

DECIDING EQUATIONS IN THE TIME WARP ALGEBRA

SAM VAN GOOL ^a, ADRIEN GUATTO ^b, GEORGE METCALFE ^c, AND SIMON SANTSCHI ^d

^a IRIF, Université Paris Cité, France
e-mail address: vangool@irif.fr

^b IRIF, Université Paris Cité, France
e-mail address: guatto@irif.fr

^c Mathematical Institute, University of Bern, Switzerland
e-mail address: george.metcalfe@unibe.ch

^d Mathematical Institute, University of Bern, Switzerland
e-mail address: simon.santschi@unibe.ch

ABSTRACT. Join-preserving maps on the discrete time scale ω^+ , referred to as time warps, have been proposed as graded modalities that can be used to quantify the growth of information in the course of program execution. The set of time warps forms a simple distributive involutive residuated lattice—called the time warp algebra—that is equipped with residual operations relevant to potential applications. In this paper, we show that although the time warp algebra generates a variety that lacks the finite model property, it nevertheless has a decidable equational theory. We also describe an implementation of a procedure for deciding equations in this algebra, written in the OCaml programming language, that makes use of the Z3 theorem prover.

1. INTRODUCTION

Graded modalities [FKM16, GKO⁺16] provide a unified setting for describing effects of a program, such as which parts of memory it modifies [LG88], or the resources it consumes, such as how long it takes to run [GS14]. Given a *type* A and *grading* f , the new type $\Box_f A$ represents a modification of A that incorporates the behavior prescribed by f . The gradings themselves are often equipped with an ordered algebraic structure that is relevant for programming applications. Typically, they form a monoid, relating $\Box_{gf} A$ to $\Box_f \Box_g A$, or admit a precision order along which the graded modality acts contravariantly, i.e., for $f \leq g$, there exists a generic program of type $\Box_g A \rightarrow \Box_f A$ that allows movement from more to

Key words and phrases: Residuated lattices, Universal algebra, Decision procedures, Graded modalities, Type systems, Programming languages.

A precursor to this paper, reporting preliminary results, appeared in the proceedings of RAMiCS 2021 [vGGMS21].

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101007627, and Swiss National Science Foundation grant 200021_215157.

less precise types. In this case, the order may yield further operations on types; e.g., if the infimum $f \wedge g$ of f and g exists, it allows values of types $\Box_f A$ and $\Box_g B$ to be converted into a value of type $\Box_{f \wedge g}(A \times B)$.

In this paper, we study an algebraic structure induced by the class of graded modalities known as *time warps*: join-preserving maps on the discrete time scale $\omega^+ = \omega \cup \{\omega\}$ [Gua18]. Informally, such maps describe the growth of data during program execution, whereby a type A describes a sequence of sets A_n of values classified by A at execution step $n \in \omega$, and the type $\Box_f A$ classifies the set of values of $A_{f(n)}$ at this step. This approach generalizes a long line of work on programming languages for embedded systems (see, e.g., [CP96]) and type theories with modal recursion operators (see, e.g., [Nak00, BMSS12]).

Let us denote the set of time warps by W . Then $\langle W, \circ, id \rangle$ is a monoid, where $f \circ g$ (often shortened to fg) is the composition of $f, g \in W$, and id is the identity map. Equipping W with the pointwise order defined by $f \leq g : \iff f(n) \leq g(n)$, for all $n \in \omega^+$, yields a complete distributive lattice $\langle W, \wedge, \vee \rangle$ satisfying for all $f, g_1, g_2, h \in W$,

$$f(g_1 \vee g_2)h = fg_1h \vee fg_2h \quad \text{and} \quad f(g_1 \wedge g_2)h = fg_1h \wedge fg_2h.$$

Equivalently, the operation \circ is a *double quasi-operator* on $\langle W, \wedge, \vee \rangle$ [GP07a, GP07b] or the algebraic structure $\langle W, \wedge, \vee, \circ, id \rangle$ is a *distributive ℓ -monoid* [GJ17, CGMS22]. Moreover, since time warps are join-preserving, there exist binary operations $\backslash, /$ on W , called *residuals*, satisfying for all $f, g, h \in W$,

$$f \leq h/g \iff fg \leq h \iff g \leq f \backslash h.$$

That is, $\langle W, \wedge, \vee, \circ, \backslash, /, id \rangle$ is a *residuated lattice* [BT03, MPT23]. From a programming perspective, residuals play a role similar to that of weakest preconditions in deductive verification. The time warp h/g may be viewed as the most general (largest) time warp g' such that $\Box_h A$ can be sent generically to $\Box_g \Box_{g'} A$. In other words, $\Box_{h/g} A$ is the most precise type B such that $\Box_g B$ is a supertype of $\Box_h A$. Similarly, $f \backslash h$ is the most general time warp f' such that $\Box_h A$ can be sent generically to $\Box_{f'} \Box_f A$. Such questions of genericity arise naturally when programming in a modular way [Gua18], justifying the consideration of residuals in gradings.

In order to benefit from the extra flexibility and descriptive power of adopting graded modalities in a programming language, the language implementation should be able to decide the order between gradings to distinguish between well-typed and ill-typed programs. That is, an algorithm is required that decides equations between terms interpreted over the algebraic structure of time warps. In a precursor to this paper [vGGMS21], decidability of the equational theory of the bounded residuated lattice $\langle W, \wedge, \vee, \circ, \backslash, /, id, \perp, \top \rangle$, with least and greatest time warps \perp and \top , respectively, was established. A key ingredient of the proof, however, involves translating terms into a rather awkward normal form which both complicates the decision procedure and obscures the algebraic meaning of the equations.

In this paper, following ideas of [EGGHK18, San20] on quantales of join-preserving maps on a complete lattice, we obtain an elegant normal form for time warp terms by introducing a constant for the ‘predecessor’ time warp p that maps each $m \in \omega \setminus \{0\}$ to $m - 1$ and is constant on $\{0, \omega\}$. More precisely, the *pointed* residuated lattice $\langle W, \wedge, \vee, \circ, \backslash, /, id, p \rangle$ is term-equivalent to the involutive residuated lattice $\langle W, \wedge, \vee, \circ, *, id \rangle$, where $f^* := f \backslash p$ is an *involution* on the lattice $\langle L, \wedge, \vee \rangle$. For convenience, we refer to both of these structures as the *time warp algebra* \mathbf{W} . Because the monoid operation distributes over both joins and meets, and meets distribute over joins, every term is equivalent in \mathbf{W} to a meet of

joins of *basic terms*, constructed using just the monoid operations and involution. It follows that the equational theory of \mathbf{W} is decidable if there exists an algorithm that decides $\mathbf{W} \models e \leq t_1 \vee \cdots \vee t_n$ for arbitrary basic terms t_1, \dots, t_n , where e is interpreted as the identity map. We provide such an algorithm by relating the existence of a counterexample to $e \leq t_1 \vee \cdots \vee t_n$ to the satisfiability of a corresponding first-order formula in ω^+ , understood as an ordered structure with a decidable first-order theory, and describe an implementation, written in the OCaml programming language, that makes use of the Z3 theorem prover.

Overview of the paper. In Section 2, we introduce the time warp algebra \mathbf{W} as an example of an involutive residuated lattice consisting of join-preserving maps on a complete chain (totally ordered set), and establish some of its elementary properties. In particular, we show that \mathbf{W} is simple, has no finite subalgebras, and generates a variety (equational class) that lacks the finite model property. We also describe a subalgebra \mathbf{R} of \mathbf{W} consisting of time warps that are ‘regular’ in the sense that they are eventually either constant or linear, and provide an explicit description of the involution operation.

In Section 3, we describe the existence of a potential counterexample to $e \leq t_1 \vee \cdots \vee t_n$, where t_1, \dots, t_n are basic terms, using the notion of a *diagram*, which provides a finite partial description of an evaluation of terms as time warps. As a consequence, an equation is satisfied by \mathbf{W} if, and only if, it is satisfied by \mathbf{R} (Theorem 3.11). Finally, we use the resulting description to reduce the existence of a counterexample to the satisfiability of a first-order formula in ω^+ , understood as an ordered structure with a decidable first-order theory, thereby establishing the decidability of the equational theory of \mathbf{W} (Theorem 3.12).

In Section 4, we describe an implementation of our decision procedure for the equational theory of \mathbf{W} , written in the OCaml programming language [LDF⁺22], that makes use of the Z3 theorem prover [dMB08]. Finally, in Section 5, we consider some potential avenues for further research; in particular, we explain how to adapt the decision procedure to deal with extra constants for first-order definable time warps such as \perp and \top , and explore a relational approach to the study of the time warp algebra and related structures.

2. THE TIME WARP ALGEBRA

In this section, we introduce and establish some elementary properties of the time warp algebra \mathbf{W} in the general framework of involutive residuated lattices. In particular, we show that the variety generated by \mathbf{W} lacks the finite model property and that ‘regular’ time warps—those that are eventually either constant or linear—form a subalgebra of \mathbf{W} . We also provide an explicit description of the involution operation on time warps.

A *pointed residuated lattice* (also known as an *FL-algebra*) is an algebraic structure $\mathbf{L} = \langle L, \wedge, \vee, \cdot, \backslash, /, e, f \rangle$ of signature $\langle 2, 2, 2, 2, 2, 0, 0 \rangle$ such that $\langle L, \cdot, e \rangle$ is a monoid, $\langle L, \wedge, \vee \rangle$ is a lattice with an order defined by $a \leq b : \iff a \wedge b = a$, and, for all $a, b, c \in L$,

$$a \leq c/b \iff ab \leq c \iff b \leq a \backslash c,$$

where the residual operators are given explicitly for $a, b \in L$ by

$$a \backslash b = \bigvee \{c \in L \mid ac \leq b\} \quad \text{and} \quad b/a = \bigvee \{c \in L \mid ca \leq b\}.$$

Moreover, it follows from the existence of residual operations \backslash and $/$ that the operation \cdot preserves existing joins in both coordinates; in particular, for all $a, b_1, b_2, c \in L$,

$$a(b_1 \vee b_2)c = ab_1c \vee ab_2c.$$

The structure \mathbf{L} is said to be *distributive* if its lattice reduct $\langle L, \wedge, \vee \rangle$ is distributive and *fully distributive* if it is distributive and, for all $a, b_1, b_2, c \in L$,

$$a(b_1 \wedge b_2)c = ab_1c \wedge ab_2c.$$

The element $f \in L$ is called *cyclic* if $f/a = a \setminus f$, for all $a \in L$, and *dualizing* if, for all $a \in L$,

$$f/(a \setminus f) = a = (f/a) \setminus f.$$

If f is both cyclic and dualizing, then the map $' : L \rightarrow L; x \mapsto x \setminus f$ is an *involution* on the lattice $\langle L, \wedge, \vee \rangle$, satisfying for all $a, b \in L$,

$$a \leq b \iff b' \leq a', \quad a'' = a, \quad (a \wedge b)' = a' \vee b', \quad \text{and} \quad (a \vee b)' = a' \wedge b'.$$

The class of pointed residuated lattices such that f is cyclic and dualizing forms a variety (equational class) that is term-equivalent to the variety of *involutive residuated lattices*: algebraic structures $\langle L, \wedge, \vee, \cdot, ', e \rangle$ of signature $\langle 2, 2, 2, 1, 0 \rangle$ such that $\langle L, \cdot, e \rangle$ is a monoid, $\langle L, \wedge, \vee \rangle$ is a lattice, $'$ is an involution on $\langle L, \wedge, \vee \rangle$, and for all $a, b, c \in L$,

$$b \leq (c'a)' \iff ab \leq c \iff a \leq (bc)'$$

The term-equivalence is implemented by defining $x' := x \setminus f$ and, conversely, $x \setminus y := (y'x)'$, $y/x := (xy)'$, and $f := e'$ (see [MPT23] for further details). For any element a of an involutive residuated lattice, we also define inductively $a^0 := e$ and $a^{k+1} := a^k \cdot a$ ($k \in \mathbb{N}$).

Now let \mathbf{Tm} denote the *term algebra* of the language of involutive residuated lattices defined over a countably infinite set of variables Var , and call a term $t \in \mathbf{Tm}$ *basic* if it is constructed using the operation symbols $\cdot, ',$ and e . As usual, an *equation* is an ordered pair of terms $s, t \in \mathbf{Tm}$, denoted by $s \approx t$, and $s \leq t$ abbreviates $s \wedge t \approx s$. An involutive residuated lattice \mathbf{L} *satisfies* an equation $s \approx t$, denoted by $\mathbf{L} \models s \approx t$, if $\llbracket s \rrbracket_h = \llbracket t \rrbracket_h$ for every homomorphism $h : \mathbf{Tm} \rightarrow \mathbf{L}$, where $\llbracket u \rrbracket_h := h(u)$ for $u \in \mathbf{Tm}$.

In a fully distributive involutive residuated lattice, every term is equivalent to a meet of joins of basic terms. More precisely:

Lemma 2.1. *There exists an algorithm that produces for any term $t \in \mathbf{Tm}$ positive integers m, n_1, \dots, n_m and basic terms $t_{i,j}$ for each $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n_i\}$ such that for any fully distributive involutive residuated lattice \mathbf{L} ,*

$$\mathbf{L} \models t \approx \bigwedge_{i=1}^m \bigvee_{j=1}^{n_i} t_{i,j}.$$

Proof. The desired basic terms are obtained by iteratively distributing joins over meets and the monoid multiplication over both meets and joins, and pushing the involution inwards using the De Morgan laws. \square

Next, let \mathbf{C} be any complete chain with a least element 0 and a greatest element ∞ . The set $\text{Res}(\mathbf{C})$ of maps on \mathbf{C} that preserve arbitrary joins forms a fully distributive pointed residuated lattice $\mathbf{Res}(\mathbf{C}) = \langle \text{Res}(\mathbf{C}), \wedge, \vee, \circ, \setminus, /, id, p \rangle$, where \wedge, \vee are defined pointwise, \circ is functional composition, id is the identity map, and p is the join-preserving map

$$p : C \rightarrow C; \quad x \mapsto \bigvee \{y \in C \mid y < x\}.$$

The least and greatest elements of $\mathbf{Res}(\mathbf{C})$ are the maps \perp, \top satisfying, respectively, $\perp(x) = 0$, for all $x \in C$, and $\top(0) = 0$ and $\top(x) = \infty$ for all $x \in C \setminus \{0\}$. Moreover, it follows from [EGGHK18, Proposition 2.6.18] and [San20, Section 4]) that p is the unique cyclic and

dualizing element of $\mathbf{Res}(\mathbf{C})$ and hence that $\mathbf{Res}(\mathbf{C})$ is term-equivalent to the involutive residuated lattice $\langle \mathbf{Res}(\mathbf{C}), \wedge, \vee, \circ, *, id \rangle$, where $f^* := f \setminus p = p / f$ for each $f \in \mathbf{Res}(\mathbf{C})$.

Recall that an algebraic structure \mathbf{A} is *simple* if $\text{Con}(\mathbf{A}) = \{\Delta_A, \nabla_A\}$, where $\text{Con}(\mathbf{A})$ denotes the set of congruences of \mathbf{A} , $\Delta_A = \{\langle a, a \rangle \mid a \in A\}$, and $\nabla_A = A \times A$.

Proposition 2.2. *For any complete chain \mathbf{C} , the pointed residuated lattice $\mathbf{Res}(\mathbf{C})$ is simple.*

Proof. Consider any $\Theta \in \text{Con}(\mathbf{Res}(\mathbf{C})) \setminus \{\Delta_{\mathbf{Res}(\mathbf{C})}\}$. Since $\Theta \neq \Delta_{\mathbf{Res}(\mathbf{C})}$ and a congruence of a pointed residuated lattice is fully determined by the congruence class of its multiplicative unit (see [BT03]), there exists an $f \in \mathbf{Res}(\mathbf{C}) \setminus \{id\}$ such that $f \Theta id$. Choose any $c \in C$ such that $f(c) \neq c$. We consider the cases $f(c) < c$ and $c < f(c)$. Suppose first that $f(c) < c$ and define the map $g \in \mathbf{Res}(\mathbf{C})$ such that for each $d \in C$,

$$g(d) := \begin{cases} 0 & \text{if } d \leq f(c); \\ \infty & \text{if } d > f(c). \end{cases}$$

Since $f \Theta id$ and Θ is a congruence, also $((g \circ f)/g) \wedge id \Theta ((g \circ id)/g) \wedge id$. Observe that

$$((g \circ f)/g)(\infty) = ((g \circ f)/g)(g(c)) = (((g \circ f)/g) \circ g)(c) \leq (g \circ f)(c) = g(f(c)) = 0,$$

so $\perp = (g \circ f)/g = ((g \circ f)/g) \wedge id$. Also, $((g \circ id)/g) \wedge id = (g/g) \wedge id = id$, so $\perp \Theta id$ and $\top = \perp \setminus id \Theta id \setminus id = id$. Hence $\top \Theta \perp$. But Θ is a lattice congruence, so its congruence classes are convex and $\Theta = \nabla_{\mathbf{Res}(\mathbf{C})}$. Suppose next that $c < f(c)$. Then $f((f \setminus id)(c)) = (f \circ (f \setminus id))(c) \leq c < f(c)$, and, since f is order-preserving and \mathbf{C} is a chain, $(f \setminus id)(c) < c$. But then, since $f \setminus id \Theta id \setminus id = id$, as in the previous case, $\Theta = \nabla_{\mathbf{Res}(\mathbf{C})}$. \square

Let us focus our attention now on the special case of the successor ordinal $\omega^+ = \omega \cup \{\omega\}$. For convenience, we call both the pointed residuated lattice $\mathbf{Res}(\omega^+)$ and the corresponding term-equivalent involutive residuated lattice, the *time warp algebra* \mathbf{W} , referring to members of W , i.e., join-preserving maps on ω^+ , as *time warps*. Clearly, a map $f: \omega^+ \rightarrow \omega^+$ is a time warp if, and only if, it is order-preserving and satisfies $f(0) = 0$ and $f(\omega) = \bigvee \{f(n) \mid n \in \omega\}$. As motivation for subsequent sections, we will show below that equational reasoning in \mathbf{W} cannot be checked by considering finite members of the variety generated by this algebra.

Observe first that $p := id^* \in W$ is the ‘predecessor’ time warp satisfying for $m \in \omega^+$,

$$p(m) = \bigvee \{n \in \omega \mid n < m\} = \begin{cases} m & \text{if } m \in \{0, \omega\}; \\ m - 1 & \text{otherwise.} \end{cases}$$

Clearly, $p < id$, so \mathbf{W} satisfies the equation $e' \leq e$.

The involution operation can be described using p as follows.

Lemma 2.3. *For any time warp f and $m \in \omega^+$,*

$$f^*(m) = \bigwedge \{p(n) \mid n \in \omega^+ \text{ and } m \leq f(n)\}.$$

Proof. Let h be the function defined by $h(m) := \bigwedge \{p(n) \mid n \in \omega^+ \text{ and } m \leq f(n)\}$ for each $m \in \omega^+$. Clearly, h is order-preserving and satisfies $h(0) = 0$. Hence, to show that h is a time warp, it remains to check that $h(\omega) = \bigvee \{h(m) \mid m \in \omega\}$. To this end, suppose first that $h(\omega) = \omega$. Then $f(n) < \omega$, for all $n \in \omega$. If also $f(\omega) < \omega$, then $h(m) = \omega$, for each $m > f(\omega)$, yielding $h(\omega) = \omega = \bigvee \{h(m) \mid m \in \omega\}$. If $f(\omega) = \omega$, then, since $\omega = f(\omega) = \bigvee \{f(n) \mid n \in \omega\}$, there exists, for each $m \in \omega$, an $n \in \omega$ with $m \leq f(n) < \omega$, yielding again $h(\omega) = \omega = \bigvee \{h(m) \mid m \in \omega\}$. Suppose now that $h(\omega) < \omega$. Then there exists

a minimal $n \in \omega \setminus \{0\}$ such that $f(n) = \omega$, and hence $h(m) = p(n)$, for each $m > f(n-1)$, yielding $h(\omega) = p(n) = \bigvee \{h(n) \mid n \in \omega\}$.

Moreover, $hf \leq p$, so $h \leq p/f = f^*$. Now consider any $m \in \omega \setminus \{0\}$. If $n \in \omega^+$ satisfies $m \leq f(n)$, then $f^*(m) \leq f^*(f(n)) \leq p(n)$. So $f^*(m) \leq h(m)$. Hence $h = f^*$, since time warps are determined by their values on $\omega \setminus \{0\}$. \square

Observe next that $p^{k+1} < p^k$ for each $k \in \mathbb{N}$ and hence that the \emptyset -generated subalgebra of \mathbf{W} is infinite. So \mathbf{W} has no finite subalgebras. Moreover, \mathbf{W} is simple, by Proposition 2.2, so any quotient of \mathbf{W} is either trivial or isomorphic to \mathbf{W} . Indeed, we do not know if the variety generated by \mathbf{W} contains *any* non-trivial finite algebra, but can show at least that it cannot contain $\mathbf{Res}(\mathbf{C})$ for any finite chain \mathbf{C} and does not have the finite model property.¹

To this end, observe first that the element p has a right inverse; that is, $p \circ s = id$, where s is the ‘successor’ time warp satisfying for each $m \in \omega^+$,

$$s: \omega^+ \rightarrow \omega^+; \quad m \mapsto \begin{cases} m & \text{if } m \in \{0, \omega\}; \\ m+1 & \text{otherwise.} \end{cases}$$

Note that $s = p \setminus id$, since $p \setminus id = (id^* \circ p)^* = (p \circ p)^*$ and, by Lemma 2.3, for each $m \in \omega^+$,

$$\begin{aligned} (p \circ p)^*(m) &= \bigwedge \{p(n) \mid n \in \omega^+ \text{ and } m \leq (p \circ p)(n)\} \\ &= \bigwedge \{n \mid n \in \omega^+ \text{ and } m \leq p(n)\} \\ &= s(m). \end{aligned}$$

Hence $p \circ (p \setminus id) = id$, and \mathbf{W} satisfies the equation $e \approx e' \cdot (e' \setminus e)$.

Proposition 2.4. *The variety generated by \mathbf{W} does not have the finite model property and does not contain $\mathbf{Res}(\mathbf{C})$ for any finite chain \mathbf{C} .*

Proof. Let $\langle L, \wedge, \vee, \cdot, ', e \rangle$ be any finite member of the variety generated by \mathbf{W} . Then, since $e' \leq e$ and \mathbf{L} is finite, there exists a $k \in \mathbb{N}$ such that $(e')^k = (e')^{k+1}$. Moreover, $e' = e'e = (e')^2 \cdot (e' \setminus e)$, since $e = e' \cdot (e' \setminus e)$ and e is the multiplicative unit of \mathbf{L} , and hence, iterating this step,

$$e' = (e')^{k+1}(e' \setminus e)^k = (e')^k(e' \setminus e)^k = e.$$

The equation $e' \approx e$ is therefore satisfied by all the finite members of the variety generated by \mathbf{W} , but not by \mathbf{W} itself, since $p < id$. Moreover, since $\mathbf{Res}(\mathbf{C})$ is finite for any finite chain \mathbf{C} and does not satisfy $e' \approx e$, it cannot belong to the variety generated by \mathbf{W} . \square

Although the variety generated by \mathbf{W} does not have the finite model property, we will show in this paper that it is generated by a subalgebra of \mathbf{W} consisting of time warps that have a simple finite description. Let us call a time warp $f \in W$ *eventually constant* if there exists an $m \in \omega$ such that $f(n) = f(m)$ for all $n \in \omega$ with $n \geq m$, *eventually linear* if there exist an $m \in \omega$ and a $k \in \mathbb{Z}$ such that $f(n) = n + k$ for all $n \in \omega$ with $n \geq m$, and *regular* if it is eventually constant or eventually linear. Equivalently, a time warp $f \in W$ is regular if, and only if, there exist $m \in \omega$, $l \in \{0, 1\}$, and $k \in \mathbb{Z} \cup \{\omega\}$ such that $f(n) = ln + k$ for all $n \in \omega$ with $n \geq m$.

Proposition 2.5. *The set of regular time warps forms a subalgebra \mathbf{R} of \mathbf{W} .*

¹Note that the involution-free reduct of \mathbf{W} generates the variety of *distributive ℓ -monoids*, which does have the finite model property; indeed, an equation is satisfied by this reduct of \mathbf{W} if, and only if, it is satisfied by the distributive ℓ -monoid of join-preserving maps on any finite chain [CGMS22].

Proof. Clearly, id is eventually linear, and hence regular. Suppose that $f, g \in W$ are regular. It is easy to see that $f \wedge g$ and $f \vee g$ are then also regular. If g is eventually constant, then there exists an $m \in \omega$ such that $g(n) = g(m)$ for all $n \in \omega$ with $n \geq m$, and hence also $f(g(n)) = f(g(m))$ for all $n \in \omega$ with $n \geq m$, so $f \circ g$ is eventually constant. Suppose then that g is eventually linear, that is, there exist an $m \in \omega$ and a $k \in \mathbb{Z}$ such that $g(n) = n + k$ for all $n \in \omega$ with $n \geq m$. If f is eventually constant, then there exists an $l \in \omega$ such that $f(n) = f(l)$ for all $n \in \omega$ with $n \geq l$, and hence $f(g(n)) = f(l)$ for all $n \in \omega$ with $n \geq \max(m, l - k)$, that is, $f \circ g$ is eventually constant. If f is eventually linear, then there exist an $l \in \omega$ and a $r \in \mathbb{Z}$ such that $f(n) = n + r$ for all $n \in \omega$ with $n \geq l$, and hence $f(g(n)) = n + k + r$ for all $n \in \omega$ with $n \geq \max(m, l + k)$.

It remains to show that f^* is regular for f regular. Suppose first that f is eventually constant, that is, there exists an $m \in \omega$ such that $f(n) = f(m)$ for all $n \in \omega$ with $n \geq m$. If $f(m) < \omega$, then for each $k \in \omega$ with $k \geq f(m) + 1$, the set $\{p(n) \mid n \in \omega^+ \text{ and } k \leq f(n)\}$ is empty and $f^*(k) = \omega$, by Lemma 2.3, i.e., f^* is eventually constant. Otherwise, $f(m) = \omega$ and for all $l \geq 1 + \max\{f(n) \mid n \in \omega, f(n) < \omega\}$, Lemma 2.3 yields

$$f^*(l) = \bigwedge \{p(n) \mid n \in \omega^+ \text{ and } l \leq f(n)\} = \bigwedge \{p(n) \mid n \in \omega^+ \text{ and } f(n) = \omega\},$$

i.e., f^* is eventually constant. Finally, suppose that f is eventually linear, that is, there exists an $m \in \omega$ and a $k \in \mathbb{Z}$ such that $f(n) = n + k$ for all $n \in \omega$ with $n \geq m$. Then every $l \in \omega$ with $l \geq m + k + 1$ lies in the image of f and $n = l - k$ is the unique solution for the equation $l = f(n)$, so $f^*(l) = l - (k + 1)$, by Lemma 2.3, i.e., f^* is eventually linear. \square

We conclude this section by providing an explicit description of the involution operation on time warps that will play a crucial role in the decidability proof in the next section. For any time warp $f \in W$, let $\text{last}(f)$ denote the smallest $m \in \omega^+$ such that f takes the same value on m (and hence all elements greater than m) as ω . More formally:

$$\text{last}(f) := \min\{m \in \omega^+ \mid f(m) = f(\omega)\}.$$

Observe that $\text{last}(f) < \omega$ if, and only if, f is eventually constant. Moreover, we have $\text{last}(f) = (id/f)f(\omega)$, since $id/f = (f \circ id^*)^* = (f \circ p)^*$ and, by Lemma 2.3,

$$(f \circ p)^*(f(\omega)) = \bigwedge \{p(n) \mid n \in \omega^+ \text{ and } f(\omega) \leq f(p(n))\} = \min\{m \in \omega^+ \mid f(m) = f(\omega)\}.$$

The behaviour of this operation with respect to compositions and involutions of time warps is easily described as follows.

Lemma 2.6. *For any time warps f and g ,*

$$\begin{aligned} \text{last}(fg) = \omega &\iff \text{last}(g) = \text{last}(f) = \omega \\ \text{last}(f^*) = \omega &\iff \text{last}(f) = \omega. \end{aligned}$$

Proof. For the first equivalence, observe that $\text{last}(fg) = \omega$ if, and only if, $fg(m) < fg(\omega)$ for all $m \in \omega$. However, $fg(m) = fg(\omega)$ for some $m \in \omega$ if, and only if, either $g(m) = g(\omega)$ for some $m \in \omega$, or $g(m) < g(\omega)$ for all $m \in \omega$ and $f(k) = f(\omega)$ for some $k \in \omega$, which is equivalent to $\text{last}(g) < \omega$ or $\text{last}(f) < \omega$. For the second equivalence, observe that $\text{last}(f^*) = \omega$ if, and only if, $f^*(m) < f^*(\omega)$ for all $m \in \omega$. But $f^*(m) = f^*(\omega)$ for some $m \in \omega$ if, and only if, $f(m) = f(\omega)$ for some $m \in \omega$, by Lemma 2.3, which is equivalent to $\text{last}(f) < \omega$. \square

Finally, we are able to provide the promised explicit description of the involution operation on time warps.

Proposition 2.7. *For any time warp f , $n \in \omega \setminus \{0\}$, and $m \in \omega$,*

$$\begin{aligned} f^*(n) = m &\iff f(m) < n \leq f(m+1) \\ f^*(\omega) = m &\iff f(m) < \omega = f(m+1) \\ f^*(n) = \omega &\iff f(\omega) < n \\ f^*(\omega) = \omega &\iff f(\omega) < \omega \text{ or } \text{last}(f) = \omega. \end{aligned}$$

Proof. Consider any $n \in \omega^+ \setminus \{0\}$ and $m \in \omega$. For the first two equivalences, observe that, using Lemma 2.3 and the assumption that $n \neq 0$,

$$f^*(n) = m \iff m = \bigwedge \{p(k) \mid k \in \omega^+ \text{ and } n \leq f(k)\} \iff f(m) < n \leq f(m+1).$$

For the second two equivalences, observe first that, using Lemma 2.3,

$$f^*(n) = \omega \iff \omega = \bigwedge \{p(k) \mid k \in \omega^+ \text{ and } n \leq f(k)\} \iff f(k) < n \text{ for all } k \in \omega.$$

It follows that $f^*(n) = \omega \iff f(k) < n \text{ for all } k \in \omega \iff f(\omega) < n$, for any $n \in \omega \setminus \{0\}$, and $f^*(\omega) = \omega \iff f(k) < \omega \text{ for all } k \in \omega \iff f(\omega) < \omega \text{ or } \text{last}(f) = \omega$. \square

3. DECIDABILITY VIA DIAGRAMS

In this section, we turn our attention to the problem of deciding equations in the time warp algebra \mathbf{W} . Observe first that for any equation $s \approx t$ with $s, t \in \text{Tm}$,

$$\mathbf{W} \models s \approx t \iff (\mathbf{W} \models s \leq t \text{ and } \mathbf{W} \models t \leq s) \iff \mathbf{W} \models e \leq (s \setminus t) \wedge (t \setminus s).$$

Moreover, by Lemma 2.1, each $u \in \text{Tm}$ is equivalent in \mathbf{W} to a meet of joins of basic terms. Hence, since $\mathbf{W} \models e \leq u_1 \wedge \dots \wedge u_m$ if, and only if, $\mathbf{W} \models e \leq u_i$, for each $i \in \{1, \dots, m\}$, we may restrict our attention to deciding equations of the form $e \leq t_1 \vee \dots \vee t_n$, where $\{t_1, \dots, t_n\}$ is any finite non-empty set of basic terms. More precisely:

Proposition 3.1. *The equational theory of \mathbf{W} is decidable if, and only if, there exists an algorithm that decides $\mathbf{W} \models e \leq t_1 \vee \dots \vee t_n$ for any basic terms t_1, \dots, t_n .*

To address this problem, we relate the existence of a counterexample to $\mathbf{W} \models e \leq t_1 \vee \dots \vee t_n$, where t_1, \dots, t_n are basic terms, to the satisfiability of a corresponding first-order formula in ω^+ , understood as an ordered structure with a decidable first-order theory. Such a counterexample is given by assigning time warps to the variables in t_1, \dots, t_n to obtain time warps $\hat{t}_1, \dots, \hat{t}_n$ satisfying $id \not\leq \hat{t}_1 \vee \dots \vee \hat{t}_n$, that is, by a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ and element $k \in \omega^+$ satisfying $k > \llbracket t_i \rrbracket_h(k)$, for each $i \in \{1, \dots, n\}$. The translation into a first-order formula is obtained in two steps. First, a ‘time variable’ κ is introduced to stand for the unknown $k \in \omega^+$ and finitely many ‘samples’ are generated that correspond to other members of ω^+ used in the computation of $\llbracket t_1 \rrbracket_h(k), \dots, \llbracket t_n \rrbracket_h(k)$. Second, finitely many quantifier-free formulas are defined that describe the relationships between samples according to the semantics of \mathbf{W} . The required formula is then the existential closure of the conjunction of these quantifier-free formulas and a formula that expresses $\kappa > \llbracket t_i \rrbracket_h(\kappa)$ for each $i \in \{1, \dots, n\}$.

Let us fix a countably infinite set \mathcal{T}_V of *time variables*, denoted by the (possibly indexed) symbol κ . We call any member of the language generated by the following grammar a *sample*, where $t \in \text{Tm}$ is any basic term and $\kappa \in \mathcal{T}_V$ is any time variable:

$$\mathcal{S} \ni \alpha ::= \kappa \mid t[\alpha] \mid \mathbf{s}(\alpha) \mid \text{last}(t).$$

Now let \rightsquigarrow be the binary relation defined on the set of all samples by

$$\begin{array}{ll} t[\alpha] \rightsquigarrow \alpha & tu[\alpha] \rightsquigarrow t[u[\alpha]] \\ \mathbf{s}(\alpha) \rightsquigarrow \alpha & t'[\alpha] \rightsquigarrow t[t'[\alpha]] \\ t[\alpha] \rightsquigarrow t[\mathbf{last}(t)] & t'[\alpha] \rightsquigarrow t[\mathbf{s}(t'[\alpha])], \end{array}$$

and let \rightsquigarrow^* denote the reflexive transitive closure of this relation. We call a sample set (i.e., set of samples) Δ *saturated* if whenever $\alpha \in \Delta$ and $\alpha \rightsquigarrow \beta$, also $\beta \in \Delta$, and define the *saturation* of a sample set Δ (analogously to the Fischer-Ladner closure of formulas in Propositional Dynamic Logic [FL79]) as

$$\Delta^{\rightsquigarrow} := \{\beta \in \mathcal{S} \mid \alpha \rightsquigarrow^* \beta \text{ for some } \alpha \in \Delta\}.$$

Crucially, this definition yields the following property:

Lemma 3.2. *The saturation of a finite sample set is finite.*

Proof. It suffices to consider a sample set with one element, so suppose that $\Delta = \{\alpha_0\}$. Let $K = \{\kappa_\alpha \mid \alpha \in \mathcal{S}\} \subseteq \mathcal{T}_V$ be the set of time variables not occurring in α_0 indexed by the set of all samples \mathcal{S} , and consider a further binary relation \dashv defined on \mathcal{S} by

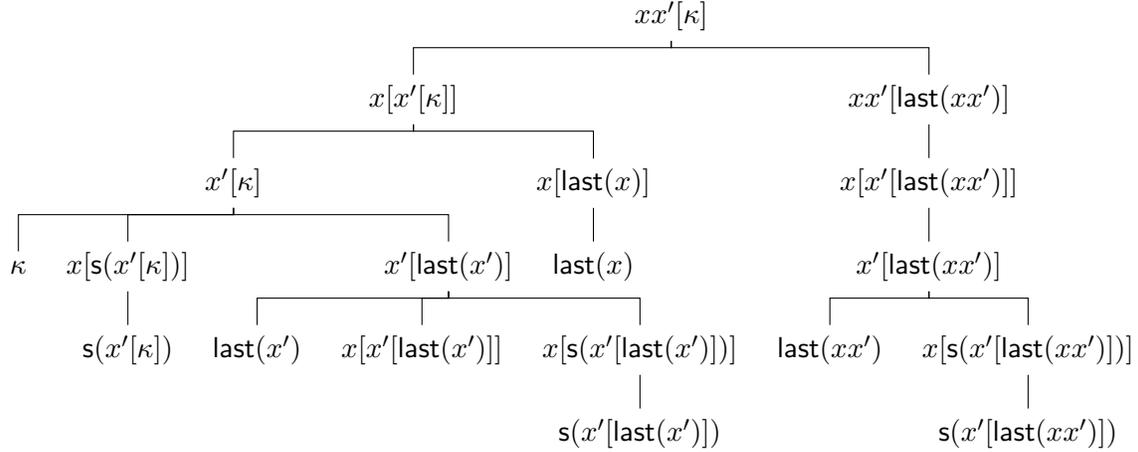
$$\begin{array}{ll} t[\alpha] \dashv \alpha & tu[\alpha] \dashv t[u[\alpha]] \\ \mathbf{s}(\alpha) \dashv \alpha & t'[\alpha] \dashv t[\kappa_{t'[\alpha]}] \\ t[\alpha] \dashv t[\mathbf{last}(t)] & t'[\alpha] \dashv t[\kappa_{\mathbf{s}(t'[\alpha])}]. \end{array}$$

Now let $\Delta^{\dashv} := \{\beta \in \mathcal{S} \mid \alpha \dashv^* \beta \text{ for some } \alpha \in \Delta\}$, where \dashv^* is the reflexive transitive closure of \dashv . For a sample α , denote by $v(\alpha)$ the set of variables from Var that occur in α , and by $\ell(\alpha)$ the number of occurrences in α of variables from Var and symbols in $\{., ', \mathbf{e}, \mathbf{s}\}$ that do not occur in the scope of an occurrence of \mathbf{last} . Note that, by definition, \dashv does not increase v and ℓ , i.e., if $\alpha \dashv \beta$, then $v(\beta) \subseteq v(\alpha)$, $\ell(\beta) \leq \ell(\alpha)$, and if $\alpha \dashv \beta$ is not of the form $t[\gamma] \dashv t[\mathbf{last}(t)]$, then $\ell(\beta) < \ell(\alpha)$. In particular, $v(\beta) \subseteq v(\alpha_0)$ and $\ell(\beta) \leq \ell(\alpha_0)$, for each $\beta \in \Delta^{\dashv}$. But clearly there are only finitely many samples of the form $t[\mathbf{last}(t)]$ such that $v(t[\mathbf{last}(t)]) \subseteq v(\alpha_0)$ and $\ell(t[\mathbf{last}(t)]) \leq \ell(\alpha_0)$. Hence Δ^{\dashv} is finite.

Now define the map $\sigma: \Delta^{\dashv} \rightarrow \Delta^{\rightsquigarrow}$ recursively by $\sigma(\kappa) = \kappa$ for $\kappa \in \mathcal{T}_V \setminus K$, $\sigma(\mathbf{last}(t)) = \mathbf{last}(t)$, $\sigma(\kappa_\alpha) = \sigma(\alpha)$ for $\kappa_\alpha \in K$, $\sigma(\mathbf{s}(\alpha)) = \mathbf{s}(\sigma(\alpha))$, $\sigma(t[\alpha]) = t[\sigma(\alpha)]$. It is straightforward to check that σ is well-defined and surjective. So $\Delta^{\rightsquigarrow}$ is finite. \square

Running example part 1. Let us consider the equation $\mathbf{e} \leq xx'$ as a running example throughout this section. In this case, the sample set of interest is $\{xx'[\kappa]\}$. In Figure 1, its saturation $\{xx'[\kappa]\}^{\rightsquigarrow}$ is visualized as a tree, where each parent node is related to its successors by \rightsquigarrow and redundant samples are omitted.

Consider next a first-order signature $\tau = \{\preceq, \mathcal{S}, 0, \omega\}$ with a binary relation symbol \preceq , a unary function symbol \mathcal{S} , and two constants 0 and ω . We denote by \mathfrak{A} the τ -structure with universe ω^+ and natural order $\preceq^{\mathfrak{A}}$, defining $\mathcal{S}^{\mathfrak{A}}(n) := n + 1$, for each $n \in \omega$, $\mathcal{S}^{\mathfrak{A}}(\omega) := \omega$, $\omega^{\mathfrak{A}} := \omega$, and $0^{\mathfrak{A}} := 0$. Since no other τ -structure will be considered in this section, we will omit the superscripts from now on. We use the symbols \neg , \wedge , \vee , \Rightarrow , and \Leftrightarrow to denote the logical connectives ‘not’, ‘and’, ‘or’, ‘implies’, and ‘if, and only if’, respectively, and let $a \prec b$ stand for $a \preceq b \wedge \neg(b \preceq a)$.

Figure 1: Visualization of the saturation of $\{xx'[\kappa]\}$

Let us fix now a saturated sample set Δ and consider its members as first-order variables. We define the following sets of first-order quantifier-free τ -formulas over Δ :

$$\text{struct}_\Delta := \{\alpha \preceq \beta \Rightarrow t[\alpha] \preceq t[\beta] \mid t[\alpha], t[\beta] \in \Delta\} \cup \quad (3.1)$$

$$\{\alpha \approx 0 \Rightarrow t[\alpha] \approx 0 \mid t[\alpha] \in \Delta\} \cup \quad (3.2)$$

$$\{s(\alpha) \approx \mathcal{S}(\alpha) \mid s(\alpha) \in \Delta\} \cup \quad (3.3)$$

$$\{\text{last}(t) \preceq \alpha \Leftrightarrow t[\text{last}(t)] \approx t[\alpha] \mid t[\alpha] \in \Delta\} \cup \quad (3.4)$$

$$\{\text{last}(t) \approx \omega \Rightarrow t[\text{last}(t)] \approx \omega \mid t[\text{last}(t)] \in \Delta\} \quad (3.5)$$

$$\text{log}_\Delta := \{e[\alpha] \approx \alpha \mid e[\alpha] \in \Delta\} \cup \quad (3.6)$$

$$\{\text{last}(e) \approx \omega \mid \text{last}(e) \in \Delta\} \cup \quad (3.7)$$

$$\{tu[\alpha] \approx t[u[\alpha]] \mid tu[\alpha] \in \Delta\} \cup \quad (3.8)$$

$$\{\text{last}(tu) \approx \omega \Rightarrow (\text{last}(t) \approx \omega \wedge \text{last}(u) \approx \omega) \mid \text{last}(tu), \text{last}(t), \text{last}(u) \in \Delta\} \quad (3.9)$$

$$\text{inv}_\Delta := \{(0 < \alpha \wedge \alpha < \omega) \Rightarrow t[t'[\alpha]] < \alpha \mid t'[\alpha] \in \Delta\} \cup \quad (3.10)$$

$$\{t'[\alpha] < \omega \Rightarrow \alpha \preceq t[s(t'[\alpha])] \mid t'[\alpha] \in \Delta\} \cup \quad (3.11)$$

$$\{\text{last}(t') \approx \omega \Rightarrow \text{last}(t) \approx \omega \mid \text{last}(t'), \text{last}(t) \in \Delta\} \quad (3.12)$$

$$\Sigma_\Delta := \text{struct}_\Delta \cup \text{log}_\Delta \cup \text{inv}_\Delta. \quad (3.13)$$

We call a \mathfrak{A} -valuation $\delta: \Delta \rightarrow \omega^+$ a Δ -*diagram* if $\mathfrak{A}, \delta \models \Sigma_\Delta$.²

²The term ‘diagram’ recalls a similar concept used to prove the decidability of the equational theory of lattice-ordered groups in [HM79].

Running example part 2. For $\Delta = \{xx'[\kappa]\}^{\rightsquigarrow}$, consider the map $\delta: \Delta \rightarrow \omega^+$ defined by

$$\begin{aligned} 0 &= \delta(x'[\kappa]) = \delta(x[x'[\kappa]]) = \delta(xx'[\kappa]), \\ 1 &= \delta(x[s(x'[\kappa])]) = \delta(s(x'[\kappa])) = \delta(\kappa), \\ \omega &= \delta(\text{last}(x)) = \delta(\text{last}(x')) = \delta(\text{last}(xx')) = \delta(x[\text{last}(x)]) = \delta(x'[\text{last}(x')]) \\ &= \delta(x[x'[\text{last}(x')]]) = \delta(x[s(x'[\text{last}(x')])]) = \delta(s(x'[\text{last}(x')])) = \delta(xx'[\text{last}(xx')]) \\ &= \delta(x[x'[\text{last}(xx')]]) = \delta(x'[\text{last}(xx')]) = \delta(x[s(x'[\text{last}(xx')])]) = \delta(s(x'[\text{last}(xx')])). \end{aligned}$$

It is readily checked that δ is a Δ -diagram; e.g., δ satisfies (3.2), since $\delta(x'[\kappa]) = 0$ and $\delta(x[x'[\kappa]]) = 0$.

Fixing a set of basic terms T and time variable κ , the following proposition relates any homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ and $n \in \omega^+$ to the existence of a Δ -diagram δ that maps κ to n and $t[\kappa]$ to $\llbracket t \rrbracket_h(n)$ for all $t \in T$, where Δ is the saturation of the sample set $\{t[\kappa] \mid t \in T\}$.

Proposition 3.3. *Let T be a set of basic terms, κ a time variable, and Δ the saturation of the sample set $\{t[\kappa] \mid t \in T\}$. Then for any homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ and $n \in \omega^+$, there exists a Δ -diagram δ such that $\delta(\kappa) = n$ and $\delta(t[\kappa]) = \llbracket t \rrbracket_h(n)$ for all $t \in T$.*

Proof. We define the map $\delta: \Delta \rightarrow \omega^+$ recursively by

$$\begin{aligned} \delta(\kappa) &:= n \\ \delta(t[\alpha]) &:= \llbracket t \rrbracket_h(\delta(\alpha)) && \text{for each } t[\alpha] \in \Delta \\ \delta(\text{last}(t)) &:= \text{last}(\llbracket t \rrbracket_h) && \text{for each } \text{last}(t) \in \Delta \\ \delta(s(\alpha)) &:= \mathcal{S}(\delta(\alpha)) && \text{for each } s(\alpha) \in \Delta. \end{aligned}$$

The map δ is well-defined since $\alpha \in \Delta$ if, and only if, there exist samples $\alpha_1, \dots, \alpha_n$ such that $\alpha_1 = t[\kappa]$ for some $t \in T$, $\alpha_n = \alpha$, and $\alpha_j \rightsquigarrow \alpha_{j+1}$ for each $j \in \{1, \dots, n-1\}$. So δ is a \mathfrak{A} -valuation. It remains to prove that δ is a diagram, i.e., that δ satisfies (3.1)-(3.12). Let us assume for convenience without further mention that all samples used are in Δ , and agree to write $\llbracket t \rrbracket$ for $\llbracket t \rrbracket_h$.

(3.1) If $\delta(\alpha) \leq \delta(\beta)$, then, by the definition of δ , using the fact that time warps are order-preserving, $\delta(t[\alpha]) = \llbracket t \rrbracket(\delta(\alpha)) \leq \llbracket t \rrbracket(\delta(\beta)) = \delta(t[\beta])$.

(3.2) If $\delta(\alpha) = 0$, then $\delta(t[\alpha]) = \llbracket t \rrbracket(\delta(\alpha)) = 0$.

(3.3) Follows directly from the definition of δ .

(3.4) Using the definition of δ ,

$$\delta(\text{last}(t)) = \text{last}(\llbracket t \rrbracket) = \min\{n \in \omega^+ \mid \llbracket t \rrbracket(n) = \llbracket t \rrbracket(\omega)\}.$$

Hence, clearly, for each $k \in \omega^+$,

$$\text{last}(\llbracket t \rrbracket) \leq k \iff \llbracket t \rrbracket(\text{last}(\llbracket t \rrbracket)) = \llbracket t \rrbracket(k),$$

and therefore, for all $t[\alpha] \in \Delta$,

$$\delta(\text{last}(t)) \leq \delta(\alpha) \iff \delta(t[\text{last}(t)]) = \delta(t[\alpha]).$$

(3.5) If $\text{last}(\llbracket t \rrbracket) = \delta(\text{last}(t)) = \omega$, then $\llbracket t \rrbracket(n) < \llbracket t \rrbracket(\omega)$ for all $n \in \omega$, and it follows that $\delta(t[\text{last}(t)]) = \llbracket t \rrbracket(\omega) = \bigvee\{\llbracket t \rrbracket(n) \mid n \in \omega\} = \omega$.

(3.6) $\delta(e[\alpha]) = id(\delta(\alpha)) = \delta(\alpha)$.

(3.7) $\delta(\text{last}(e)) = \text{last}(id) = \omega$.

(3.8) $\delta(tu[\alpha]) = \llbracket tu \rrbracket(\delta(\alpha)) = \llbracket t \rrbracket(\llbracket u \rrbracket(\delta(\alpha))) = \delta(t[u[\alpha]])$.

- (3.9) If $\delta(\text{last}(tu)) = \omega$, then $\text{last}(\llbracket t \rrbracket \llbracket u \rrbracket) = \omega$ and, using Lemma 2.6, it follows that $\delta(\text{last}(t)) = \text{last}(\llbracket t \rrbracket) = \omega$ and $\delta(\text{last}(u)) = \text{last}(\llbracket u \rrbracket) = \omega$.
- (3.10) If $0 < \delta(\alpha) < \omega$, then Proposition 2.7 yields $\delta(t[t'[\alpha]]) = \llbracket t \rrbracket(\llbracket t \rrbracket^*(\delta(\alpha))) < \delta(\alpha)$.
- (3.11) If $\llbracket t \rrbracket^*(\delta(\alpha)) = \delta(t'[\alpha]) < \omega$, then either $\delta(\alpha) = 0 \leq \delta(t[s(t'[\alpha])])$ or $\delta(\alpha) > 0$ and Proposition 2.7 yields $\delta(\alpha) \leq \llbracket t \rrbracket(\llbracket t \rrbracket^*(\delta(\alpha)) + 1) = \delta(t[s(t'[\alpha])])$.
- (3.12) It follows directly from Lemma 2.6 that $\text{last}(\llbracket t \rrbracket^*) = \delta(\text{last}(t')) = \omega$ if, and only if, $\delta(\text{last}(t)) = \text{last}(\llbracket t \rrbracket) = \omega$. \square

Running example part 3. Consider again $\Delta = \{xx'[\kappa]\}^{\sim}$. Applied to a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ with $h(x) = id$ and $n = 1$, the construction in the above proof yields exactly the diagram $\delta: \Delta \rightarrow \omega^+$ defined in part 2 of the running example.

Next, we provide an algorithm that converts any Δ -diagram δ for a finite saturated sample set Δ into an algorithmic description of a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ satisfying $\llbracket t \rrbracket_h(\delta(\alpha)) = \delta(t[\alpha])$ for all $t[\alpha] \in \Delta$. Indeed, we show that this homomorphism can be described by mapping each $x \in \text{Var}$ occurring in Δ to a *regular* time warp, i.e., a time warp that is either eventually constant or eventually linear. The main challenge here is to show that each partial map consisting of ordered pairs $(\delta(\alpha), \delta(x[\alpha]))$ with $x[\alpha] \in \Delta$ extends to a time warp and that a composition of these time warps given by a basic term t extends the set of ordered pairs $(\delta(\alpha), \delta(t[\alpha]))$ with $t[\alpha] \in \Delta$.

For any finite saturated sample set Δ , Δ -diagram δ , and basic term t , let

$$\llbracket t \rrbracket_\delta := \{(\delta(\alpha), \delta(t[\alpha])) \mid t[\alpha] \in \Delta\}.$$

We will say that a time warp f

- *extends* $\llbracket t \rrbracket_\delta$ if $f(i) = j$ for all $(i, j) \in \llbracket t \rrbracket_\delta$;
- *strongly extends* $\llbracket t \rrbracket_\delta$ if it extends $\llbracket t \rrbracket_\delta$ and also

$$\llbracket t \rrbracket_\delta \neq \emptyset \text{ and } \delta(\text{last}(t)) = \omega \implies \text{last}(f) = \omega.$$

For convenience, let us fix for the following five lemmas — which collectively describe how to obtain a time warp that strongly extends $\llbracket t \rrbracket_\delta$ based on the structure of t — a finite saturated sample set Δ and Δ -diagram δ .

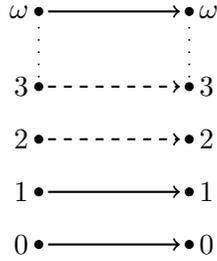
Lemma 3.4. *There exists an algorithm that produces an algorithmic description, for any variable $x \in \text{Var}$, of a time warp f that strongly extends $\llbracket x \rrbracket_\delta$.*

Proof. If $\llbracket x \rrbracket_\delta = \emptyset$, then any time warp strongly extends it, so we may assume that $\llbracket x \rrbracket_\delta \neq \emptyset$. By (3.1), if $x[\alpha], x[\beta] \in \Delta$ with $\delta(\alpha) \leq \delta(\beta)$, then also $\delta(x[\alpha]) \leq \delta(x[\beta])$. Hence $\llbracket x \rrbracket_\delta$ is a partial order-preserving map from ω^+ to ω^+ . Moreover, since Δ is saturated, it follows that $(\delta(\text{last}(x)), \delta(x[\text{last}(x)])) \in \llbracket x \rrbracket_\delta$. Now for $x[\alpha] \in \Delta$, if $\delta(\alpha) = 0$, then, by (3.2), $\delta(x[\alpha]) = 0$, and if $\delta(\alpha) \geq \delta(\text{last}(x))$, then, by (3.4), $\delta(x[\alpha]) = \delta(x[\text{last}(x)])$. Hence the relation $X := \llbracket x \rrbracket_\delta \cup \{(0, 0), (\omega, \delta(x[\text{last}(x)]))\}$ is also a partial order-preserving map. For each $i \in \omega$, there exists a unique pair $(i_1, j_1), (i_2, j_2) \in X$ such that $i_1 \leq i < i_2$ and there is no $(i_3, j_3) \in X$ with $i_1 < i_3 < i_2$, and we define

$$f(i) := \min(j_2, j_1 \oplus (i - i_1)),$$

where $n \oplus m := \min\{\omega, n + m\}$. Let also $f(\omega) := \delta(x[\text{last}(x)])$.

Clearly, since X is order-preserving, also f is order-preserving. Moreover, it extends $\llbracket x \rrbracket_\delta$, since $i = i_1 < \omega$ implies $f(i_1) = \min(j_2, j_1) = j_1$. In particular, $f(0) = 0$. To confirm that f is a time warp, it remains to show that $f(\omega) = \bigvee \{f(i) \mid i \in \omega\}$. If $\delta(x[\text{last}(x)]) = f(\omega) < \omega$, then, by (3.5), $\delta(\text{last}(x)) < \omega$ and, since f is order-preserving, $f(i) = f(\omega)$ for each $i \geq$

Figure 2: Visualization of the extension of $[x]_\delta$

$\delta(\text{last}(x))$ and $f(\omega) = f(\delta(\text{last}(x))) = \bigvee\{f(i) \mid i \in \omega\}$. If $f(\omega) = \omega$, then for each $j \in \omega$, there exists an $i \in \omega$ such that $f(i) > j$, and hence $\bigvee\{f(i) \mid i \in \omega\} = \omega = f(\omega)$.

Finally, suppose that $\delta(\text{last}(x)) = \omega$. Then (3.5) yields $(\omega, \omega) \in [x]_\delta$ and, by (3.4), for any $(i, j) \in [x]_\delta$, if $i \in \omega$, then also $j \in \omega$. Hence, $\text{last}(f) = \omega$, by the definition of f . So f strongly extends $[x]_\delta$. \square

Running example part 4. Consider again $\Delta = \{xx'[\kappa]\}^{\sim}$ and the Δ -diagram $\delta: \Delta \rightarrow \omega^+$ from parts 2 and 3 of the running example. The construction in the above proof yields the strong extension of $[x]_\delta$ depicted in Figure 2, where the added relations are indicated with dashed arrows. Indeed, the obtained extension is the time warp id , which we already know to be a suitable extension.

Lemma 3.5. *Let t_1 and t_2 be basic terms. If f_1 and f_2 strongly extend $[t_1]_\delta$ and $[t_2]_\delta$, respectively, then $f_1 f_2$ strongly extends $[t_1 t_2]_\delta$.*

Proof. Suppose that f_1 and f_2 strongly extend $[t_1]_\delta$ and $[t_2]_\delta$, respectively. Then for all $t_1 t_2[\alpha] \in \Delta$,

$$\begin{aligned}
 f_1 f_2(\delta(\alpha)) &= f_1(f_2(\delta(\alpha))) && \text{(by definition)} \\
 &= f_1(\delta(t_2[\alpha])) && \text{(since } f_2 \text{ extends } [t_2]_\delta) \\
 &= \delta(t_1[t_2[\alpha]]) && \text{(since } f_1 \text{ extends } [t_1]_\delta) \\
 &= \delta(t_1 t_2[\alpha]) && \text{(by (3.8)).}
 \end{aligned}$$

So $f_1 f_2$ extends $[t_1 t_2]_\delta$, and it remains to show that the extension is strong. We can assume that $[t_1 t_2]_\delta$ is non-empty, since otherwise there is nothing to prove. Suppose that $\delta(\text{last}(t_1 t_2)) = \omega$. Then $\delta(\text{last}(t_1)) = \delta(\text{last}(t_2)) = \omega$, by (3.9), and, since f_1 and f_2 strongly extend $[t_1]_\delta$ and $[t_2]_\delta$, respectively, also $\text{last}(f_1) = \text{last}(f_2) = \omega$. Hence $\text{last}(f_1 f_2) = \omega$, by Lemma 2.6. \square

Lemma 3.6. *Let t be a basic term. If f strongly extends $[t]_\delta$, then f^* strongly extends $[t']_\delta$.*

Proof. Let $t'[\alpha] \in \Delta$. Note that, by (3.2), if $\delta(\alpha) = 0$, then $\delta(t'[\alpha]) = 0 = f^*(0)$. Hence we can assume that $\delta(\alpha) > 0$. Suppose first that $\delta(t'[\alpha]) < \omega$. Then, by (3.11), $\delta(\alpha) \leq \delta(t[\mathfrak{s}(t'[\alpha])])$. Since f extends $[t]_\delta$, it follows that $\delta(\alpha) \leq \delta(t[\mathfrak{s}(t'[\alpha])]) = f(\delta(\mathfrak{s}(t'[\alpha])))$, and, by (3.3), also $f(\delta(\mathfrak{s}(t'[\alpha]))) = f(\delta(t'[\alpha]) + 1)$, i.e., $\delta(\alpha) \leq f(\delta(t'[\alpha]) + 1)$. We have two cases:

- (1) $\delta(\alpha) < \omega$. Then, by (3.10) and since f extends $[t]_\delta$, it follows that $f(\delta(t'[\alpha])) = \delta(t[t'[\alpha]]) < \delta(\alpha)$. So $\delta(t'[\alpha]) = f^*(\delta(\alpha))$, by Proposition 2.7.
- (2) $\delta(\alpha) = \omega$. Then, since $\delta(t'[\alpha]) < \omega$, by (3.4) and (3.5), it follows that $\text{last}(t') < \omega$ and $\delta(t'[\text{last}(t')]) = \delta(t'[\alpha])$. Hence, as in the first case, $f(\delta(t'[\text{last}(t')])) < \delta(\text{last}(t'))$. So $f(\delta(t'[\alpha])) < \omega = \delta(\alpha) \leq f(\delta(t'[\alpha]) + 1)$, and $\delta(t'[\alpha]) = f^*(\delta(\alpha))$, by Proposition 2.7.

Otherwise $\delta(t'[\alpha]) = \omega$ and we again have two cases:

- (1) $\delta(\alpha) < \omega$. Then, since $f(\delta(t'[\alpha])) = \delta(t[t'[\alpha]]) < \delta(\alpha)$ by (3.10), it follows that $f(\omega) = f(\delta(t'[\alpha])) < \delta(\alpha)$. Hence $f^*(\delta(\alpha)) = \omega$, by Proposition 2.7.
- (2) $\delta(\alpha) = \omega$. Then there are two cases. If $\delta(\text{last}(t')) < \omega$, then, using the previous cases, $f^*(\delta(\text{last}(t'))) = \omega$, and hence $f^*(\omega) = \omega$. Otherwise $\delta(\text{last}(t')) = \omega$. Then $\delta(\text{last}(t)) = \omega$, by (3.12), and, since f strongly extends $[t]_\delta$, also $\text{last}(f) = \omega$. But then $f^*(\omega) = \omega$, by Proposition 2.7.

It remains to show that the extension of $[t']_\delta$ to f^* is strong. We can assume that $[t']_\delta$ is non-empty. Suppose that $\delta(\text{last}(t')) = \omega$. Then $\delta(\text{last}(t)) = \omega$, by (3.12), and, since f strongly extends $[t]_\delta$, also $\text{last}(f) = \omega$. Hence $\text{last}(f^*) = \omega$, by Lemma 2.6. \square

Lemma 3.7. *The time warp id strongly extends $[e]_\delta$.*

Proof. The extension property follows from (3.6); the fact that it is strong follows from the fact that $\text{last}(id) = \omega$. \square

Lemma 3.8. *For every basic term t and homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$, if $h(x)$ strongly extends $[x]_\delta$ for each $x \in \text{Var}$ occurring in t , then $\llbracket t \rrbracket_h$ strongly extends $[t]_\delta$.*

Proof. By induction on t . The case $t = x$ is immediate and the other cases follow from Lemmas 3.5–3.7, and the induction hypothesis. \square

The next proposition is then a direct consequence of Lemmas 3.4 and 3.8.

Proposition 3.9. *There exists an algorithm that produces for any finite saturated sample set Δ and Δ -diagram δ , an algorithmic description of a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ satisfying $\llbracket t \rrbracket_h(\delta(\alpha)) = \delta(t[\alpha])$ for all $t[\alpha] \in \Delta$.*

We have now assembled all the necessary ingredients to relate the validity of an equation in the time warp algebra to the existence of a corresponding diagram.

Proposition 3.10. *Let t_1, \dots, t_n be basic terms, κ a time variable, and Δ the saturation of the sample set $\{t_1[\kappa], \dots, t_n[\kappa]\}$. Then $\mathbf{W} \not\models id \leq t_1 \vee \dots \vee t_n$ if, and only if, there exists a Δ -diagram δ such that $\delta(\kappa) > \delta(t_i[\kappa])$ for each $i \in \{1, \dots, n\}$.*

Proof. Suppose first that $\mathbf{W} \not\models id \leq t_1 \vee \dots \vee t_n$. Then there exist a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ and an $m \in \omega^+$ such that $m = id(m) > \llbracket t_i \rrbracket_h(m)$ for each $i \in \{1, \dots, n\}$. Hence, by Proposition 3.3, there exists a Δ -diagram δ such that $\delta(\kappa) = m > \llbracket t_i \rrbracket_h(m) = \delta(t_i[\kappa])$ for each $i \in \{1, \dots, n\}$.

For the converse, suppose that there exists a Δ -diagram δ such that $\delta(\kappa) > \delta(t_i[\kappa])$ for each $i \in \{1, \dots, n\}$. Then, by Proposition 3.9, there exists a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ such that $\llbracket t_i \rrbracket_h(\delta(\kappa)) = \delta(t_i[\kappa])$ for each $i \in \{1, \dots, n\}$. So $id(\delta(\kappa)) = \delta(\kappa) > \llbracket t_i \rrbracket_h(\delta(\kappa))$ for each $i \in \{1, \dots, n\}$. Hence $\mathbf{W} \not\models id \leq t_1 \vee \dots \vee t_n$. \square

Running example part 5. The Δ -diagram $\delta: \Delta \rightarrow \omega^+$ for $\Delta = \{xx'[\kappa]\}^{\rightsquigarrow}$ from parts 2 and 3 of the running example witnesses the failure of the equation $e \leq xx'$ in \mathbf{W} , since $\delta(\kappa) = 1 > 0 = \delta(xx'[\kappa])$.

Observe now that the time warps constructed in the proof of Lemma 3.4 are regular. Hence, since, by Proposition 2.5, the set of regular time warps forms a subalgebra \mathbf{R} of \mathbf{W} , we obtain the following correspondence:

Theorem 3.11. *For any $s, t \in \mathbf{Tm}$, $\mathbf{W} \models s \approx t$ if, and only if, $\mathbf{R} \models s \approx t$.*

Finally, let t_1, \dots, t_n be basic terms, κ a time variable, and Δ the saturation of the sample set $\{t_1[\kappa], \dots, t_n[\kappa]\}$. Define also $\text{fail} := \{t_1[\kappa] \prec \kappa, \dots, t_n[\kappa] \prec \kappa\}$ and let ψ_Δ be the existential closure of the τ -formula $\bigwedge(\Sigma_\Delta \cup \text{fail})$. Then, by Proposition 3.10, it follows that $\mathbf{W} \not\models id \leq t_1 \vee \dots \vee t_n$ if, and only if, there exists a \mathfrak{A} -valuation δ such that $\mathfrak{A}, \delta \models \Sigma_\Delta \cup \text{fail}$ if, and only if, $\mathfrak{A} \models \psi_\Delta$. Since the first-order theory of $\langle \omega^+, \preceq \rangle$ is decidable (see, e.g., [LL66]) and the operations 0 , ω , and \mathcal{S} are definable in $\langle \omega^+, \preceq \rangle$, we obtain the following result:

Theorem 3.12. *The equational theory of \mathbf{W} is decidable.*

4. IMPLEMENTATION

We have implemented the decision procedure described in Section 3 as a software tool. The tool is written in the OCaml programming language [LDF⁺22] and makes use of the Z3 theorem prover [dMB08]. It reads an equation between terms in the language $\{\wedge, \vee, \cdot, \backslash, /, ', e\}$ from the command line and translates it into a Satisfiability Modulo Theory (SMT) query. This query expresses the existence of corresponding diagrams for the equation, constructed as in Section 3, and is unsatisfiable if, and only if, the equation is satisfied by the time warp algebra. If the query is satisfiable, the tool translates the model provided by Z3 into a homomorphism from \mathbf{Tm} to \mathbf{W} falsifying the initial problem.

Outline. We summarize the main steps of the decision procedure below, mentioning details of the implementation in passing.

- (1) The input problem is normalized into an inequation of the form $e \leq t$. This is always possible since:
 - any inequation $u \leq v$ can be rewritten into the inequation $e \leq u \backslash v$;
 - an equation $u \approx v$ can be rewritten into the inequation $e \leq (u \backslash v) \wedge (v \backslash u)$.
- (2) The inequation $e \leq t$ is rewritten, as per Lemma 2.1, into an inequation of the form

$$e \leq \bigwedge_{i=1}^m \bigvee_{j=1}^{n_i} t_{i,j},$$

where each $t_{i,j}$ is a basic term, by first iteratively replacing subterms of the form $r \backslash s$ and s/t with $(s'r)'$ and $(rs')'$, respectively, then iteratively distributing joins over meets and the monoid multiplication over meets and joins, pushing the involution inwards using the De Morgan laws. Each inequation $e \leq \bigvee_{j=1}^{n_i} t_{i,j}$ can then be checked independently for $i \in \{1, \dots, m\}$.

- (3) Optionally, each basic term $t_{i,j}$ can be rewritten using simple algebraic rules. In effect, this step computes normal forms for the following convergent rewrite system:

$$t \cdot e \rightarrow t \qquad e \cdot t \rightarrow t \qquad t'' \rightarrow t.$$

This simplification step is optional but easy to implement and can drastically reduce the cost of the rest of the decision procedure.

- (4) The tool computes for each i in $\{1, \dots, m\}$, the saturated sample set

$$\Delta_i := \{t_{i,1}[\kappa_i], \dots, t_{i,n_i}[\kappa_i]\}^{\sim\sim},$$

where $\kappa_1, \dots, \kappa_n$ are distinct sample variables. Samples are represented as first-order terms with maximal sharing [FC06], so that sample equations can be tested in constant time. This accelerates the membership tests used to check whether saturation is complete.

- (5) For each i in $\{1, \dots, m\}$, the set Δ_i is traversed to generate an SMT query built from the quantifier-free first-order formulas described in Section 3. These first-order formulas are expressed over the theory known as *Integer Difference Logic* in the SMT literature [BFT16], which corresponds to the first-order theory of the structure $\mathcal{Z} = \langle \mathbb{Z}, \leq^{\mathcal{Z}}, 0^{\mathcal{Z}}, \mathcal{S}^{\mathcal{Z}} \rangle$, where $\leq^{\mathcal{Z}}$ is the natural order on \mathbb{Z} , $0^{\mathcal{Z}} = 0$, and $\mathcal{S}^{\mathcal{Z}}(m) = m + 1$, for each $m \in \mathbb{Z}$. We encode \mathfrak{A} in this structure by representing $\omega^{\mathfrak{A}}$ as $0^{\mathcal{Z}}$ and $n \in \omega$ as $\mathcal{S}^{\mathcal{Z}}n$. More precisely, writing $\langle - \rangle$ for the encoding map, we let

$$\begin{aligned} \langle 0^{\mathfrak{A}} \rangle &:= \mathcal{S}^{\mathcal{Z}}0^{\mathcal{Z}} \\ \langle \omega^{\mathfrak{A}} \rangle &:= 0^{\mathcal{Z}} \\ \langle \xi \approx \mathcal{S}^{\mathfrak{A}}t \rangle &:= (\langle t \rangle \approx 0^{\mathcal{Z}} \wedge \xi \approx 0^{\mathcal{Z}}) \vee (\mathcal{S}^{\mathcal{Z}}0^{\mathcal{Z}} \leq^{\mathcal{Z}} \langle t \rangle \wedge \xi \approx \mathcal{S}^{\mathcal{Z}}\langle t \rangle) \\ \langle t \leq^{\mathfrak{A}} u \rangle &:= (\langle u \rangle \approx 0^{\mathcal{Z}}) \vee (\mathcal{S}^{\mathcal{Z}}0^{\mathcal{Z}} \leq^{\mathcal{Z}} \langle t \rangle \wedge \langle t \rangle \leq^{\mathcal{Z}} \langle u \rangle), \end{aligned}$$

where ξ denotes an arbitrary first-order variable. Since the operation $\mathcal{S}^{\mathfrak{A}}$ is encoded as a formula rather than a term in the target theory, we have assumed without loss of generality that it only appears at the root of an equation. Additionally, for each first-order variable ξ , the tool generates the formula $0^{\mathcal{Z}} \leq^{\mathcal{Z}} \xi$.

- (6) Finally, the i th query is passed to the SMT solver. It is satisfiable if, and only if, $\mathbf{W} \models e \leq \bigvee_{j=1}^{n_i} t_{i,j}$. In this case, the tool uses the diagram returned by the SMT solver to construct a homomorphism $h: \mathbf{Tm} \rightarrow \mathbf{W}$ and element $k \in \omega^+$ satisfying $k > \llbracket t_{i,j} \rrbracket_h(k)$, for each $j \in \{1, \dots, n_i\}$, following the approach described in the proof of Lemma 3.4.

Examples. The running example $e \leq xx'$ considered in Section 3 can be handled by our implementation. The generated SMT query contains 139 clauses. Z3 checks its satisfiability and constructs the diagram

$$\begin{aligned} 0 &= \delta(xx'[\kappa]) = \delta(x[x'[\kappa]]) = \delta(xx'[\text{last}(xx')]) = \delta(x[x'[\text{last}(xx')]]) = \delta(x'[\text{last}(xx')]) \\ &= \delta(x[x'[\text{last}(x')]]) = \delta(x[\text{s}(x'[\text{last}(xx')])]) = \delta(\text{last}(xx')) \\ 1 &= \delta(\text{s}(x'[\text{last}(xx')])) = \delta(\text{last}(x')) \\ 2 &= \delta(x[\text{s}(x'[\kappa])]) = \delta(x'[\kappa]) = \delta(x[\text{s}(x'[\text{last}(x')])]) = \delta(x'[\text{last}(x')]) = \delta(\kappa) \\ 3 &= \delta(\text{s}(x'[\text{last}(x')])) = \delta(\text{s}(x'[\kappa])) = \delta(x[\text{last}(x)]) \\ 4 &= \delta(\text{last}(x)), \end{aligned}$$

which expresses that any time warp f satisfying $f(n) = 1$ for $n < 3$, $f(3) = 2$, and $f(4) = 3$, satisfies $f(f^*(1)) < 1$.

As a further example, consider the inequation $e \leq x$, which is small enough to make each step of the decision procedure explicit. The corresponding saturated sample set is

$$\Delta := \{\kappa, x[\kappa], \text{last}(x), x[\text{last}(x)]\}.$$

The tool traverses Δ to generate a first-order formula expressing the existence of a diagram. Expressed in mathematical notation, this formula is

$$\begin{aligned}
& \exists \xi_\kappa, \xi_{x[\kappa]}, \xi_{\text{last}(x)}, \xi_{x[\text{last}(x)]}, \bigwedge_{\alpha \in \Delta} (0^{\mathcal{Z}} \leq^{\mathcal{Z}} \alpha) \\
& \wedge (\xi_\kappa \preceq^{\mathfrak{A}} \xi_{\text{last}(x)} \Rightarrow \xi_{x[\kappa]} \preceq^{\mathfrak{A}} \xi_{x[\text{last}(x)]}) \\
& \wedge (\xi_{\text{last}(x)} \preceq^{\mathfrak{A}} \xi_\kappa \Rightarrow \xi_{x[\text{last}(x)]} \preceq^{\mathfrak{A}} \xi_{x[\kappa]}) \\
& \wedge \bigwedge_{\alpha \in \{\kappa, \text{last}(x)\}} (\xi_\alpha \approx 0^{\mathfrak{A}} \Rightarrow \xi_{x[\alpha]} \approx 0^{\mathfrak{A}}) \\
& \wedge \bigwedge_{\alpha \in \{\kappa, \text{last}(x)\}} (\xi_{\text{last}(x)} \preceq^{\mathfrak{A}} \xi_\alpha \Rightarrow \xi_{x[\alpha]} \approx \xi_{x[\text{last}(x)]}) \\
& \wedge \xi_{\text{last}(x)} \approx \omega^{\mathfrak{A}} \Rightarrow \xi_{x[\text{last}(x)]} \approx \omega^{\mathfrak{A}} \\
& \wedge \xi_{x[\kappa]} \prec^{\mathfrak{A}} \xi_\kappa,
\end{aligned}$$

where the symbols $\preceq^{\mathfrak{A}}$, $0^{\mathfrak{A}}$, $\omega^{\mathfrak{A}}$, and $\mathcal{S}^{\mathfrak{A}}$ stand for their actual encodings in terms of $\leq^{\mathcal{Z}}$, $0^{\mathcal{Z}}$, and $\mathcal{S}^{\mathcal{Z}}$. Figure 3 displays the same formula as a query in concrete SMT-LIB syntax. The final step is to check its validity using Z3. The SMT solver builds the diagram

$$\begin{aligned}
0 &= \delta(x[\text{last}(x)]) = \delta(x[\kappa]) \\
1 &= \delta(\text{last}(x)) \\
2 &= \delta(\kappa),
\end{aligned}$$

expressing in particular that mapping x to the constant time warp \perp falsifies $e \leq x$ in \mathbf{W} .

Computational complexity. The worst-case computational complexity of the decision procedure is exponential in the size of the initial equation. This complexity breaks down as follows. Steps 2 and 4 have a running time that is exponential in the size of their input data. In particular, the size of the saturated sample set for a given basic term $t_{i,j}$ is exponential in the size of $t_{i,j}$ in general. The size of the SMT query generated during step 5 is polynomial in the number of samples in the input sample set. Finally, the SMT solver has a running time that is exponential in the size of the SMT query in the worst case, as the conjunctive fragment of difference logic is decidable in polynomial time [Dil89].

Software complexity. The entirety of the implementation, not counting Z3 and its OCaml interface, takes slightly less than 1900 lines of code. The core procedure and data structures represent around 550 lines of code. The code is available online as libre software.³

³<https://github.com/adrieng/twmc>

```

(declare-fun !0 () Int)
(declare-fun !1 () Int)
(declare-fun !2 () Int)
(declare-fun !3 () Int)
(assert (<= 0 !0))
(assert (<= 0 !1))
(assert (<= 0 !2))
(assert (<= 0 !3))
(assert (let ((a!1 (or (= !2 0) (and (not (= !3 0)) (<= !3 !2))))
              (a!2 (or (= !0 0) (and (not (= !1 0)) (<= !1 !0)))))
        (=> a!1 a!2)))
(assert (let ((a!1 (or (= !3 0) (and (not (= !2 0)) (<= !2 !3))))
              (a!2 (or (= !1 0) (and (not (= !0 0)) (<= !0 !1)))))
        (=> a!1 a!2)))
(assert (=> (= !3 1) (= !1 1)))
(assert (let ((a!1 (or (= !3 0) (and (not (= !2 0)) (<= !2 !3))))
              (and (=> a!1 (= !0 !1)) (=> (= !0 !1) a!1))))
        (=> (= !2 1) (= !0 1)))
(assert (let ((a!1 (or (= !2 0) (and (not (= !2 0)) (<= !2 !2))))
              (and (=> a!1 (= !0 !0)) (=> (= !0 !0) a!1))))
        (=> (= !2 0) (= !0 0)))
(assert (let ((a!1 (or (= !3 0) (and (not (= !1 0)) (<= !1 !3))))
              (and a!1 (not (= !1 !3)))))
        (=> (= !2 0) (= !0 0)))

```

Figure 3: SMT query whose models correspond to diagrams for $e \leq x$

5. CONCLUDING REMARKS

The main contribution of this paper is a proof that the equational theory of the time warp algebra \mathbf{W} is decidable, supported by an implementation of a corresponding decision procedure. There remain, however, many open problems regarding computational and structural properties of \mathbf{W} and related algebras of join-preserving maps on complete chains. Below, we briefly discuss some of these problems and potential avenues for further research.

An axiomatic description. Although we have provided an algorithm to decide equations in the time warp algebra, an axiomatic description of its equational theory is still lacking. The algebra generates a variety of fully distributive involutive residuated lattices satisfying the equations $e' \leq e$ and $e \approx e' \cdot (e' \setminus e)$, but not much more is known. An axiomatization could also provide the basis for developing a sequent calculus for reasoning in \mathbf{W} along the lines of Yetter’s cyclic linear logic without exponentials [Yet90].

Computational complexity. As explained in Section 4, the decision procedure described in this paper provides an exponential upper bound for the computational complexity of deciding equations in the time warp algebra. A lower bound for this problem follows from the fact that the involution-free reduct of \mathbf{W} generates the variety of distributive ℓ -monoids, which has a co-NP-complete equational theory [CGMS22]. Tight complexity bounds have yet to be determined, however.

Applications to graded modalities. The implementation described in Section 4 can be readily integrated in a type checker for a programming language with a graded modality whose gradings are first-order terms over \mathbf{W} , generalizing previous work [Gua18]. In broad strokes, such a type checker reduces each subtyping problem arising during type checking to a finite set of inequalities in the time warp algebra. The simplest interesting case is that of checking whether the type $\Box_t A$ is a subtype of $\Box_u A$, which reduces to checking $\mathbf{W} \models u \leq t$. The other cases in reducing type-checking problems to universal-algebraic problems depend on the exact language of types under consideration, including its subtyping relationship, which is beyond the scope of the current paper. Let us mention, however, that certain programming language features, such as the modeling of fixed computation rates, may require extending the language of gradings, and hence also the decision procedure. We discuss a simple example of such an extension in the next paragraph.

Extending the language. For certain applications, it may be profitable to extend the language of time warps with further operations. To illustrate, let us describe how ‘definable’ time warps can be added as constants to the language, while preserving decidability of the resulting equational theory. Let X be a countably infinite set of fresh variables for building first-order formulas. We call a time warp f *definable* if there exist first-order formulas $\phi_f(x, y)$ and $\psi_f(z)$ in the signature $\tau = \{\preceq, \mathcal{S}, 0, \omega\}$ with free variables x, y , and z , respectively such that for every \mathfrak{A} -valuation $v: X \rightarrow \omega^+$,

$$\mathfrak{A}, v \models \phi_f(x, y) \iff f(v(x)) = v(y) \quad \text{and} \quad \mathfrak{A}, v \models \psi_f(z) \iff v(z) = \text{last}(f).$$

For example, the time warp \top satisfying $\top(0) = 0$ and $\top(n) = \omega$, for each $n \in \omega^+ \setminus \{0\}$, is definable by $\phi_\top(x, y) := (x \approx 0 \Rightarrow y \approx 0) \wedge (0 \prec x \Rightarrow y \approx \omega)$ and $\psi_\top(z) := z \approx \mathcal{S}(0)$. Similarly, the time warp \perp satisfying $\perp(n) = 0$, for each $n \in \omega^+$, is definable by $\phi_\perp(x, y) := y \approx 0$ and $\psi_\perp(z) := z \approx 0$.

Let \mathcal{D} denote the set of definable time warps. Given any $F = \{f_1, \dots, f_n\} \subseteq \mathcal{D}$, let $\mathbf{W}[F]$ denote the algebra $\langle \text{Res}(\omega^+), \wedge, \vee, \circ, *, \text{id}, f_1, \dots, f_n \rangle$. We call a term in the signature $\{\cdot, ', e, f_1, \dots, f_n\}$ *F-basic* and define *F-samples* analogously to samples by allowing *F*-basic terms in the construction. The notion of saturation is extended to *F*-sample sets, and following the proof of Lemma 3.2 shows that also the saturation of a finite *F*-sample set is finite. The notion of a diagram is then extended to saturated *F*-sample sets Δ by defining the sets struct_Δ , log_Δ , inv_Δ as before, and also

$$\text{const}_\Delta := \{\phi_f(\alpha, f[\alpha]) \mid f[\alpha] \in \Delta, f \in F\} \cup \{\psi_f(\text{last}(f)) \mid \text{last}(f) \in \Delta, f \in F\}.$$

Let $\Sigma_\Delta := \text{struct}_\Delta \cup \text{log}_\Delta \cup \text{inv}_\Delta \cup \text{const}_\Delta$. A Δ -*diagram* is a \mathfrak{A} -valuation $\delta: \Delta \cup X \rightarrow \omega^+$ such that $\mathfrak{A}, \delta \models \Sigma_\Delta$. It is clear that Proposition 3.3 extends to saturated *F*-sample sets, using the definition of a definable time warp and the fact that homomorphisms from the term algebra into $\mathbf{W}[F]$ map f to f , for each $f \in F$. Similarly, since every $f \in F$ is completely described by the formulas $\phi_f(x, y)$ and $\psi_f(z)$, also Proposition 3.9, extends to saturated *F*-sample sets. Hence, arguing as in Section 3, the equational theory of $\mathbf{W}[F]$ is decidable.

A relational approach. We end this paper by discussing a different, relational point of view on time warps, which we believe will play a role in generalizing and extending the methods of this paper. It is well-known that, if \mathbf{L} is a sufficiently well-behaved lattice, then any completely join-preserving function on \mathbf{L} can be uniquely described via a certain binary relation on the set P of completely join-prime elements of \mathbf{L} . The binary relations that

occur in this way are exactly the *distributors* on P , when viewed as a posetal category. This raises the question as to whether the decision procedure we give here in the case $\mathbf{L} = \omega^+$ can be generalized to the algebraic structure of distributors on an object of a sufficiently well-behaved category P . In order to illustrate the idea, we explain here how the results of the paper can be alternatively understood in this relational language.

Let us write $P = \{1, 2, \dots\}$ for the set of positive natural numbers, so that ω^+ is isomorphic to the lattice of downward closed subsets of P , via the function that sends $x \in \omega^+$ to $\downarrow x \cap P$. A time warp $f: \omega^+ \rightarrow \omega^+$ may then be viewed alternatively as a binary relation $R \subseteq P \times P$ that is *monotone* (also called a *weakening relation* in, e.g., [GJ20]): that is, for any $x, x', y, y' \in P$, if $x' \leq x$, xRy , and $y \leq y'$, then $x'Ry'$. Indeed, given a time warp f , the binary relation R_f defined by

$$R_f := \{(x, y) \in P^2 \mid x \leq f(y)\}$$

is clearly monotone, and conversely, given a monotone binary relation R on P , the function $f: P \rightarrow \omega^+$ defined, for $y \in P$, by

$$f(y) := \bigvee \{x \in P \mid xRy\}$$

is order-preserving and hence extends uniquely to a time warp by setting $f(0) := 0$ and $f(\omega) := \bigvee_{p \in P} f(p)$. These assignments $f \mapsto R_f$ and $R \mapsto f_R$ yield an order-isomorphism between the lattice of time warps and the lattice $\mathcal{M}(P)$ of monotone binary relations on P , ordered by inclusion. Let us denote by $*$ the associative operation of relational composition on $\mathcal{M}(P)$, that is, $R * S := \{(x, z) \in P^2 \mid xRy \text{ and } ySz \text{ for some } y \in P\}$. Note that the natural order relation, \leq_P , on P is a neutral element in $\mathcal{M}(P)$ for this composition operation $*$. Straightforward computations show that, for any $f, g \in W$,

$$R_{f \circ g} = R_f * R_g \quad \text{and} \quad R_{f^*} = \{(x, y) \in P^2 \mid f(x) \leq y - 1\} = \{(x, y) \in P^2 \mid y \not\leq f(x)\},$$

where we note that the latter can also be written as the complementary relation of the converse of the relation R_f . Hence, the lattice isomorphism between the lattice of time warps and $\mathcal{M}(P)$ extends to an isomorphism between involutive residuated lattices, by equipping the lattice $\mathcal{M}(P)$ with intersection, union, relational composition $*$, the unary operation $(-)^*$ given by $R^* := P^2 \setminus (R^{-1})$, and the neutral element \leq_P .

From the above considerations, since the isomorphisms are effective, it follows that questions about the theory of the structure \mathbf{W} may be effectively translated into questions about the structure $\mathcal{M}(P)$. Suppose that t is a term in the language of involutive residuated lattices that uses variables x_1, \dots, x_n . The above isomorphism allows us to translate the term t into a formula $T(k, k', X_1, \dots, X_n)$ of second-order logic over the structure $\langle P, \leq \rangle$, where the X_i are binary predicate symbols and k and k' are first-order variables.

REFERENCES

- [BFT16] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). <https://smtlib.cs.uiowa.edu/>, 2016.
- [BMSS12] Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. First steps in synthetic guarded domain theory: step-indexing in the topos of trees. *Log. Methods Comput. Sci.*, 8(4), 2012. doi:10.2168/LMCS-8(4:1)2012.
- [BT03] Kevin Blount and Constantine Tsinakis. The structure of residuated lattices. *Int. J. Algebr. Comput.*, 13(4):437–461, 2003. doi:10.1142/S0218196703001511.

- [CGMS22] Almudena Colacito, Nikolaos Galatos, George Metcalfe, and Simon Santachi. From distributive ℓ -monoids to ℓ -groups, and back again. *J. Algebra*, 601:129–148, 2022. doi:10.1016/j.jalgebra.2022.02.012.
- [CP96] Paul Caspi and Marc Pouzet. Synchronous Kahn networks. In Robert Harper and Richard L. Wexelblat, editors, *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming, ICFP 1996, Philadelphia, Pennsylvania, USA, May 24-26, 1996*, pages 226–238. ACM, 1996. doi:10.1145/232627.232651.
- [Dil89] David L. Dill. Timing assumptions and verification of finite-state concurrent systems. In Joseph Sifakis, editor, *Proc. Automatic Verification Methods for Finite State Systems, International Workshop, Grenoble, France, June 12-14, 1989*, volume 407 of *LNCS*, pages 197–212. Springer, 1989. doi:10.1007/3-540-52148-8_17.
- [dMB08] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Z3: an efficient SMT solver. In C. R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008. doi:10.1007/978-3-540-78800-3_24.
- [EGGHK18] Patrik Eklund, Javier Gutiérrez García, Ulrich Höhle, and Jari Kortelainen. *Semigroups in complete lattices. Quantales, modules and related topics*, volume 54 of *Dev. Math.* Cham: Springer, 2018. doi:10.1007/978-3-319-78948-4.
- [FC06] Jean-Christophe Filliâtre and Sylvain Conchon. Type-safe modular hash-consing. *Proc. 2006 workshop on ML*, pages 12–19, 2006. doi:10.1145/1159876.1159880.
- [FKM16] Soichiro Fujii, Shin-ya Katsumata, and Paul-André Melliès. Towards a formal theory of graded monads. In Bart Jacobs and Christof Löding, editors, *Foundations of Software Science and Computation Structures - 19th International Conference, FOSSACS 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016. Proceedings*, volume 9634 of *Lecture Notes in Computer Science*, pages 513–530. Springer, 2016. doi:10.1007/978-3-662-49630-5_30.
- [FL79] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *J. Comput. System Sci.*, 18(2):194–211, 1979. doi:10.1016/0022-0000(79)90046-1.
- [GJ17] Nikolaos Galatos and Peter Jipsen. Distributive residuated frames and generalized bunched implication algebras. *Algebra Universalis*, 78:303–336, 2017. doi:10.1007/s00012-017-0456-x.
- [GJ20] Nikolaos Galatos and Peter Jipsen. The structure of generalized BI-algebras and weakening relation algebras. *Algebra Universalis*, 81(3):35, 2020. doi:10.1007/s00012-020-00663-9.
- [GKO⁺16] Marco Gaboardi, Shin-ya Katsumata, Dominic Orchard, Flavien Breuvar, and Tarmo Uustalu. Combining effects and coeffects via grading. *ACM SIGPLAN Notices*, 51(9):476–489, 2016. doi:10.1145/2951913.2951939.
- [GP07a] Mai Gehrke and Hilary A. Priestley. Canonical extensions of double quasioperator algebras: An algebraic perspective on duality for certain algebras with binary operations. *J. Pure Appl. Algebra*, 209(1):269–290, 2007. doi:10.1016/j.jpaa.2006.06.001.
- [GP07b] Mai Gehrke and Hilary A. Priestley. Duality for double quasioperator algebras via their canonical extensions. *Studia Logica*, 86(1):31–68, 2007. doi:10.1007/s11225-007-9045-x.
- [GS14] Dan R. Ghica and Alex I. Smith. Bounded linear types in a resource semiring. In Zhong Shao, editor, *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014. Proceedings*, volume 8410 of *Lecture Notes in Computer Science*, pages 331–350. Springer, 2014. doi:10.1007/978-3-642-54833-8_18.
- [Gua18] Adrien Guatto. A generalized modality for recursion. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 482–491. ACM, 2018. doi:10.1145/3209108.3209148.
- [HM79] W. Charles Holland and Stephen H. McCleary. Solvability of the word problem in free lattice-ordered groups. *Houston J. Math.*, 5(1):99–105, 1979.
- [LDF⁺22] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. The OCaml system release 4.14, 2022. URL: <https://v2.ocaml.org/manual/>.

- [LG88] John M. Lucassen and David K. Gifford. Polymorphic effect systems. In Jeanne Ferrante and Peter Mager, editors, *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages, San Diego, California, USA, January 10-13, 1988*, pages 47–57. ACM Press, 1988. doi:10.1145/73560.73564.
- [LL66] Hans Läuchli and Joseph Leonard. On the elementary theory of linear order. *Fund. Math.*, 59:109–116, 1966. doi:10.4064/FM-59-1-109-116.
- [MPT23] George Metcalfe, Francesco Paoli, and Constantine Tsinakis. *Residuated Structures in Algebra and Logic*, volume 277 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2023.
- [Nak00] Hiroshi Nakano. A modality for recursion. In *15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000*, pages 255–266. IEEE Computer Society, 2000. doi:10.1109/LICS.2000.855774.
- [San20] Luigi Santocanale. The involutive quantaloid of completely distributive lattices. In Uli Fahrenberg, Peter Jipsen, and Michael Winter, editors, *Relational and Algebraic Methods in Computer Science - 18th International Conference, RAMiCS 2020, Palaiseau, France, April 8-11, 2020, Proceedings [postponed]*, volume 12062 of *Lecture Notes in Computer Science*, pages 286–301. Springer, 2020. doi:10.1007/978-3-030-43520-2_18.
- [vGGMS21] Sam van Gool, Adrien Guatto, George Metcalfe, and Simon Santschi. Time warps, from algebra to algorithms. In Uli Fahrenberg, Mai Gehrke, Luigi Santocanale, and Michael Winter, editors, *Relational and Algebraic Methods in Computer Science - 19th International Conference, RAMiCS 2021, Marseille, France, November 2-5, 2021, Proceedings*, volume 13027 of *Lecture Notes in Computer Science*, pages 309–324. Springer, 2021. doi:10.1007/978-3-030-88701-8_19.
- [Yet90] David N. Yetter. Quantales and (noncommutative) linear logic. *J. Symbolic Logic*, 55(1):41–64, 1990. doi:10.2307/2274953.