Synthesis for Temporal Logic over the Reals

Tim French John M^{c}Cabe-Dansted Mark Reynolds [1]

School of Computer Science and Software Engineering, The University of Western Australia
35 Stirling Highway, Crawley WA 6009
Perth, Australia.

Abstract

We develop the notion of synthesizing, or constructing, a temporal structure over the real numbers flow of time, from a given temporal or first-order specification. We present a new notation for giving a manageable description of the compositional construction of such a model and an efficient procedure for finding it from the specification.

Keywords: Compositional Models, RTL, Continuous Time.

1 Introduction

Standard temporal logics are based on a discrete, natural numbers model of time [8]. However, a dense, continuous or specifically real-numbers model of time may be better for many applications, ranging from philosophical, natural language and AI modelling of human reasoning to computing and engineering applications of concurrency, refinement, open systems, analogue devices and metric information.

The most natural and well-established such temporal logic is RTL, propositional temporal logic over real-numbers time using the Until and Since connectives introduced in [5]. We know from [5] that, as far as defining properties is concerned, this logic is as expressive as the first-order monadic logic of the real numbers order, and so RTL is at least as expressive as any other standard temporal logic which could be defined over real-numbers time.

Reasoning in RTL is fairly well understood: complete, Hilbert-style axioms systems for RTL are given in [4] and [11]. Satisfiability and validity in RTL is decidable [1]. However, the decision procedure in [1] uses Rabin's non-elementarily complex decision procedure for the second-order monadic logic of two successors, and so is far from practical. Furthermore, deciding validity in the equally expressive first-order monadic logic of the real order is a non-elementary problem [17]. More recently, there has been some more pos-

[1] The work was partially supported by the Australian Research Council.

itive news as [15] showed that deciding (validity or satisfiability in) RTL is PSPACE-complete.

Our task here, synthesis, is a harder problem than satisfiability checking. It requires an algorithm which can output a complete description of a specific model of the input formula, whenever the input formula is satisfiable. In this paper we give the first synthesis result for RTL.

Towards this end, we first present a new suitable notation for describing models in a concrete way. The compositional approach presented here was hinted at in [12], and traces back to pioneering work in [6] and [1]. It uses a small number of distinct operations for putting together a larger model from one or more smaller ones, or copies thereof. For example, the *shuffle* construct makes a new linear structure from a dense mixture of copies of a finite number of simpler ones. A good overview of the mathematics of linear orders may be found in [16].

We introduce a formal model expression language for defining a model via these inductive operations. In fact, we first give a language for making general linear structures in this way and then define a restricted sub-language (the real model expression language) capable of specifying structures with the real number flow of time. Having a formal model building language opens up the possibility for workable definitions of such tasks as synthesis and model checking for real-flowed structures. It also allows us to formalise questions of expressibility and to assess the computational complexity of these reasoning tasks.

One of our results here, echoing the earlier work of [6], [1] and others is that the real model expression language is able to describe some real-flowed model of every satisfiable RTL formula.

The major advance of this paper on that previously-known expressiveness result, is that we also present an EXPTIME procedure for finding the real model expression of a model from any given satisfiable RTL formula. EXPTIME is best possible. The real model expression tells us exactly how to make a specific real-flowed model of the formula [2]. This is our synthesis result.

Some of the proofs here use the mosaic techniques for temporal logic developed in [15]. These mosaics are small pieces of a real-flowed structure. In that paper we try to find a finite set of small pieces which is sufficient to be used to build a real-numbers model of a given formula. We showed that if a formula was satisfiable then we could find a sufficient set of mosaics. In this paper, we go further and show how to build a compositional model (i.e. one corresponding to an expression in our model language) when there is such a set of mosaics.

The extension here is built on a series of lemmas mirroring some of the earlier results but keeping a much closer track on a relationship between mosaics and some compositionally built structures which witness them. Thus there is a series of quite long and detailed lemmas needed for this result. In the short

[2] Of course, any isomorphic real-flowed structure will also be a model

version of this paper we only sketch some of the new proofs but the full proofs can be found in [2].

In section 2 we present RTL and monadic logic. In section 3 we introduce the compositional approach to building linear models. In section 4 we remind ourselves of useful properties of mosaics from [15]. In section 5 we use mosaics to give a proof of the previously-known expressiveness result for RTL: satisfiability implies satisfiability in a compositional model. We apply this technique to present the new synthesis result and we conclude in section 6.

2 The logic

Fix a countable set $\mathbf{L}$ of atoms. Here, frames $(T, <)$, or flows of time, will be irreflexive linear orders. Structures $\mathbf{T} = (T, <, h)$ will have a frame $(T, <)$ and a valuation h for the atoms i.e. for each atom $p \in \mathbf{L}$, $h(p) \subseteq T$. Of particular importance will be *real* structures $\mathbf{T} = (\mathbb{R}, <, h)$ which have the real numbers flow (with their usual irreflexive linear ordering). We will also introduce structures over sub-orders of this standard model using $<$ to mean the usual ordering. For example, $(]0, 1[, <, h)$ is some structure over the open unit interval of the reals.

The language $L(U, S)$ is generated by the 2-place connectives U and S along with classical $\neg$ and $\wedge$. That is, we define the set of formulas recursively to contain the atoms and for formulas α and β we include $\neg\alpha$, $\alpha \wedge \beta$, $U(\alpha, \beta)$ and $S(\alpha, \beta)$.

Formulas are evaluated at points in structures $\mathbf{T} = (T, <, h)$. We write $\mathbf{T}, x \models \alpha$ when α is true at the point $x \in T$. This is defined recursively as follows. Suppose that we have defined the truth of formulas α and β at all points of $\mathbf{T}$. Then for all points x:

$\mathbf{T}, x \models p$	iff	$x \in h(p)$, for p atomic;
$\mathbf{T}, x \models \neg\alpha$	iff	$\mathbf{T}, x \not\models \alpha$;
$\mathbf{T}, x \models \alpha \wedge \beta$	iff	both $\mathbf{T}, x \models \alpha$ and $\mathbf{T}, x \models \beta$;
$\mathbf{T}, x \models U(\alpha, \beta)$	iff	there is $y > x$ in T such that $\mathbf{T}, y \models \alpha$ and for all $z \in T$ such that $x < z < y$ we have $\mathbf{T}, z \models \beta$; and
$\mathbf{T}, x \models S(\alpha, \beta)$	iff	there is $y < x$ in T such that $\mathbf{T}, y \models \alpha$ and for all $z \in T$ such that $y < z < x$ we have $\mathbf{T}, z \models \beta$.

The logic is discussed more fully in [13], [15] and [14], for example. See those references for investigations of the "strict" versus "non-strict" connectives, infix versus postfix operators, etc. We use the following abbreviations in illustrating the logic: $F\alpha = U(\alpha, \top)$, "alpha will be true (sometime in the future)"; $G\alpha = \neg F(\neg\alpha)$, "alpha will always hold (in the future)"; and their mirror images P and H. Particularly for dense time applications we also have: $\Gamma^+\alpha = U(\top, \alpha)$, "alpha will be constantly true for a while after now"; and $K^+\alpha = \neg\Gamma^+\neg\alpha$,

"alpha will be true arbitrarily soon". They have mirror images Γ^- and K^-.

2.1 Reasoning with RTL

A formula ϕ is $\mathbb{R}$-*satisfiable* if it has a real model: i.e. there is a real structure $\mathbf{S} = (\mathbb{R}, <, h)$ and $x \in \mathbb{R}$ such that $\mathbf{S}, x \models \phi$. A formula is $\mathbb{R}$-*valid* iff it is true at all points of all real structures. Of course, a formula is $\mathbb{R}$-valid iff its negation is not $\mathbb{R}$-satisfiable. We will refer to the logic of L(U,S) over real structures as RTL.

Let RTL-SAT be the problem of deciding whether a given formula of $L(U, S)$ is $\mathbb{R}$-satisfiable or not. The main result of [15] is:

Theorem 2.1 *RTL-SAT is PSPACE-complete.*

In order to help get a feel for the sorts of formulas which are valid in RTL it is worth considering a few formulas in the language. $U(\top, \bot)$ is a formula which only holds at a point with a discrete successor point so $G\neg U(\top, \bot)$ is valid in RTL. $Fp \to FFp$ is a formula which can be used as an axiom for density and is also a valid in RTL.

$(\Gamma^+ p \wedge F\neg p) \to U(\neg p \vee K^+(\neg p), p)$ was used as an axiom for Dedekind completeness (in [11]) and is valid. Recall that a linear order is Dedekind complete if and only if each non-empty subset which has an upper bound has a least upper bound. The formula says that if p is true constantly for a while but not forever then there is an upper bound on the interval in which it remains true. This formula is not valid in the temporal logic with until and since over the rational numbers flow of time.

One of the most interesting valid formulas of RTL is Hodkinson's axiom "Sep" (see [11]). It is

$$K^+ p \wedge \neg K^+(p \wedge U(p, \neg p)) \to K^+(K^+ p \wedge K^- p).$$

This can be used in an axiomatic completeness proof to enforce the *separability* of the linear order:

Definition 2.2 *A linear order is separable iff it has a countable suborder which is spread densely throughout the order: i.e. between every two elements of the order lies an element of the suborder.*

The fact that the rationals are dense in them shows that the reals are separable. There are dense, Dedekind complete linear orders with end points which are not separable (e.g. , see [11]). The negation of Sep will be satisfiable over them but not over the reals.

As we have noted in the introduction, there are complete axiom systems for RTL in [4] and in [11]: the former using a special rule of inference and the latter just using orthodox rules.

Rabin's decision procedure for the second-order monadic logic of two successors [10] is used in [1] to show that RTL is decidable. One of the two decision procedures in that paper just gives us a non-elementary upper bound on the complexity of RTL-SAT.

2.2 Monadic Logic

The first-order monadic language of order, FOMLO, is a first order language which can describe the structures we are dealing with and it is useful to translate between it and the temporal language.

The relation symbols of FOMLO are 2-ary $<$ and 1-ary, or monadic, P_0, P_1, P_2, $\dots$ each corresponding respectively to the atoms $p_0, p_1, p_2, \dots$ of L. So atomic propositions are $x_i < x_j$ and $P_k(x_j)$ for each variable symbol x_i and each 1-ary relation symbol P_k. Formulas of the language are built up from the atoms as follows: $\neg\alpha$, $\alpha \wedge \beta$, and $\forall x_i \alpha$.

The notions of free and bound variables and sentences are as usual.

Given a temporal structure $(T, <, g)$ we can evaluate monadic formulas in it by interpreting the 1-ary predicates P_i as 1-ary relations on (i.e. subsets of) T using the valuation $g(p_i)$ to tell us where the interpretation of P_i holds as follows:

$(T, <, g), \mu \models P_i(x_j)$ iff $t_j \in g(p_i)$
$(T, <, g), \mu \models x_i < x_j$ iff $t_i < t_j$
$(T, <, g), \mu \models \neg\alpha$ iff it is not the case that $(T, <, g), \mu \models \alpha$
$(T, <, g), \mu \models \alpha_1 \wedge \alpha_2$ iff $(T, <, g), \mu \models \alpha_1$ and $(T, <, g), \mu \models \alpha_2$
$(T, <, g), \mu \models \forall x_i \alpha$ iff for every $d \in T$, $(T, <, g), \mu[x_i \mapsto d] \models \alpha$

Here μ is a (possibly partial) map from $\{x_1, x_2, \dots\}$ to T and $\mu[x_i \mapsto d]$ is the map which is the same as μ except that x_i is mapped to d. We require that μ is defined on all the free variables of α. The truth of $(T, <, g), \mu \models \alpha$ does not depend on the value of $\mu(x_i)$ if x_i is not free in α.

Definition 2.3 *We say that the temporal language $L(B)$ is* expressively complete *over class K of linear orders iff for every FOMLO formula $\alpha(t)$, there is some ϕ of the temporal language such that ϕ is equivalent to α over K.*

Kamp showed in [5] that $L(U, S)$ is expressively complete over $\mathbb{R}$ and over $\mathbb{N}$.

2.3 Isomorphisms

An isomorphism is a bijective mapping from one structure to another that preserves the temporal relation $<$ and the valuation h. This is an important notion of equivalence for us, as we will show that equivalent structures satisfy the same set of formulas in $L(U, S)$.

Definition 2.4 *We say two structures $\mathbf{T} = (T, <, h)$ and $\mathbf{T}' = (T', <', h')$ are isomorphic (written $\mathbf{T} \cong \mathbf{T}'$) if and only if there is a bijection $f : T \longrightarrow T'$ where for all $x, y \in T$ $x < y$ if and only if $f(x) <' f(y)$, and for all $p \in P$ $x \in h(p)$ if and only if $f(x) \in h'(p)$.*

It is well known that isomorphisms between structures preserve the truth of formulas of temporal logic.

Lemma 2.5 *Suppose that $\mathbf{T} = (T, <, h)$, $\mathbf{T}' = (T', <', h')$ and $\mathbf{T} \cong \mathbf{T}'$, via the bijection f. Then for any $\alpha \in L(U, S)$, for any $t \in T$, $\mathbf{T}, t \models \alpha$ if and only if $\mathbf{T}', f(t) \models \alpha$.*

3 Building Structures

We introduce a notation which allows the description of a temporal structure in terms of simple basic structures via a small number of ways of putting structures together to form larger ones.

The general idea is simple: using singleton structures (the flow of time is one point), we build up to more complex structures by the recursive application of four operations. They are:

- concatenation or sum of two structures, consisting of one followed by the other;
- ω repeats of some structure laid end to end towards the future;
- ω repeats laid end to end towards the past;
- and making a densely thorough *shuffle* of copies from a finite set of structures.

These operations are well-known from the study of linear orders (see, for example, [6,16,1]).

Model Expressions are an abstract syntax for defining models that are constructed using the follow set of primitive operators:

$$\mathcal{I} ::= \; a \mid \mathcal{I} + \mathcal{J} \mid \overleftarrow{\mathcal{I}} \mid \overrightarrow{\mathcal{I}} \mid \langle \mathcal{I}_0, \ldots, \mathcal{I}_n \rangle$$

where $a \in \Sigma$, where Σ is some alphabet [3]. We refer to these operators, respectively, as *a letter*, *concatenation*, *lead*, *trail*, and *shuffle*.

Definition 3.1 (Correspondence) *Given* $\Sigma = 2^P$*, a model expression* $\mathcal{I}$ corresponds *to a structure as follows. A letter* a *corresponds to any single point model* $(\{x\}, <, h)$ *where* $<$ *is the empty relation and* $h(p) = \{x\}$ *if and only if* $p \in a$*. For the inductive cases we require the notion of an isomorphism (Definition 2.4). Then:*

- $\mathcal{I} + \mathcal{J}$ *corresponds to a structure* $(T, <, h)$ *if and only if* T *is the disjoint union of two sets* U *and* V *where* $\forall u \in U$*,* $\forall v \in V$*,* $v < V$ *and* $\mathcal{I}$ *corresponds to* $(U, <^U, h^U)$ *and* $\mathcal{J}$ *corresponds to* $(V, <^V, h^V)$*. (*$<^U, h^U$ *refers to the restriction of the relations* $<$ *and* h *to apply only to elements of* U*).*
- $\overleftarrow{\mathcal{I}}$ *corresponds to the structure* $(T, <, h)$ *if and only if* T *is the disjoint union of sets* $\{U_i | i \in \omega\}$ *where for all* i*, for all* $u \in U_i$*, for all* $v \in U_{i+1}$*,* $v < u$*, and* $\mathcal{I}$ *corresponds to* $(U_i, <^{U_i}, h^{U_i})$*.*
- $\overrightarrow{\mathcal{I}}$ *corresponds to the structure* $(T, <, h)$ *if and only if* T *is the disjoint union of sets* $\{U_i | i \in \omega\}$ *where for all* i*, for all* $u \in U_i$*, for all* $v \in U_{i+1}$*,* $u < v$*, and* $\mathcal{I}$ *corresponds to* $(U_i, <^{U_i}, h^{U_i})$*.*
- $\langle \mathcal{I}_0, \ldots, \mathcal{I}_n \rangle$ *corresponds to the structure* $(T, <, h)$ *if and only if* T *is the disjoint union of sets* $\{U_i | i \in \mathbb{Q}\}$ *where*

(i) *for all* $i \in \mathbb{Q}$ $(U_i, <^{U_i}, h^{U_i})$ *corresponds to some* $\mathcal{I}_j$ *for* $j \leq n$*,*

---

[3] Typically, we will let $\Sigma = \wp(\mathbf{L})$ so the letter indicates the atoms true at a point

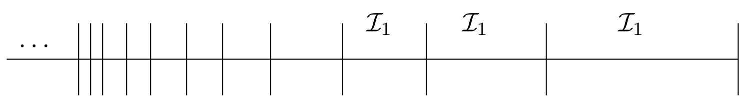

Fig. 1. The lead operation, where $\mathcal{I} = \overleftarrow{\mathcal{I}_1}$

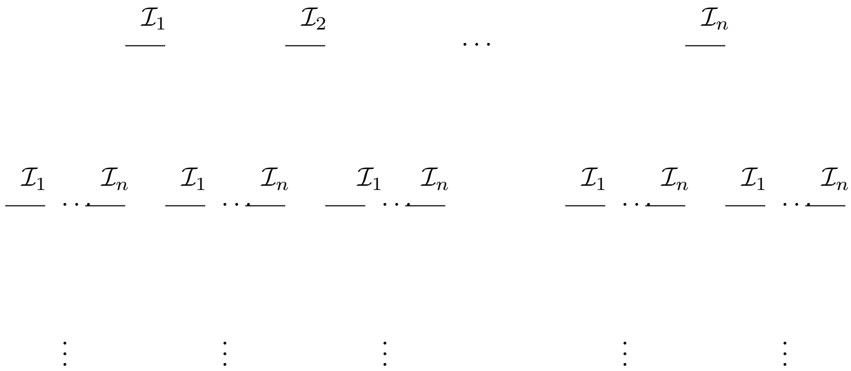

Fig. 2. The shuffle operation, where $\mathcal{I} = \langle \mathcal{I}_1, \ldots, \mathcal{I}_n \rangle$

(ii) *for every* $j \leq n$*, for every* $a \neq b \in \mathbb{Q}$*, there is some* $k \in (a, b)$ *where* $\mathcal{I}_j$ *corresponds to* $(U_k, <^{U_k}, h^{U_k})$*,*

(iii) *for every* $a < b \in \mathbb{Q}$ *for all* $u \in U_a$*, for all* $v \in U_b$*,* $u < v$*.*

We will give an illustration of the non-trivial operations below. The *lead* operation, $\mathcal{I} = \overleftarrow{\mathcal{I}_1}$ corresponds to ω submodels, each corresponding to $\mathcal{I}$, and each preceding the last, as illustrated in Figure 1.

The *trail* operator is the mirror image of the *lead* operation, whereby $\mathcal{I} = \overrightarrow{\mathcal{I}_1}$ corresponds to ω structures, each corresponding to $\mathcal{I}_1$ and each proceeding the earlier structures.

The *shuffle* operator is harder to represent with a diagram. The model expression $\mathcal{I} = \langle \mathcal{I}_1, \ldots \mathcal{I}_n \rangle$ corresponds to a dense, thorough mixture of intervals corresponding to $\mathcal{I}_1$, ..., $\mathcal{I}_n$, without endpoints.

Model expressions give us a grammar that corresponds to structures over general linear frames in a similar manner to the way regular expressions corresponds to words over a given alphabet. Our particular interest in this paper is for frames that are isomorphic to the real numbers so we are required to identify a sublanguage of model expressions. However, the recursive definition of correspondence given in Definition 3.1 is restricted to countable frames. To address this we:

(i) define a Dedekind closure of a structure.

(ii) show that there is a sublanguage of model expressions that correspond to dense, separable structures without endpoints, which agree with their Dedekind closures on the interpretation of $L(U, S)$ formulas.

As every real valued structure is isomorphic to a dense, separable, Dedekind complete structure without end-points, and vice-versa (see [11]) this is sufficient to justify the use of (some) model expressions as the base artefact for synthesis and model checking results.

To define a sublanguage of separable, dense structures without endpoints, we must address the fact that some of the operators of model expressions, such as concatenation, naturally imply a discrete gap in the linear order. We build *real model expressions* via induction using the definitions above:

$$\mathcal{R} ::= \langle a_0, \ldots, a_m, \mathcal{R}_1, \ldots, \mathcal{R}_n \rangle \mid \mathcal{R}_0 + a + \mathcal{R}_1 | \overleftarrow{a + \mathcal{R}} \mid \overrightarrow{\mathcal{R} + a}$$

where $a, a_i \in \Sigma$, and $m, n \geq 0$. The letter a_0 is used as a sort of background filler to ensure that the shuffle is Dedekind complete. The abstract syntax for real model expressions is a direct sub-language for the abstract syntax for general model expressions. We note that their syntax will always define open intervals and that the base element of this recursion is a shuffle containing only points. This will define a dense, separable linear order with all the letters homogeneously distributed across the linear order.

Such a sublanguage is suggested in [1] where similar refinements of the [6] operations were applied to provide a decidability result for the monadic theory of the reals. The following lemmas are implicit in that work, but we include them here for completeness.

Lemma 3.2 *Every real model expression corresponds to some structure whose frame is dense, separable and without end-points.*

It is important to note that the corresponding structures are *not* based on a real frame. In fact, any structure corresponding to a model expression $\mathcal{I}$ must be countable and therefore cannot be isomorphic to the reals. However, real model expressions are sufficient for our purposes as the set of formulas satisfiable over the reals is exactly the set of formulas satisfiable over dense, separable, Dedekind complete linear orders without endpoints [3]. As real model expressions correspond to dense, separable linear orders without endpoints, we must show that we can take a further step to a related Dedekind complete order without affecting the interpretation of $L(U, S)$ formulas.

To address this we define a *Dedekind* closure of a structure, and show that any model corresponding to a real model expression agrees with its Dedekind closures on the interpretation of $L(U, S)$ formulae.

Definition 3.3 *Given a structure* $\mathbf{T} = (T, <, h)$*, we say a* Dedekind gap *is pair of sequences in* T*,* $\ell_0, \ell_1, \ldots$ *and* $u_0, u_1, \ldots$*, where for all* i*,* $\ell_i < \ell_{i+1}$*,* $u_i > u_{i+1}$ *and for all* j*,* $\ell_i < u_j$*, and there is no point* $x \in T$ *where for all* i*,* $\ell_i < x < u_i$*.*
Given a point $x \in T$*, we say a* context of x *is the triple* (a, A, B) *where* $a \in \Sigma$ *and* $A, B \subseteq \Sigma$ *where:*

(i) $x \in h(a)$,

(ii) $A = \{b \mid \forall t < x \exists u, t < u < x, \text{ and } u = h(b)\}$

(iii) $B = \{b \mid \forall t > x \exists u, x < u < t, \text{ and } u = h(b)\}$.

A Dedekind gap is curable *if there is a context* (a, A, A) *such that for all* i *there are points* x_i *and* y_i *where for all* j, $\ell_i < x_i < u_j$, $\ell_j < y_i < u_i$, *and the context of* x_i *and* y_i *is* (a, A, A).

If every Dedekind gap in $\mathbf{T}$ *is curable, the* Dedekind closure *of* $\mathbf{T}$ *is the structure* $\mathbf{T}^* = (T \cup X, <^*, h^*)$ *where:*

(i) X *is a set of new points, one for each Dedekind gap of* $\mathbf{T}$,

(ii) $<^*$ *is the extension of* $<$ *such that if* x *is a new point corresponding to a gap defined by* $\ell_0, \ell_1, \ldots$ *and* $u_0, u_1, \ldots$ *then for all* i, $\ell_i <^* x$ *and* $x <^* u_i$.

(iii) h^* *is the extension of* h *such that for each new point* x, $x \in h^*(a)$ *where* (a, A, A) *is some context for* x.

Note that not every structure has a Dedekind closure. However, we have defined real model expressions in such a way that they guarantee that every Dedekind gap will be curable, and furthermore, the cure will not affect the interpretation of any formula.

Lemma 3.4 *Every structure corresponding to a real model expression agrees with its Dedekind closures on the interpretation of* $L(U, S)$ *formulae.*

Proof. (Sketch)

We show a structure corresponding to a real model expression agrees with its Dedekind closure on the interpretation of $L(U, S)$ formulae by induction over the complexity of formulae. We must show that the addition of the new points in a Dedekind closure does not affect to interpretation of any $L(U, S)$ formula at any of the original points of the structure. This is clearly only relevant in the case of the U and S operators.

Suppose $\mathbf{T} = (T, <, h)$ is a structure and $\mathbf{T}^* = (T \dot{\cup} X, <^*, h^*)$ is its Dedekind closure. Let $U(\alpha, \beta)$ be given, and suppose that for all $x \in T$ $\mathbf{T}^*, x \models \alpha$ if and only if $\mathbf{T}, x \models \alpha$, and for all $x \in X$, $\mathbf{T}^*, x \models \alpha$ if and only if both for all $t < x$ there is some $\ell \in T$ where $t < l < x$ and $\mathbf{T}, \ell \models \alpha$ and for all $t > x$ there is some $u \in T$ where $x < u < t$ and $\mathbf{T}, u \models \alpha$, (and the same conditions hold for β). Where α and β are propositional, these conditions follow from the definition of Dedekind closure.

Suppose that $x \in T$. If $\mathbf{T}^*, x \models U(\alpha, \beta)$, then there is some point $y > x$ and $\mathbf{T}^*, y \models \alpha$, and for all z where $x < z < y$, $\mathbf{T}^*, z \models \beta$. Therefore, there must be some point $y' \in T$ where $x < y' < z$ where $\mathbf{T}, y' \models \alpha$, and for every point, $t \in T$, where $x < t < y'$ where must have $\mathbf{T}, t \models \beta$ by the induction hypothesis. Conversely, suppose that $\mathbf{T}, x \models U(\alpha, \beta)$. Therefore there is some $y \in T$ where $y > x$ and $\mathbf{T}, y \models \alpha$, and for all $z \in T$ where $x < z < y$, $\mathbf{T}, z \models \beta$. By the induction hypothesis, we have $\mathbf{T}^*, y \models \alpha$ and for all $z \in T \cup X$ where $x <^* z <^* y$, we have $\mathbf{T}^*, z \models \beta$, so $\mathbf{T}^*, x \models U(\alpha, \beta)$.

Suppose now that $x \notin T$. If $\mathbf{T}^*, x \models U(\alpha, \beta)$, then there is some point

$y > x$ and $\mathbf{T}^*, y \models \alpha$, and for all z where $x < z < y$, $\mathbf{T}^*, z \models \beta$. By the induction hypothesis there must be some $x' \in T$ where $x' <^* x$ and for all z where $x' <^* z < y$, $\mathbf{T}^*, z \models \beta$. Hence there must be some $y' \in T$, where $y' <^* y$, $\mathbf{T}, y \models \alpha$ and for all $z \in T$ where $x' \leq z' < y'$, $\mathbf{T}, z \models U(\alpha, \beta)$. Conversely, if for all $t < x$ there is some $\ell \in T$ where $\mathbf{T}, \ell \models U(\alpha, \beta)$ then there must be some $y > x$ where $\mathbf{T}, y \models \alpha$ and for all $z \in T$, $\ell < z < y$, $\mathbf{T}, z \models \beta$. By the induction hypothesis, $\mathbf{T}^*, y \models \alpha$, and for all $z \in T \cup X$, $\ell < z < y$, $\mathbf{T}^*, z \models \beta$ so we must have $\mathbf{T}^*, x \models U(\alpha, \beta)$.

As the case for S is symmetric, we can see that all points $x \in T$ maintain their interpretation of $L(U, S)$ formulas in the Dedekind closure of $\mathbf{T}$, as required. □

Finally we must show that every real model expression corresponds to some real valued structure.

Lemma 3.5 *Every structure corresponding to a real model expression is dense, separable, without endpoints and agrees with its Dedekind closure on the interpretation of $L(U, S)$ formulae.*

Proof. We can see every structure, $\mathbf{T}$, corresponding to a real model expression has a Dedekind closure by construction. Every concatenation, lead and trail operation in a real model expression explicitly includes a single point between the two sub-expressions, so the only place a Dedekind gap may occur is in the shuffle operation. As every shuffle must include at least one single point structure, and the shuffle is dense, then there is a dense set of points in a structure corresponding to the shuffle, where each point has a consistent context. These points can be used to cure all Dedekind defects in $\mathbf{T}$ without affecting the interpretation of any $L(U, S)$ formulae. From Lemma 3.2 we have that $\mathbf{T}$ is dense separable and without end-points so the result follows. □

It is straightforward to make the notation completely formal in the case of a finite set of atoms, and this is the case when we are considering a particular temporal formula. For example, let $[p, \neg q]$ represent a singleton structure with the obvious valuation. We might then suggest $\langle\langle([p, q])\rangle\rangle + [p, q] + \langle\langle([p, q], [p, \neg q])\rangle\rangle + [p, q] + \langle\langle([p, q])\rangle\rangle$, as a model expression for $Gp \wedge U(q, \neg U(q, \neg q) \wedge \neg U(q, q))$.

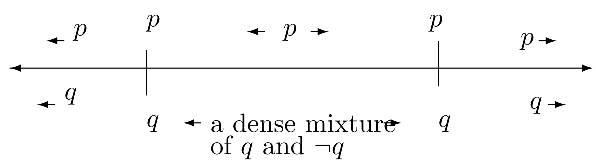

Fig. 3. Diagram representing: $\langle\langle([p, q])\rangle\rangle + [p, q] + \langle\langle([p, q], [p, \neg q])\rangle\rangle + [p, q] + \langle\langle([p, q])\rangle\rangle$

Definition 3.6 *We say that a real-flowed structure* $(\mathbb{R}, <, h)$ *is a* compositional real structure *(or* model*) iff it is isomorphic to the Dedekind closure of a structure which corresponds to some real model expression. In that case, we say that it* realizes *the expression.*

Thus, a compositional real structure is real-flowed by definition. Note also that the real model expression tells us exactly what the model looks like (up to isomorphism).

One of our main results in this paper is that an RTL formula has a real-flowed model iff it has a compositional real model. In the next two sections we briefly describe the proof which uses the mosaic technique for temporal logics.

Before we do so it may be worth noting that there is a similar sort of result in [1] where it is shown that a RTL formula has a real-flowed model iff it has a model with the valuation of each atom being a Borel set, i.e. one obtained from open sets by iterated application of complementation and countable union.

The second half of that paper [1] presents a series of operations corresponding to those of the real model expressions, to show the decidability of the monadic theory of the reals.

The important advantages of our new result are that we provide an explicit notation that is adequate for representing real structures, we are able to give a finite representation in this notation for a model that supports a given satisfiable $L(U, S)$ formula, and we are able to give an efficient means for finding it.

4 Mosaics for U and S

Much of the hard work for us is done by a theorem (4.15 below) from [15] which, unfortunately, must come after a large host of definitions. Our new work also needs these definitions. Due to space considerations the definitions (from that 40 page paper) are selectively presented and only sketched.

In [15], we decided the satisfiability of formulas by considering sets of small pieces of real structures. The idea is based on the mosaics seen in [7] and applied to modal logics. Satisfiability can be decided by checking to see if there exists a finite set of mosaics sufficient to build a model of the formula.

For us, a mosaic is a small piece of a model consisting of three sets of formulas representing those true at each of two points (called the start and end of the mosaic) and those true at all points in-between (called the cover of the mosaic). There are *coherence* conditions on the mosaic which are necessary for it to be part of a model. Note that in the context of a particular formula, ϕ say, (whose satisfiability we might be investigating) we can limit our attention to a finite closure set of formulas and so make these mosaics finite in size. The set of subformulas of ϕ and their negations are a sufficient closure set and will be denoted $\mathrm{Cl}\phi$.

Definition 4.1 *Suppose ϕ is from $L(U, S)$. A ϕ-*mosaic *is a triple (A, B, C) of subsets of the closure set of ϕ such that A and C are maximally propositionally consistent, and B is closed under adding or removing double negations (within the closure set) and the following four* coherency *conditions hold:*

C1. *if* $\neg U(\alpha, \beta) \in A$ *and* $\beta \in B$ *then we have:*
C1.1. $\neg\alpha \in C$;
C1.2. either $\neg\beta \in C$ *or* $\neg U(\alpha, \beta) \in C$ *(or both);*
C1.3. $\neg\alpha \in B$; *and*
C1.4. $\neg U(\alpha, \beta) \in B$.
C2. *if* $U(\alpha, \beta) \in A$ *and* $\neg\alpha \in B$ *then we have:*
C2.1. $\alpha \notin C$ *implies both* $\beta \in C$ *and* $U(\alpha, \beta) \in C$;
C2.2. $\beta \in B$; *and*
C2.3. $U(\alpha, \beta) \in B$.
C3-4 *mirror images of C1-C2.*

Conceptually, A represents the start of the mosaic, C represents the end, and B represents all the points between. We want to show the equivalence of the existence of a model to the existence of a certain set of mosaics: enough mosaics to build a whole model. So the whole set of mosaics also has to obey some conditions. Such conditions are often called *saturation* conditions.

We say that two mosaics compose if the end of the first is the same set as the start of the second. We can then define their composition as the mosaic corresponding to the first point of the first mosaic through to the second point of the second mosaic. Composition is associative.

Definition 4.2 *We say that* ϕ*-mosaics* (A', B', C') *and* (A'', B'', C'') compose *iff* $C' = A''$. *In that case, their* composition *is* $(A', B' \cap C' \cap B'', C'')$.

4.1 Defects

Most of the hard work of finding a saturated set of mosaics is done by breaking up, or decomposing, mosaics into composing sequences of other mosaics. We use a notion of a full decomposition which means that the sequence includes witnesses to all the appropriate formulas in the starts and end of the decomposed mosaic. Eg, if $U(\alpha, \beta)$ is in the start of a mosaic but β is not in its cover then in any full decomposition of the mosaic we should find an initial sequence of mosaics with β in their covers and ends followed by a mosaic with β in its cover and α in its end.

Definition 4.3 *A* defect *in a mosaic* (A, B, C) *is either*

1. *a formula* $U(\alpha, \beta) \in A$ *with either*
   1.1 $\beta \notin B$,
   1.2 ($\alpha \notin C$ *and* $\beta \notin C$), *or*
   1.3 ($\alpha \notin C$ *and* $U(\alpha, \beta) \notin C$);
2. *mirror for* $S(\alpha, \beta) \in C$; *or*
3. *a formula* $\beta \in \mathrm{Cl}\phi$ *with* $\neg\beta \notin B$.

We refer to defects of type 1 to 3 (as listed here). Note that the same formula may be both a type 1 or 2 defect and a type 3 defect in the same mosaic. In that case we count it as two separate defects.

We can talk of sequences of mosaics composing and then find their composition. We define the composition of a sequence of length one to be just the mosaic itself. We leave the composition of an empty sequence undefined.

Definition 4.4 *A* decomposition *for a mosaic* (A, B, C) *is any finite sequence of mosaics* $(A_1, B_1, C_1), (A_2, B_2, C_2), \ldots, (A_n, B_n, C_n)$ *which composes to* (A, B, C).

It will be useful to introduce an idea of fullness of decompositions. This is intended to be a decomposition which provides witnesses to the cure of every defect in the decomposed mosaic.

Definition 4.5 *The decomposition above is* full *iff the following three conditions all hold:*

1. *for all* $U(\alpha, \beta) \in A$ *we have*
   1.1. $\beta \in B$ *and either*
   ($\beta \in C$ *and* $U(\alpha, \beta) \in C$) *or* $\alpha \in C$*,*
   1.2. *or there is some* i *with* $1 \leq i < n$ *such that*
   $\alpha \in C_i$*,* $\beta \in B_j$ *(all* $j \leq i$*)*
   and $\beta \in C_j$ *(all* $j < i$*);*
2. *the mirror image of 1.; and*
3. *for each* $\beta \in \mathrm{Cl}\phi$ *such that* $\neg\beta \notin B$ *there is some* i
   such that $1 \leq i < n$
   and $\beta \in C_i$.

If 1.2 above holds in the case that $U(\alpha, \beta) \in A$ is a type 1 defect in (A, B, C) then we say that *a cure for the defect is witnessed* (in the decomposition) by the end of (A_i, B_i, C_i) (or equivalently by the start of $(A_{i+1}, B_{i+1}, C_{i+1})$). Similarly for the mirror image for $S(\alpha, \beta) \in C$. If $\beta \in C_i$ is a type 3 defect in (A, B, C) then we also say that *a cure for this defect is witnessed* (in the decomposition) by the end of (A_i, B_i, C_i). If a cure for any defect is witnessed then we say that the defect is cured.

Lemma 4.6 *If* $m_1, \ldots, m_n$ *is a full decomposition of* m *then every defect in* m *is cured in the decomposition.*

4.2 Tactics

We would like to introduce an idea of mosaics being fully decomposed in terms of simpler ones. However, sometimes it is allowed that mosaics are decomposed in terms of themselves or equally complicated mosaics. Some recursion is allowed. In order to specify which types of recursion are allowed, we introduce several "tactics" as a sort of meta-level description of how mosaics can be decomposed. There are three tactics with the familiar names of lead, trail and shuffle.

We will see in the next section that by repeated use of a particular tactic to decompose a mosaic we end up showing more or less that it has a model built compositionally via the construction technique with the same name. For example, a mosaic which can be decomposed via a lead tactic into simpler mosaics has a model which is built from simpler structures via a lead construction.

We shall write $\langle p_1, \ldots, p_n \rangle$ for the sequence of mosaics containing $p_1, \ldots, p_n$ in that order. We shall write $\pi^\wedge\rho$ for the sequence resulting from the concatenation of sequences π and ρ in that order. Sequences will always be finite.

Definition 4.7 *We say that m is fully decomposed by the* tactic lead σ, *for some sequence σ of mosaics iff $\langle m \rangle ^\wedge \sigma$ is a full decomposition of m.*

The trail σ tactic is mirror.

4.3 Shuffles

The term shuffle has been used in the literature (see, for example, [6] or [1]) to refer to a certain method of constructing a linear structure (often a monadic one) from a thorough mixture of smaller linear structures. A formal definition was given in Section 3.

The intention here is similar except now we need to deal with mosaics corresponding to linear structures. We consider a shuffle S of linear structures $U_0, U_1, \ldots, U_s, V_1, V_2, \ldots, V_r$ where each U_i is a singleton structure and each V_i is a non-singleton structure consisting of a finite sequence of other structures. Thus, we actually only consider an MPC set P_i instead of U_i and a non-empty composing sequence λ_i of mosaics instead of V_i. In this case it is possible to construct a certain set $\{o, m', m'', x_0, \ldots, x_r, y_0, \ldots, y_s\}$ of mosaics such that one, o, corresponds to S and each one in the set has a full decomposition in terms of others in the set and/or the mosaics which decompose each λ_i.

The mosaic m' corresponds to any proper initial interval of S ending at a copy of U_0 while m'' corresponds to any proper final interval of S beginning at a copy of U_0. Each x_i is satisfied by any interior interval of S starting at the end of a copy of V_i (or a copy of U_0 if $i = 0$) and ending at the start of a copy of V_{i+1} (or a copy of U_0 if $i = r$). There are a lot of intervals of this form, for each i, but each one satisfies x_i. Each y_i is satisfied by any interior interval of S starting at a copy of U_i and ending at a copy of U_{i+1} (or a copy of U_0 if $i = s$).

It can be shown that these mosaics can be used to mutually decompose each other according to the patterns described in F1 to F6 in the following definition. We do not need to prove this so we will not. However, this observation does provide the intuition behind this rather involved construct. The fact that we have these mutual full decompositions, means that the mosaic o can be fully decomposed by the others mentioned and they in turn can be fully decomposed and so on: i.e. this provides a tactic for iterative full decompositions.

Definition 4.8 *Suppose $0 \leq r$, each $\lambda_i (1 \leq i \leq r)$ is a non-empty composing sequence of ϕ-mosaics, and $P_0, \ldots, P_s$ $(0 \leq s)$ are maximally propositionally consistent subsets of* Clϕ.

Suppose ϕ-mosaic $o = (A, B, C)$ and:

$$m' = (A, B, P_0);$$
$$y_i = (P_i, B, P_{i+1}) \ (0 \leq i \leq s-1);$$
$$y_s = (P_s, B, P_0);$$
$$m'' = (P_0, B, C); \text{ and}$$
$$\mu = \langle y_0, \ldots, y_s \rangle.$$

If $r = 0$ suppose $\lambda = \langle\rangle$, the empty sequence, but otherwise, if $r > 0$, suppose:

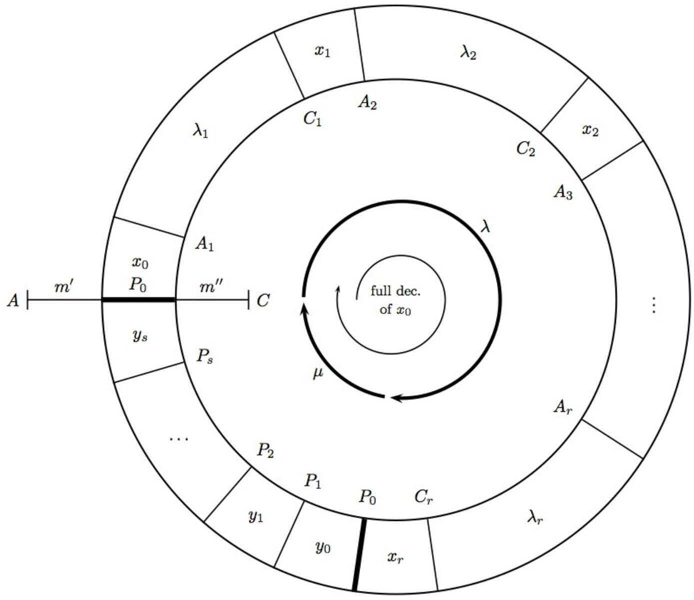

Fig. 4. Full decompositions in a shuffle

A_i is the start of the first mosaic in $\lambda_i (1 \leq i \leq r)$;
C_i is the end of the last mosaic in $\lambda_i (1 \leq i \leq r)$;
$x_0 = (P_0, B, A_1)$;
$x_i = (C_i, B, A_{i+1})$, $(1 \leq i \leq r-1)$;
$x_r = (C_r, B, P_0)$;
$\lambda = \langle x_0 \rangle ^\wedge \lambda_1^\wedge \langle x_1 \rangle ^\wedge \ldots ^\wedge \lambda_r^\wedge \langle x_r \rangle$.

Further suppose that m', m'', and each y_i and x_i are mosaics.

Then, we say that o is fully decomposed by the tactic shuffle *$(\langle P_0, \ldots, P_s \rangle, \langle \lambda_1, \ldots, \lambda_r \rangle)$ iff the following conditions all hold:*

- *F1.* *o is fully decomposed by $\langle m' \rangle ^\wedge \lambda ^\wedge \mu ^\wedge \langle m'' \rangle$;*
- *F2.* *if $r > 0$, x_0 is fully decomposed by $\lambda ^\wedge \mu ^\wedge \langle x_0 \rangle$;*
- *F3.* *if $0 < i < r$, x_i is fully decomposed by $\langle x_i \rangle ^\wedge \lambda_{i+1}^\wedge \langle x_{i+1} \rangle ^\wedge \ldots ^\wedge \lambda_r^\wedge \langle x_r \rangle ^\wedge \mu ^\wedge \langle x_0 \rangle ^\wedge \lambda_1^\wedge \langle x_1 \rangle ^\wedge \ldots ^\wedge \lambda_i^\wedge \langle x_i \rangle$;*
- *F4.* *if $r > 0$, x_r is fully decomposed by $\langle x_r \rangle ^\wedge \mu ^\wedge \lambda$;*
- *F5.* *if $0 \leq i < s$, y_i is fully decomposed by $\langle y_i, y_{i+1}, \ldots, y_s \rangle ^\wedge \lambda ^\wedge \langle y_0, \ldots, y_i \rangle$;*
- *F6.* *y_s is fully decomposed by $\langle y_s \rangle ^\wedge \lambda ^\wedge \mu$.*

The order of the mutual full decompositions specified in F1-F6 in the definition is illustrated in the circular diagram in figure 4 (borrowed from [15]).

Note that as $s \geq 0$ there is at least one P_i involved in the shuffle. In a general linear order setting we could define a shuffle with no P_is (provided that then $r > 0$) but over the reals it turns out to be crucial to require at least one

P_i. This is because, as it is not too hard to see, a shuffle of only non-singleton closed intervals of the reals cannot be both Dedekind complete and separable (i.e. having a countable dense suborder).

4.4 Real Mosaic Systems

Finally (in [15]) we are able to introduce the notion of a real mosaic system (RMS): one in which each mosaic can be fully decomposed in terms of the three tactics and of simpler mosaics. This is our version of a saturated set of mosaics.

Definition 4.9 *For $\phi \in L(U, S)$, suppose S is a set of ϕ-mosaics and $n \geq 0$.*

A ϕ-mosaic m is a level n^+ member of S *iff m is the composition of a sequence of mosaics, each of them being either a level n member of S or fully decomposed by the tactics* lead σ *or* trail σ *with each mosaic in σ being a level n member of S.*

A ϕ-mosaic m is a level $(n+1)^-$ member of S *iff m is the composition of a sequence of mosaics, each of them being either a level n^+ member of S or fully decomposed by the tactics* lead σ *or* trail σ *with each mosaic in σ being a level n^+ member of S.*

A ϕ-mosaic $m \in S$ is a level n member of S *iff m is the composition of a sequence of mosaics with each of them being either a level n^- member of S or a mosaic which is fully decomposed by the tactic* shuff $\langle P_0, \ldots, P_s \rangle, \langle \sigma_1, \ldots, \sigma_r \rangle$ *with each mosaic in each σ_i being a level n^- member of S.*

Note that it is generally possible for mosaics to be level 0 members of some S provided that they are compositions of mosaics which can be fully decomposed by shuffles in which there are no sequences (i.e. , $r = 0$).

Definition 4.10 *For $\phi \in L(U, S)$, a* real mosaic system *of ϕ-mosaics is a set S of ϕ-mosaics such that for every $m \in S$ there exists some n such that m is a level n member of S. For any n we say that S is a real mosaic system of depth n iff every $m \in S$ is a level n member of S.*

4.5 Relativization

We will now relate the satisfiability of a formula ϕ to that of certain mosaics. Obviously, a formula will be satisfiable over the reals iff it is satisfiable over the $]0, 1[$ flow. Furthermore, this happens iff a relativized version of the formula is satisfiable somewhere in the interior of a model over $[0, 1]$. To define this relativization we need to use a new atom to indicate points in the interior. Hence the next few definitions.

Definition 4.11 *Given ϕ and an atom q which does not appear in ϕ, we define a map $* = *_q^\phi$ on formulas in $\mathrm{Cl}\phi$ recursively:*

1. $*p = p \wedge q$,
2. $*\neg\alpha = \neg(*\alpha) \wedge q$,
3. $*(\alpha \wedge \beta) = *(\alpha) \wedge *(\beta) \wedge q$,
4. $*U(\alpha, \beta) = U(*\alpha, *\beta) \wedge q$, *and*
5. $*S(\alpha, \beta) = S(*\alpha, *\beta) \wedge q$.

So $*^{\phi}_{q}(\phi)$ will be a formula using only q and atoms from ϕ.

Lemma 4.12 *$*^{\phi}_{q}(\phi)$ is at most 3 times as long as ϕ.*

With the relativization machinery we can then define a relativized mosaic to be one which could correspond to the whole of a $[0,1]$ structure in which q is true of exactly the interior $]0,1[$ and the interior is a model of ϕ.

Definition 4.13 *We say that a $*^{\phi}_{q}(\phi)$-mosaic (A,B,C) is (ϕ,q)*-relativized *iff*

1. *$\neg q$ is in A and no $S(\alpha,\beta)$ is in A;*
2. *$q \in B$ and $\neg *^{\phi}_{q}(\phi) \notin B$; and*
3. *$\neg q \in C$ and no $U(\alpha,\beta)$ is in C.*

Here we confirm that ϕ is satisfiable over the reals exactly when we can find such a relativized mosaic.

Lemma 4.14 *Suppose that ϕ is a formula of $L(U,S)$ and q is an atom not appearing in ϕ. Then ϕ is $\mathbb{R}$-satisfiable iff there is some fully $[0,1]$-satisfiable (ϕ,q)-relativized $*^{\phi}_{q}(\phi)$-mosaic.*

Proof. Let $* = *^{\phi}_{q}$ and let $\zeta :]0,1[\to \mathbb{R}$ be any order preserving bijection.

Suppose that ϕ is $\mathbb{R}$-satisfiable. Say that $\mathbf{S} = (\mathbb{R},<,g)$, $s_0 \in \mathbb{R}$ and $\mathbf{S}, s_0 \models \phi$. Let $\mathbf{T} = ([0,1],<,h)$ where for atom $p \neq q$, $h(p) = \{t \in]0,1[| \zeta(t) \in g(p)\}$; and $h(q) =]0,1[$. An easy induction on the construction of formulas in $\mathrm{Cl} * \phi$ shows that $\mathbf{T}, \zeta^{-1}(s_0) \models *\phi$ and so $\mathrm{mos}^{*\phi}_{\mathbf{T}}(0,1)$ is the right mosaic.

Suppose mosaic $(A,B,C) = \mathrm{mos}(0,1)$ from structure $\mathbf{T} = ([0,1], <, h)$ is a (ϕ,q)-relativized $*(\phi)$-mosaic. Thus $q \in B$ and $\neg q \in A \cap C$. Define $\mathbf{S} = (\mathbb{R},<,g)$ via $s \in g(p)$ iff $\zeta^{-1}(s) \in h(p)$ for any atom p (including $p = q$). As $\neg * \phi \notin B$, there is some z such that $0 < z < 1$ and $\mathbf{T}, z \models *\phi$. It is easy to show that $\mathbf{S}, \zeta(z) \models \phi$. □

Our satisfiability procedure in [15] is to guess a relativized mosaic (A,B,C) and then check that (A,B,C) is fully $[0,1]$-satisfiable. Thus we now turn to the question of deciding whether a relativized mosaic is satisfiable.

4.6 Satisfiability of Mosaics

For us, in this paper, the important result from [15] is that which shows the equivalence of satisfiability of a formula to the fact of its negation not being in the cover of some mosaic in some RMS. The negation is not in the cover of a mosaic if the mosaic represents a pair of points in a model with a witness to ϕ in between. There is actually a slight further complication in that we need to deal with mosaics (relativized ones) which (by virtue of the formulas contained within each end) must be covering the whole of the structure of which they are part. This idea of relativization allows us to take care of formulas such as $U(\alpha,\beta)$ lying in the end of a mosaic: this cannot happen in a relativized mosaic.

Theorem 4.15 *([15] Theorem 72) Suppose ϕ is a formula of $L(U,S)$ and q is an atom not appearing in ϕ. Suppose $\psi = *^{\phi}_{q}(\phi)$ has length N.*

Then the following are equivalent:

1. ϕ *is* $\mathbb{R}$*-satisfiable;*
2. *there is a* (ϕ, q)*-relativized* ψ*-mosaic which is a level* N *member of some RMS.*

5 Expressiveness

In this section we use mosaics to show that if a formula is satisfiable over the reals then it is satisfiable in a compositional real model. This shows that the compositional model building method has adequate expressiveness for describing models of temporal specifications. Despite this first result not being surprising given earlier precedents [6], the new technique here allows us to to establish the new synthesis theorem later in the section.

In order to work with open intervals of the reals and match mosaics to such structures, we need to rework some definitions and several of the lemmas from [15]. The first definition gives us something like a Hintikka structure to witness a mosaic, at least as far as possible internal structure is concerned.

Definition 5.1 *Suppose* ϕ *is a formula of* $L(U, S)$ *and* $m = (A, B, C)$ *is a* ϕ*-mosaic.*

Suppose $x, y \in \mathbb{R}$ *with* $x < y$.

Say that structure $(]x, y[, <, h)$ supports m *iff there is a map* $\mu :]x, y[\to \wp(\mathrm{Cl}\phi)$ *satisfying:*

R0. *for each* $z \in]x, y[$, $p \in \mu(z)$ *iff* $z \in h(p) \cap \mathrm{Cl}\phi$ *;*

R1. *for each* $z \in]x, y[$, $\mu(z)$ *is a maximally propositionally consistent subset of* $\mathrm{Cl}\phi$*;*

R2. *Suppose* $z \in]x, y[$. *Then* $U(\alpha, \beta) \in \mu(z)$ *iff either*
- *R2.1,* *there is* u *such that* $z < u < y$ *and* $\alpha \in \mu(u)$ *and for all* v, *if* $z < v < u$ *then* $\beta \in \mu(v)$ *or*
- *R2.2,* $\alpha \in C$ *and for all* v, *if* $z < v < y$ *then* $\beta \in \mu(v)$ *or*
- *R2.3,* $\beta \in C$, $U(\alpha, \beta) \in C$ *and for all* v, *if* $z < v < y$, *then* $\beta \in \mu(v)$*;*

R3. *the mirror image of R2 for* $S(\alpha, \beta)$*;*

R4. $U(\alpha, \beta) \in A$ *iff either*
- *R4.1,* *there is* u *such that* $x < u < y$ *and* $\alpha \in \mu(u)$ *and for all* v, *if* $x < v < u$ *then* $\beta \in \mu(v)$ *or*
- *R4.2,* $\alpha \in C$ *and for all* v, *if* $x < v < y$ *then* $\beta \in \mu(v)$ *or*
- *R4.3,* $\beta \in C$, $U(\alpha, \beta) \in C$ *and for all* v, *if* $x < v < y$, *then* $\beta \in \mu(v)$*;*

R5. *the mirror image of R4 for* $S(\alpha, \beta)$*; and*

R6. *for each* $\beta \in \mathrm{Cl}\phi$, β *is in the cover of* m *iff for all* u, *if* $x < u < y$, $\beta \in \mu(u)$.

(Also say that m *is supported by* $(]x, y[, <, h)$, *via* μ.*)*

A straightforward induction on the construction of formulas tells us the following.

Lemma 5.2 *If m is supported by* $\mathbf{S}$ *and that is isomorphic to* $\mathbf{T}$ *then m is also supported by* $\mathbf{T}$.

As we have seen, the class of compositional real structures is closed under appropriate semantic sum, trail, lead and shuffle operations. Thus we can proceed through an induction on mosaics in an RMS, relating them to successively more complex compositional structures.

For the full proofs see [2].

Lemma 5.3 *If m is the composition of m' and m'' with each of m' and m'' being supported by a compositional real structure then m is supported by a compositional real structure.*

Lemma 5.4 *If m is fully decomposed by the tactic* lead (σ) *with each mosaic in σ being supported by a compositional real structure, then m is supported by a compositional real structure. There is a mirror image result for* trail (σ).

Lemma 5.5 *If m is fully decomposed by the tactic*

$$\text{shuff}\ (\langle P_0, \ldots, P_s\rangle, \langle \lambda_1, \ldots, \lambda_r\rangle)$$

with each mosaic in each λ_i being supported by a compositional real structure, then m is supported by a compositional real structure.

We have shown that each of the operations of composition and using lead, trail or shuffle tactics preserves compositional supportedness of mosaics. A simple induction allows us to conclude that any mosaic in a real mosaic system is supported by a compositional real structure.

Lemma 5.6 *Suppose ϕ is a formula of $L(U,S)$ and m is a ϕ-mosaic.*

If m appears in an RMS then m is supported by a compositional real structure.

Proof. Given the real mosaic system S of ϕ-mosaics, we can easily proceed by induction on k to show that any level k member $m' \in S$ is supported by a compositional real structure. Each step of the induction is just a use of one or two of the preceding lemmas 5.3, 5.4, its mirror image and 5.5.

Suppose this is true for $k \geq -1$: it is true for $k = -1$ because there are no level -1 members of S. Lemma 5.4 tells us that any mosaic which can be fully decomposed by lead (σ) is supported if each of the mosaics in σ are. By lemma 5.3 this (and its mirror image) means that any level k^+ member of S is supported. By lemma 5.3 and lemma 5.4 and its mirror image we have that any level $(k+1)^-$ member of S is supported. By lemmas 5.3 and 5.5 it follows that any level $k+1$ member of S is supported as required. □

5.1 Relativized and Supported means Modeled

Lemma 5.7 *Suppose ϕ is a formula of $L(U,S)$ and q is an atom not appearing in ϕ. Suppose $\psi = *^{\phi}_{q}(\phi)$ has length N.*

Suppose that ψ-mosaic m is (ϕ, q)-relativized and is supported by a real structure $\mathbf{T} = (\mathbb{R}, <, h)$.

Then there is $r \in \mathbb{R}$ such that $\mathbf{T}, r \models \phi$.

Proof. Suppose ϕ is a formula of $L(U, S)$ and q is an atom not appearing in ϕ. Suppose $\psi = *_q^\phi(\phi)$ has length N.

Suppose that ψ-mosaic $m = (A, B, C)$ is (ϕ, q)-relativized and is supported by a real structure $(\mathbb{R}, <, h)$.

By lemma 5.2, m is also supported by $\mathcal{T} = (]0, 1[, <, h')$ for some h'. Let $f : \mathbb{R} \to]0, 1[$ be the bijection such that for all $p \in P$, for all $t \in]0, 1[$, $t \in h'(p)$ iff $f(t) \in h(p)$.

Let $\mu :]0, 1[\to \wp(\mathrm{Cl}\phi)$ satisfy R1-R6 in the definition of supporting.

First, we claim that for all $\alpha \leq \phi$, for all $x \in]0, 1[$, $\mathbf{T}, x \models \alpha$ iff $*\alpha \in \mu(x)$. The proof is by induction on construction of α, and has a series of cases (each with forward and converse arguments) but is straightforward.

Now as m is (ϕ, q)-relativized, $\neg * \phi \notin B$.

By R6, there must be some $u \in]0, 1[$ such that $*\phi \in \mu(u)$. By our induction result, $\mathbf{T}, u \models \phi$.

By lemma 2.5, $(\mathbb{R}, <, h), f(u) \models \phi$ as required. □

5.2 Finding a model of a formula

The main result is that we can use such an RMS to describe a model in our new notation.

Theorem 5.8 *A formula ϕ from $L(U, S)$ is $\mathbb{R}$-satisfiable iff there is a compositional real model of ϕ. There is some c such that, for $\mathbb{R}$-satisfiable ϕ, a satisfying model can be described by an expression of shuffles, leads, trails and sums of length $< 2^{c|\phi|^2}$ (this bound is best possible).*

Furthermore, there is an EXPTIME procedure for finding such an expression.

Proof. Suppose ϕ is a formula of $L(U, S)$ and q is an atom not appearing in ϕ. Suppose $\psi = *_q^\phi(\phi)$ has length N.

One direction of the theorem is immediate: by Definition 3.6, any compositional real model of ϕ is a real model of ϕ.

For the other direction assume ϕ is $\mathbb{R}$-satisfiable.

By theorem 4.15, there is a (ϕ, q)-relativized ψ-mosaic m which is a level N member of some RMS.

By lemma 5.6, m is supported by a compositional real structure $\mathbf{T}$, say, corresponding to compositional real expression $\mathcal{I}$.

By lemma 5.7, there is $r \in \mathbb{R}$ such that $\mathbf{T}, r \models \phi$.

The bound follows from consideration of the level of the relativized mosaic in the RMS, a level which from consideration of arguments in [15] can be shown to be at most six times the length of ϕ. We also show that the length of decomposition sequences at each level is bounded by an exponential in $|\phi|$.

We can also show that bound is best possible by considering a formula describing a binary counter. Given n, use n atoms to describe a counter which increases at discrete intervals. A formula of quadratic length (in n) can be used to specify such a model but a description of the model in our construction notation needs to be of exponential length.

Suppose that the atom p marks discrete points (so $\delta = G(p \rightarrow U(p, \neg p)$ is true), $\overline{b} = b_1, \ldots, b_n$ are the n bits of a counter that counts modulo 2^n and:

$$\phi_i = (p \wedge \bigwedge_{j=1}^{i-1} b_j) \rightarrow (b_i \rightarrow U(\neg b_i, \neg p) \wedge \neg b_i \rightarrow U(b_i, \neg p))$$

$$\psi_i = (p \wedge \neg \bigwedge_{j=1}^{i-1} b_j) \rightarrow (b_i \rightarrow U(b_i, \neg p) \wedge \neg b_i \rightarrow U(\neg b_i, \neg p))$$

The formula ϕ_i specifies if all bits less significant than b_i are true, then b_i will invert its valuation the next time p is true, and ψ_i specifies that if any less significant bit is not true, b_i will retain its current valuation at the next point that p is true.

Clearly $\delta \wedge G \bigwedge_{i=1}^{n} (\phi_i \wedge \psi_i)$ is satisfiable, and is of length $O(n^2)$. However, any real model expression describing a structure that satisfies this formula would have to contain at least 2^n distinct letters.

□

Note that thanks to the expressive completeness result in [5], we know that any satisfiable sentence of the first-order monadic logic of the reals also has a compositional real model. To find a description of a model from the sentence must be a hard problem as deciding validity in this logic is non-elementarily complex [17].

Theorem 5.9 *There is an EXPTIME procedure which given a formula ϕ from $L(U, S)$ will decide whether ϕ is $\mathbb{R}$-satisfiable or not and if so will provide an expression for a compositional model of ϕ.*

Proof. Finally, we can give an EXPTIME procedure for finding and printing out a model of any satisfiable RTL formula. The set of all ϕ-mosaics is of size exponential in the length of ϕ. There is a fairly straightforward procedure (in the style of [9]) for going through the set repeatedly and removing mosaics which can not be fully decomposed in terms of other simpler ones in the set. If ϕ is satisfiable we will eventually end up with an RMS and another straightforward EXPTIME procedure reads out the description of a model of ϕ. By repeatedly decomposing mosaics as specified in the RMS we can produce the expression in a top-down manner. □

Other results from [15] allow us to conclude that if we find all possible starting points (i.e. relativized mosaics in the RMS) and follow all possible ways of decomposing the mosaics (as given in the RMS) then we will eventually output a list of possible models of the formula which is in a certain sense exhaustive. Any real model of ϕ will be back-and-forth equivalent to one of the compositional models which is listed.

6 Conclusion

We have investigated a compositional approach to building linear temporal structures as a way of working with models on a continuous (real numbers) flow of time. Structures are built by putting together smaller structures in a recursive way, with copies of the smaller ones occupying successive intervals of time. We have formalised the approach so that such models can be described clearly and efficiently.

We have identified a sub-language of the formal compositional model building language which can be used (in a slightly modified way) to build real-flowed structures. Any RTL formula satisfiable in the reals is satisfiable in such a compositional real-flowed model. We presented an efficient method for building a real-flowed model of any given a satisfiable formula.

The approaches here may generalise to general linear models of time [2].

References

[1] Burgess, J. P. and Y. Gurevich, *The decision problem for linear temporal logic*, Notre Dame J. Formal Logic **26** (1985), pp. 115–128.

[2] French, T., J. McCabe-Dansted and M. Reynolds, *Synthesis and model checking for continuous time: Long version*, Technical report, CSSE, UWA (April 2012), "http://www.csse.uwa.edu.au/∼mark/research/Online/sctm.htm".

[3] Gabbay, D., I. Hodkinson and M. Reynolds, "Temporal Logic: Mathematical Foundations and Computational Aspects, Volume 1," Oxford University Press, 1994.

[4] Gabbay, D. M. and I. M. Hodkinson, *An axiomatisation of the temporal logic with until and since over the real numbers*, J. Logic and Computation **1** (1990), pp. 229 – 260.

[5] Kamp, H., "Tense logic and the theory of linear order," Ph.D. thesis, U. Calif. (1968).

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

[7] Németi, I., *Decidable versions of first order logic and cylindric-relativized set algebras*, in: *Logic Colloquium '92* (1995), pp. 171–241.

[8] Pnueli, A., *The temporal logic of programs*, in: *Proceedings of the Eighteenth Symposium on Foundations of Computer Science*, 1977, pp. 46–57.

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

[10] Rabin, M. O., *Decidability of second order theories and automata on infinite trees*, American Mathematical Society Transactions **141** (1969), pp. 1–35.

[11] Reynolds, M., *An axiomatization for Until and Since over the reals without the IRR rule*, Studia Logica **51** (1992), pp. 165–193.

[12] Reynolds, M., *Continuous temporal models*, in: M. Stumptner, D. Corbett and M. J. Brooks, editors, *Australian Joint Conference on Artificial Intelligence*, Lecture Notes in Computer Science **2256** (2001), pp. 414–425.

[13] Reynolds, M., *The complexity of the temporal logic with "until" over general linear time*, J. Comput. Syst. Sci. **66** (2003), pp. 393–426.

[14] Reynolds, M., *Dense time reasoning via mosaics*, in: *Proceedings of the 2009 16th International Symposium on Temporal Representation and Reasoning* (2009), pp. 3–10.

[15] Reynolds, M., *The complexity of the temporal logic over the reals*, Annals of Pure and Applied Logic **161** (2010), pp. 1063–1096, online at doi:10.1016/j.apal.2010.01.002.

[16] Rosenstein, J. G., "Linear Orderings," Academic Press, New York, 1982.

[17] Stockmeyer, L., "The complexity of decision problems in automata and logic," Ph.D. thesis, M.I.T. (1974).