}} z \wedge y \preceq_{\text{desc}} x) \right)$$

and for $\mu(x, y)$, by

$$(y \preceq_{\text{desc}} x) \wedge \forall z \left((z \preceq_{\text{desc}} x) \rightarrow \exists x (x = z \wedge x \leq y) \right).$$

Thus, it suffices to prove the three-variable property for such ordered forests, which will be referred to as \mathcal{A} , \mathcal{B} , etc. We shall use pebble games. Let $\mathbf{G}_m^v(\mathcal{A}, a_1, b_1, \mathcal{B}, b_1, b_2)$ be the m -move, v -pebble game on structures \mathcal{A} and \mathcal{B} where initially

pebbles x_i are placed on a_i in \mathcal{A} and b_i in \mathcal{B} . Player II has a winning strategy for $\mathbf{G}_m^v(\mathcal{A}, a_1, b_1, \mathcal{B}, b_1, b_2)$ iff \mathcal{A}, a_1, a_2 and \mathcal{B}, b_1, b_2 agree on all formulas with at most v variables and quantifier-depth m . We know from [13] that to prove Theorem 6.1, it suffices to show the following,

Claim A.9 *For all k , if Player II has a winning strategy for the game $\mathbf{G}_{3k+2}^3(\mathcal{A}, a_1, a_2; \mathcal{B}, b_1, b_2)$. Then she also has a winning strategy for the game $\mathbf{G}_k^k(\mathcal{A}, a_1, a_2; \mathcal{B}, b_1, b_2)$.*

We will show how Player II can win the k -pebble game by maintaining a set of 3-pebble sub-games on which she will copy Player I's moves and decide on good responses using her winning strategy for these smaller 3-pebble games. The choice of these sub-games will partition the universe $|\mathcal{A}| \cup |\mathcal{B}|$ so that each play by Player I in the k -pebble game will be answered in one 3-pebble game. This is similar to the proof that linear orderings have the 3-variable property [13].

The subgames, $\mathbf{G}_m^3(\mathcal{A}, a_1, a_2; \mathcal{B}, b_1, b_2)$, that Player II maintains will all be *vertical* in which $a_2 \preceq_{\text{desc}} a_1$ and $b_2 \preceq_{\text{desc}} b_1$ hold, or *horizontal* in which $a_1 \prec_{\text{sib}} a_2$ and $b_1 \prec_{\text{sib}} b_2$ hold.

The following lemma gives the beginning strategy of Player II in which she replaces an arbitrary game configuration with a set of configurations each of which is vertical or horizontal.

Lemma A.10 *If Player II wins $\mathbf{G}_{m+4}^3(\mathcal{A}, a_1, a_2; \mathcal{B}, b_1, b_2)$. Then there are points a'_1, a'_2 from \mathcal{A} and b'_1, b'_2 from \mathcal{B} such that Player II wins the horizontal game $\mathbf{G}_{m+2}^3(\mathcal{A}, a'_1, a'_2; \mathcal{B}, b'_1, b'_2)$ and the vertical games $\mathbf{G}_{m+2}^3(\mathcal{A}, a'_i, a_i; \mathcal{B}, b'_i, b_i)$ for $i = 1, 2$.*

Proof: For this proof since \mathcal{A} and \mathcal{B} are fixed, we will describe a game only by listing the chosen points, e.g., $(a_1, a_2; b_1, b_2)$. We simulate two moves of the game, $\mathbf{G}_{m+4}^3(a_1, a_2; b_1, b_2)$, in which we choose Player I's moves and then Player II answers according to her winning strategy. Let $u + v$ denote the least common ancestor of u and v . First, we have Player I place pebble x_3 on a'_1 , the unique child of $a_1 + a_2$ that is an ancestor of a_1 . (Note that if $a'_1 = a_1$ then this move can be skipped and similarly for the second move if $a'_2 = a_2$.) Player II answers by placing x_3 on some point b'_1 . Second, Player I should move pebble x_1 from a_1 to a'_2 , the unique child of $a_1 + a_2$ that is an ancestor of a_2 . Player II moves x_1 to some point b'_2 .

Since Player II has moved according to her winning strategy, we have that she still has a winning strategy for the three games in the statement of the lemma. Furthermore, since a'_1 and a'_2 are siblings and we have two remaining moves, b'_1 and b'_2 must be siblings as well. \square

Using Lemma A.10 we initially partition the universe according to four subgames:

- $(a_r, a_p; b_r, b_p)$ with domain everything not below a_p or b_p . Here $a_p = a_1 + a_2$, i.e., the parent of a'_1 , $b_p = b_1 + b_2$, i.e., the parent of b'_1 and a_r and b_r are the roots of \mathcal{A} and \mathcal{B} , (the roots are not necessary but then the subgames are all on horizontal or vertical pairs), or
- $(a'_1, a_1; b'_1, b_1)$ with domain everything below a'_1 or b'_1 ,
- $(a'_2, a_2; b'_2, b_2)$, with domain everything below a'_2 or b'_2 ,
- $(a'_1, a'_2; b'_1, b'_2)$, with the remaining domain.

We now have to explain, inductively, how all moves of Player I in the k -pebble game are answered by Player II and how, in the process, the universe is further partitioned. We inductively assume that Player II has a winning strategy for each of the 3-pebble, m -move sub-games. There are two cases:

Vertical: Player I places a new pebble on a point a that is in the domain of a vertical game: $(a_1, a_2; b_1, b_2)$. We thus know that a_1 is a proper ancestor of a . The interesting case is where neither of a and a_2 is above the other so, without loss of generality, assume that $a < a_2$. We place x_3 on a'_2 , the child of $a + a_2$ that is above a_2 . Let Player II move according to her winning strategy, placing x_3 on some point b'_2 . We split the original game into $(a_1, a'_2; b_1, b'_2)$ and $(a'_2, a_2; b'_2, b_2)$ so Player II has a winning strategy for these 3-pebble, $m - 1$ move sub-games. Next, in the $(a_1, a'_2; b_1, b'_2)$ game we place x_3 on a_p , the parent of a'_2 and we let Player II answer according to her winning strategy, placing x_3 on some point, b_p . We then split off the game $(a_1, a_p; b_1, b_p)$.

Returning to the game $(a_1, a'_2; b_1, b'_2)$, we have Player I place x_3 on a' , the sibling of a'_2 above a , and let Player II answer according to her winning strategy, placing x_3 on some point, b' .

Finally, we let Player I move x_1 to a , and let Player II reply with x_1 on some point b .

The sub-games are thus: $(a_1, a_p; b_1, b_p)$, $(a', a'_2; b', b'_2)$, $(a', a; b', b)$, and $(a'_2, a_2; b'_2, b_2)$ and Player II has winning strategies for the \mathbf{G}_{m-3}^3 game on all of them.

Horizontal: In this case, we have the configuration, $(a_1, a_2; b_1, b_2)$, consisting of a pair of siblings. The only interesting case occurs when Player I puts a new pebble on some vertex, a , s.t. $a_1 < a < a_2$. In this case, we have Player I place pebble x_3 on a' , the sibling of a_1 above a . Player II will place pebble x_3 on some vertex, b' , which must be a sibling of b_1 and b_2 .

Next, in the game below a' and b' , we let Player I place pebble x_2 on a and we let Player II answer according to her winning strategy in this game, placing x_2 on some vertex, b . The domain of the original configuration is thus split into domains for three sub-games: $(a_1, a'; b_1, b')$, $(a', a_2; b', b_2)$, and $(a', a; b', b)$. On each of these, Player II has a winning strategy for the 3-pebble, $m - 2$ move game.

We now complete the proof that Player II wins $\mathbf{G}_k^k(a_1, a_2; b_1, b_2)$. Whenever Player I places a new pebble on some point, say a , in the original game, Player II will answer as described above, i.e., in one of the little games we will have Player II wins $\mathbf{G}_{3r}^3(a, a'; b, b')$ where there are r moves remaining in the big game.

Player II then answers in the big game by placing the corresponding pebble on b . To see that the resulting moves are a win for Player II, we must just consider any two pebbled points, $a_i, a_j \in \mathcal{A}$, and $b_i, b_j \in \mathcal{B}$. If they came from the same sub-game, then they agree on relations $\preceq_{\text{desc}}, \prec_{\text{sib}}$ because Player II wins the sub-game. Otherwise, a_i, b_i came from one sub-game, G_i , and a_j, b_j came from another sub-game, G_j . By our choice of the domains and transitivity of $\preceq_{\text{desc}}, \prec_{\text{sib}}$, it thus follows that a_i, a_j stand in the same relation with respect to $\preceq_{\text{desc}}, \prec_{\text{sib}}$ as b_i, b_j do. \square

Proof of Theorem 6.2

The translation from unary-NWTL into FO^2 is standard and can be done with negligible blow-up in the size of the formula, so we concentrate on the other direction. The proof generalizes the proof of an analogous result for unary temporal logic over words from [9].

Given an FO^2 formula $\varphi(x)$ the translation procedure works as follows. When $\varphi(x)$ is atomic, i.e., of the form $a(x)$, it outputs a . 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

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

where β is a propositional formula, each formula χ_i is an atomic order formula, each formula ξ_i is an atomic or existential FO^2 formula with $\text{qdp}(\xi_i) < \text{qdp}(\varphi)$, and each formula ζ_i is an atomic or existential FO^2 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. We obtain the following equivalent formulation for φ :

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

We proceed by a case distinction on which order relation holds between x and y , where $x \leq y$. We consider mutually exclusive cases, determined by the following formulas, which we call *order types*.

- Ψ_0 is $x = y$.
- For each $t \in T$, Ψ_t is $S^t(x, y)$.
- For each $t \in T$, Φ_t is $\exists z (S^t(x, z) \wedge z <^t y)$.
- Let $o = t_1, t_2, \dots, t_k$ be a sequence over T such that $2 \leq k \leq 5$, all t_i 's are distinct, and a call never appears before return (that is, if $t_i = c$ then $t_j \neq r$ for $j > i$). Then Ψ_o stands for

$$\exists z_1, z'_1, z_2, z'_2, \dots, z_k (S^{t_1}(x, z_1) \wedge z_1 \leq^{T_1} z'_1 \wedge S^{t_2}(z'_1, z_2) \wedge z_2 \leq^{T_2} z'_2 \wedge \dots \wedge z_k \leq^{T_k} y)$$

where for $1 \leq i \leq k$, the set T_i equals the set $\{t_1, t_2, \dots, t_i\}$, but with r removed if both c and r belong to this set.

We claim that these order types are mutually exclusive and complete, and are expressible in unary-NWTL (and hence, in FO^2). First, let us show that the order types form a disjoint partition, meaning for all pairs (x, y) such that $x \leq y$, we have exactly one of these relationships holding true. To see this, suppose $x < y$. Then either $S^t(x, y)$ holds for some type t (and the successor relations S^t 's are disjoint), or there is a path from x to y that uses at least two edges. The key observation is that a path from x to y is a summary path iff the path does not contain a call edge followed later by a return edge. Also, there is a unique summary path from x to y . We can now classify the paths by the edge types that this unique summary path contains, and the order in which they first appear in the path. For example, $\Phi_c(x, y)$ holds when there is a path from x to y using 2 or more call edges; $\Phi_{c,cir}(x, y)$ holds when there is a path from x to y which begins with a call edge, uses at least one call-interior-return summary edge, and uses only these two types of edges; $\Phi_{r,i,c}(x, y)$ holds when there is a path from x to y that can be split into three consecutive parts: a part containing only return edges, a part containing at least one internal and only internal and return edges, and a part containing at least one call and only call and internal edges. Note that some of these order types are empty: for example, two summary edges can never follow one another, and hence $\Phi_{cr}(x, y)$ can never hold. Emptiness of some of the order types is not relevant to the proof.

When we assume that one of these order types is true, each atomic order formula evaluates to either \top or \perp , in particular, each of the χ_i 's evaluates to either \top or \perp ; we will denote this truth value by χ_i^τ . For example, when $\Psi_{cr}(x, y)$ holds then (1) $S^t(x, y)$ is true for $t = cr$ and false for $t \neq cr$, and (2) \leq^Γ is true if Γ contains cr or if Γ contains both c and r , and false otherwise.

We can finally rewrite φ as follows, where Υ stands for the set of all order types:

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

If τ is an order type, $\psi(x)$ an FO^2 formula, and ψ' an equivalent unary-NWTL formula, there is a way to obtain a unary-NWTL formula $\tau\langle\psi\rangle$ equivalent to $\exists y(\tau \wedge \psi(y))$, as follows. Assume that $x \leq y$.

- For the order type Ψ_0 , $\tau\langle\psi'\rangle$ is ψ' itself.
- For each $t \in T$, for the order type Ψ_t , $\tau\langle\psi'\rangle$ is $\bigcirc^t \psi'$.
- For each $t \in T$, for the order type Φ_t , $\tau\langle\psi'\rangle$ is $\bigcirc^t \bigcirc^t \diamond^{\{t\}} \psi'$.
- For order type Ψ_o , where $o = t_1, t_2, \dots, t_k$ is a sequence over T , $\tau\langle\psi'\rangle$ is $\bigcirc^{t_1} \diamond^{T_1} \bigcirc^{t_2} \dots \diamond^{T_k} \psi'$, where for $1 \leq i \leq k \leq 5$, the set T_i equals the set $\{t_1, t_2, \dots, t_i\}$, but with r removed if both c and r belong to this set.

The case corresponding to past operators is analogous. Our procedure will therefore recursively compute ξ'_i for $i < s$ and $\zeta'_i(x)$ for $i < t$ and output

$$\bigvee_{\bar{\gamma} \in \{\top, \perp\}^s} \left(\bigwedge_{i < s} (\xi'_i \leftrightarrow \gamma_i) \wedge \bigvee_{\tau \in \Upsilon} \tau \langle \beta(\chi_0^\tau, \dots, \chi_{r-1}^\tau, \bar{\gamma}, \zeta'_0(x), \dots, \zeta'_{t-1}(x)) \rangle \right) . \quad (2)$$

Now we verify that $|\varphi'|$ and $\text{odp}(\varphi')$ are bounded as stated in the theorem. Note that the size $|\varphi'|$ is measured by viewing the unary-NWTL formula as a DAG, i.e., sharing identical subformulas. That $\text{odp}(\varphi') \leq 10 \text{qdp}(\varphi)$ is easily seen from the operator depth in the translation table above. The proof that $|\varphi'| \leq 2^{c|\varphi|(\text{qdp}(\varphi)+1)}$ for some constant c is inductive on the quantifier depth of φ . The base case is trivial, and the only interesting case in the inductive step is when φ is of the form $\exists y \varphi^*(x, y)$ as above. In this case, we have to estimate the length of (2). There are $2^s \leq 2^{|\varphi|}$ possibilities for $\bar{\gamma}$ in (2), and each disjunct in (2) has length at most $d|\varphi| \max_{i < s, j < t} (|\xi'_i|, |\zeta'_j|)$ for some constant d . By induction hypothesis, the latter is bounded by $d|\varphi| 2^{c|\varphi|\text{qdp}(\varphi)}$, which implies the claim, provided c is chosen large enough.

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

Model checking of unary-NWTL can be achieved with the same complexity as for NWTL using a variant of the tableaux construction in Section 5. ■