Complexity of monodic guarded fragments over linear and real time

Ian Hodkinson [1]

Department of Computing, Imperial College London, London SW7 2AZ, UK

1 Introduction

Propositional temporal logic is now well established, both in its theory and its utility for practical applications in computing. In contrast, predicate, or first-order, temporal logic has been less studied. Unpublished results of Lindström and Scott in the 1960s showed that even weak fragments of first-order temporal logic are highly undecidable, and these and later similar results (see, e.g., [6,15]) may have suggested that other areas were more profitable to work on.

Recently, however, some decidable fragments of first-order temporal logic have been found. The so-called *monodic fragments,* originating in [15], have now been quite extensively investigated. In these fragments, formulas beginning with a temporal operator are required to have at most one free variable. Also, the 'first-order part' of formulas must lie in some decidable fragment of first-order logic with very mild closure properties. Suitable fragments include the monadic fragment (with only unary relation symbols), the one- and two-variable fragments, and the guarded, loosely guarded, and packed fragments. [15] showed that the monodic fragments based on all these first-order fragments are decidable over a wide range of linear flows of time. (We use 'flow of time' synonymously with 'strict partial order'.) A notable case left open in [15] is for real

Email address: `imh@doc.ic.ac.uk` (Ian Hodkinson).
URL: `www.doc.ic.ac.uk/~imh` (Ian Hodkinson).

[1] Research partially supported by U.K. EPSRC grant GR/S19905/01. I thank Mark Reynolds and the staff of the School of Information Technology, Murdoch University, Western Australia, for helpful conversations and kind hospitality during my visit in winter 2003, when much of this paper was written. I am grateful to Rod Downey and Robert Goldblatt for inviting me to submit this paper, and to the referee for helpful comments.

Preprint submitted to Annals of Pure and Applied Logic *17 October 2004*

numbers flow of time: while it was shown that decidability does hold for structures with finite first-order domains, without this restriction it remains an open problem whether any non-trivial monodic fragment is decidable over $(\mathbb{R}, <)$. Results for branching time have also been established — some positive, some negative — and axiomatisations and resolution and tableau procedures have been developed and implemented. For some of this work see [2,4,5,7,12,17,19,28]. The expressive power of monodic fragments, though obviously less than that of full first-order temporal logic, is still sufficient to encode many other logics of interest; see [16] for some examples and applications.

Even more recently, the examination of the computational complexity of monodic fragments has begun. In [13], it was shown that the one-variable fragment of linear first-order temporal logic, even with sole temporal operator $\Box$ (standing for 'always in the future'), is EXPSPACE-hard over natural numbers time; consequently, so are the one-variable, two-variable and monadic monodic fragments with the temporal connectives Until and Since. This EXPSPACE lower bound actually applies to any class of linear flows of time containing a flow that embeds the natural numbers, and it holds both for structures with arbitrary first-order domains and for structures where the domain is finite.

Some upper bounds were established in [13], again both for arbitrary and for finite first-order domains, but only for flow of time the natural numbers. Roughly speaking, it was shown that the complexity of a monodic fragment over natural numbers time is at most the maximum of EXPSPACE and the complexity of the underlying first-order fragment. In particular, it follows that over natural numbers time, the monodic packed fragment (a generalised guarded fragment) has the same complexity as its pure first-order part — 2EXPTIME-complete. The proof is similar to the one in [27] giving a PSPACE algorithm to decide propositional temporal logic over natural numbers time.

For many, especially in computing, the only linear flow of time of importance is the natural numbers, and [13] will be the start and end of the story. But other linear flows of time have received attention, for example in modal logic and philosophy — see, for example, Kamp's expressive completeness results [18] and Burgess and Gurevich's decision procedures [3]. Recently, Reynolds established that the satisfiability problem for propositional temporal logic with Until and Since over arbitrary linear time and over the real numbers is PSPACE-complete [25,24]. The proofs often require powerful methods and cast new light on the inner workings of the logics. Proving results for the real numbers is especially useful, as results for other flows of time can often be derived as corollaries. It therefore seems (to us) natural to investigate the complexity of monodic fragments over linear flows of time other than the natural numbers.

Here, we contribute to this investigation by showing that the monodic packed fragment with Until and Since is 2EXPTIME-complete over quite a wide range of linear flows of time. For structures with arbitrary first-order domains, we prove in theorem 4.8 that satisfiability for monodic packed sentences is 2EXPTIME-complete over the class of all linear flows of time, dense flows, discrete flows, the rationals, and some others. With finite first-order domains, we can do better: in theorem 5.14, we show that the monodic packed fragment is 2EXPTIME-complete over the real numbers, and hence over essentially all commonly-used linear flows of time — for example, the natural numbers (known from [13]), integers, rationals, all linear flows, all dense flows, all discrete flows, and indeed any first-order definable class of linear flows. The same results hold for the monodic guarded and loosely guarded fragments: see remark 6.1.

We concentrate on the monodic packed fragment here for three reasons. First, it is a rather attractive monodic fragment, because equality can be included; for other monodic fragments, adding equality can result in loss of decidability [5]. To include equality is desirable in any case, but its presence also aids the handling of constants, and so makes the proofs simpler. Second, it is a fragment for which the complexity of our algorithm is theoretically optimal, because the first-order packed fragment is already 2EXPTIME-complete. Third, by the same token, our results show that the monodic packed fragment, a full-blown first-order temporal logic, is computationally no more expensive at all to use than the non-temporal first-order packed fragment. We find this very striking.

Our proofs start off by applying the 'quasimodel' technique of [15], but in the main they are similar to, and parts of them are actually borrowed from, the recent mosaic-based work of Reynolds on complexity of propositional temporal logic with Until and Since over linear and real time [25,24]. Reynolds established PSPACE-completeness by a sophisticated argument which we have not yet been able to generalise to the monodic case; our proof for arbitrary domains follows the simpler proof of EXPTIME upper bounds given early in [25]. The proof for finite domains, over the real numbers, extends the argument in [15, §7] and involves ideas of [3,20] — in particular, the second, 'model-theoretic' proof in [3] of decidability of propositional temporal logic with Until and Since over the real numbers.

We would like to stress that because the methods in this paper only give 2EXPTIME upper bounds, they are not entirely satisfactory. 2EXPTIME is a very high complexity, and while the monodic packed fragment is 2EXPTIME-complete, for many monodic fragments only EXPSPACE lower complexity bounds are known. Examples include the monodic 2-variable fragment, the monadic monodic fragment, and monodic guarded and packed fragments with bounded number of variables. Generalising Reynolds's more powerful arguments to monodic fragments might be a way to obtain EXPSPACE upper bounds; this would certainly improve the results presented here.

Outline of paper In section 2, we define the monodic packed fragment formally. Section 3 establishes some results on quasimodels, mosaics, and bags of mosaics that will be needed in both the main proofs. Section 4 contains the proof for arbitrary domains and section 5 the proof for finite domains over the real numbers. We conclude in section 6 with some remarks about possible extensions to our results.

Notation and conventions A *linear order,* or *linear flow of time,* is a pair $(I, <)$, where I is a non-empty set and $<$ an irreflexive transitive relation on I satisfying trichotomy: $\forall xy(x < y \vee x = y \vee y < x)$. We write $x \leq y$ to abbreviate $x < y \vee x = y$, and $x > y$ means $y < x$. We usually identify (notationally) a linear order or other classical first-order structure with its domain. We use standard notation for intervals of a linear order I: for $x, y \in I$, we let $(x, y) = \{z \in I : x < z < y\}$, $[x, y] = \{z \in I : x \leq z \leq y\}$, and define $[x, y)$ and $(x, y]$ similarly. We often let such an interval $[x, y]$ also denote the induced linear order $([x, y], < \cap ([x, y] \times [x, y]))$. A linear flow is *dense* if it has at least two elements[2] and satisfies $\forall xy(x < y \rightarrow \exists z(x < z < y))$.

For a structure M and a constant c of its signature, c^M denotes the interpretation of c in M. For a set I and sets S_i $(i \in I)$, $\prod_{i \in I} S_i$ denotes as usual the set of all maps $a : I \rightarrow \bigcup_{i \in I} S_i$ such that $a(i) \in S_i$ for each $i \in I$. When I is finite, we sometimes write maps in $\prod_{i \in I} S_i$ as the sequence of their values. We write $\bar{\mathsf{x}}, \bar{a}$, etc., for *tuples* of variables, elements, etc. For an equivalence relation $\sim$ on a set S, we write $S/{\sim}$ for the set of $\sim$-classes, and $s/{\sim}$ for the $\sim$-class of an element $s \in S$.

We use the term *'mirror image'* in the usual way in temporal logic: the mirror image of a condition is obtained by swapping past and future notions within it. For example, we would exchange $<$ with $>$, Until with Since, left with right, initial points with endpoints, start with end, etc.

2 Monodic packed fragment

The *guarded fragment* was introduced by Andréka, van Benthem, and Németi in [1] as a fragment of first-order logic with the 'nice' properties of modal logic (in particular, decidability with reasonable complexity). All quantifiers in guarded formulas are relativised to an atomic formula (the 'guard'). The *packed fragment,* introduced by Marx in [21], extends the guarded fragment by allowing weaker guards, and has the same complexity (2EXPTIME-complete), so we will use the packed fragment here. The *clique-guarded fragment* defined in [8] is a syntactic variant of it.

The following combines the definition of *monodic* formulas from [15] with a minor

[2] To avoid considering degenerate cases, in this paper we take 'dense' to imply that there are at least two elements.

modification of Marx's definition of the packed fragment from [21].

Definition 2.1 (monodic packed fragment) Let L be a signature with at most constants and relation symbols (no function symbols). An L-formula γ is said to be a *packing guard* if γ is a conjunction of atomic and existentially-quantified atomic formulas (i.e., of the form α or $\exists \bar{\mathsf{x}}\alpha$, where α is atomic, possibly an equality) such that for any two distinct free variables of γ, there is a conjunct of γ in which they both occur free. The *monodic packed fragment* of predicate temporal logic in signature L consists of the following formulas:

- Any atomic L-formula (which can be an equality, $\top$, or $\bot$) is monodic packed.
- Boolean combinations of monodic packed formulas are monodic packed.
- If x is a variable and φ, ψ are monodic packed formulas with free variable at most x, then $\mathsf{U}(\varphi, \psi)$ and $\mathsf{S}(\varphi, \psi)$ (*Until* and *Since,* respectively) are monodic packed formulas.
- If γ is a packing guard, φ is a monodic packed formula, every free variable of φ is free in γ, and $\bar{\mathsf{x}}$ is a tuple of variables, then $\exists \bar{\mathsf{x}}(\gamma \wedge \varphi)$ is a monodic packed formula.

Any propositional temporal formula is monodic packed, since we allow nullary relation symbols in L. So is any packed first-order formula — the first-order packed fragment is contained in the monodic packed fragment.

Definition 2.2 For a formula φ, we let $|\varphi|$ denote the length of (i.e., the number of symbols in) φ.

Our complexity results will be stated in terms of $|\varphi|$. For formal complexity purposes, algorithms must encode φ, subformulas of φ, etc., using symbols from a fixed finite alphabet. This has the consequence that relation symbols, constants, and variables are not stored as single symbols but as strings, such as R_0, c_{16}, x_{141}. So storing φ takes space $\mathcal{O}(|\varphi| \log |\varphi|)$. This increase is at most quadratic and is never a problem in practice, but it should be borne in mind.

Definition 2.3 Let L be as above and let $T = (T, <)$ be a linear flow of time. A *temporal L-structure with flow of time T* is a triple $\mathcal{M} = (T, D, (M_t : t \in T))$, where D is a non-empty set (the *domain* of $\mathcal{M}$), each M_t is an L-structure with domain D, and $c^{M_t} = c^{M_u}$ for all constants $c \in L$ and all $t, u \in T$.

Given $\mathcal{M}$, $t \in T$, and an assignment h of variables into D, we define the semantics of monodic packed formulas written with variables from the domain of h as follows:

- For atomic α, we let $\mathcal{M}, t, h \models \alpha$ iff $M_t, h \models \alpha$.
- $\mathcal{M}, t, h \models \neg\varphi$ iff $\mathcal{M}, t, h \not\models \varphi$, and $\mathcal{M}, t, h \models \varphi \wedge \psi$ iff $\mathcal{M}, t, h \models \varphi$ and $\mathcal{M}, t, h \models \psi$.

- $\mathcal{M}, t, h \models \exists \bar{\mathsf{x}} \varphi$ iff $\mathcal{M}, t, g \models \varphi$ for some assignment g that agrees with h except perhaps on the variables in $\bar{\mathsf{x}}$.
- $\mathcal{M}, t, h \models \mathsf{U}(\varphi, \psi)$ iff there is $u \in T$ with $u > t$, $\mathcal{M}, u, h \models \varphi$, and $\mathcal{M}, v, h \models \psi$ for all $v \in T$ with $t < v < u$.
- $\mathcal{M}, t, h \models \mathsf{S}(\varphi, \psi)$ iff there is $u \in T$ with $u < t$, $\mathcal{M}, u, h \models \varphi$, and $\mathcal{M}, v, h \models \psi$ for all $v \in T$ with $u < v < t$ (i.e., the mirror image of the clause for Until).

Occasionally, for a formula φ with free variable x, if $h(\mathsf{x}) = a$, we write $\mathcal{M}, t \models \varphi(a)$ instead of $\mathcal{M}, t, h \models \varphi$. For a sentence φ, we drop the assignment h and just write $\mathcal{M}, t \models \varphi$.

Note that our temporal structures have constant domains, and the interpretations of constants and assignments to variables are rigid. Also, we use the 'strict' semantics of Until (U) and Since (S) ($u > t$ rather than $u \geq t$, etc), this being more suitable for the flows of time considered here. Strict Until and Since can express connectives such as the non-strict and weak Until and Since, Tomorrow, etc. Indeed, [15] proves expressive completeness results for monodic formulas written with them, over Dedekind complete linear flows of time.

Abbreviations $\Diamond\psi$ abbreviates $\psi \vee \mathsf{U}(\psi, \top) \vee \mathsf{S}(\psi, \top)$, and $\Box\psi$ abbreviates $\neg\Diamond\neg\psi$. The abbreviations $\vee, \rightarrow, \forall$ are defined as usual.

Definition 2.4 A monodic packed L-sentence φ is said to be *satisfiable* in a temporal L-structure $\mathcal{M}$ with flow of time T if there is $t \in T$ with $\mathcal{M}, t \models \varphi$. In this case, we say that φ is satisfiable in $\mathcal{M}$, and that $\mathcal{M}$ is a model of φ.

We will be interested in satisfiability in temporal structures with arbitrary first-order domains, as in the definition, and also in temporal structures whose first-order domains are finite. These problems are different. The domain of any temporal structure in which the monodic packed sentence $\Box\exists\mathsf{x}(P(\mathsf{x}) \wedge \neg\mathsf{S}(P(\mathsf{x}), \top))$ (given in [15]) is satisfiable is at least as large as its flow of time. As we are mainly interested in infinite flows of time, this shows that satisfiability for arbitrary domains does not imply satisfiability for finite domains.

The satisfiability problem for the packed fragment of first-order logic is 2EXPTIME-complete [8,9,21]. Moreover, the packed fragment has the *finite model property* — any satisfiable packed sentence has a finite model [11,14]. So the satisfiability problem for the packed fragment over finite structures is the same as the satisfiability problem over arbitrary structures, and is also 2EXPTIME-complete. Since the first-order packed fragment is a subfragment of the monodic packed fragment, we deduce:

Proposition 2.5 *Any satisfiability problem for the monodic packed fragment, whether* *over arbitrary temporal structures or over temporal structures with finite domains, is* 2EXPTIME-*hard.*

3 Preliminaries

For the rest of the paper, we fix a packed monodic sentence φ with no function symbols.

There are two complexity proofs in this paper; this section develops some material needed in both of them. The core idea in both proofs is to condense an arbitrary temporal structure $\mathcal{M}$ into some finite object, preserving enough information to tell whether φ was satisfiable in $\mathcal{M}$. We do this using *quasimodels* and *mosaics*. Here, we recall the technique of quasimodels from [15]; a quasimodel is the result of condensing (essentially by filtration) the first-order part of a temporal structure but leaving the flow of time intact. Then we introduce the mosaics (similar to those in [25,24]) and the 'bags of mosaics' that we will use to collapse the temporal part of a quasimodel into a finite object and so complete the condensation process.

3.1 Types and state candidates

We let L be the finite signature consisting of the relation symbols and constants that occur in φ. We also fix a variable x not occurring in φ.

Definition 3.1

(1) Define $\mathrm{sub}_{\mathsf{x}}\varphi$ to be the finite set

$$\{\psi(\mathsf{x}/\mathsf{y}),\ \neg\psi(\mathsf{x}/\mathsf{y}),\ \mathsf{x}=c,\ \mathsf{x}\neq c\ :\ \psi(\mathsf{y})\text{ a subformula of }\varphi,\ c\text{ a constant in }L\}.$$

Here, $\psi(\mathsf{y})$ denotes that ψ is a sentence or has a single free variable, y, and $\psi(\mathsf{x}/\mathsf{y})$ denotes the result of substituting x for all free occurrences of y in ψ.

(2) A *type for* φ is a maximal boolean consistent subset $p \subseteq \mathrm{sub}_{\mathsf{x}}\varphi$. That is:
- for all $\neg\psi \in \mathrm{sub}_{\mathsf{x}}\varphi$, we have $\neg\psi \in p \iff \psi \notin p$,
- for all $\psi \wedge \chi \in \mathrm{sub}_{\mathsf{x}}\varphi$, we have $\psi \wedge \chi \in p \iff \psi \in p$ and $\chi \in p$.

(3) We let $\mathcal{T}(\varphi)$ denote the (finite) set of types for φ.

For example, let $\mathcal{M} = ((I,<), D, (M_t : t \in I))$ be a temporal structure, and for $t \in I$ and $a \in D$ write $\mathrm{tp}_t(a) = \{\psi \in \mathrm{sub}_{\mathsf{x}}\varphi : \mathcal{M}, t \models \psi(a)\}$. This is a type for φ.

Definition 3.2

(1) For each subformula ψ of φ of the form $\mathsf{U}(\alpha,\beta)$ or $\mathsf{S}(\alpha,\beta)$ with one (respectively zero) free variable(s), we introduce a new unary (respectively, nullary) relation symbol R_ψ. The *surrogate* of ψ is $R_\psi(\mathsf{y})$ if ψ has free variable y, and R_ψ if ψ is a sentence.
(2) For any subformula ψ of φ, the formula $\overline{\psi}$ is obtained by replacing all maximal subformulas $\mathsf{U}(\alpha,\beta)$ and $\mathsf{S}(\alpha,\beta)$ of ψ by their surrogates. For $p \subseteq \mathrm{sub}_{\mathsf{x}}\varphi$, we write $\overline{p}$ for $\{\overline{\psi} : \psi \in p\}$.

Definition 3.3 A *state candidate (for φ)* is a non-empty set of types for φ. A *realisable state candidate* (respectively, *finitely realisable state candidate*) is a state candidate Σ such that for some (respectively, finite) structure M we have

$$M \models \underbrace{\bigwedge_{p\in\Sigma} \exists \mathsf{x} \bigwedge \overline{p} \ \wedge \ \forall \mathsf{x} \bigvee_{p\in\Sigma} \bigwedge \overline{p}}_{\alpha_\Sigma} .$$

Observe that the sentence α_Σ here is (up to logical equivalence) in the packed fragment, since we can guard the $\exists\mathsf{x}$ and $\forall\mathsf{x}$ by $\mathsf{x}=\mathsf{x}$. This has some consequences:

Lemma 3.4 *A state candidate is realisable iff it is finitely realisable.*

Proof. By [11,14], the packed fragment has the finite model property. So Σ is a realisable state candidate iff α_Σ has a model, iff it has a finite model, iff Σ is a finitely realisable state candidate. □

In spite of this lemma, we still wish to preserve the distinction between the two notions, because our work may be applicable to other logics that do not have the finite model property.

We will need to know how hard it is to decide whether a given set Σ of types is a realisable state candidate. If we simply apply a decision procedure for the packed fragment to α_Σ, it could take treble exponential time in $|\varphi|$, since the length of α_Σ is potentially exponential in $|\varphi|$. Fortunately, we can do better:

Proposition 3.5 *Let Σ be a set of types for φ. It is decidable in* 2EXPTIME *in $|\varphi|$ (i.e., in time at most $2^{2^{f(|\varphi|)}}$, for some fixed polynomial $f(n)$) whether Σ is a realisable state candidate.*

Proof (sketch). The result can be seen quite easily by inspection of the proofs in [9,8] that the satisfiability problem for the clique-guarded fragment (even with fixed point operators) is in 2EXPTIME. We will outline a similar argument.

Let us first review how satisfiability of a first-order packed sentence σ (with no function symbols) is decided. One way to do it is by considering 'pieces' of a potential model of σ. We formalise this as follows. All formulas below are implicitly put in negation-normal form, with all negations pushed next to atomic formulas and double negations removed. So we must take $\wedge, \vee, \neg, \forall$, and $\exists$ as primitive symbols. Assume that σ is written with k variables. Introduce new constants $w_1, \ldots, w_k$, called *witnesses.* Let $\mathcal{W}$ be the set consisting of these witnesses together with the constants in σ. Below, c, c', etc., will range over $\mathcal{W}$; $\bar{c}$ denotes a tuple of such c. Let $\mathcal{S}(\sigma)$ be the set of subformulas of σ and their negations, together with $\mathsf{x} = \mathsf{y}$ and $\mathsf{x} \neq \mathsf{y}$ (where x and y are distinct variables). Let $\mathcal{C}(\sigma)$ be the set of all sentences of the form $\psi(\bar{c})$, where $\psi(\bar{\mathsf{x}}) \in \mathcal{S}(\sigma)$. A *condition* is a maximal boolean consistent subset C of $\mathcal{C}(\sigma)$, containing all sentences of the form $c = c$, such that if $\psi(\bar{c}, c_1) \in C$ and $c_1 = c_2 \in C$ then $\psi(\bar{c}, c_2) \in C$, and if $\forall\bar{\mathsf{x}}\psi(\bar{x}) \in C$ then $\psi(\bar{c}) \in C$ for all $\bar{c}$. A condition can be thought of as describing the 'piece' of a model defined by the interpretations of the constants c.

The notion that we have enough pieces to form a model of σ is formalised by a *saturated set $\mathcal{N}$ of conditions:* one such that

(1) there is $C \in \mathcal{N}$ with $\sigma \in C$,
(2) for each $C \in \mathcal{N}$ and each sentence $\exists\bar{\mathsf{x}}\psi(\bar{\mathsf{x}}, \bar{c}) \in C$, where $\exists\bar{\mathsf{x}}\psi(\bar{\mathsf{x}}, \bar{\mathsf{y}}) \in \mathcal{S}(\sigma)$, there is $D \in \mathcal{N}$ such that (i) $\psi(\bar{c}', \bar{c}) \in D$ for some $\bar{c}'$, and (ii) $\chi \in C \iff \chi \in D$ for each $\chi \in \mathcal{C}(\sigma)$ that only involves witnesses occurring in $\bar{c}$.

If σ has a model, say M, let $\mathcal{E}$ be the set of all expansions of M to interpret the witnesses. Then for each $M^+ \in \mathcal{E}$, the set $C(M^+) = \{\psi \in \mathcal{C}(\sigma) : M^+ \models \psi\}$ is a condition, and $\mathcal{N} = \{C(M^+) : M^+ \in \mathcal{E}\}$ is a saturated set of conditions. Conversely, a model M of σ can be constructed from a saturated set of conditions in a step-by-step fashion (cf. [9, theorem 3.7]). M satisfies each condition in $\mathcal{N}$ if the witnesses are interpreted appropriately, and every element of M is named by some $c \in \mathcal{W}$ under some such interpretation. So σ has a model iff there exists a saturated set of conditions. But the existence of such a set can be decided in 2EXPTIME in $|\sigma|$, by a method of Pratt [22]. The algorithm initialises $\mathcal{N}$ to be the set of *all* conditions; the size of this set is at most doubly exponential. Then it loops, in each iteration deleting all conditions C that fail property 2 above. On termination, the algorithm states that a saturated set exists iff some C in the final $\mathcal{N}$ contains σ. (This is easily seen to be correct; for more details, see theorem 4.7. If k is fixed in advance, the procedure only takes EXPTIME.)

A simple modification of this decides realisability of a given set Σ of types for φ. The set $\mathcal{S}(\sigma)$ above is replaced by the set of all formulas $\overline{\psi}$ and their negations, where ψ is a subformula of φ, together with $\mathsf{x} = \mathsf{y}$ and $\mathsf{x} \neq \mathsf{y}$. Conditions are then defined as before. Notice that a condition C determines the type $\mathrm{tp}_C(c)$ of any $c \in \mathcal{W}$, via $\mathrm{tp}_C(c) = \{\psi(\mathsf{x}) \in \mathrm{sub}_{\mathsf{x}}\varphi : \overline{\psi}(c) \in C\}$. We can then show that Σ is realisable iff there is a non-empty set $\mathcal{N}$ of conditions satisfying property 2 above, and such that

(3) all conditions in $\mathcal{N}$ agree on all L-sentences,
(4) $\Sigma = \{\mathrm{tp}_C(c) : C \in \mathcal{N},\ c \in \mathcal{W}\}$.

These requirements ensure that all and only the types in Σ are realised in the model built 'step by step' from $\mathcal{N}$. The existence of such an $\mathcal{N}$ can be decided much as before. Requirement (3) is handled by adding an outer loop running over all sets S of L-sentences that can occur in conditions; the inner loop initialises $\mathcal{N}$ to the set of all conditions containing S, and proceeds as before. The whole process takes in 2Exptime in $|\varphi|$ (and only Exptime if k is fixed in advance). □

3.2 Runs and (pre-)quasimodels

Definition 3.6 Let $I = (I, <)$ be a linear order. A map $r : I \to \mathcal{T}(\varphi)$ is said to be a *run (over I)* if the following three conditions hold:

(1) for each constant $c \in L$ and each $t, u \in I$, we have $\mathsf{x} = c \in r(t)$ iff $\mathsf{x} = c \in r(u)$,
(2) for each $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$ and each $t \in I$, we have $\mathsf{U}(\alpha, \beta) \in r(t)$ iff:
  (a) there is $u \in I$ with $t < u$, $\alpha \in r(u)$, and $\beta \in r(v)$ for all $v \in I$ with $t < v < u$, or
  (b) I contains a maximal element y (say), $\mathsf{U}(\alpha, \beta) \in r(y)$, and $\beta \in r(v)$ for all $v \in I$ with $v > t$,
(3) the mirror image condition for Since.

We say that r is a *full run*[3] if it satisfies condition 1 and the (stronger) forms of conditions 2 and 3 in which part (b) is deleted.

Continuing the example in section 3.1, for any $a \in D$ the map $r_a : t \mapsto \mathrm{tp}_t(a)$ is a full run over I, and its restriction to any closed interval J of I is a run over J. The reader should check that any full run is a run, so that deleting part (b) does indeed strengthen conditions 2 and 3.

The following definition is crucial to the paper.

[3] In [12,15], 'full runs' were simply called 'runs'.

Definition 3.7 Let I be a linear order. A *pre-quasimodel (for φ over I)* is a triple

$$\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R}),$$

where each Σ_t is a realisable state candidate and $\mathcal{R} \subseteq \prod_{t \in I} \Sigma_t$ is a set of runs such that $\Sigma_t = \{r(t) : r \in \mathcal{R}\}$ for each $t \in I$. $\mathcal{Q}$ is said to be a *quasimodel (for φ)* if each $r \in \mathcal{R}$ is a full run, and $\varphi \in r(t)$ for some $r \in \mathcal{R}$ and $t \in I$. $\mathcal{Q}$ is said to be *finitary* if each Σ_t ($t \in I$) is a finitely realisable state candidate, and $\mathcal{R}$ is finite.

Definition 3.8 If $\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R})$ is a pre-quasimodel for φ, and $x \leq y$ in I, write $\mathcal{Q} \restriction [x, y]$ for

$$\big([x, y], (\Sigma_t : t \in [x, y]), \{r \restriction [x, y] : r \in \mathcal{R}\}\big).$$

(Recall that we identify the interval $[x, y]$ with the induced suborder on it.)

Lemma 3.9 *$\mathcal{Q} \restriction [x, y]$ is a pre-quasimodel. If $\mathcal{Q}$ is finitary then so is $\mathcal{Q} \restriction [x, y]$.*

Proof. Straightforward. □

The following result from [12] will be the starting point for our complexity results.

Fact 3.10 ([12, theorem 3]) *Let I be any linear order. Then φ has a model with flow of time I (and finite domain) iff there is a (respectively, finitary) quasimodel for φ over I.*

The proof of '$\Rightarrow$' is straightforward: each element a of the domain D of a model $\mathcal{M}$ of φ with flow of time T gives rise to a full run r_a given by $r_a(t) = \{\psi \in \mathrm{sub}_{\mathsf{x}}\varphi : \mathcal{M}, t \models \psi(a)\}$, and we let $\Sigma_t = \{r_a(t) : a \in D\}$. To prove the converse, given a quasimodel $\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R})$ of φ, we first take a model $M_t \models \alpha_{\Sigma_t}$ for each $t \in T$. We make a suitable number of copies of runs in $\mathcal{R}$, resulting in a multiset D, say; this is arranged to be finite if $\mathcal{Q}$ is finitary. Then, in a 'model-theoretic' step which can be carried through for the packed fragment even though equality is present, we manipulate the M_t so that they all have domain D, and for all $r \in D$, $t \in I$, and $\psi \in \mathrm{sub}_{\mathsf{x}}\varphi$, we have $M_t \models \overline{\psi}(r)$ iff $\psi \in r(t)$. We then form the temporal structure $\mathcal{M} = (I, D, (M_t : t \in I))$, and show by induction on subformulas ψ of φ that for any assignment h of the variables into D, and any $t \in I$, $\mathcal{M}, t, h \models \psi$ iff $M_t, h \models \overline{\psi}$. It follows that φ is satisfiable in $\mathcal{M}$.

3.3 Mosaics

We will use 'mosaics' to describe runs. The following definition is based on one in [24]. For convenience, we have retained only the conditions we need here, and we have dualised the notion of 'cover' from [24] by couching it in terms of satisfiability rather than validity.

Definition 3.11 A *mosaic* (for φ) is a triple (A, B, C), where A and C are types for φ, $B \subseteq \mathrm{sub}_{\mathsf{x}}\varphi$, and:

(1) For each formula $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$ with $\neg\beta \notin B$, we have

$$\mathsf{U}(\alpha, \beta) \in A \iff (\alpha \in B \cup C \text{ or } \beta, \mathsf{U}(\alpha, \beta) \in C).$$

Note that all formulas mentioned are in $\mathrm{sub}_{\mathsf{x}}\varphi$.

(2) A mirror image condition for Since.

(3) For each constant $c \in L$, the conditions $\mathsf{x} = c \in A$, $\mathsf{x} = c \in B$, $\mathsf{x} \neq c \notin B$, and $\mathsf{x} = c \in C$ are equivalent.

For a mosaic $m = (A, B, C)$ we write $\mathrm{st}(m) = A$, $\mathrm{end}(m) = C$ and $\mathrm{cov}(m) = B$ (these stand for start, end, cover).

The motivation for this definition is that runs can be represented by mosaics.

Definition 3.12 Let I be a linear order with endpoints $x < y$ and with at least 3 points, and let $r : I \to \mathcal{T}(\varphi)$. Define

$$\mathrm{mos}(r) = \Big(r(x),\ \bigcup_{z \in (x,y)} r(z),\ r(y)\Big).$$

Lemma 3.13 *If r is a run over I, then $m = \mathrm{mos}(r)$ is a mosaic.*

Proof. Let $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$, and assume that $\neg\beta \notin \mathrm{cov}(m) = \bigcup_{z \in (x,y)} r(z)$. As the $r(z)$ are types, β or $\neg\beta$ is in each of them, so $\beta \in \bigcap_{z \in (x,y)} r(z)$. Hence, as r is a run, we have $\mathsf{U}(\alpha, \beta) \in r(x)$ iff there is $z \in I$ with $z > x$ and $\alpha \in r(z)$, or $\beta, \mathsf{U}(\alpha, \beta) \in r(y)$. This is iff $\alpha \in \mathrm{cov}(m) \cup \mathrm{end}(m)$ or $\beta, \mathsf{U}(\alpha, \beta) \in \mathrm{end}(m)$. The checks for Since are similar. The checks for constants are easy and left to the reader. □

There is a partial converse to this:

Lemma 3.14 *If $r : I \to \mathcal{T}(\varphi)$ and $\mathrm{mos}(r)$ is a mosaic then r satisfies condition 1 of definition 3.6.*

Proof. Let c be a constant in L, and let $x, y \in I$. Write m for $\mathrm{mos}(r)$. If $\mathsf{x} = c \in r(x)$, then $\mathsf{x} = c \in \mathrm{st}(m) \cup \mathrm{cov}(m) \cup \mathrm{end}(m)$. Now m is a mosaic, so this implies that $\mathsf{x} \neq c \notin \mathrm{st}(m) \cup \mathrm{cov}(m) \cup \mathrm{end}(m) \supseteq r(y)$; and this implies that $\mathsf{x} = c \in r(y)$. □

Just as mosaics will be used to describe runs, so sets of mosaics, or 'bags', will be used to describe pre-quasimodels, which are at heart sets of runs. As a bag is a finite object, this will complete the 'condensation' process discussed at the start of the section.

Definition 3.15 A *bag* (respectively, *finitary bag)* (for φ) is a set μ of mosaics such that the sets $\mathrm{st}(\mu) = \{\mathrm{st}(m) : m \in \mu\}$ and $\mathrm{end}(\mu) = \{\mathrm{end}(m) : m \in \mu\}$ are (respectively, finitely) realisable state candidates.

Definition 3.16 Let I be a linear order with endpoints $x < y$ and with at least 3 points. Let $\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R})$ be a pre-quasimodel over I. Define $\mathrm{bag}(\mathcal{Q}) = \{\mathrm{mos}(r) : r \in \mathcal{R}\}$.

Lemma 3.17 *For $I, \mathcal{Q}$ as above,* $\mathrm{bag}(\mathcal{Q})$ *is a bag. If $\mathcal{Q}$ is finitary then* $\mathrm{bag}(\mathcal{Q})$ *is a finitary bag.*

Proof. Trivial, by lemma 3.13 and since $\mathrm{st}(\mathrm{bag}(\mathcal{Q})) = \{\mathrm{st}(m) : m \in \mathrm{bag}(\mathcal{Q})\} = \{\mathrm{st}(\mathrm{mos}(r)) : r \in \mathcal{R}\} = \{r(x) : r \in \mathcal{R}\} = \Sigma_x$, a (finitely) realisable state candidate, and similarly for end, with y. $\square$

We now show how to tell from its bag whether a pre-quasimodel is a quasimodel.

Definition 3.18 A bag μ is said to be *perfect* if

(1) there is no formula $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$ with $\mathsf{U}(\alpha, \beta) \in \bigcup \mathrm{end}(\mu)$,
(2) there is no formula $\mathsf{S}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$ with $\mathsf{S}(\alpha, \beta) \in \bigcup \mathrm{st}(\mu)$,
(3) $\varphi \in \bigcup_{m \in \mu} \Big(\mathrm{st}(m) \cup \mathrm{cov}(m) \cup \mathrm{end}(m) \Big)$.

Lemma 3.19 *Let $\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R})$ be a pre-quasimodel for φ, where I has endpoints and at least three points. Then* $\mathrm{bag}(\mathcal{Q})$ *is perfect iff $\mathcal{Q}$ is a quasimodel for φ.*

Proof. Suppose that the endpoints of I are x, y with $x < y$. We have $\mathrm{end}(\mathrm{bag}(\mathcal{Q})) = \{r(y) : r \in \mathcal{R}\}$. Assume first that $\mathrm{bag}(\mathcal{Q})$ is perfect. Let $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$. So $\mathsf{U}(\alpha, \beta) \notin r(y)$ for all $r \in \mathcal{R}$. As $\mathcal{Q}$ is a pre-quasimodel, if $r \in \mathcal{R}$, $t \in I$ and $\mathsf{U}(\alpha, \beta) \in r(t)$, there must be $u > t$ in I with $\beta \in r(u)$ and $\alpha \in r(v)$ for all $v \in (t, u)$. The case of Since is similar. This shows that each $r \in \mathcal{R}$ is a full run. Now take $m \in \mathrm{bag}(\mathcal{Q})$ with $\varphi \in \mathrm{st}(m) \cup \mathrm{cov}(m) \cup \mathrm{end}(m)$, and $r \in \mathcal{R}$ with $\mathrm{mos}(r) = m$. So $\varphi \in r(t)$ for some $t \in I$. So $\mathcal{Q}$ is a quasimodel for φ.

Conversely, assume that $\mathcal{Q}$ is a quasimodel for φ, and suppose for contradiction that there is some $\mathsf{U}(\alpha, \beta) \in \bigcup \mathrm{end}(\mathrm{bag}(\mathcal{Q}))$, so $\mathsf{U}(\alpha, \beta) \in r(y)$ for some $r \in \mathcal{R}$. As r is a full run, there is $u \in I$ with $u > y$, $\alpha \in r(u)$, etc., which is a contradiction. Similarly, we can show that there is no formula $\mathsf{S}(\alpha, \beta) \in \bigcup \mathrm{st}(\mathrm{bag}(\mathcal{Q}))$. Moreover, as $\mathcal{Q}$ is a quasimodel for φ, we have $\varphi \in r(u)$ for some $r \in \mathcal{R}$ and $u \in I$. Let $m = \mathrm{mos}(r) \in \mathrm{bag}(\mathcal{Q})$. Then $\varphi \in \bigcup_{t \in I} r(t) = \mathrm{st}(m) \cup \mathrm{end}(m) \cup \mathrm{cov}(m)$. So $\mathrm{bag}(\mathcal{Q})$ is perfect. □

3.5 Numbers and sizes

We will need bounds on the number of types, mosaics, and bags, and the space they take up. The following is clear:

Lemma 3.20 *Let $|\varphi| = k$, say.*

(1) $|\mathrm{sub}_{\mathsf{x}}\varphi| \leq 2k + 2 \leq 4k$ (since $k \geq 1$).
(2) The number $|\mathcal{T}(\varphi)|$ of types for φ is at most $\natural(\varphi) = 2^{2k}$ (since no type contains a formula and its negation).
(3) The number of mosaics is at most $\sharp(\varphi) = \natural(\varphi)^3 = 2^{8k}$.
(4) Any formula in $\mathrm{sub}_{\mathsf{x}}\varphi$ takes at most space k^2 to write (see the comments following definition 2.2), and a mosaic involves at most $8k$ such formulas. So any mosaic can be written in space $8k^3$.
(5) Any bag can therefore be written in space $2^{8k} \cdot 8k^3$.
(6) The number of bags is at most $\flat(\varphi) = 2^{2^{8k}}$ (this bounds the number of sets of mosaics).

3.6 Sums

Both our proofs will work by decomposing pre-quasimodels into 'sums' of smaller ones. So we have to introduce sums of mosaics and bags.

Definition 3.21 For mosaics m, n_i for $i \leq k$ for some $k < \omega$, we write $m = \sum_{i \leq k} n_i$ if

(1) $\mathrm{st}(m) = \mathrm{st}(n_0)$, and $\mathrm{end}(m) = \mathrm{end}(n_k)$,
(2) $\mathrm{st}(n_{i+1}) = \mathrm{end}(n_i)$ for each $i < k$,
(3) $\mathrm{cov}(m) = \mathrm{cov}(n_0) \cup \bigcup_{1 \leq i < k}(\mathrm{st}(n_i) \cup \mathrm{cov}(n_i) \cup \mathrm{end}(n_i)) \cup \mathrm{cov}(n_k)$.

Note that m is unique (if it exists). So we can write $\sum_{i \leq k} n_i$ for m, if condition 2 holds.

Definition 3.22 For bags μ, ν_i for $i \leq k$ for some $k < \omega$, write $\mu \equiv \sum_{i \leq k} \nu_i$ if

(1) for all $m \in \mu$ there are $n_i \in \nu_i$ (for each $i \leq k$) with $m = \sum_{i \leq k} n_i$,

(2) for all $i \leq k$ and $n_i \in \nu_i$, there are $n_j \in \nu_j$ for each $j \leq k$ with $j \neq i$, and $m \in \mu$, such that $m = \sum_{j \leq k} n_j$.

We write '$\equiv$' rather than '$=$' because there can be several (or no) μ with $\mu \equiv \sum_{i \leq k} \nu_i$. Of course, if $k = 1$ we write $m = n_0 + n_1$ and $\mu \equiv \nu_0 + \nu_1$.

Lemma 3.23 *Suppose that* $\mu, \nu_0, \ldots, \nu_k$ *are bags, and* $\mu \equiv \sum_{i \leq k} \nu_i$. *Then* $\mathrm{st}(\mu) = \mathrm{st}(\nu_0)$, $\mathrm{end}(\mu) = \mathrm{end}(\nu_k)$, *and* $\mathrm{end}(\nu_i) = \mathrm{st}(\nu_{i+1})$ *for each* $i < k$.

Proof. Purely routine applications of the definitions. □

4 Complexity of monodic fragments over linear time

In this section, we show that the satisfiability problem for the packed monodic fragment over dense linear time with endpoints (and hence over several other kinds of linear flow of time) is in 2EXPTIME. The method is an adaptation of [25, theorem 3]. The rough idea is as follows. In any full run r in a quasimodel, all Until-formulas $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$ are 'witnessed'. Whenever such a formula appears in a type $r(t)$, there will be a 'witness' $u > t$ with $\alpha \in r(u)$ and $\beta \in r(v)$ for all v between t and u. All Since-formulas in $\mathrm{sub}_{\mathsf{x}}\varphi$ are similarly witnessed. So a mosaic representing such a run should be decomposable into smaller mosaics whose endpoints witness all Until and Since formulas, and all formulas in its cover. Taking a run for each mosaic in the bag describing the quasimodel suggests that the bag itself should also be decomposable into smaller 'witnessing bags' which in turn should be decomposable. We will show that there is a quasimodel for φ iff there is a perfect bag that can be hereditarily decomposed in this way. This criterion can be decided in 2EXPTIME.

Recall that φ is a fixed monodic packed sentence; all pre-quasimodels, types, mosaics, bags, and so on, are for φ.

4.1 Defects and full decompositions

Definition 4.1 Let $m = (A, B, C)$ be a mosaic. A *defect in m* is either

(1) (cover defect) a formula in B,
(2) (future defect) a formula $\mathsf{U}(\alpha, \beta) \in A$ with either
  (a) $\neg\beta \in B$, or
  (b) $\alpha, \beta \notin C$, or
  (c) $\alpha, \mathsf{U}(\alpha, \beta) \notin C$.

(3) (past defect) the mirror image of condition 2.

Definition 4.2 Let $k < \omega$.

(1) For mosaics m, n_i for $i \leq k$, we write $m \stackrel{\text{full}}{=} \sum_{i \leq k} n_i$ (full decomposition) if $m = \sum_{i \leq k} n_i$ and
   (a) for each cover defect $\beta \in \text{cov}(m)$, there is $i < k$ with $\beta \in \text{end}(n_i)$,
   (b) for each future defect $\mathsf{U}(\alpha, \beta) \in \text{st}(m)$, there is $i < k$ with $\alpha \in \text{end}(n_i)$ and $\beta \in \text{end}(n_j)$ for all $j < i$ and $\neg\beta \notin \text{cov}(n_j)$ for all $j \leq i$.
   (c) a mirror image condition for past defects.

(2) For bags μ, ν_i for $i \leq k$, we write $\mu \stackrel{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$, and say that $\langle \nu_0, \ldots, \nu_k \rangle$ is a *full decomposition* of μ, if
   (a) for all $m \in \mu$ there are $n_i \in \nu_i$ (each $i \leq k$) with $m \stackrel{\text{full}}{=} \sum_{i \leq k} n_i$,
   (b) for all $i \leq k$ and $n_i \in \nu_i$, there are $n_j \in \nu_j$ for each $j \leq k$ with $j \neq i$, and $m \in \mu$, such that $m = \sum_{j \leq k} n_j$. (We do *not* require $\stackrel{\text{full}}{=}$ here.)

Obviously, if $\mu \stackrel{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$ then $\mu \equiv \sum_{i \leq k} \nu_i$.

4.2 Decomposition trees

By a *tree,* we will mean a non-empty partially-ordered set $(T, <)$ such that for each 'node' $t \in T$, the set $\{u \in T : u < t\}$ is well-ordered. We write $ht(t)$ (the *height* of t) for the order-type of this set; it is a unique ordinal. The *height* $ht(T)$ *of* T is the least ordinal α such that no node of T has height α. We will only consider *rooted trees* — those with a unique node (the *root*) of height 0. We often write simply T for $(T, <)$. A *successor* of a node t is a node $u > t$; such a u is a *child* of t if $ht(u) = ht(t) + 1$. Two distinct nodes are *siblings* if they are both children of the same node. A *leaf* is a node with no children. A tree is *finitely branching* if every node has finitely many children.

We will consider trees T of height ω, endowed with a binary relation $\prec$ such that for any $t \in T$, the restriction of $\prec$ to the children of t (if any) is a linear ordering. We extend $\prec$ 'lexicographically' to the whole of T by: for any $t, u \in T$, let $t \prec u$ iff $t < u$ or there are siblings t', u' with $t' \leq t$, $u' \leq u$, and $t' \prec u'$. It can be checked that this defines a linear order, called here an *earlier-later ordering,* on T.

Definition 4.3 Let μ_0 be a bag. A *decomposition tree* for μ_0 is a pair (T, μ), where T is a finitely branching rooted tree of height ω such that each node has at least two children, endowed with a earlier-later ordering $\prec$ as above, and μ is a map associating a bag $\mu(t)$ with each node $t \in T$, satisfying

(1) $\mu(t_0) = \mu_0$, where t_0 is the root of T,
(2) for each $t \in T$, if the children of t are $t_0, \ldots, t_k$ with $t_0 \prec \cdots \prec t_k$, then we have $\mu(t) \stackrel{\text{full}}{\equiv} \sum_{i \leq k} \mu(t_i)$.

We often write T for (T, μ).

Lemma 4.4 *Let $\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R})$ be a pre-quasimodel for φ, where I is a dense order with endpoints. Then* $\text{bag}(\mathcal{Q})$ *has a decomposition tree, each node of which has at most $4|\varphi| \cdot 2^{8|\varphi|}$ children.*

Proof. Assume that the endpoints of I are x, y, with $x < y$. We build a decomposition tree (T, μ) for $\text{bag}(\mathcal{Q})$ by induction, so that condition 2 of definition 4.3 holds for each $t \in T$.

The nodes of T will be pairs of elements of I. We define the root of T to be $\langle x, y \rangle$, and set $\mu(\langle x, y \rangle) = \text{bag}(\mathcal{Q})$.

Suppose we have built the nodes of T of height $\leq n$, for some $n \geq 0$, and defined μ on them to satisfy condition 2 of definition 4.3. Pick a node t of height n, and assume inductively that $t = \langle u, v \rangle$, where $u < v$ in I, and that $\mu(t) = \text{bag}(\mathcal{Q} \restriction [u, v])$. (By lemma 3.9, $\mathcal{Q} \restriction [u, v]$ is a pre-quasimodel.) For each mosaic $m \in \mu(t)$, pick $r \in \mathcal{R}$ with $\text{mos}(r \restriction [u, v]) = m$. Then, for each cover defect $\beta \in \text{cov}(m) = \bigcup_{z \in (u,v)} r(z)$, pick a single $z \in (u, v)$ with $\beta \in r(z)$. Similarly, for each future defect $\mathsf{U}(\alpha, \beta)$ of m, pick $z \in (u, v)$ with $\alpha \in r(z)$ and $\beta \in \bigcap_{t \in (u,z)} r(t)$. The definitions of 'run' and 'defect' (definitions 3.6 and 4.1) allow us to do this. Pick similar witnesses for past defects. Do this for all $m \in \mu(t)$.

Let the chosen points of (u, v) be $d_1, \ldots, d_k$ with $u = d_0 < d_1 < \cdots < d_{k+1} = v$. As I is dense, there are cover defects in each $\text{mos}(r \restriction [u, v])$, so $k \geq 1$. To bound k above, we observe that there are at most $\sharp(\varphi) = 2^{8|\varphi|}$ mosaics in $\mu(t)$, and we added at most one point d_i for each defect in each of them. Given a mosaic m, each formula in $\text{sub}_{\mathsf{x}}\varphi$ could be a cover defect of m, but for any constant c, $\mathsf{x} = c$ and $\mathsf{x} \neq c$ cannot both be cover defects. So there are at most $3|\varphi|$ cover defects of m. Each formula of the form $\mathsf{U}(\alpha, \beta)$ or $\mathsf{S}(\alpha, \beta)$ in $\text{sub}_{\mathsf{x}}\varphi$ could additionally be a future or a past defect of m, respectively, and there are $< |\varphi|$ of these. Hence, the number of defects of m is $< 4|\varphi|$. So by lemma 3.20, $k < 4|\varphi| \cdot 2^{8|\varphi|}$.

Define $\nu_i = \text{bag}(\mathcal{Q} \restriction [d_i, d_{i+1}])$ for each $i \leq k$. By lemma 3.9, $\mathcal{Q} \restriction [d_i, d_{i+1}]$ is a pre-quasimodel, so by lemma 3.17, ν_i is a bag. We claim that $\mu(t) \stackrel{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$. Let $m \in \mu(t)$. Suppose $r \in \mathcal{R}$ was the run we chose with $\text{mos}(r \restriction [u, v]) = m$. For each $i \leq k$ let

$n_i = \mathrm{mos}(r \restriction [d_i, d_{i+1}]) \in \nu_i$. Since we picked enough witnesses for r, we have

$$m = \mathrm{mos}(r) \stackrel{\text{full}}{=} \sum_{i \leq k} \mathrm{mos}(r \restriction [d_i, d_{i+1}]) = \sum_{i \leq k} n_i.$$

Conversely, for each $i \leq k$ and $n_i \in \nu_i$, there is $r \in \mathcal{R}$ with $\mathrm{mos}(r \restriction [d_i, d_{i+1}]) = n_i$. Let $n_j = \mathrm{mos}(r \restriction [d_j, d_{j+1}]) \in \nu_j$ for each $j \leq k$ with $j \neq i$, and $m = \mathrm{mos}(r \restriction [u, v]) \in \mu(t)$. Then $m = \sum_{j \leq k} n_j$ (we may not have $\stackrel{\text{full}}{=}$ here). Hence, $\mu(t) \stackrel{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$, as claimed.

So we add to T the children $t_0 \prec \cdots \prec t_k$ of t with $t_i = \langle d_i, d_{i+1} \rangle$ for each $i \leq k$. This is at least two and at most $4|\varphi| \cdot 2^{8|\varphi|}$ children. We define $\mu(t_i) = \nu_i$ for each i. We have $\mu(t) \stackrel{\text{full}}{\equiv} \sum_{i \leq k} \mu(t_i)$. Doing this for all nodes t of height n in T completes the induction. The result is a decomposition tree T for $\mathrm{bag}(\mathcal{Q})$ with the required features. □

Lemma 4.5 *Any bag with a decomposition tree is of the form* $\mathrm{bag}(\mathcal{Q})$ *for some pre-quasimodel* $\mathcal{Q}$ *over a dense linear order with endpoints.*

Proof. Let (T, μ) be a decomposition tree for the bag μ_0. We use it to build a pre-quasimodel $\mathcal{Q} = (I, (\Sigma_t : t \in I), \mathcal{R})$ with $\mathrm{bag}(\mathcal{Q}) = \mu_0$. First we define the linear order I. For each $t \in T$, we will associate rational numbers t^- and t^+, with $t^- < t^+$. With the root, we associate $0, 1$. Inductively, if we have associated t^-, t^+ with t, and the children of t are $t_0, \ldots, t_k$, say, with $t_0 \prec \cdots \prec t_k$, we choose rationals $t^- = q_0 < q_1 < \cdots < q_{k+1} = t^+$ and set $t_i^- = q_i$ and $t_i^+ = q_{i+1}$ for each $i \leq k$. We let I be the suborder of $(\mathbb{Q}, <)$ consisting of all rationals associated with nodes of T. Clearly, I is dense and has endpoints 0,1 — indeed, we could easily arrange that $I = [0, 1] \cap \mathbb{Q}$. Notice above that $t^- = t_0^-$ and $t^+ = t_k^+$. Hence, each point in I associated with a node $t \in T$ is also associated with successors of t of all heights $> ht(t)$.

For $t \in T$, we define $\Sigma_{t^-} = \mathrm{st}(\mu(t))$ and $\Sigma_{t^+} = \mathrm{end}(\mu(t))$. Lemma 3.23 shows that this determines a well-defined set Σ_x for each $x \in I$; it is a realisable state candidate since each $\mu(t)$ is a bag.

It remains to define the set $\mathcal{R}$ of runs of $\mathcal{Q}$.

Notation For $l < \omega$ let $T(l)$ denote the (non-empty) set of nodes of T of height l, and suppose that $T(l) = \{t_0^l, \ldots, t_{h(l)}^l\}$, where $t_0^l \prec \cdots \prec t_{h(l)}^l$. For each $i \leq h(l)$, we suppose that the children of t_i^l in T are $t_{i_-}^{l+1} \prec \cdots \prec t_{i_+}^{l+1}$, where $0 \leq i_- < i_+ \leq h(l+1)$.

A *walk along* $T(l)$ is a sequence $\langle m_0, \ldots, m_{h(l)} \rangle \in \prod_{i \leq h(l)} \mu(t_i^l)$ such that $\mathrm{end}(m_i) = \mathrm{st}(m_{i+1})$ for all $i < h(l)$, and $\sum_{i \leq h(l)} m_i \in \mu_0$. For each $l < \omega$, we will define a set $\mathcal{W}_l$ of walks along $T(l)$ so that 'every mosaic in every bag from $T(l)$ is hit by a walk in $\mathcal{W}_l$'

— formally, that

$$\{m_i : \langle m_0, \ldots, m_{h(l)} \rangle \in \mathcal{W}_l\} = \mu(t^l_i) \quad \text{for every } i \leq h(l). \tag{4.1}$$

We define $\mathcal{W}_0$ as simply the set $\{\langle m \rangle : m \in \mu_0\}$. Let $l < \omega$ and assume inductively that $\mathcal{W}_l$ is defined, satisfying (4.1). We obtain $\mathcal{W}_{l+1}$ as follows.

(1) *Fully decompose each walk in* $\mathcal{W}_l$. Let $w = \langle m_0, \ldots, m_{h(l)} \rangle \in \mathcal{W}_l$. For each $i \leq h(l)$, choose a mosaic $n_j \in \mu(t^{l+1}_j)$ for each j with $i_- \leq j \leq i_+$, so that $m_i \overset{\text{full}}{\equiv} \sum_j n_j$. This is possible because T is a decomposition tree, so $\mu(t^l_i) \overset{\text{full}}{\equiv} \sum_j \mu(t^{l+1}_j)$. Then $\sum_{s \leq h(l+1)} n_s = \sum_{i \leq h(l)} m_i \in \mu_0$, so $w' = \langle n_0, \ldots, n_{h(l+1)} \rangle$ is a walk along $T(l+1)$. Such a w' is called a *full decomposition* of w. We put it into $\mathcal{W}_{l+1}$. We do this for each $w \in \mathcal{W}_l$.

(2) *Ensure that every mosaic in every bag from* $T(l+1)$ *is hit by a walk.* Let $j \leq h(l+1)$ and let n be any mosaic in $\mu(t^{l+1}_j)$. Assume that t^{l+1}_j is a child of t^l_i, say, so that $i_- \leq j \leq i_+$. Since $\mu(t^l_i) \overset{\text{full}}{\equiv} \sum_{i_- \leq s \leq i_+} \mu(t^{l+1}_s)$, we may choose mosaics $n_s \in \mu(t^{l+1}_s)$ for each s with $i_- \leq s \leq i_+$, such that $n_j = n$ and $m = \sum_{i_- \leq s \leq i_+} n_s \in \mu(t^l_i)$. By the inductive hypothesis, there is $\langle m_0, \ldots, m_{h(l)} \rangle \in \mathcal{W}_l$ such that $m_i = m$. As T is a decomposition tree, for each $k \leq h(l)$ with $k \neq i$, we may choose mosaics $n_s \in \mu(t^{l+1}_s)$ for each $k_- \leq s \leq k_+$ with $m_k = \sum_{k_- \leq s \leq k_+} n_s$. Then $\sum_{s \leq h(l+1)} n_s = \sum_{s \leq h(l)} m_s \in \mu_0$, so $\langle n_0, \ldots, n_{h(l+1)} \rangle$ is a walk along $T(l+1)$. Put it into $\mathcal{W}_{l+1}$. We have $n_j = n$ as required. Doing this for each mosaic $n \in \mu(t)$ for each $t \in T(l+1)$ completes the construction of $\mathcal{W}_{l+1}$ and preserves (4.1).

So we have defined a set $\mathcal{W}_l$ of walks along $T(l)$ for each $l < \omega$. Let $w = \langle m_0, \ldots, m_{h(l)} \rangle \in \mathcal{W}_l$. Suppose $j \leq h(l)$ and that the rationals $x < y$ are associated with t^l_j. Then we define $w[x] = \text{st}(m_j) \in \Sigma_x$ and $w[y] = \text{end}(m_j) \in \Sigma_y$. If $j < h(l)$ then $\text{end}(m_j) = \text{st}(m_{j+1})$, so this is well-defined.

A *run sequence* is a sequence $\rho = \langle w_l, w_{l+1}, \ldots \rangle$ for some $l < \omega$, where for each $i \geq l$, $w_i \in \mathcal{W}_i$ and w_{i+1} is a full decomposition of w_i. For such a ρ, we define $r_\rho \in \prod_{x \in I} \Sigma_x$ by letting $r_\rho(x) = w_i[x]$ for any $i \geq l$ such that x is associated with a node in $T(i)$. Since x is associated with a node of $T(l)$ for all but finitely many $l < \omega$, lemma 3.23 can be used to check that r_ρ is thereby well-defined and total. It also follows that if $\rho' = \langle w_{l+1}, w_{l+2}, \ldots \rangle$ then $r_\rho = r_{\rho'}$.

Claim. Let $\rho = \langle w_{l_0}, w_{l_0+1}, \ldots \rangle$ be a run sequence. Then r_ρ is a run on I.

Proof of claim. Because for each $l \geq l_0$, w_l is fully decomposed by w_{l+1}, it follows that if $w_l = \langle m_0, \ldots, m_{h(l)} \rangle$ then for each $i \leq j \leq h(l)$ we have

$$\mathrm{mos}(r_\rho \restriction [(t^l_i)^-, (t^l_j)^+]) = \sum_{i \leq k \leq j} m_k. \tag{4.2}$$

To see this, note first that if $(t^l_i)^- < z < (t^l_j)^+$ then any formula in $r_\rho(z)$ is in the cover of $\sum_{i \leq k \leq j} m_k$. Conversely, any formula in this cover either is in the start or end of, or is a cover defect of, some m_k. In either case, it is in the start or end of a mosaic refining m_k in the full decomposition w_{l+1} of w_l, and hence is in $r_\rho(z)$ for some $z \in ((t^l_i)^-,\ (t^l_j)^+)$.

Now we prove the claim. We require that condition 2 of definition 3.6 holds: for each $x \in I$ and $\mathsf{U}(\alpha, \beta) \in \mathrm{sub}_\mathsf{x}\varphi$, we have $\mathsf{U}(\alpha, \beta) \in r_\rho(x)$ iff:

(a) there is $v \in I$ with $x < v$, $\alpha \in r_\rho(v)$, and $\beta \in r_\rho(z)$ for all $z \in I$ with $x < z < v$, or
(b) $\mathsf{U}(\alpha, \beta) \in r_\rho(1)$, and $\beta \in r_\rho(z)$ for all $z \in I$ with $z > x$.

Let $x \in I$. If $x = 1$, there is nothing to prove. Assume $x < 1$. We may choose $t^l_i \in T(l)$ for some $l \geq l_0$ (where $i \leq h(l)$) such that x is the smaller of the two points of I associated with t^l_i. Let $w_l = \langle m_0, \ldots, m_{h(l)} \rangle$. So $r_\rho(x) = \mathrm{st}(m_i)$.

By (4.2), for any $j \leq h(l)$, if $\beta \in \mathrm{cov}(m_j)$ then $\beta \in r_\rho(z)$ for some $z \in I$ lying strictly between the points associated with t^l_j.

Assume that $\mathsf{U}(\alpha, \beta) \in r_\rho(x)$. Take the greatest integer j with $i \leq j \leq h(l)$ and such that $\mathsf{U}(\alpha, \beta) \in \mathrm{st}(m_j)$ and $\neg\beta \notin \mathrm{cov}(m_k) \cup \mathrm{end}(m_k)$ for all $i \leq k < j$. (Since i has this property, the set of such j is non-empty.)

- If $\mathsf{U}(\alpha, \beta) \in \mathrm{end}(m_j)$ and $\neg\beta \notin \mathrm{cov}(m_j) \cup \mathrm{end}(m_j)$, then by maximality of j we must have $j = h(l)$, and it follows from (4.2) that $\neg\beta \notin \mathrm{cov}(\sum_{i \leq k \leq j} m_k) = \bigcup_{x<y<1} r(y)$, so (b) above holds.
- If $\alpha \in \mathrm{end}(m_j)$ and $\neg\beta \notin \mathrm{cov}(m_j)$, then using (4.2), we see that (a) holds with 'v' as $(t^l_j)^+$.
- Otherwise, $\mathsf{U}(\alpha, \beta)$ is a future defect in m_j. Now w_{l+1} fully decomposes w_l. So if $w_{l+1} = \langle n_0, \ldots, n_{h(l+1)} \rangle$, then using the earlier notation, there is $j_- \leq k < j_+$ with $\alpha \in \mathrm{end}(n_k)$, $\beta \in \mathrm{end}(n_s)$ for all $j_- \leq s < k$, and $\neg\beta \notin \mathrm{cov}(n_s)$ for all $j_- \leq s \leq k$. Let v be the greater point associated with t^{l+1}_k. Then $\alpha \in r_\rho(v)$, and using (4.2) we see that $\beta \in r_\rho(z)$ for all $z \in I$ with $x < z < v$, as required.

Conversely, if (a) or (b) holds, then by increasing l if need be, we may assume that the 'v' in (a), if applicable, is associated with a node in $T(l)$ — say, $v = (t^l_j)^+$, and $x = (t^l_i)^-$, where $i \leq j \leq h(l)$. If (b) holds, take $v = 1$, $j = h(l)$, and again assume $x = (t^l_i)^-$. Let $w_l = \langle m_0, \ldots, m_{h(l)} \rangle$. By (4.2), $\neg\beta \notin \mathrm{cov}(m_j)$ and either $\alpha \in \mathrm{end}(m_j)$ (if (a) holds) or $\beta, \mathsf{U}(\alpha, \beta) \in \mathrm{end}(m_j)$ (if (b) holds). As m_j is a mosaic, $\mathsf{U}(\alpha, \beta) \in \mathrm{st}(m_j)$. If $j > i$

then $\mathsf{U}(\alpha, \beta) \in \mathrm{end}(m_{j-1})$, and by (4.2), $\beta \in \mathrm{end}(m_{j-1})$ and $\neg\beta \notin \mathrm{cov}(m_{j-1})$; so as m_{j-1} is a mosaic, $\mathsf{U}(\alpha, \beta) \in \mathrm{st}(m_{j-1})$. Continuing inductively in this way, we obtain $\mathsf{U}(\alpha, \beta) \in \mathrm{st}(m_i) = r_\rho(x)$, as required.

The check for Since is similar. Finally let $c \in L$ be a constant. To show that $\mathsf{x} = c \in r_\rho(t)$ iff $\mathsf{x} = c \in r_\rho(u)$ for all $t, u \in I$, it suffices to show that for each $l \geq l_0$, if $w_l = \langle m_0, \ldots, m_{h(l)} \rangle$, then $\mathsf{x} = c \in \mathrm{st}(m_i)$ iff $\mathsf{x} = c \in \mathrm{end}(m_i)$ for each $i \leq h(l)$. But this is immediate since m_i is a mosaic. (Alternatively, we could apply (4.2) and lemma 3.14.) This proves the claim.

We set $\mathcal{R} = \{r_\rho : \rho \text{ a run sequence}\}$.

Claim. $\mathcal{Q} = (I, (\Sigma_x : x \in I), \mathcal{R})$ is a pre-quasimodel with $\mathrm{bag}(\mathcal{Q}) = \mu_0$.

Proof of claim. For arbitrary $x \in I$, it is required that $\Sigma_x = \{r(x) : r \in \mathcal{R}\}$. Take $t \in T(l)$ (for some $l < \omega$) such that x is associated with t, and take a type $p \in \Sigma_x$. We have $p = \mathrm{st}(m)$ or $p = \mathrm{end}(m)$ for some $m \in \mu(t)$. Now m is hit by a walk in $\mathcal{W}_l$: there is $w = \langle m_0, \ldots, m_{h(l)} \rangle \in \mathcal{W}_l$ and $i \leq h(l)$ with $m_i = m$. As each walk in each $\mathcal{W}_{l'}$ is fully decomposed in $\mathcal{W}_{l'+1}$, there is a run sequence $\rho = \langle w_l, w_{l+1}, \ldots \rangle$ with $w_l = w$. Then $r_\rho \in \mathcal{R}$ and $r_\rho(x) = p$. Since p was arbitrary, we obtain $\Sigma_x \subseteq \{r(x) : r \in \mathcal{R}\}$, and the converse inclusion is trivial.

A similar argument shows that for each $m \in \mu_0$, there is a run sequence $\rho = \langle w_0, w_1, \ldots \rangle$ with $w_0 = \langle m \rangle$. By (4.2), $\mathrm{mos}(r_\rho) = m$. Conversely, if $\rho = \langle w_l, w_{l+1}, \ldots \rangle$ is a run sequence with $w_l = \langle m_0, \ldots, m_{h(l)} \rangle$, say, then (4.2) and the definition of 'walk' yields $\mathrm{mos}(r_\rho) = \sum_{i \leq h(l)} m_i \in \mu_0$. It is now plain that $\mathrm{bag}(\mathcal{Q}) = \mu_0$. □

Corollary 4.6 *The following are equivalent:*

(1) φ has a model whose flow of time is linear, dense, and with endpoints,
(2) there is a set $\mathcal{B}$ of bags such that:
- *(a) some bag $\mu_0 \in \mathcal{B}$ is perfect,*
- *(b) for each $\mu \in \mathcal{B}$ there are $k \leq 4|\varphi| \cdot 2^{8|\varphi|}$ and $\nu_0, \ldots, \nu_k \in \mathcal{B}$ with $\mu \overset{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$.*

Proof. By fact 3.10, φ has a model with flow of time a dense linear order I with endpoints, iff there is a quasimodel for φ over I. By lemma 3.19, this is iff there is a pre-quasimodel $\mathcal{Q}$ over I with $\mathrm{bag}(\mathcal{Q})$ perfect. By lemmas 4.4 and 4.5, this is iff there are a perfect bag μ_0, and a decomposition tree T for μ_0 with each node having at most $4|\varphi| \cdot 2^{8|\varphi|}$ children.

If there are such T and μ_0, then $\mathcal{B} = \{\mu(t) : t \in T\}$ satisfies the conditions of part 2 of the corollary. But given $\mathcal{B}, \mu_0$ satisfying these conditions, we may easily construct a decomposition tree for μ_0 with the stated bound on numbers of children by induction on heights of nodes, using bags from $\mathcal{B}$ to label nodes. □

Theorem 4.7 *The problem of whether a monodic packed sentence is satisfiable in some temporal structure (with arbitrary first-order domain and) with a dense linear flow of time with endpoints is* 2EXPTIME*-complete.*

Proof. By proposition 2.5, the problem is 2EXPTIME-hard. So it remains only to determine in 2EXPTIME whether φ is satisfiable over dense linear time with endpoints. To do this, we will specify an algorithm that checks the criterion provided by corollary 4.6, using an approach originating in [22] (we used it already in proposition 3.5).

The algorithm first constructs the set $\mathcal{B}_0$ of all bags, by enumerating all sets of mosaics and placing each one in $\mathcal{B}_0$ iff its start and end are realisable state candidates; by proposition 3.5, whether a set of types for φ is a realisable state candidate can be tested in 2EXPTIME in $|\varphi|$. There are at most $\flat(\varphi)$ bags (see lemma 3.20), so the construction of $\mathcal{B}_0$ takes 2EXPTIME.

Now the algorithm loops to construct a chain $\mathcal{B}_0 \supseteq \mathcal{B}_1 \supseteq \cdots$. Let $n \geq 0$ and assume that $\mathcal{B}_n$ has been constructed. For each $\mu \in \mathcal{B}_n$ the algorithm searches for bags $\nu_0, \ldots, \nu_k \in \mathcal{B}_n$ for some $k \leq 4|\varphi| \cdot 2^{8|\varphi|}$, with $\mu \overset{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$. If such bags can be found, the algorithm includes μ in $\mathcal{B}_{n+1}$; otherwise, μ is not included. There are at most double-exponentially many sequences $\nu_0, \ldots, \nu_k$ to check, for each of at most double-exponentially many bags μ in $\mathcal{B}_n$; and given $\nu_0, \ldots, \nu_k$, whether or not $\mu \overset{\text{full}}{\equiv} \sum_{i \leq k} \nu_i$ can easily be determined in 2EXPTIME. So the construction of $\mathcal{B}_{n+1}$ from $\mathcal{B}_n$ takes 2EXPTIME.

Once $\mathcal{B}_{n+1}$ is constructed, the algorithm checks (in 2EXPTIME) whether there is a perfect bag in $\mathcal{B}_{n+1}$. If not, the algorithm terminates with result 'φ `is unsatisfiable`'. Otherwise, if $\mathcal{B}_{n+1} = \mathcal{B}_n$, the algorithm terminates with result 'φ `is satisfiable`'. Otherwise, the next iteration begins. The number of iterations is therefore bounded by $|\mathcal{B}_0|$, and so is at most double exponential. This algorithm therefore terminates in 2EXPTIME.

If the algorithm claims that φ is satisfiable, then clearly the final $\mathcal{B}_n$ it constructed satisfies the conditions of corollary 4.6. Conversely, if there is a set $\mathcal{B}$ of bags as in the corollary, then a simple induction shows that $\mathcal{B} \subseteq \mathcal{B}_n$ for all n. The algorithm can therefore never claim that φ is unsatisfiable, and hence it will eventually state that φ is satisfiable. So the algorithm is correct. □

We can easily use theorem 4.7 to obtain complexity results over additional classes of flows of time.

Theorem 4.8 *The satisfiability problem for monodic packed sentences in temporal structures with arbitrary first-order domains, over each of the following (classes of) flows of time, is* 2EXPTIME*-complete.*

(1) All linear flows.
(2) Discrete linear flows — those satisfying $\forall x(\exists y(x < y) \to \exists z(x < z \wedge \neg\exists t(x < t < z)))$ *and its mirror image obtained by replacing '<' by '>'.*
(3) Dense linear flows.
(4) More generally, any class $\mathcal{C}$ *such that for some propositional temporal formula* ξ *written with Until and Since,* $\mathcal{C}$ *is the class of linear flows in which* ξ *is satisfiable.*
(5) The rationals, $(\mathbb{Q}, <)$*.*

Proof. By proposition 2.5, we already have 2EXPTIME-hardness, so we only need prove that the above problems are in 2EXPTIME. First, we show that all the problems reduce to problems of type 1. Let φ be a monodic packed sentence. Clearly, φ is satisfiable in discrete time iff $\varphi \wedge \Box((\mathsf{U}(\top,\top) \to \mathsf{U}(\top,\bot)) \wedge (\mathsf{S}(\top,\top) \to \mathsf{S}(\top,\bot)))$ is satisfiable in linear time, and satisfiable in dense time iff $\varphi \wedge \Box(\mathsf{U}(\top,\top) \to \neg\mathsf{U}(\top,\bot))$ is satisfiable in linear time. We next reduce each type 4 problem to linear satisfiability. Suppose that ξ is a propositional formula and that $\mathcal{C}$ is the class of all linear flows of time in which ξ is satisfiable. Given φ, we may rename its non-logical symbols so that none of them occur in ξ. It is now clear that φ is satisfiable over a flow of time in $\mathcal{C}$ iff $\varphi \wedge \Diamond\xi$ is satisfiable over linear time.

It remains to deal with case 5: $(\mathbb{Q}, <)$. It is not hard to see that φ is satisfiable over $\mathbb{Q}$ iff it is satisfiable in a dense linear flow of time without endpoints. This can be shown using the downward Löwenheim–Skolem theorem (see, e.g., [10]); similar results for monodic fragments were shown in [15, theorem 15] and [12, theorem 1]. But φ is satisfiable over such a flow iff $\varphi \wedge \Box(\mathsf{U}(\top,\top) \wedge \mathsf{S}(\top,\top) \wedge \neg\mathsf{U}(\top,\bot))$ is satisfiable over linear time.

So all the problems reduce (in logarithmic space) to case 1. The theorem is therefore established if we show that satisfiability over linear time is decidable in 2EXPTIME. Let φ be a monodic packed sentence, and let q be a propositional atom not occurring in φ. Define the *temporal relativisation* φ^q of φ to q by induction on φ. If φ is atomic we let $\varphi^q = \varphi$. We let $(\varphi \wedge \psi)^q = \varphi^q \wedge \psi^q$, $(\neg\varphi)^q = \neg\varphi^q$, and $(\exists\bar{\mathsf{x}}(\gamma \wedge \varphi))^q = \exists\bar{\mathsf{x}}(\gamma \wedge \varphi^q)$; the main cases are $\mathsf{U}(\varphi,\psi)^q = \mathsf{U}(q \wedge \varphi^q, q \to \psi^q)$, and $\mathsf{S}(\varphi,\psi)^q = \mathsf{S}(q \wedge \varphi^q, q \to \psi^q)$. It should be clear that for any temporal structure $\mathcal{M} = ((T,<), D, (M_t : t \in T))$, if $(I,<)$ is the suborder of $(T,<)$ with $I = \{t \in T : M_t \models q\}$, then for any $u \in I$,

$$\mathcal{M}, u \models \varphi^q \iff ((I,<), D, (M_t : t \in I)), u \models \varphi. \tag{4.3}$$

Finally we let

$$\varphi^\delta = (q \wedge \varphi^q) \wedge \underbrace{\mathsf{U}(\top,\top)}_{\geq 2 \text{ elements}} \wedge \underbrace{\Box\neg\mathsf{U}(\top,\bot)}_{\text{dense}} \wedge \underbrace{\Diamond\neg\mathsf{U}(\top,\top) \wedge \Diamond\neg\mathsf{S}(\top,\top)}_{\text{with endpoints}}.$$

Clearly, φ^δ is in the monodic packed fragment and is constructible in logarithmic space in the size of φ. It is routine to check, using (4.3), that a sentence φ is satisfiable over linear time iff φ^δ has a model with dense flow of time with distinct endpoints. So by theorem 4.7, we may decide in 2EXPTIME whether φ is satisfiable over linear time. □

We are unable to handle classes defined by arbitrary first-order sentences, as Until and Since are not expressively complete over linear time (see [6] for background on this notion). This could perhaps be remedied by adding the so-called *Stavi connectives;* we conjecture that the methods described here generalise to handle these. When we come to the real numbers, next, there will be no such limitation, as Until and Since are expressively complete over this flow. We are unable to provide any complexity results for temporal structures with flow of time $(\mathbb{R},<)$ and with arbitrary domains — in this context, it is an open problem whether any non-trivial monodic fragment is even decidable.

5 Complexity with finite domains over real numbers

Now we show that the satisfiability problem for the monodic packed fragment over $(\mathbb{R},<)$ with finite first-order domains is also 2EXPTIME-complete. We can then show by reduction that satisfiability with finite first-order domains over a wide range of linear flows of time has the same complexity. We adapt the decision procedure for monodic sentences over $\mathbb{R}$ with finite domains given in [15] (the 3-theories and characters of that paper are replaced by mosaics and bags here). In §4, we *decomposed* a quasimodel or perfect bag; now, we will *construct* a perfect bag from simpler bags. We will use three kinds of construction operation: taking a sum of two bags, summing ω copies of a single bag, either forwards or backwards, and 'shuffling' finitely many bags densely together, with some realisable state candidates mixed in (this is how the construction must start off). Our aim is to show that all and only the bags of pre-quasimodels of φ can be built in this way. We prove by induction on the construction that any bag constructed like this is the bag associated with a pre-quasimodel. Conversely, we use Ramsey's theorem and properties of $\mathbb{R}$ in a way familiar from [3,20] to show that the bag of any pre-quasimodel is constructible. It is then easy to write down a 2EXPTIME algorithm to test whether there exists a perfect constructible bag, and hence a quasimodel of φ.

In this section, *all bags, realisable state candidates, pre-quasimodels, etc., are implicitly finitary,* though sometimes we say 'finitary' explicitly, for emphasis. We do not use full decompositions here. All intervals (e.g., [0,1]) are intervals in $(\mathbb{R}, <)$; recall that we also write $[0,1]$ for the linear order $([0,1], <)$.

Definition 5.1 A bag μ is said to be *realisable* if there is a finitary pre-quasimodel $\mathcal{Q}$ for φ over $[0,1]$ with $\text{bag}(\mathcal{Q}) = \mu$.

5.1 Sums of bags

Recall from definition 3.22 the notion of $\mu \equiv \nu_0 + \nu_1$, for bags μ, ν_0, ν_1.

Lemma 5.2 *Suppose that μ, ν_0, ν_1 are bags with $\mu \equiv \nu_0 + \nu_1$. If ν_0 and ν_1 are realisable, then so is μ.*

Proof. For each i, let $\mathcal{Q}_i = ([0,1], (\Sigma^i_t : t \in [0,1]), \mathcal{R}_i)$ be a pre-quasimodel with $\text{bag}(\mathcal{Q}_i) = \nu_i$. We define $\mathcal{Q} = ([0,1], (\Sigma_t : t \in [0,1]), \mathcal{R})$ as follows. We let

$$\Sigma_t = \begin{cases} \Sigma^0_{2t}, & \text{if } t \in [0, 1/2], \\ \Sigma^1_{2t-1}, & \text{if } t \in [1/2, 1]. \end{cases}$$

This is well-defined when $t = 1/2$, by lemma 3.23. For $r_i \in \mathcal{R}_i$ $(i = 0, 1)$ such that $r_0(1) = r_1(0)$, define $r_0 + r_1$ by

$$(r_0 + r_1)(t) = \begin{cases} r_0(2t), & \text{if } t \in [0, 1/2], \\ r_1(2t-1), & \text{if } t \in [1/2, 1]. \end{cases}$$

It is easily checked that $r_0 + r_1$ is a well-defined run in $\prod_{t \in [0,1]} \Sigma_t$ and $\text{mos}(r_0 + r_1) = \text{mos}(r_0) + \text{mos}(r_1)$. Now we define $\mathcal{R}$.

(1) Let $m \in \mu$. There are $n_0 \in \nu_0$ and $n_1 \in \nu_1$ with $m = n_0 + n_1$. Pick $r_i \in \mathcal{R}_i$ with $\text{mos}(r_i) = n_i$ $(i = 0, 1)$, and put $r_0 + r_1$ into $\mathcal{R}$. Note that $\text{mos}(r_0 + r_1) = m$. Do this for each $m \in \mu$.
(2) Let $i < 2$ and $r_i \in \mathcal{R}_i$. Let $n_i = \text{mos}(r_i) \in \nu_i$. There is $n_{1-i} \in \nu_{1-i}$ with $n_0 + n_1 \in \mu$. Pick $r_{1-i} \in \mathcal{R}_{1-i}$ with $\text{mos}(r_{1-i}) = n_{1-i}$, and put $r_0 + r_1$ into $\mathcal{R}$. Again, $\text{mos}(r_0 + r_1) \in \mu$. Do this for each $i < 2$ and each of the finitely many $r_i \in \mathcal{R}_i$ (recall that $\mathcal{Q}_i$ is finitary).

Clearly, $\mathcal{R}$ is finite. We check that $\mathcal{Q}$ is a finitary pre-quasimodel. Clearly, $r(t) \in \Sigma_t$ for all $r \in \mathcal{R}$ and $t \in [0,1]$. Let $t \in [0,1]$ and $p \in \Sigma_t$. We require $p = r(t)$ for some $r \in \mathcal{R}$. Suppose that $t \in [0, 1/2]$; the other case is similar. As $p \in \Sigma^0_{2t}$ and $\mathcal{Q}_0$ is a pre-quasimodel, there is $r_0 \in \mathcal{R}_0$ with $r_0(2t) = p$. By the second clause defining $\mathcal{R}$, there is $r_1 \in \mathcal{R}_1$ with $r_0 + r_1 \in \mathcal{R}$; then $(r_0 + r_1)(t) = r_0(2t) = p$ as required.

Now we check that $\mathrm{bag}(\mathcal{Q}) = \mu$. It is clear by construction that $\mathrm{mos}(r) \in \mu$ for every $r \in \mathcal{R}$. Conversely, let $m \in \mu$. By the first clause defining $\mathcal{R}$, there are $n_0 \in \nu_0$, $n_1 \in \nu_1$ with $m = n_0 + n_1$, and $r_i \in \mathcal{R}_i$ $(i = 0, 1)$ with $\mathrm{mos}(r_i) = n_i$ and $r_0 + r_1 \in \mathcal{R}$. Clearly, $\mathrm{mos}(r) = m$. □

5.2 Iteration

Now we consider sums of ω copies of a fixed mosaic and bag.

Definition 5.3

(1) Let m, n be mosaics. We write $m \equiv n \cdot \omega$ if $\mathrm{st}(m) = \mathrm{st}(n) = \mathrm{end}(n)$, $\mathrm{cov}(m) = \mathrm{st}(n) \cup \mathrm{cov}(n)$, and for each formula $\mathsf{S}(\alpha, \beta) \in \mathrm{end}(m)$, we have $\neg\beta \notin \mathrm{cov}(m)$.
We write $m \equiv n \cdot \omega^*$ if the mirror image of this holds.

(2) Let μ, ν be bags. We write $\mu \equiv \nu \cdot \omega$ if
  (a) for each $m \in \mu$ there is $n \in \nu$ with $m \equiv n \cdot \omega$,
  (b) for each $n \in \nu$ there is $m \in \mu$ with $m \equiv n \cdot \omega$.
  We define $\mu \equiv \nu \cdot \omega^*$ similarly.

Observe that if $\mu \equiv \nu \cdot \omega$ or $\mu \equiv \nu \cdot \omega^*$ then $\mathrm{st}(\nu) = \mathrm{end}(\nu)$.

Lemma 5.4 *Let μ, ν be bags with $\mu \equiv \nu \cdot \omega$. If ν is realisable then so is μ. The same holds when $\mu \equiv \nu \cdot \omega^*$.*

Proof. Let $\mathcal{Q} = ([0,1], (\Sigma_t : t \in [0,1]), \mathcal{R})$ be a pre-quasimodel with $\mathrm{bag}(\mathcal{Q}) = \nu$. We will define a pre-quasimodel $\mathcal{Q}^* = ([0,1], (\Sigma^*_t : t \in [0,1]), \mathcal{R}^*)$ with $\mathrm{bag}(\mathcal{Q}^*) = \mu$. Pick real numbers $0 = x_0 < x_1 < \cdots < 1$ with $\sup\{x_i : i < \omega\} = 1$, and an order-isomorphism $\theta_i : [x_i, x_{i+1}] \to [0,1]$ for each $i < \omega$. For $t \in [0,1]$, define

$$\Sigma^*_t = \begin{cases} \Sigma_{\theta_i(t)}, & \text{if } i < \omega \text{ and } t \in [x_i, x_{i+1}], \\ \mathrm{end}(\mu), & \text{if } t = 1. \end{cases}$$

This is well-defined, since $\mathrm{st}(\nu) = \mathrm{end}(\nu)$ and so $\Sigma_0 = \Sigma_1$; it is a realisable state candidate since $\mathcal{Q}$ is a pre-quasimodel and μ a bag. For each $m \in \mu$ and $n \in \nu$ with

$m \equiv n \cdot \omega$, and $r \in \mathcal{R}$ with $\text{mos}(r) = n$, define $r^m \in \prod_{t \in [0,1]} \Sigma_t^*$ by

$$r^m(t) = \begin{cases} r(\theta_i(t)), & \text{if } i < \omega \text{ and } t \in [x_i, x_{i+1}], \\ \text{end}(m), & \text{if } t = 1. \end{cases}$$

Since $m \equiv n \cdot \omega$, $r(0) = \text{st}(m) = \text{end}(m) = r(1)$, so this is well-defined.

Claim. r^m is a run and $\text{mos}(r^m) = m$.

Proof of claim. We have $r^m(0) = r(0) = \text{st}(n) = \text{st}(m)$, and $r^m(1) = \text{end}(m)$. Also, $\bigcup_{t \in (0,1)} r^m(t) = \bigcup_{t \in [0,1]} r(t) = \text{st}(n) \cup \text{end}(n) \cup \text{cov}(n) = \text{cov}(m)$. Hence, $\text{mos}(r^m) = m$.

Now we check that r^m is a run. As m is a mosaic, we know by lemma 3.14 that for each constant $c \in L$ and each $t, u \in [0,1]$, $\mathsf{x} = c \in r^m(t)$ iff $\mathsf{x} = c \in r^m(u)$. Now let $t \in [0,1)$; let $i < \omega$ be such that $t \in [x_i, x_{i+1})$. Let $\mathsf{U}(\alpha, \beta) \in \text{sub}_{\mathsf{x}}\varphi$.

First assume that $\mathsf{U}(\alpha, \beta) \in r^m(t)$ and that there is no $u \in (t, 1]$ with $\alpha \in r^m(u)$ and $\beta \in r^m(v)$ for all $v \in (t, u)$. So $\mathsf{U}(\alpha, \beta) \in r(\theta_i(t))$, and there is no $u \in (\theta_i(t), 1]$ with $\alpha \in r(u)$ and $\beta \in r(v)$ for all $v \in (t, u)$. As r is a run, we obtain $\mathsf{U}(\alpha, \beta) \in r(1) = \text{end}(n) = \text{st}(n) = r(0) = \text{st}(m)$, and $\beta \in \bigcap_{v \in (\theta_i(t), 1]} r(v)$. So $\mathsf{U}(\alpha, \beta) \in r^m(x_{i+1})$. Now if there is $u \in (x_{i+1}, x_{i+2}]$ with $\alpha \in r^m(u)$ and $\beta \in r^m(v)$ for all $v \in (x_{i+1}, u)$, then $\beta \in r^m(v)$ for all $v \in (t, u)$ and we contradict our assumption on r^m. So there is no $u \in (0, 1]$ with $\alpha \in r(u)$ and $\beta \in r(v)$ for all $v \in (0, u)$. As $\mathsf{U}(\alpha, \beta) \in r(0)$ and r is a run, we deduce that $\beta, \neg\alpha \in \bigcap_{v \in (0,1]} r(v)$. As the $r(v)$ are types, we obtain $\alpha, \neg\beta \notin \bigcup_{v \in (0,1]} r(v) = \text{cov}(n) \cup \text{end}(n) = \text{cov}(m)$.

Now m is a mosaic and $\mathsf{U}(\alpha, \beta) \in \text{st}(m)$: so $\alpha \in \text{cov}(m) \cup \text{end}(m)$ or $\beta, \mathsf{U}(\alpha, \beta) \in \text{end}(m)$. Since $\alpha \notin \text{cov}(m)$, we get $\alpha \in \text{end}(m) = r^m(1)$ or $\beta, \mathsf{U}(\alpha, \beta) \in \text{end}(m) = r^m(1)$, and we already have $\beta \in r^m(u)$ for all $u \in (t, 1)$. We are done.

Next, suppose that there is $u \in (t, 1)$ with $\alpha \in r^m(u)$ and $\beta \in r^m(v)$ for all $v \in (t, u)$. Let $i \leq j < \omega$ with $u \in (x_j, x_{j+1}]$. If $j = i$ we obtain $\mathsf{U}(\alpha, \beta) \in r^m(t)$ since r is a run. If $j > i$, we obtain $\mathsf{U}(\alpha, \beta) \in r^m(x_j) = r(1)$ and $\beta \in r(v)$ for all $v \in (\theta_i(t), 1]$, so again $\mathsf{U}(\alpha, \beta) \in r(\theta_i(t)) = r^m(t)$ since r is a run.

Now assume that $\alpha \in r^m(1)$ or $\beta, \mathsf{U}(\alpha, \beta) \in r^m(1)$, and $\beta \in r^m(v)$ for all $v \in (t, 1)$. Then $\neg\beta \notin \text{cov}(n) \cup \text{end}(n) = \text{cov}(m)$. As $r^m(1) = \text{end}(m)$ and m is a mosaic, $\mathsf{U}(\alpha, \beta) \in \text{st}(m) = r(0)$. Hence also, $\mathsf{U}(\alpha, \beta) \in \text{st}(m) = \text{end}(n) = r(1)$. As clearly $\beta \in \bigcap_{v \in (0,1]} r(v)$, and r is a run, we obtain $\mathsf{U}(\alpha, \beta) \in \bigcap_{v \in (0,1)} r(v)$. It follows that $\mathsf{U}(\alpha, \beta) \in r^m(t)$.

Now let $t \in (0, 1]$ and $\mathsf{S}(\alpha, \beta) \in \text{sub}_{\mathsf{x}}\varphi$. Most checks are similar to the Until case, except when $t = 1$. We consider only this case. Assume that $\mathsf{S}(\alpha, \beta) \in r^m(1) = \text{end}(m)$. Because $m \equiv n \cdot \omega$, we have $\neg\beta \notin \text{cov}(m)$. As m is a mosaic, $\alpha \in \text{st}(m) \cup \text{cov}(m) = \text{st}(n) \cup \text{cov}(n)$, or $\beta, \mathsf{S}(\alpha, \beta) \in \text{st}(m) = \text{st}(n)$. We know from $\neg\beta \notin \text{cov}(m)$ that $\beta \in r^m(t)$ for all $t < 1$. So $(*)$ $\alpha \in r^m(t)$ for some $t \in [0, 1)$ with $\beta \in r^m(u)$ for all $u \in (t, 1)$, or $\mathsf{S}(\alpha, \beta) \in r^m(0)$ and $\beta \in r^m(u)$ for all $u \in [0, 1)$, as required. Conversely, if $(*)$ holds, then we easily obtain $\neg\beta \notin \text{cov}(m)$ and, because r is a run, $\beta, \mathsf{S}(\alpha, \beta) \in r(1) = \text{st}(m)$. So $\mathsf{S}(\alpha, \beta) \in \text{end}(m)$ as m is a mosaic. This proves the claim.

Now we define $\mathcal{R}^*$.

(1) For each $m \in \mu$, pick $n \in \nu$ with $m \equiv n \cdot \omega$, and $r \in \mathcal{R}$ with $\text{mos}(r) = n$, and put r^m into $\mathcal{R}^*$.
(2) For each $r \in \mathcal{R}$, let $n = \text{mos}(r) \in \nu$, pick $m \in \mu$ with $m \equiv n \cdot \omega$, and put r^m into $\mathcal{R}^*$.

Clearly, $\mathcal{R}^*$ is finite (since $\mathcal{R}$ is). It can be checked that $\mathcal{Q}^* = ([0, 1], (\Sigma^*_t : t \in [0, 1]), \mathcal{R}^*)$ is a finitary pre-quasimodel with $\text{bag}(\mathcal{Q}^*) = \mu$. The proof when $\mu \equiv \nu \cdot \omega^*$ is the mirror image. □

5.3 Shuffles

This is the most interesting step. We 'shuffle' finitely many mosaics and types, and finitely many bags and realisable state candidates, densely together.

Definition 5.5 Let σ be a set of mosaics, let τ be a non-empty set of types for φ, and let m be a mosaic. We write $m \equiv \text{sh}(\sigma, \tau)$ ('shuffle') if

(1) $\text{cov}(m) = \bigcup \tau \cup \bigcup_{n \in \sigma} \Big(\text{st}(n) \cup \text{cov}(n) \cup \text{end}(n) \Big)$,
(2) for each formula $\mathsf{U}(\alpha, \beta) \in \text{sub}_{\mathsf{x}}\varphi$, the following are equivalent:
 (a) $\mathsf{U}(\alpha, \beta) \in \text{st}(m)$,
 (b) $\neg\beta \notin \text{cov}(m)$ and $\neg\mathsf{U}(\alpha, \beta) \notin \text{cov}(m)$,
 (c) $\mathsf{U}(\alpha, \beta) \in \bigcup \tau \cup \bigcup_{n \in \sigma} \text{end}(n)$.
(3) the mirror image condition for Since-formulas.

For example, imagine that we replace every rational q in $(0, 1)$ by a copy I_q of $[0, 1]$ (this process is formalised in the proof of proposition 5.8 below). The resulting flow of time remains isomorphic to $[0, 1]$. Imagine a run r over it, whose restriction r_q to each I_q satisfies $\text{mos}(r_q) \in \sigma$, and each mosaic in σ is equal to $\text{mos}(r_q)$ for a dense set of rationals q; and similarly, each $r(i)$ for irrational $i \in (0, 1)$ is in τ, and each type in τ is the value of $r(i)$ for a dense set of i. Then we will have $\text{mos}(r) \equiv \text{sh}(\sigma, \tau)$.

Definition 5.6 Let $\mathcal{B}$ be a set of bags and $\mathcal{C}$ a non-empty set of realisable state candidates. Let μ be a bag. We write $\mu \equiv \text{sh}(\mathcal{B}, \mathcal{C})$ if

(1) for each $m \in \mu$, there are $\sigma \subseteq \bigcup \mathcal{B}$ with $\sigma \cap \nu \neq \emptyset$ for all $\nu \in \mathcal{B}$, and $\tau \subseteq \bigcup \mathcal{C}$ with $\tau \cap \Sigma \neq \emptyset$ for all $\Sigma \in \mathcal{C}$, such that $m \equiv \text{sh}(\sigma, \tau)$,

(2) for each $X \in \bigcup \mathcal{B} \cup \bigcup \mathcal{C}$, there are $m \in \mu$, $\sigma \subseteq \bigcup \mathcal{B}$ with $\sigma \cap \nu \neq \emptyset$ for all $\nu \in \mathcal{B}$, and $\tau \subseteq \bigcup \mathcal{C}$ with $\tau \cap \Sigma \neq \emptyset$ for all $\Sigma \in \mathcal{C}$, such that $X \in \sigma \cup \tau$ and $m \equiv \text{sh}(\sigma, \tau)$.

Though there is no prior restriction on the sizes of $\mathcal{B}, \mathcal{C}$ in shuffles, we can easily obtain one:

Lemma 5.7 *Let $\mathcal{B}$ be a set of bags, $\mathcal{C}$ a non-empty set of realisable state candidates, and μ a bag. Suppose that $\mu \equiv \text{sh}(\mathcal{B}, \mathcal{C})$. Then there are $\mathcal{B}_0 \subseteq \mathcal{B}$ and non-empty $\mathcal{C}_0 \subseteq \mathcal{C}$ with $|\mathcal{B}_0| \leq \sharp(\varphi)$ and $|\mathcal{C}_0| \leq \natural(\varphi)$, and with $\mu \equiv \text{sh}(\mathcal{B}_0, \mathcal{C}_0)$. (Here, $\sharp, \natural$ are as in lemma 3.20.)*

Proof. $\bigcup \mathcal{B}$ is a set of mosaics, so $|\bigcup \mathcal{B}|$ is bounded by the number of mosaics for φ, which is at most $\sharp(\varphi)$ (see lemma 3.20). So we can choose $\mathcal{B}_0 \subseteq \mathcal{B}$ with $|\mathcal{B}_0| \leq \sharp(\varphi)$ and $\bigcup \mathcal{B}_0 = \bigcup \mathcal{B}$. Similarly, we may choose $\mathcal{C}_0 \subseteq \mathcal{C}$ with $|\mathcal{C}_0| \leq \natural(\varphi)$ and $\bigcup \mathcal{C}_0 = \bigcup \mathcal{C}$. It is plain that $\mathcal{C}_0 \neq \emptyset$ and $\mu \equiv \text{sh}(\mathcal{B}_0, \mathcal{C}_0)$. □

The conditions in definition 5.6 above are what is needed to ensure that — roughly speaking — if $\mu \equiv \text{sh}(\mathcal{B}, \mathcal{C})$ then a pre-quasimodel with bag μ can be obtained by replacing each point in $[0, 1]$ by either a pre-quasimodel with bag in $\mathcal{B}$, or a realisable state candidate from $\mathcal{C}$, always replacing irrationals by realisable state candidates, all to be done densely and bringing in the whole of $\mathcal{B}$ and $\mathcal{C}$. This is the content of the next proposition.

Proposition 5.8 *Let $\mathcal{B}$ be a (possibly empty) set of realisable bags and $\mathcal{C}$ a non-empty set of realisable state candidates. Let μ be a bag with $\mu \equiv \text{sh}(\mathcal{B}, \mathcal{C})$. Then μ is realisable.*

Proof. For each $\nu \in \mathcal{B}$, let

$$\mathcal{Q}_\nu = ([0, 1], (\Sigma_t^\nu : t \in [0, 1]), \mathcal{R}_\nu)$$

be a pre-quasimodel such that $\text{bag}(\mathcal{Q}_\nu) = \nu$. Choose a map $\xi : (0, 1) \to \mathcal{B} \cup \mathcal{C}$ such that $\xi^{-1}(X) = \{t \in (0, 1) : \xi(t) = X\}$ is dense in $(0, 1)$ for each $X \in \mathcal{B} \cup \mathcal{C}$ and countable for each $X \in \mathcal{B}$; we can do this because $\mathcal{C} \neq \emptyset$. Let

$$S = \{t \in (0, 1) : \xi(t) \in \mathcal{B}\}.$$

S is either empty or a countable dense subset of $[0, 1]$. Define a linear order I_t for each

$t \in [0, 1]$, by

$$I_t = \begin{cases} [0, 1], & \text{if } t \in S, \\ [0, 0], & \text{otherwise.} \end{cases}$$

Define I to be the linear order $\sum_{t\in[0,1]} I_t$; formally this is

$$I = \bigcup \Big\{ I_t \times \{t\} : t \in [0, 1] \Big\},$$

endowed with the lexicographic ordering $\langle x, t\rangle < \langle y, u\rangle$ iff $t < u$ or ($t = u$ and $x < y$). I has endpoints and is dense, Dedekind complete, and separable; so (see, e.g., [26, theorem 2.30]) it is isomorphic to $[0, 1]$. So it suffices to provide a finitary pre-quasimodel $\mathcal{Q} = (I, (\Sigma_i : i \in I), \mathcal{R})$ over I with $\text{bag}(\mathcal{Q}) = \mu$.

First we define the realisable state candidates $\Sigma_{\langle x,t\rangle}$ for $\langle x, t\rangle \in I$:

$$\Sigma_{\langle x,t\rangle} = \begin{cases} \text{st}(\mu), & \text{if } t = 0, \\ \text{end}(\mu), & \text{if } t = 1, \\ \Sigma_x^{\xi(t)}, & \text{if } t \in S, \\ \xi(t), & \text{if } t \in (0, 1) \setminus S. \end{cases}$$

We now attempt to define $\mathcal{R}$. This will take some time. First, choose an equivalence relation $\sim$ on $(0, 1)$ refining $\ker(\xi)$ (i.e., with $t \sim u \;\Rightarrow\; \xi(t) = \xi(u)$), such that each $\sim$-class is dense in $(0, 1)$, and with $|\xi^{-1}(X)/{\sim}| = |X|$ for each $X \in \mathcal{B} \cup \mathcal{C}$. Note that there are finitely many $\sim$-classes. Let $\mathcal{Z}$ be the set of all maps ζ defined on $(0, 1)$ and such that

$$\begin{aligned} &\zeta(t) = \zeta(u) \text{ for all } t, u \in (0, 1) \text{ with } t \sim u, \\ &\zeta(t) \in \mathcal{R}_{\xi(t)} \text{ if } t \in S, \\ &\zeta(t) \in \xi(t) \quad \text{if } t \in (0, 1) \setminus S. \end{aligned}$$

Note that $\mathcal{Z}$ is finite. Each $\zeta \in \mathcal{Z}$ picks a type or a run for each $\sim$-class. It induces a choice of type for each interior point of I, and we will extend this to a run over I. For each $\zeta \in \mathcal{Z}$, let

$$\begin{aligned} \sigma_\zeta &= \{\text{mos}(\zeta(t)) : t \in S\} \quad \subseteq \bigcup \mathcal{B}, \\ \tau_\zeta &= \{\zeta(t) : t \in (0, 1) \setminus S\} \subseteq \bigcup \mathcal{C}. \end{aligned}$$

Note that $\tau_\zeta \neq \emptyset$.

Lemma 5.9 *Suppose that* $\sigma \subseteq \bigcup \mathcal{B}$ *with* $\sigma \cap \nu \neq \emptyset$ *for all* $\nu \in \mathcal{B}$*, and* $\tau \subseteq \bigcup \mathcal{C}$ *with* $\tau \cap \Sigma \neq \emptyset$ *for all* $\Sigma \in \mathcal{C}$*.*

(1) Let $t_0 \in S$ *and* $s \in \mathcal{R}_{\xi(t_0)}$ *be such that* $\operatorname{mos}(s) \in \sigma$*. Then there is* $\zeta \in \mathcal{Z}$ *with* $\zeta(t_0) = s$*,* $\sigma_\zeta = \sigma$*, and* $\tau_\zeta = \tau$*.*
(2) Let $t_1 \in (0,1) \setminus S$ *and* $p \in \xi(t_1) \cap \tau$*. Then there is* $\zeta \in \mathcal{Z}$ *with* $\zeta(t_1) = p$*,* $\sigma_\zeta = \sigma$*, and* $\tau_\zeta = \tau$*.*

Proof. For each $\nu \in \mathcal{B}$, we have $|\xi^{-1}(\nu)/\sim| = |\nu| \geq |\sigma \cap \nu| > 0$, so since $\nu = \operatorname{bag}(\mathcal{Q}_\nu) = \{\operatorname{mos}(r) : r \in \mathcal{R}_\nu\}$, we may choose a map $\theta_\nu : \xi^{-1}(\nu)/\sim \to \mathcal{R}_\nu$ such that

$$\Big\{ \operatorname{mos}(\theta_\nu(t/\sim)) : t \in \xi^{-1}(\nu) \Big\} = \sigma \cap \nu. \tag{5.1}$$

In part 1 of the lemma, we may assume that $\theta_{\xi(t_0)}(t_0/\sim) = s$. Similarly, for each $\Sigma \in \mathcal{C}$, we have $|\xi^{-1}(\Sigma)/\sim| = |\Sigma| \geq |\tau \cap \Sigma| > 0$, so we may choose a surjection $\theta_\Sigma : \xi^{-1}(\Sigma)/\sim \to \tau \cap \Sigma$. In part 2 of the lemma, we may assume that $\theta_{\xi(t_1)}(t_1/\sim) = p$.

Now define $\zeta(t) = \theta_{\xi(t)}(t/\sim)$, for each $t \in (0,1)$. Clearly, $\zeta \in \mathcal{Z}$. In part 1 of the lemma, we have $\zeta(t_0) = s$, and in part 2 we have $\zeta(t_1) = p$. By (5.1), for each $\nu \in \mathcal{B}$ we have $\{\operatorname{mos}(\zeta(t)) : t \in (0,1),\ \xi(t) = \nu\} = \sigma \cap \nu$. It follows that

$$\sigma_\zeta = \{\operatorname{mos}(\zeta(t)) : t \in S\} = \bigcup_{\nu \in \mathcal{B}} (\sigma \cap \nu) \quad = \quad \sigma.$$

Similarly, $\tau_\zeta = \tau$. □

Now, for each $\zeta \in \mathcal{Z}$ and $m \in \mu$ with $m \equiv \operatorname{sh}(\sigma_\zeta, \tau_\zeta)$, define $r = r^m_\zeta \in \prod_{i \in I} \Sigma_i$ by

$$r(x,t) = \begin{cases} \operatorname{st}(m), & \text{if } t = 0, \\ \operatorname{end}(m), & \text{if } t = 1, \\ (\zeta(t))(x), & \text{if } t \in S, \\ \zeta(t), & \text{if } t \in (0,1) \setminus S. \end{cases}$$

(Here and below, we write $r(x,t)$ instead of $r(\langle x, t\rangle)$.)

Lemma 5.10 *Let* $m \in \mu$*,* $\zeta \in \mathcal{Z}$*, and suppose that* $m \equiv \operatorname{sh}(\sigma_\zeta, \tau_\zeta)$*. Then* $r = r^m_\zeta$ *is a run over* I*, and* $\operatorname{mos}(r) = m$*.*

Proof. That $\operatorname{mos}(r) = m$ is easily seen. We have $r(0,0) = \operatorname{st}(m)$, $r(0,1) = \operatorname{end}(m)$, and

$$
\begin{aligned}
\bigcup_{\langle x,t\rangle\in I, t\in(0,1)} r(x,t) &= \Big(\bigcup_{t\in(0,1)\setminus S}\zeta(t)\Big)\cup\Big(\bigcup_{t\in S, x\in[0,1]}\zeta(t)(x)\Big)\\
&= \bigcup\tau_\zeta \cup \bigcup_{t\in S}\Big(\mathrm{st}(\mathrm{mos}(\zeta(t)))\cup\mathrm{cov}(\mathrm{mos}(\zeta(t)))\cup\mathrm{end}(\mathrm{mos}(\zeta(t)))\Big)\\
&= \bigcup\tau_\zeta \cup \bigcup_{n\in\sigma_\zeta}\Big(\mathrm{st}(n)\cup\mathrm{cov}(n)\cup\mathrm{end}(n)\Big)\\
&= \mathrm{cov}(m).
\end{aligned}
$$

Now we check that r is a run. As $\mathrm{mos}(r)$ is a mosaic, for any constant c in L, and any $\langle x,t\rangle, \langle y,u\rangle \in I$, if $\mathsf{x} = c \in r(x,t)$ then by lemma 3.14, $\mathsf{x} = c \in r(y,u)$. Next, let $\mathsf{U}(\alpha,\beta) \in \mathrm{sub}_{\mathsf{x}}\varphi$ and $\langle x,t\rangle \in I$ with $t < 1$.

(1) Assume that $\mathsf{U}(\alpha,\beta) \in r(x,t)$, and that there is no 'witness' $\langle y,u\rangle > \langle x,t\rangle$ in I with $u < 1$, $\alpha \in r(y,u)$, and $\beta \in r(z,v)$ for all $\langle z,v\rangle \in (\langle x,t\rangle,\langle y,u\rangle)$. We show first that $\mathsf{U}(\alpha,\beta) \in \mathrm{st}(m)$.
  (a) If $t = 0$, then obviously $\mathsf{U}(\alpha,\beta) \in r(0,0) = \mathrm{st}(m)$.
  (b) Assume now that $t \in S$. We have $\mathsf{U}(\alpha,\beta) \in \zeta(t)(x)$, and there is no witness to it in $\zeta(t)$. As $\zeta(t)$ is a run, we must have $\mathsf{U}(\alpha,\beta) \in \zeta(t)(1) = \mathrm{end}(\mathrm{mos}(\zeta(t))) \subseteq \bigcup_{n\in\sigma_\zeta}\mathrm{end}(n)$. As $m \equiv \mathrm{sh}(\sigma_\zeta,\tau_\zeta)$, we obtain $\mathsf{U}(\alpha,\beta) \in \mathrm{st}(m)$.
  (c) If $t \in (0,1)\setminus S$, then $\mathsf{U}(\alpha,\beta) \in r(x,t) = \zeta(t) \subseteq \bigcup\tau_\zeta$, so by the shuffle conditions we again obtain $\mathsf{U}(\alpha,\beta) \in \mathrm{st}(m)$.

  Since $m \equiv \mathrm{sh}(\sigma_\zeta,\tau_\zeta)$, it now follows from the shuffle conditions that $\neg\beta \notin \mathrm{cov}(m)$. So $\beta \in r(z,v)$ for all $\langle z,v\rangle \in (\langle x,t\rangle,\langle 1,0\rangle)$. By density and lack of a witness, we also have $\alpha \notin \mathrm{cov}(m)$. As m is a mosaic, either $\alpha \in \mathrm{end}(m) = r(0,1)$, or else $\beta, \mathsf{U}(\alpha,\beta) \in \mathrm{end}(m) = r(0,1)$, as required.

(2) Next suppose that there is $\langle y,u\rangle > \langle x,t\rangle$ in I with $\alpha \in r(y,u)$ and $\beta \in r(z,v)$ for all $\langle z,v\rangle \in (\langle x,t\rangle,\langle y,u\rangle)$. We require that $\mathsf{U}(\alpha,\beta) \in r(x,t)$.
  (a) If $t = u$, then $t \in S$, so $\alpha \in \zeta(t)(y)$ and $\beta \in \zeta(t)(z)$ for all $z \in I_t$ with $x < z < y$. As $\zeta(t) \in \mathcal{R}_{\xi(t)}$ is a run, we obtain $\mathsf{U}(\alpha,\beta) \in \zeta(t)(x) = r(x,t)$.
  (b) Assume then that $t < u$. So by density, $\neg\beta \notin \mathrm{cov}(m)$, and also $\alpha \in \mathrm{cov}(m) \cup \mathrm{end}(m)$. As m is a mosaic, we have $\mathsf{U}(\alpha,\beta) \in \mathrm{st}(m)$. By shuffle conditions, this implies that $\neg\mathsf{U}(\alpha,\beta) \notin \mathrm{st}(m) \cup \mathrm{cov}(m) \supseteq r(x,t)$, so that $\mathsf{U}(\alpha,\beta) \in r(x,t)$.

(3) Finally suppose that $\mathsf{U}(\alpha,\beta) \in r(0,1) = \mathrm{end}(m)$ and $\beta \in r(z,v)$ for all $\langle z,v\rangle > \langle x,t\rangle$. Again we obtain $\neg\beta \notin \mathrm{cov}(m)$ by density conditions, and clearly, $\beta \in \mathrm{end}(m)$. As m is a mosaic, $\mathsf{U}(\alpha,\beta) \in \mathrm{st}(m)$. So by shuffle conditions, $\neg\mathsf{U}(\alpha,\beta) \notin \mathrm{cov}(m)$. As $r(x,t) \subseteq \mathrm{st}(m) \cup \mathrm{cov}(m)$, it follows that $\mathsf{U}(\alpha,\beta) \in r(x,t)$ as before.

The proof for Since is a mirror image argument. □

We now define

$$\mathcal{R} = \Big\{ r_\zeta^m : \zeta \in \mathcal{Z},\ m \in \mu,\ m \equiv \mathrm{sh}(\sigma_\zeta, \tau_\zeta) \Big\}.$$

By lemma 5.10, $\mathcal{R}$ is a finite set of runs.

Claim 1. $\{\mathrm{mos}(r) : r \in \mathcal{R}\} = \mu$.

Proof of claim. If $r = r_\zeta^m \in \mathcal{R}$ then by lemma 5.10, $\mathrm{mos}(r) = m \in \mu$. Conversely, let $m \in \mu$. Since $\mu \equiv \mathrm{sh}(\mathcal{B}, \mathcal{C})$, we have $m \equiv \mathrm{sh}(\sigma, \tau)$ for some $\sigma \subseteq \bigcup \mathcal{B}$ with $\sigma \cap \nu \neq \emptyset$ for all $\nu \in \mathcal{B}$, and some $\tau \subseteq \bigcup \mathcal{C}$ with $\tau \cap \Sigma \neq \emptyset$ for all $\Sigma \in \mathcal{C}$. By lemma 5.9, there is $\zeta \in \mathcal{Z}$ with $\sigma_\zeta = \sigma$ and $\tau_\zeta = \tau$. So $m \equiv \mathrm{sh}(\sigma_\zeta, \tau_\zeta)$, so $r_\zeta^m \in \mathcal{R}$, and by lemma 5.10, $\mathrm{mos}(r_\zeta^m) = m$. This proves the claim.

Claim 2. $\Sigma_i = \{r(i) : r \in \mathcal{R}\}$ for each $i \in I$.

Proof of claim. It follows from the definitions that $r(x,t) \in \Sigma_{\langle x,t\rangle}$ for all $\langle x,t\rangle \in I$ and $r_\zeta^m \in \mathcal{R}$. For the converse, take $\langle x,t\rangle \in I$ and a type $p \in \Sigma_{\langle x,t\rangle}$. We have to produce $r \in \mathcal{R}$ with $r(x,t) = p$. There are three cases.

(1) If $t \in \{0,1\}$, then $p = \mathrm{st}(m)$ or $p = \mathrm{end}(m)$ for some $m \in \mu$. By claim 1, there is $r \in \mathcal{R}$ with $\mathrm{mos}(r) = m$, and it follows that $r(x,t) = p$.
(2) If $t \in S$, let $\nu = \xi(t)$. Then $p \in \Sigma_x^\nu$, so as $\mathcal{Q}_\nu$ is a pre-quasimodel, there is $s \in \mathcal{R}_\nu$ with $s(x) = p$. Let $\mathrm{mos}(s) = n \in \nu$. Since $\mu \equiv \mathrm{sh}(\mathcal{B}, \mathcal{C})$, there are $\sigma \subseteq \bigcup \mathcal{B}$, $\tau \subseteq \bigcup \mathcal{C}$, and $m \in \mu$, with $n \in \sigma$, $\sigma \cap \nu' \neq \emptyset$ for each $\nu' \in \mathcal{B}$, $\tau \cap \Sigma \neq \emptyset$ for each $\Sigma \in \mathcal{C}$, and $m \equiv \mathrm{sh}(\sigma, \tau)$. By lemma 5.9, there is $\zeta \in \mathcal{Z}$ with $\zeta(t) = s$, $\sigma_\zeta = \sigma$, and $\tau_\zeta = \tau$. So $r_\zeta^m \in \mathcal{R}$ and $r_\zeta^m(x,t) = \zeta(t)(x) = s(x) = p$.
(3) Now assume that $t \in (0,1) \setminus S$. As $\mu \equiv \mathrm{sh}(\mathcal{B}, \mathcal{C})$, there are $m \in \mu$, $\sigma \subseteq \bigcup \mathcal{B}$ with $\sigma \cap \nu \neq \emptyset$ for all $\nu \in \mathcal{B}$, and $\tau \subseteq \bigcup \mathcal{C}$ with $\tau \cap \Sigma \neq \emptyset$ for all $\Sigma \in \mathcal{C}$, such that $p \in \tau$ and $m \equiv \mathrm{sh}(\sigma, \tau)$. By lemma 5.9, there is $\zeta \in \mathcal{Z}$ with $\zeta(t) = p$, $\sigma_\zeta = \sigma$, and $\tau_\zeta = \tau$. Then $r_\zeta^m \in \mathcal{R}$ and $r_\zeta^m(x,t) = \zeta(t) = p$.

This proves the claim. We can now see that $\mathcal{Q} = (I, (\Sigma_i : i \in I), \mathcal{R})$ is a finitary pre-quasimodel with $\mathrm{bag}(\mathcal{Q}) = m$. Proposition 5.8 is now proved. □

5.4 Soundness and completeness

We now show that our syntactic constructions of bags exactly match pre-quasimodels. We repeat that all bags, pre-quasimodels, realisable state candidates, etc., are implicitly assumed to be *finitary.*

Definition 5.11 Define sets $\mathcal{S}_n$ ($n < \omega$) of bags by induction as follows. Let $\mathcal{S}_0 = \emptyset$. Given $\mathcal{S}_{3n}$ for $n < \omega$, define:

- $\mathcal{S}_{3n+1}$ is the union of $\mathcal{S}_{3n}$ and the set of all bags μ such that $\mu \equiv \mathrm{sh}(\mathcal{B}, \mathcal{C})$ for some $\mathcal{B} \subseteq \mathcal{S}_n$ and some non-empty set $\mathcal{C}$ of realisable state candidates,
- $\mathcal{S}_{3n+2}$ is the union of $\mathcal{S}_{3n+1}$ and the set of all bags μ such that there are bags $\nu_0, \nu_1 \in \mathcal{S}_{3n+1}$ with $\mu \equiv \nu_0 + \nu_1$,
- $\mathcal{S}_{3n+3}$ is the union of $\mathcal{S}_{3n+2}$ and the set of all bags μ such that $\mu \equiv \nu \cdot \omega$ or $\mu \equiv \nu \cdot \omega^*$ for some $\nu \in \mathcal{S}_{3n+2}$.

A bag μ is said to be *constructible* if $\mu \in \bigcup_{n<\omega} \mathcal{S}_n$. That is, μ is obtainable from realisable state candidates by finitely many applications of shuffle, $\cdot\omega$, $\cdot\omega^*$, and $+$ (see definitions 5.3, 5.6, and 3.22).

Proposition 5.12 *A bag is constructible iff it is realisable.*

Proof. A simple induction on n, using lemmas 5.2 and 5.4 and proposition 5.8, shows that every bag in $\mathcal{S}_n$ is realisable.

The converse is similar to proofs in [15,3,20], so we will be brief. Take a finitary pre-quasimodel $\mathcal{Q} = ([0,1], (\Sigma_t : t \in [0,1]), \mathcal{R})$. We wish to show that $\mathrm{bag}(\mathcal{Q})$ is constructible. Define a binary relation $\sim$ on $[0,1]$, by $x \sim y$ iff $x = y$, or $x < y$ and $\mathrm{bag}(\mathcal{Q} \restriction [z,t])$ is constructible for all $x \leq z < t \leq y$, or $x > y$ and $\mathrm{bag}(\mathcal{Q} \restriction [z,t])$ is constructible for all $y \leq z < t \leq x$. It is easily checked that if $x < y < z$ in $[0,1]$ then

$$\mathrm{bag}(\mathcal{Q} \restriction [x,z]) \equiv \mathrm{bag}(\mathcal{Q} \restriction [x,y]) + \mathrm{bag}(\mathcal{Q} \restriction [y,z]). \tag{5.2}$$

It follows that $\sim$ is transitive, and hence an equivalence relation. Clearly, each $\sim$-class E is convex. We claim it is a closed interval in $\mathbb{R}$. Suppose for example that $y = \sup(E)$. We show that $y \in E$. Of course, $y \in [0,1]$, so it suffices to check that for each $x \in E$ with $x < y$, $\mathrm{bag}(\mathcal{Q} \restriction [x,y])$ is constructible. Choose $x < x_0 < x_1 < \cdots$ in E with $\sup\{x_i : i < \omega\} = y$. For a non-singleton closed interval $I \subseteq [0,1]$, let

$$\chi(I) = \langle \mathrm{mos}(r \restriction I) : r \in \mathcal{R} \rangle.$$

Since $\mathcal{R}$ is finite, χ is finitely-valued. So we may suppose by Ramsey's theorem [23] that $\chi([x_i, x_j])$ is constant for all $i < j < \omega$. It can now be checked that $\mathrm{bag}(\mathcal{Q} \restriction [x_0, y]) \equiv \mathrm{bag}(\mathcal{Q} \restriction [x_0, x_1]) \cdot \omega$. As $x_0 \sim x_1$, $\mathrm{bag}(\mathcal{Q} \restriction [x_0, x_1])$ is constructible, so $\mathrm{bag}(\mathcal{Q} \restriction [x_0, y])$ is constructible too. It follows by (5.2) that $\mathrm{bag}(\mathcal{Q} \restriction [x,y])$ is constructible as required. The case of $y = \inf(E)$ is similar, using $\cdot\omega^*$.

Hence, $\mathrm{bag}(\mathcal{Q} \restriction E)$ is constructible for each $\sim$-class E with $|E| > 1$. So it suffices to show that $[0,1]$ is a $\sim$-class.

Assume for contradiction that there are at least two $\sim$-classes. Since $\sim$-classes are closed intervals, the condensation ordering $[0,1]/\sim$ is dense. Extend the aforementioned map χ to singleton intervals by

$$\chi([t,t]) = \langle r(t) : r \in \mathcal{R}\rangle.$$

Choose a non-empty open interval $O \subseteq [0,1]/\sim$ such that $|\{\chi(E) : E \in O\}|$ is least possible. So each value that χ takes on O is its value on a dense set of elements of O. Let

$$\mathcal{B} = \{\text{bag}(\mathcal{Q} \restriction E) : E \in O,\ |E| > 1\},$$
$$\mathcal{C} = \{\Sigma_t : [t,t] \in O\}.$$

Each $\nu \in \mathcal{B}$ is constructible. Standard arguments using separability of $\mathbb{R}$ show that there is a dense set of singleton $\sim$-classes in O, so $\mathcal{C} \neq \emptyset$.

Now take any $x, y \in \bigcup O$ with $x < y$. We claim that $\text{bag}(\mathcal{Q} \restriction [x,y])$ is constructible.

(1) If $x \sim y$, this is clear.
(2) Assume that $x \not\sim y$ and suppose that x is maximal and y minimal in their respective $\sim$-classes. It can be checked that for any $r \in \mathcal{R}$ we have

$$\text{mos}(r \restriction [x,y]) \equiv \text{sh}\Big(\{\text{mos}(r \restriction E) : E \in O,\ |E| > 1\},\ \{r(t) : [t,t] \in O\}\Big).$$

Hence, $\text{bag}(\mathcal{Q} \restriction [x,y]) \equiv \text{sh}(\mathcal{B}, \mathcal{C})$, so is constructible.
(3) Assume that $x \not\sim y$ and suppose that x is maximal in its $\sim$-class but $y > y^- = \inf(y/\sim)$. Then $y \sim y^-$, so $\text{bag}(\mathcal{Q} \restriction [y^-, y])$ is constructible. As above, $\text{bag}(\mathcal{Q} \restriction [x, y^-])$ is constructible. Then by (5.2), $\text{bag}(\mathcal{Q} \restriction [x,y])$ is constructible.
(4) The other two cases are proved by combinations of similar arguments.

This proves the claim. It follows that $x \sim y$ for any $x, y \in \bigcup O$, contradicting that O contains more than one $\sim$-class. $\square$

5.5 *The algorithm*

Theorem 5.13 *The problem of whether a monodic packed sentence is satisfiable in a temporal structure with finite first-order domain and flow of time* $[0,1]$ *is* 2EXPTIME-*complete.*

Proof. The problem is certainly 2EXPTIME-hard, by proposition 2.5. We turn to finding an 2EXPTIME algorithm for satisfiability. Take a monodic packed sentence φ. The following are equivalent:

(1) φ is satisfiable over $[0,1]$,

(2) there exists a finitary quasimodel for φ over $[0,1]$ (by fact 3.10),
(3) there exists a perfect realisable finitary bag (by lemma 3.19),
(4) there exists a perfect constructible finitary bag (by proposition 5.12).

But the latter condition is easy to decide by a 2EXPTIME algorithm. We simply construct the sets $\mathcal{S}_n$ of definition 5.11 by induction, stopping when either a perfect bag is found in $\mathcal{S}_n$, or when $\mathcal{S}_n = \mathcal{S}_{n+3}$ and no perfect bag has been found. Let us briefly check that this can be done in 2EXPTIME. By lemma 5.7, $\mathcal{S}_{3n+1}$ is the union of $\mathcal{S}_{3n}$ and the set of all bags μ such that $\mu \equiv \mathrm{sh}(\mathcal{B}, \mathcal{C})$ for some $\mathcal{B} \subseteq \mathcal{S}_n$ with $|\mathcal{B}| \leq \sharp(\varphi)$, and some non-empty set $\mathcal{C}$ of realisable state candidates with $|\mathcal{C}| \leq \natural(\varphi)$. These bounds are exponential (see lemma 3.20); and since $\mathcal{S}_{3n}$ is a set of bags and so is at most double-exponential in size, there are at most doubly exponentially many sets $\mathcal{B} \subseteq \mathcal{S}_{3n}$ and sets $\mathcal{C}$ of realisable state candidates to consider. Thus, $\mathcal{S}_{3n+1}$ is constructible from $\mathcal{S}_{3n}$ in 2EXPTIME. So are $\mathcal{S}_{3n+2}$ and $\mathcal{S}_{3n+3}$, as is easily seen. So, much as in theorem 4.7, it is plain that the chain of $\mathcal{S}_n$ can be constructed in 2EXPTIME. □

5.6 Corollaries

The following is easily obtainable by the techniques of proposition 4.8 and the expressive completeness of Until and Since over $\mathbb{R}$ ([18]; other proofs are given in [6]). Similar reductions were given in [15, corollary 37]. The case of $(\mathbb{N}, <)$ was first proved in [13].

Theorem 5.14 *The satisfiability problem for monodic packed sentences in temporal structures with finite first-order domains, and with flow of time (in) any of the following, is* 2EXPTIME*-complete:*

(1) $(\mathbb{R}, <)$
(2) $(\mathbb{N}, <)$
(3) $(\mathbb{Z}, <)$
(4) the class of all finite linear flows of time
(5) the class of all linear flows of time
(6) $(\mathbb{Q}, <)$
(7) any given first-order-definable class of linear flows of time.

Proof. By proposition 2.5, we only have to show that the satisfiability problems are solvable in 2EXPTIME. Consider the case of $(\mathbb{R}, <)$. For any monodic packed L-sentence φ, let

$$\varphi^\rho = (q \wedge \varphi^q) \wedge \Diamond\Big(\neg q \wedge \neg \mathsf{S}(\top, \top) \wedge \mathsf{U}(\neg q \wedge \neg \mathsf{U}(\top, \top), q)\Big),$$

where the relativisation φ^q is as in theorem 4.8. If this formula is true at some time in a temporal structure $\mathcal{M} = ([0,1], D, (M_t : t \in [0,1]))$ with finite domain D, then

q is true in $\mathcal{M}$ at just the points in $(0, 1)$, and φ is true at some time point in the restricted structure $((0, 1), D, (M_t : t \in (0, 1)))$. The flow of time of this structure is isomorphic to $(\mathbb{R}, <)$. Conversely, if φ is satisfiable in a temporal structure with flow of time $\mathbb{R}$ and finite domain, then it is true at some time u in a temporal structure of the form $((0, 1), D, (M_t : t \in (0, 1)))$ with finite D, since $((0, 1), <) \cong (\mathbb{R}, <)$. Let N_t be an expansion of M_t making q true (for each $t \in (0, 1)$), and let N_0, N_1 be arbitrary L-structures with domain D and making q false. Then $([0, 1], D, (N_t : t \in [0, 1])), u \models \varphi^\rho$. We conclude that φ has a model with flow of time $(\mathbb{R}, <)$ iff φ^ρ has a model with flow of time $[0, 1]$. So part 1 follows by reduction from theorem 5.13. The remaining cases are covered by reductions given in [15, corollary 37]. □

6 Concluding remarks

We have established 2EXPTIME-completeness of the satisfiability problem for the monodic packed fragment over temporal structures with arbitrary and with finite first-order domains, over a range of (classes of) linear flows of time.

Remark 6.1 Theorems 4.8 and 5.14 also hold for the monodic guarded and monodic loosely guarded fragments, since the first-order guarded and loosely guarded fragments are 2EXPTIME-complete subfragments of the packed fragment.

Various issues remain outstanding:

(1) It may be of interest to extend the results of section 4 to classes of linear flows of time defined by arbitrary first-order sentences.
(2) Though modifications would be needed because equality may have to be omitted, our algorithms may work for other monodic fragments for which the problem of deciding whether a state candidate is realisable is 2EXPTIME-hard. If this problem is, say, 2EXPSPACE-complete, then we would expect a 2EXPSPACE algorithm to decide the monodic fragment. We would also expect our methods to provide 2EXPTIME algorithms to decide the monadic and 2-variable monodic fragments over the same flows of time.
(3) As we said earlier, it would be highly desirable to give (if possible) an EXPSPACE algorithm to decide monodic fragments of which the first-order part is decidable in EXPSPACE, over the flows of time considered here.
(4) Among other important open problems, chief is whether any non-trivial monodic fragment is even decidable in the case of arbitrary first-order domains with flow of time the real numbers.
(5) It may be of interest to devise tableau- or resolution-style algorithms to decide the

problems addressed in this paper.

(6) Complexity results for decidable monodic fragments over branching time also need to be established.

References

[1] H. Andréka, J. van Benthem, and I. Németi, *Modal logics and bounded fragments of predicate logic*, J. Philosophical Logic 27 (1998), 217–274.

[2] S. Bauer, I. Hodkinson, F. Wolter, and M. Zakharyaschev, *On non-local propositional and weak monodic quantified* CTL^*, J. Logic Computat. 14 (2004), 3–22.

[3] J. P. Burgess and Y. Gurevich, *The decision problem for linear temporal logic*, Notre Dame J. Formal Logic 26 (1985), no. 2, 115–128.

[4] A. Degtyarev, M. Fisher, and B. Konev, *Monodic temporal resolution*, Proceedings of CADE'19: International Conference on Automated Deduction (Berlin) (F. Baader, ed.), Lecture Notes in Computer Science, Springer-Verlag, 2003, pp. 397–411.

[5] A. Degtyarev, M. Fisher, and A. Lisitsa, *Equality and monodic first-order temporal logic*, Studia Logica 72 (2002), 147–156.

[6] D. Gabbay, I. Hodkinson, and M. Reynolds, *Temporal logic: mathematical foundations and computational aspects, vol. 1*, Clarendon Press, Oxford, 1994.

[7] D. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev, *Many-dimensional modal logics: theory and applications*, Studies in Logic, vol. 148, North-Holland, Amsterdam, 2003.

[8] E. Grädel, *Decision procedures for guarded logics*, Automated Deduction - CADE16, LNCS, vol. 1632, Springer-Verlag, 1999, Proceedings of 16th International Conference on Automated Deduction, Trento, 1999, pp. 31–51.

[9] ———, *On the restraining power of guards*, J. Symbolic Logic 64 (1999), 1719–1742.

[10] W. Hodges, *Model theory*, Encyclopedia of mathematics and its applications, vol. 42, Cambridge University Press, 1993.

[11] I. Hodkinson, *Loosely guarded fragment of first-order logic has the finite model property*, Studia Logica 70 (2002), 205–240.

[12] ———, *Monodic packed fragment with equality is decidable*, Studia Logica 72 (2002), 185–197.

[13] I. Hodkinson, R. Kontchakov, A. Kurucz, F. Wolter, and M. Zakharyaschev, *On the computational complexity of decidable fragments of first-order linear temporal logics*, Proc. TIME-ICTL (M. Reynolds and A. Sattar, eds.), IEEE, 2003, pp. 91–98.

[14] I. Hodkinson and M. Otto, *Finite conformal hypergraph covers and Gaifman cliques in finite structures*, Bull. Symbolic Logic 9 (2003), 387–405.

[15] I. Hodkinson, F. Wolter, and M. Zakharyaschev, *Decidable fragments of first-order temporal logics*, Ann. Pure Appl. Logic 106 (2000), 85–134.

[16] ______, *Monodic fragments of first-order temporal logics: 2000–2001 A.D.*, Logic for Programming, Artificial Intelligence and Reasoning (R. Nieuwenhuis and A. Voronkov, eds.), LNAI, vol. 2250, Springer-Verlag, 2001, pp. 1–23.

[17] ______, *Decidable and undecidable fragments of first-order branching temporal logics*, Proc. 17th IEEE Symposium on Logic in Computer Science (LICS), IEEE Inc., 2002, pp. 393–402.

[18] H. Kamp, *Tense logic and the theory of linear order*, Ph.D. thesis, University of California, Los Angeles, 1968.

[19] B. Konev, A. Degtyarev, and M. Fisher, *Handling equality in monodic temporal resolution*, Logic for Programming and Automated Reasoning (LPAR 2003) (Berlin) (M. Vardi and A. Voronkov, eds.), Lecture Notes in Artificial Intelligence, vol. 2850, Springer-Verlag, 2003, pp. 211–225.

[20] H. Läuchli and J. Leonard, *On the elementary theory of linear order*, Fundamenta Mathematicae 59 (1966), 109–116.

[21] M. Marx, *Tolerance logic*, J. Logic, Language and Information 10 (2001), 353–374.

[22] V. R. Pratt, *Models of program logics*, Proc. 20th IEEE Symposium on Foundations of Computer Science, San Juan, 1979, pp. 115–122.

[23] F. P. Ramsey, *On a problem of formal logic*, Proc. London Math. Soc. 30 (1930), 264–286.

[24] M. Reynolds, *The complexity of temporal logic over the reals*, (2001), preprint, `www.it.murdoch.edu.au/~mark/research/online/cort.html`

[25] ______, *The complexity of the temporal logic with "until" over general linear time*, J. Comput. Systems Sci. 66 (2003), 393–426.

[26] J. G. Rosenstein, *Linear orderings*, Academic Press, New York, 1982.

[27] A. P. Sistla and E. M. Clarke, *The complexity of propositional linear temporal logics*, J. ACM 32 (1985), 733–749.

[28] F. Wolter and M. Zakharyaschev, *Axiomatizing the monodic fragment of first-order temporal logic*, Ann. Pure. Appl. Logic 118 (2002), 133–145.