The temporal logic of two-dimensional Minkowski spacetime with slower-than-light accessibility is decidable

Robin Hirsch [1] Brett McLean [2]

Department of Computer Science
University College London
Gower Street, London WC1E 6BT

Abstract

We work primarily with the Kripke frame consisting of two-dimensional Minkowski spacetime with the irreflexive accessibility relation 'can reach with a slower-than-light signal'. We show that in the basic temporal language, the set of validities over this frame is decidable. We then refine this to PSPACE-complete. In both cases the same result for the corresponding reflexive frame follows immediately. With a little more work we obtain PSPACE-completeness for the validities of the Halpern–Shoham logic of intervals on the real line with two different combinations of modalities.

Keywords: temporal logic, basic temporal language, Minkowski spacetime, frame validity, Halpern–Shoham logic.

1 Introduction

Minkowski spacetime refers to the flat spacetime of special relativity, where in any inertial coordinates $(\boldsymbol{r}, ct)$, light travels in straight lines at a 45° angle to the time axis. To view this as a Kripke frame, there are at least four natural accessibility relations to chose from: reflexive and irreflexive versions of 'can reach with a lightspeed-or-slower signal' and 'can reach with a slower-than-light signal'. See below for the formal definitions in the two-dimensional case.

For the basic *modal* language, Goldblatt found that, regardless of the number of spatial dimensions, both reflexive choices produce the logic **S4.2** (reflexivity, transitivity, confluence) [5]. Shapirovsky and Shehtman proved the irreflexive slower-than-light logic is **OI.2**—transitivity, seriality, confluence, *two-density* (see Section 2)—again regardless of dimension [15]. Both **S4.2** and **OI.2** are PSPACE-complete [13]. In contrast, Shapirovsky has shown that with irreflexive *exactly lightspeed* accessibility, validity is undecidable [14].

[1] r.hirsch@ucl.ac.uk

[2] b.mclean@cs.ucl.ac.uk

For the basic *temporal* language, the problems of axiomatising and determining the complexity of the validities of these frames had all been open for decades—Shehtman recommended an investigation of the two-dimensional case (one time and one space dimension) in the concluding remarks of [16]. In the two-dimensional case, starting with coordinates (r, ct), we may rotate the axes through 45° to get coordinates (x, y), where $x = \frac{1}{\sqrt{2}}(ct + r),\ y = \frac{1}{\sqrt{2}}(ct - r)$. With these coordinates the reflexive lightspeed-or-slower relation $\leq$, and the irreflexive slower-than-lightspeed relation $\prec$, are given by

$$(x, y) \leq (x', y') \iff x \leq x' \wedge y \leq y', \quad (x, y) \prec (x', y') \iff x < x' \wedge y < y \quad (1)$$

where $<, \leq$ on the right are the usual irreflexive/reflexive orderings of the reals, and $<, \preceq$ are obtained from $\leq, \prec$ by deleting/adding the identity, respectively.

Recently, Hirsch and Reynolds managed to show that for this two-dimensional case, with either reflexive or irreflexive lightspeed-or-slower accessibility, the validity problem is PSPACE-complete [7]. However, they were unable to obtain decidability/complexity results for slower-than-light accessibility. Indeed decidability appears in item (2) in the list of open problems at the end of their paper. In this paper we solve Hirsch and Reynolds' problem, eventually proving the following.

Theorem 6.1 *On the frame consisting of two-dimensional Minkowski spacetime equipped with the irreflexive slower-than-light accessibility relation, the set of validities of the basic temporal language is* PSPACE*-complete. The same is true with reflexive slower-than-light accessibility.*

The proof of Theorem 6.1 follows that of [7] very closely. The only additional insight needed is that all pertinent information about the behaviour of a valuation on a light-line can be captured in a finite way—see Definition 2.3—, despite all of a light-line's points being mutually inaccessible.

The proof of the main result is structured as follows.

(i) Given a fixed formula ϕ whose satisfiability is to be determined, we define the maximal consistent sets of subformulas/negated subformulas of ϕ, where consistency is with respect to the class of all temporal frames.

(ii) We define *surrectangles* by recording a maximal consistent set at each point of a rectangle, together with some information about the maximal consistent sets holding near but beyond the boundaries of the rectangle.

(iii) Starting in Section 3, we define *biboundaries*, also based on the maximal consistent sets. Biboundaries have finite specifications and the intuition for them is as a record of the information contained near the boundary of a surrectangle, plus a little from its interior. Indeed we define the biboundary ∂^s determined by a given surrectangle s.

(iv) We define three operations on biboundaries: *joins*, *limits*, and *shuffles*, and we define the set of *fabricated biboundaries* to be those biboundaries formed by iterating these operations, starting from certain basic biboundaries. This is all computable, and the iterative procedure must terminate,

for the biboundaries are finite in number.

(v) We show that every fabricated biboundary is given by ∂^s for some surrectangle s, by describing a recursive construction of s. (Section 4)

(vi) Conversely, we show that for every surrectangle s, the biboundary ∂^s is fabricated. (Section 5)

(vii) We deduce the decidability of the validity problem over our frame.

(viii) In Section 6, we give a procedure for deciding if a biboundary is fabricated using polynomial space, and also note the satisfiability task is PSPACE-hard. Consequently our result that validity is decidable is refined to validity being PSPACE-complete.

In [6], Halpern and Shoham introduced a family of modal logics in which the entities under discussion are intervals. These logics have subsequently come to be highly influential, and their axiomatisability, decidability, and complexity extensively studied [19,2,9,4,11,10,3]. In Section 7, we describe how the proof of the complexity of validity for two-dimensional Minkowski spacetime can be adapted to prove PSPACE-completeness of the temporal logic of intervals where the accessibility relation is $\texttt{overlaps} \cup \texttt{meets} \cup \texttt{before}$, or its reflexive closure.

2 Preliminaries and surrectangles

We often, but not exclusively, follow the terminology and notation of [7]. We take as primitive the propositional connectives $\neg$ and $\vee$, and modal operators $\mathbf{F}$ and $\mathbf{P}$; the usual abbreviations apply. The semantics is the usual semantics on temporal frames—Kripke frames for which the accessibility relation for $\mathbf{P}$ is the converse of that for $\mathbf{F}$. An example of a formula that is valid over two-dimensional slower-than-light frames but not over irreflexive lightspeed-or-slower frames is the two-density formula

$$(\mathbf{F}p \wedge \mathbf{F}q) \to \mathbf{F}(\mathbf{F}p \wedge \mathbf{F}q)$$

asserting that if x precedes both z_1 and z_2, then there exists y that is between x and z_1 *and* between x and z_2.

Throughout, ϕ is a fixed formula whose satisfiability is to be determined. The **closure** $\mathrm{Cl}(\phi)$ of ϕ is the set of all subformulas and negated subformulas of ϕ. A **maximal consistent set** is a subset of $\mathrm{Cl}(\phi)$ that is satisfiable in some temporal frame and is maximal with respect to the inclusion ordering, subject to the satisfiability constraint. We denote the set of maximal consistent sets by MCS. It is well known that satisfiability in a temporal frame is decidable, indeed PSPACE-complete [17]. Hence the set of all maximal consistent sets of ϕ is a (total) computable function of ϕ. The relation on MCS given by

$$\begin{aligned} m \lesssim n \iff \quad & \forall \mathbf{F}\psi \in \mathrm{Cl}(\phi)\ ((\psi \in n \to \mathbf{F}\psi \in m) \wedge (\mathbf{F}\psi \in n \to \mathbf{F}\psi \in m)) \\ & \wedge \forall \mathbf{P}\psi \in \mathrm{Cl}(\phi)\ ((\psi \in m \to \mathbf{P}\psi \in n) \wedge (\mathbf{P}\psi \in m \to \mathbf{P}\psi \in n)) \end{aligned}$$

is transitive, and so defines a preorder on reflexive elements. We call a $\lesssim$-equivalence class a **cluster** and write $\leq$ for the partial order on clusters induced

by $\lesssim$. The notation $<$ means $\leq$ but not equal. We extend these notations to compare a cluster c and a (not necessarily $\lesssim$-reflexive) maximal consistent set m as follows. Write $m \leq c$ if for all $n \in c$ we have $m \lesssim n$, and write $m < c$ if for all $n \in c$ we have $(m \lesssim n \wedge m \neq n)$. Similarly for $c \leq m$ and $c < m$.

Definition 2.1 A formula of the form $\mathbf{F}\psi$ is a **future defect** of a maximal consistent set m if $\mathbf{F}\psi \in m$. A future defect of a set S of maximal consistent sets is a future defect of any member of S *unless* we explicitly identify S as a cluster. The formula $\mathbf{F}\psi$ is a future defect of a cluster c if $\mathbf{F}\psi$ is contained in some $m \in c$, but, for all $n \in c$, we have $\psi \notin n$. A future defect $\mathbf{F}\psi$ is **passed up** to a set S of maximal consistent sets if either ψ or $\mathbf{F}\psi$ belongs to some $m \in S$. A **past defect** is defined similarly.

Recall from (1) the ordering $\prec$ and the associated 'can reach with a slower-than-light signal' Kripke frame $(\mathbb{R}^2, \prec)$. Subsets of $\mathbb{R}^2$ inherit the same ordering. The notation $\uparrow \boldsymbol{x}$ denotes the set $\{\boldsymbol{y} \mid \boldsymbol{x} \prec \boldsymbol{y}\}$. (The set $\boldsymbol{y}$ is drawn from should be clear.) We define the operations $\vee$ and $\wedge$ by

$$(x, y) \vee (x', y') = (\max\{x, x'\}, \max\{y, y'\}),$$
$$(x, y) \wedge (x', y') = (\min\{x, x'\}, \min\{y, y'\}).$$

(According to lightspeed-or-slower accessibility $\vee$ is join and $\wedge$ is meet). We also define the partial order $\triangleleft$ on $\mathbb{R}^2$ by $(x, y) \triangleleft (x', y') \iff x \leq x', y \geq y'$, and $(x, y) \neq (x', y')$. When this holds we say that (x, y) is 'northwest' of (x', y').

The **diagonal dual** of a condition c on a frame/model whose domain is a subset of $\mathbb{R}^2$ is the condition that c holds on the frame/model obtained by swapping the x- and y-axes (which will not affect accessibility). The **temporal dual** is the result of reversing the x-axis and reversing the y-axis (which reverses accessibility), and also swapping $\mathbf{F}$ and $\mathbf{P}$ in formulas. When we say 'all duals' we mean the diagonal dual, the temporal dual, and the diagonal temporal dual.

Suppose we have a preorder-preserving map f from a subset of $\mathbb{R}^2$ to MCS and that f is defined on $\mathcal{U} \cap \uparrow \boldsymbol{x}$ for some $\mathcal{U}$ an open neighbourhood (in $\mathbb{R}^2$) of $\boldsymbol{x}$. Then it is easy to see there is a unique cluster c such that there exists a neighbourhood $\mathcal{U}'$ of $\boldsymbol{x}$ such that f only takes values in c on $\mathcal{U}' \cap \uparrow \boldsymbol{x}$. We denote this cluster by $f^+(\boldsymbol{x})$. The analogous value for $\downarrow \boldsymbol{x}$ we denote $f^-(\boldsymbol{x})$.

Let $(x, y), (x', y)$ be two distinct points on a horizontal line segment in the domain of the preorder-preserving map f. Then $(x, y) \not\prec (x', y)$ so we do not know the relation between $f(x, y)$ and $f(x', y)$. However, it is easy to see that $x \leq x' \implies f^+(x, y) \lesssim f^+(x', y)$. Similarly for f^-. If f^+ is constantly equal to c on an open line segment l and defined at the left end of l then f^+ also equals c there.[3] Similarly for f^- with right ends. All diagonal duals of statements in this paragraph (that is, statements for vertical lines) hold similarly.

The following concept will be used in definitions that follow, and informs the way we think of a bi-trace (Definition 2.3) as specifying behaviour *on* a horizontal or vertical line segment.

[3] For line segments, 'open' and 'closed' have their usual meanings of 'excludes end points' and 'includes end points', respectively.

Definition 2.2 Let c^- and c^+ be clusters. The **interpolant** of c^- and c^+ is the set of all $m \in \mathsf{MCS}$ with $c^- \leq m \leq c^+$ such that all future defects of m are passed up to c^+, and all past defects of m are passed down to c^-.

Definition 2.3 A **bi-trace** (of length n) is two sequences $c_0^+ \leq \cdots \leq c_n^+$ and $c_0^- \leq \cdots \leq c_n^-$ of clusters, and one sequence $b_1, \ldots, b_n \in \mathsf{MCS}$, with the following constraints.

- For each $i \leq n$ we have $c_i^- \leq c_i^+$, and the interpolant of c_i^- and c_i^+ is nonempty.
- For each $i < n$ either $c_i^- < c_{i+1}^-$ or $c_i^+ < c_{i+1}^+$.
- For each $i < n$ we have $c_i^- \leq b_{i+1} \leq c_{i+1}^+$.
- For each $i < n$ all future defects of b_{i+1} are passed up to c_{i+1}^+ and all past defects passed down to c_i^-.

Formally, we consider a bi-trace to be the interleaving of its three sequences, with the advantage that we can use the notation $(c_0^-, c_0^+, b_1, c_1^-, \ldots, c_n^+)$ to indicate a bi-trace. We call the c_i^+'s the **upper clusters**, the c_i^-'s the **lower clusters**, c_0^+ and c_0^- the **initial clusters**, and c_n^+ and c_n^- the **final clusters**. Each pair c_i^-, c_i^+ is a **cluster pair** of the bi-trace, and each b_i a **transition value**. The top part of Figure 1 suggests how to visualise a bi-trace. There are only finitely many bi-traces (because of the second condition in their definition).

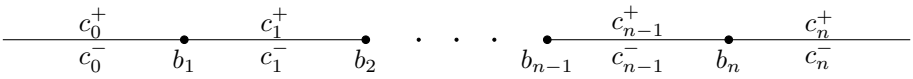

Figure 1. A bi-trace

Definition 2.4 Let $t_1 = (c_0^-, c_0^+, b_1, \ldots, c_n^+)$ and $t_2 = (e_0^-, e_0^+, d_1, \ldots, e_k^+)$ be bi-traces, and $a \in \mathsf{MCS}$. Then $t_1 + a + t_2$ is defined and equal to $(c_0^-, c_0^+, b_1, \ldots, c_n^+, a, e_0^-, e_0^+, d_1, \ldots, e_k^+)$ if this is a bi-trace. It is also defined if the final clusters of t_1 equal the initial clusters of t_2, and a is in the interpolant of c_n^- and c_n^+, in which case it equals $(c_0^-, c_0^+, b_1, \ldots, c_n^+, d_1, \ldots, e_k^+)$.

We now define surrectangles. Intuitively, the domain of a surrectangle is a rectangle plus infinitesimally more beyond any closed edges of the rectangle (a 'surreal rectangle'), and a surrectangle records a valuation on this domain.

Definition 2.5 A **rectangle** is a product of two intervals of $\mathbb{R}$ (with unbounded and single-point intervals both allowed); it is **degenerate** if either interval is a single point, otherwise it is **nondegenerate**. An **edge** of a rectangle R is an edge of the closure of R in $\mathbb{R}^2$ and is not considered to include its end points; a **closed edge** of R is an edge of R contained in R. An **upper edge** of R is either a horizontal edge with maximal vertical component, or a vertical edge with maximal horizontal component, a lower edge is defined dually. A rectangle is **open**/**closed** if it is open/closed in $\mathbb{R}^2$. The notation $[\boldsymbol{b}, \boldsymbol{t}]$, for points $\boldsymbol{b} = (b_1, b_2)$ and $\boldsymbol{t} = (t_1, t_2)$, signifies the closed rectangle $\{(x, y) \in \mathbb{R}^2 \mid b_1 \leq x \leq t_1 \text{ and } b_2 \leq y \leq t_2\}$.

A **surrectangle** consists of the following data.

(1) A preorder-preserving map $f\colon (R, \prec) \to (\mathsf{MCS}, \lesssim)$, for some *nondegenerate* rectangle R. We call R the **rectangle** and f the **core map** of the surrectangle.

(2) For each closed horizontal upper edge e of R with endpoints (x_0, y) and (x', y),[4] a finite sequence $c_0^+ \leq \cdots \leq c_n^+$ of clusters and a sequence $(x_1, y), \ldots, (x_n, y) \in e$, with $x_0 < \cdots < x_n < x_{n+1}$, defining x_{n+1} to be x'. Similar finite sequences for any closed vertical and/or lower edges.[5]

And is required to satisfy the following constraints.

(3) For each closed horizontal upper edge as above, f^- is constant on each open line segment $((x_i, y), (x_{i+1}, y))$ (let this constant cluster be c_i^-), and $(c_0^-, c_0^+, b_1, c_1^-, \ldots, c_n^-, c_n^+)$ forms a bi-trace, where b_i is defined to be $f(x_i, y)$, for each i. Also, all duals of this constraint. (Figure 2 suggests how to visualise a closed edge.)

(4) For any point $\boldsymbol{x} \in R$, if $\mathbf{F}\psi \in f(\boldsymbol{x})$ either
- **resolved internally:** there is $\boldsymbol{y} \in {\uparrow}\boldsymbol{x}$ such that $\psi \in f(\boldsymbol{y})$,
- **passed upwards:** R has a boundary point $\boldsymbol{y}$ either due north or due east of $\boldsymbol{x}$ such that $\mathbf{F}\psi$ is passed up to $f(\boldsymbol{y})$.

(5) The temporal dual of (4) holds.

In (2), we call $c_0^+, \ldots, c_n^+$ the **supplementary clusters** of e, we call $(x_1, y), \ldots, (x_n, y)$ the **transition points** of e, and we call $f(x_1, y), \ldots, f(x_n, y)$ the **transition values**. Note that the clusters c_i^- are determined by f^- in (3), and for edges not contained in R (for example if the rectangle is unbounded in the corresponding direction) supplementary clusters and transition points are not defined.

Let $\boldsymbol{b}$ and $\boldsymbol{t}$ be respectively the lower-left and upper-right corners of R. Then $f^+(\boldsymbol{b})$ and $f^-(\boldsymbol{t})$ are necessarily defined. The **height** of the surrectangle is the maximum possible length of a chain of clusters (not necessarily in the image of f) from its **lower cluster** $f^+(\boldsymbol{b})$ to its **upper cluster** $f^-(\boldsymbol{t})$. Surrectangles also inherit descriptions such as open/closed from their underlying rectangle.

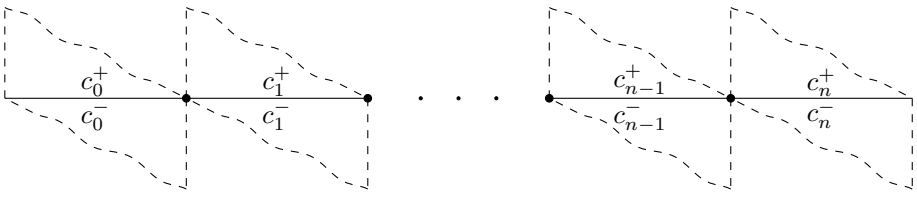

Figure 2. A closed edge of a surrectangle

[4] Here, it could be that $x_0 = -\infty$ and/or $x' = \infty$; this does not present any problems.

[5] For lower edges, these supplementary clusters are the ones denoted c_i^-.

3 Biboundaries

Definition 3.1 A **biboundary** is a partial map ∂ on $\{-, +, b, t, l, r, N, S, E, W\}$. It must be defined on $-$ and $+$, and be cluster-valued there. If ∂ is defined on N, S, E, or W, it is bi-trace-valued there, and if defined on b, t, l, or r, it is MCS-valued there. It is defined on b if and only if it is defined on S and W; similarly for t and N, E, for l and N, W, and for r and S, E. The following conditions must also be satisfied.

(i) If $\partial(b)$ is defined, then $\partial(b) \leq \partial(-)$ and every future defect of $\partial(b)$ is passed up to $\partial(-)$.

(ii) If $\partial(W)$ is defined, then the initial upper cluster of $\partial(W)$ equals $\partial(-)$.

(iii) If $\partial(l)$ is defined, then it is less than or equal to the initial upper cluster of $\partial(N)$, and every future defect of $\partial(l)$ is passed up to that cluster.

(iv) Every future defect of $\partial(+)$ is passed up to the interpolant of the final clusters of $\partial(N)$, to the interpolant of the final clusters of $\partial(E)$, or to $\partial(t)$, with ∂ defined in the places appropriate to the case.

(v) All duals of (i)–(iv) hold (with the evident meaning of duals).

Since there are only finitely many maximal consistent sets, clusters, and bi-traces, there are only finitely many biboundaries. A biboundary is **closed** if it is defined on N, S, E, and W (hence also on b, l, r, and t).

Let s be a surrectangle. The biboundary ∂^s determined by ∂ is defined in the obvious way.

We now define three types of operations on biboundaries: joins, limits, and shuffles. See Figure 3 for visual representations of limits and shuffles.

Definition 3.2 A biboundary ∂ is the **vertical join** of biboundaries ∂_1 and ∂_2, written $\partial_1 \oplus_- \partial_2$, if

- $\partial_1(N)$ and $\partial_2(S)$ are both defined and are equal,
- either $\partial_1(W)$, $\partial_2(W)$, and $\partial(W)$ are all defined, $\partial_1(l) = \partial_2(b)$, and $\partial(W) = \partial_1(W) + \partial_1(l) + \partial_2(W)$, or $\partial_1(W), \partial_2(W)$ and $\partial(W)$ are all undefined; similarly for E,
- ∂ agrees with ∂_1 on b, S, r, and $-$, and with ∂_2 on l, N, t, and $+$.

The diagonal-dual concept is a **horizontal join**, written $\partial_1 \oplus_| \partial_2$.[6]

Definition 3.3 A biboundary ∂^* is the **southeastern limit** of a biboundary ∂_0 using biboundaries $\partial_1, \partial_2, \partial_3$ if

- $\partial_0 = (\partial_2 \oplus_| \partial_3) \oplus_- (\partial_0 \oplus_| \partial_1)$,
- the lower cluster of $\partial_1(E)$ is constantly $\partial_0(+)$,
- the upper cluster of $\partial_2(S)$ is constantly $\partial_0(-)$,
- ∂^* agrees with ∂_0 over $\{-, +, l, W, N\}$,

[6] The $-$ and $|$ subscripts indicate the orientation of the shared edge.

- $\partial^*(S)$, if defined, is a bi-trace where the upper cluster is constantly $\partial_0(-)$,
- $\partial^*(E)$, if defined, is a bi-trace where the lower cluster is constantly $\partial_0(+)$.

A **northwestern limit** is defined dually.

If Δ is a set of biboundaries and there are $\partial_0, \partial_1, \partial_2, \partial_3 \in \Delta$ such that ∂^* is the southeastern limit of ∂_0 using $\partial_1, \partial_2, \partial_3$, then ∂^* is a southeastern limit over Δ. Northwestern limits are dual. We say that ∂^* is a limit over Δ if it is either a southeastern or northwestern limit over Δ.

Definition 3.4 Let Δ be a collection of *closed* biboundaries. The biboundary ∂' is a **shuffle** of Δ if there is a *nonempty* set $M \subseteq \mathsf{MCS}$ such that

(i) if $\partial'(W)$ is defined, then all upper clusters of $\partial'(W)$ equal $\partial'(-)$,

(ii) every future defect $\mathbf{F}\psi$ of $\partial'(-)$ is passed up to some $m \in M$, or there is a $\partial \in \Delta$ such that $\mathbf{F}\psi$ is passed up to either $\partial(b), \partial(l)$, $\partial(r)$, the interpolant of some cluster pair of $\partial(W)$ or $\partial(S)$, or some transition value of $\partial(W)$ or $\partial(S)$,

(iii) for all $\partial \in \Delta$, we have $\partial(t) \leq \partial'(+)$, all future defects of $\partial(t)$ are passed up to $\partial'(+)$, and all upper clusters of $\partial(N)$ equal $\partial'(+)$,

(iv) for all $m \in M$, we have $m \leq \partial'(+)$, and all future defects of every m are passed up to $\partial'(+)$,

(v) all duals of (i), (ii), (iii), and (iv) hold.

Now we are ready to define the biboundaries that we proceed to show are precisely those obtained from surrectangles.

Definition 3.5 A **ground fabricated biboundary** is a biboundary ∂ such that

(i) $\partial(-) = \partial(+)$,

(ii) if $\partial(N)$ is defined, then all lower clusters of $\partial(N)$ equal $\partial(+)$,

(iii) all duals of (ii) hold.

A **fabricated biboundary** is either a ground fabricated biboundary, or a biboundary obtained recursively as the join, limit, or shuffle of fabricated biboundaries.

4 From fabricated biboundaries to surrectangles

In this section we show that every fabricated biboundary is the biboundary obtained from some surrectangle, by describing how to construct such a surrectangle from a given biboundary. We use the recursive structure of fabricated biboundaries as given by their definition.

When we say a function f fills X **densely** with M we mean that $f(\boldsymbol{x}) \in M$ for all $\boldsymbol{x} \in X$, and for each $m \in M$ the set $f^{-1}(m)$ is dense in X. It is clear that if $\mathcal{U}$ is an open subset of $\mathbb{R}^2$ and c is a cluster, then there exists f that fills $\mathcal{U}$ densely with c (and this remains true when restrictions are placed on the behaviour of f outside of $\mathcal{U}$). Similarly when $\mathcal{U}$ is an open line segment.

Further, if for a biboundary ∂ we have a surrectangle satisfying $\partial^s = \partial$, we may assume that for every closed edge e of s and each associated cluster pair c_i^-, c_i^+ between transition points $\boldsymbol{x}_i, \boldsymbol{x}_{i+1}$, the core map of s fills $(\boldsymbol{x}_i, \boldsymbol{x}_{i+1})$ densely with the interpolant of c_i^- and c_i^+. Hence if surrectangles s_1 and s_2 have a common edge e, on which we obtain the same bi-trace and transition points from both surrectangles, then we may assume s_1 and s_2 agree on e.

Lemma 4.1 *Let s be a surrectangle with core map f and let g_1 and g_2 be order-preserving bijections $\mathbb{R} \to \mathbb{R}$. Then there is a surrectangle s' with core map $(x, y) \mapsto f(g_1(x), g_2(y))$, the same supplementary clusters and transition values as s, and a transition point $(g_1^{-1}(x), g_2^{-1}(y))$ for every transition point (x, y) of s. Moreover, s' yields the same biboundary as s.*

The proof of Lemma 4.1 is routine, and omitted. The proofs of the following four lemmas may be found in the appendix. Figure 3 illustrates the proofs of the last two.

Lemma 4.2 (ground biboundaries) *Let ∂ be a ground fabricated biboundary. Then there exists a surrectangle s such that $\partial^s = \partial$.*

Lemma 4.3 (joins) *Let ∂_1 and ∂_2 be biboundaries such that the vertical join $\partial_1 \oplus_- \partial_2$ exists, and suppose there exist surrectangles s_1 and s_2 with $\partial^{s_1} = \partial_1$ and $\partial^{s_2} = \partial_2$. Then there exists a surrectangle s such that $\partial^s = \partial_1 \oplus_- \partial_2$.*

Similarly for horizontal joins.

Lemma 4.4 (limits) *Let ∂^* be a southeastern limit of ∂_0 using $\partial_1, \partial_2, \partial_3$, and suppose there are surrectangles s_0, s_1, s_2, s_3 such that $\partial^{s_i} = \partial_i$, for $i < 4$. Then there exists a surrectangle s^* such that $\partial^{s^*} = \partial^*$. Similarly for northwestern limits.*

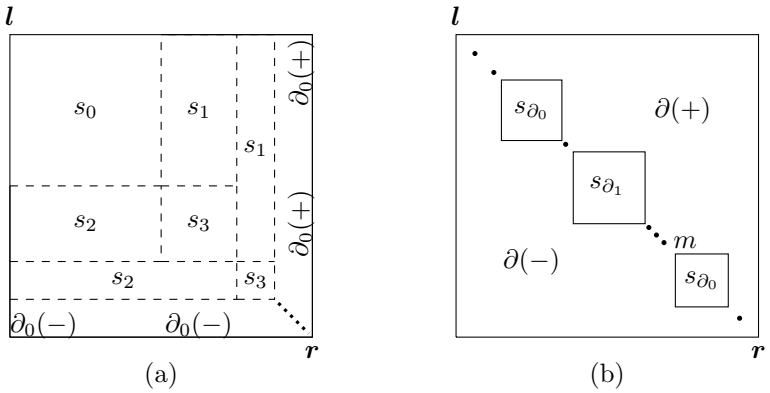

Figure 3. (a) A surrectangle for a southeastern limit of ∂_0 using $\partial_1, \partial_2, \partial_3$, and (b) a surrectangle for a shuffle of $\partial_0, \partial_1, \ldots$

Lemma 4.5 (shuffles) *Let ∂' be a shuffle of Δ, and suppose that for all $\partial \in \Delta$ there is a surrectangle s_∂ with $\partial^{s_\partial} = \partial$. Then there is a surrectangle s' with $\partial^{s'} = \partial'$.*

By Lemma 4.2, the set Δ of biboundaries that can be obtained from surrectangles contains the ground fabricated biboundaries. By Lemma 4.3, Lemma 4.4, and Lemma 4.5, it is closed under joins, limits, and shuffles. Hence Δ contains all fabricated biboundaries. This proves, as promised, the following.

Lemma 4.6 *Let ∂ be a fabricated biboundary. Then there exists a surrectangle s such that $\partial^s = \partial$.*

5 From surrectangles to fabricated biboundaries

In this section we show that every biboundary obtained from a surrectangle is a fabricated biboundary (Lemma 5.6). We do this by induction on the height of the surrectangle.

If a surrectangle s has height 0, then the upper and lower clusters of ∂^s are equal, and from that it is easy to see that all conditions for ∂^s to be a ground fabricated biboundary are satisfied.

In the remainder of this section, we assume that for every surrectangle s of height no greater than N, the biboundary ∂^s is fabricated, and aim to prove the statement for all surrectangles of height $N+1$. First, note that by Lemma 4.1, we only need prove the result for surrectangles with bounded domains. So henceforth we take bounded domains as an assumption.

Let s be a surrectangle with domain R, core map f, northwest corner $\boldsymbol{l}$, southeast corner $\boldsymbol{r}$, lower cluster c^-, and upper cluster c^+. Let $I^- = f^{-1}(c^-)$ be the subset of R that maps to the lower cluster, and $I^+ = f^{-1}(c^+)$. Define Γ^- to be the *closure* (in $\mathbb{R}^2$) of the intersection of the boundary of I^- with the interior of R. Define Γ^+ similarly. Elementary topology shows that Γ^- and Γ^+ are homeomorphic to closed (proper) line segments, meeting the boundary of R only at their endpoints, and linearly ordered by $\lhd$; see [7, Lemmas 2.10 and 2.11] for details. The notation $s{\restriction}_{[\boldsymbol{x},\boldsymbol{y}]}$ signifies the surrectangle formed in the obvious way by restriction of s to $[\boldsymbol{x},\boldsymbol{y}] \cap R$ (assuming this is nondegenerate). It is straightforward to check that $s{\restriction}_{[\boldsymbol{x},\boldsymbol{y}]}$ is indeed a surrectangle.

As before, proofs of lemmas are contained in the appendix.

Lemma 5.1 *If Γ^- and Γ^+ are disjoint then ∂^s is fabricated.*

If Γ^- and Γ^+ are *not* disjoint, define a binary relation $\equiv$ over $R \setminus (\downarrow I^- \cup \uparrow I^+)$ by letting $\boldsymbol{x} \equiv \boldsymbol{y} \iff$ for *all* nondegenerate rectangles $[\boldsymbol{w},\boldsymbol{z}] \subseteq [\boldsymbol{x} \wedge \boldsymbol{y}, \boldsymbol{x} \vee \boldsymbol{y}]$ the biboundary $\partial^{s{\restriction}_{[\boldsymbol{w},\boldsymbol{z}]}}$ is fabricated—clearly reflexive and symmetric, and also transitive (use a join of four biboundaries and the induction hypothesis), so an equivalence relation. By using joins and the induction hypothesis if necessary, we may assume $\boldsymbol{l}, \boldsymbol{r} \in R \setminus (\downarrow I^- \cup \uparrow I^+)$. We aim to show that $\boldsymbol{l} \equiv \boldsymbol{r}$.

Lemma 5.2 *Let $\boldsymbol{x}_0 \lhd \boldsymbol{x}_1 \lhd \ldots \in R \setminus (\downarrow I^- \cup \uparrow I^+)$ be an infinite sequence converging to $\boldsymbol{x}$. If for all i we have $\boldsymbol{y} \equiv \boldsymbol{x}_i$ then $\boldsymbol{y} \equiv \boldsymbol{x}$.*

Let $P = \{\boldsymbol{l}, \boldsymbol{r}\} \cup (\Gamma^- \cap \Gamma^+)$—a closed set, inheriting a linear order $\lhd$ from Γ^-, and a subset of $R \setminus (\downarrow I^- \cup \uparrow I^+)$. Define a binary relation $\approx$ over P as the smallest equivalence relation such that

(a) $\boldsymbol{p} \approx \boldsymbol{q}$ whenever $\boldsymbol{p}$ is an immediate $\lhd$-successor of $\boldsymbol{q}$,

(b) $\boldsymbol{p} \approx \boldsymbol{q}$ whenever $[\boldsymbol{p} \wedge \boldsymbol{q}, \boldsymbol{p} \vee \boldsymbol{q}]$ is degenerate,

(c) all equivalence classes are topologically closed (in $\mathbb{R}^2$, equivalently, in P).

Lemma 5.3 *$\boldsymbol{p} \approx \boldsymbol{q}$ implies $\boldsymbol{p} \equiv \boldsymbol{q}$.*

For any $\approx$-equivalence class E, let $l(E), r(E)$ be the extreme points with respect to $\lhd$ (equal if E is a singleton). As E is closed, $l(E), r(E) \in E$. Write $R(E)$ for the closed rectangle $[l(E) \wedge r(E), l(E) \vee r(E)]$ (a singleton if and only if E is a singleton). Since $\boldsymbol{p} \approx \boldsymbol{q}$ we know $l(E) \equiv r(E)$.

Lemma 5.4 *Either all elements of P are $\approx$-equivalent, or there are uncountably many singleton equivalence classes of $\approx$.*

Let $\boldsymbol{x} = r(E)$, where E is the $\approx$-equivalence class of $\boldsymbol{l}$, and let $\boldsymbol{y} = l(E')$ where E' is the $\approx$-equivalence class of $\boldsymbol{r}$. Since $\boldsymbol{x}$ is the most southeastern point in E, and has no immediate $\lhd$-successor in P, there is no point of Γ^- due south of $\boldsymbol{x}$, similarly no point of Γ^+ due east of $\boldsymbol{x}$. Dual conditions hold for $\boldsymbol{y}$. It follows that f^+ is constantly c^- on the south and west edges of $[\boldsymbol{x} \wedge \boldsymbol{y}, \boldsymbol{x} \vee \boldsymbol{y}]$, and f^- is constantly c^+ on the north and east edges.

Since $\boldsymbol{l} \approx \boldsymbol{x}$ and $\boldsymbol{y} \approx \boldsymbol{r}$ we know that $\boldsymbol{l} \equiv \boldsymbol{x}$ and $\boldsymbol{y} \equiv \boldsymbol{r}$. The next lemma shows that $\boldsymbol{x} \equiv \boldsymbol{y}$. We omit the proof, which is to check each of the conditions of Definition 3.4.

Lemma 5.5 *Suppose the upper cluster on the south and west edges of s is constantly c^-, and the lower cluster on the north and east edges of s is constantly c^+. Then ∂^s is a shuffle of $\partial^{s\restriction R(E)}$ where E ranges over $\approx$-equivalence classes.*

Using joins and the induction hypothesis, this proves $\boldsymbol{l} \equiv \boldsymbol{r}$. Hence we have obtained our goal.

Lemma 5.6 *Let s be any surrectangle. The biboundary ∂^s is fabricated.*

Combining Lemmas 4.6 and 5.6 we have the following.

Lemma 5.7 *A biboundary ∂ is of the form ∂^s for some surrectangle s if and only if ∂ is a fabricated biboundary.*

Now we are in a position to prove our first main result.

Theorem 5.8 *It is decidable whether a formula of the basic temporal language is valid on the frame consisting of two-dimensional Minkowski spacetime equipped with the irreflexive slower-than-light accessibility relation. The same is true with reflexive slower-than-light accessibility.*

Proof Decidability of validity is equivalent to decidability of satisfiability; we prove the latter.

We first show that satisfiability of ϕ on $(\mathbb{R}^2, \prec)$ is equivalent to the existence of an open surrectangle having some point assigned a maximal consistent set containing ϕ. An open surrectangle consists only of its core map f. Given a valuation v on $(\mathbb{R}^2, \prec)$ and a point $\boldsymbol{x}$ at which ϕ holds, define f by $f(\boldsymbol{y}) = \{\psi \in \mathrm{Cl}(\phi) \mid (\mathbb{R}^2, \prec), v, \boldsymbol{y} \models \psi\}$, and then f will be a surrectangle on $\mathbb{R}^2$ with $\phi \in f(\boldsymbol{x})$. Conversely, given such an f, for any propositional variable

p if p appears in ϕ, define $v(p) = \{\boldsymbol{y} \mid p \in f(\boldsymbol{y})\}$, and otherwise define $v(p)$ arbitrarily. Then ϕ holds at $\boldsymbol{x}$ under valuation v.

Let $NE(x, y) = \{(x', y') \in \mathbb{R}^2 \mid x \leq x' \text{ and } y \leq y'\}$ and define $NW(\boldsymbol{x})$, $SE(\boldsymbol{x})$, and $SW(\boldsymbol{x})$ similarly, in the evident way. The existence of an open surrectangle having some point $\boldsymbol{x}$ assigned a maximal consistent set containing ϕ is equivalent to the existence of four surrectangles with domains $NE(\boldsymbol{x})$, $NW(\boldsymbol{x})$, $SE(\boldsymbol{x})$, and $SW(\boldsymbol{x})$, agreeing on their shared edges and at $\boldsymbol{x}$, and with ϕ contained in the maximal consistent set assigned to $\boldsymbol{x}$. By Lemma 5.7, this is in turn equivalent to the existence of four fabricated biboundaries of the appropriate types that match up in the appropriate way.

Since satisfiability of ϕ is equivalent to the existence of a set of four fabricated biboundaries with properties that are easily checked, and the set of all fabricated biboundaries is finite and computable, the satisfiability problem is decidable. This completes the proof of the irreflexive case.

The reflexive case follows by the reduction given by recursively replacing subformulas of ϕ of the form $\mathbf{F}\varphi$ with $\varphi \wedge \mathbf{F}\varphi$, and similarly for $\mathbf{P}$. □

6 A PSPACE procedure for fabricated biboundaries

In this section, we refine the decidability results of Theorem 5.8 to show the validity problems are PSPACE-complete.

Theorem 6.1 *On the frame consisting of two-dimensional Minkowski spacetime equipped with the irreflexive slower-than-light accessibility relation, the set of validities of the basic temporal language is* PSPACE*-complete. The same is true with reflexive slower-than-light accessibility.*

As mentioned in the introduction, for the reflexive frame, the validities of the purely *modal* fragment of the basic temporal language form **S4.2**, and for the irreflexive frame, **OI.2**. These are both known to be PSPACE-complete [13], so the validity problems for the *entire* basic temporal language are PSPACE-hard.

We provide a nondeterministic polynomial space algorithm for satisfiability, for the irreflexive frame. Hence validity is in coNPSPACE. By the Immerman–Szelepcsényi theorem [8,18], coNPSPACE = NPSPACE, and by Savitch's theorem [12], NPSPACE = PSPACE, giving the result. The reflexive case follows by the same reduction as before.

Throughout this section, let the length of ϕ be n. By structural induction, the length of a formula bounds the number of its subformulas, so the cardinality of $\mathrm{Cl}(\phi)$ is linear in n. Hence any maximal consistent set—a subset of $\mathrm{Cl}(\phi)$—can be stored using a linear number of bits, and $|\mathsf{MCS}|$ is at most exponential in n. All pertinent information about any cluster c can also be stored in a linear number of bits, for we only need record $\{\psi \in \mathrm{Cl}(\phi) \mid \exists m \in c : \psi \in m\}$. The maximal length of a chain of distinct clusters or irreflexive members of MCS is also linear in n. Hence any biboundary can be stored using a quadratic number of bits, and the number of biboundaries is exponential in n.

Having nondeterministically chosen a bit string representing some $m \subseteq$

$\mathrm{Cl}(\phi)$—a putative maximal consistent set—the conjunct over all formulas in m is quadratic in n, so its satisfiability can be determined in polynomial space [17]. Hence we can determine if a chosen bit string represents a maximal consistent set using polynomial space.

We can similarly nondeterministically 'guess' sets of the form $\{\psi \in \mathrm{Cl}(\phi) \mid \exists m \in c : \psi \in m\}$ for some cluster c, using polynomial space. To do this, first guess a maximal consistent set m and check it is reflexive; keep m in memory. Then one-by-one for each bit string representing some $m' \subseteq \mathrm{Cl}(\phi)$, check $m' \in \mathsf{MCS}$ and that $m \lesssim m' \lesssim m$. If so, add all elements of m' to an ongoing collection of formulas, discarding m' after each iteration.

From the preceding discussion, it is clear we can also determine if a chosen string of bits represents a *biboundary* using polynomial space.

The algorithm for satisfiability of ϕ first nondeterministically chooses four bit strings and checks they represent four compatible biboundaries, with ϕ at the appropriate corners—performed using polynomial space. Then for each of these in turn, it is checked whether the biboundary is fabricated. The remainder of the proof is devoted to showing that this check can be performed using polynomial space. The procedure to do this is shown in Algorithm 1.

Algorithm 1 Nondeterministic procedure to decide whether ∂ is fabricated

procedure FABRICATED(∂) **choose** either
  option 0
    check ∂ is ground;
  option 1
    choose ∂_1, ∂_2; **check** they are biboundaries
    check their join (**choose** some direction) is ∂; **release** ∂
    check FABRICATED(∂_1); **tail-call** FABRICATED(∂_2)
  option 2
    choose $\partial_1, \partial_2, \partial_3, \partial_4$; **check** they are biboundaries
    check ∂ is the limit (**choose** direction) of ∂_1 using $\partial_2, \partial_3, \partial_4$; **release** ∂
    check FABRICATED(∂_2), FABRICATED(∂_3); **release** ∂_2, ∂_3; **check** FABRICATED(∂_4); **tail-call** FABRICATED(∂_1)
  option 3
    choose $k \in \{0, 1, \ldots, n\}$, $\partial_1, \ldots, \partial_k$; **check** they are biboundaries
    choose $m_1, \ldots, m_n$; **check** they are maximal consistent sets
    check ∂ is the shuffle of $\partial_1, \ldots, \partial_k$ using $m_1, \ldots, m_n$; **release** ∂
    for $i = 1, \ldots, k$ **do**
      check FABRICATED(∂_i)
    end for
end procedure

We assume that during execution of Algorithm 1, the formula ϕ is a global constant (and therefore is n too). All choices are made nondeterministically. At a **check** the algorithm fails if the check fails, otherwise it proceeds. A **tail-call** uses tail recursion. The algorithm succeeds if it terminates without failure.

Clearly Algorithm 1 can only succeed for a biboundary ∂ if ∂ is fabricated. Let $S(N)$ be the amount of space required by Algorithm 1 to succeed for any fabricated biboundary of height N. We prove that $S(N) = \mathcal{O}((N+1)n^3)$ (as a function of N and n) by induction on N. Hence Algorithm 1 requires space $\mathcal{O}(n^4)$ to check any biboundary.

If ∂ is a fabricated biboundary of height 0, then it is ground, so success can be achieved by taking the first branch, requiring space $\mathcal{O}(n^2)$.

For the inductive step, we show that $S(N+1) \leq S(N) + \mathcal{O}(n^3)$. The proof has a similar structure to the proof in the previous section. Let ∂ be a fabricated biboundary of height $N+1$, and s be a surrectangle such that $\partial^s = \partial$. Define Γ^-, Γ^+, $\boldsymbol{l}$, and $\boldsymbol{r}$ as in the previous section. Let δ be the space required to store any biboundary. Thus δ is a quadratic function of n, independent of N.

As always, proofs of lemmas are relegated to the appendix.

Lemma 6.2 *If Γ^- and Γ^+ are disjoint, the space required is bounded by $S(N)+\delta$.*

Lemma 6.3 *If Γ^- and Γ^+ intersect only on the boundary of s, the space required is bounded by $S(N) + 3\delta$.*

As in the previous section, let $P = \{\boldsymbol{l}, \boldsymbol{r}\} \cup (\Gamma^- \cap \Gamma^+)$. Additionally, let $\Delta = \{\partial^{s\restriction_{[\boldsymbol{x}\wedge\boldsymbol{y}, \boldsymbol{x}\vee\boldsymbol{y}]}} \mid \boldsymbol{x}, \boldsymbol{y} \in P,$ and $\boldsymbol{y}$ is an immediate $\lhd$-successor of $\boldsymbol{x}\}$.

Lemma 6.4 *If the biboundary of s is a shuffle of a subset Δ' of Δ, the space required is bounded by $S(N) + 3\delta + n\delta$.*

If $\boldsymbol{l}$ and $\boldsymbol{r}$ are each the limit of elements of P in the interior of s, then the biboundary of s is the shuffle of Δ', using M, where $\Delta' = \{\partial^{s\restriction_{[\boldsymbol{x}\wedge\boldsymbol{y}, \boldsymbol{x}\vee\boldsymbol{y}]}} \mid \boldsymbol{x}, \boldsymbol{y} \in P \setminus \{\boldsymbol{l}, \boldsymbol{r}\},$ and $\boldsymbol{y}$ is the successor of $\boldsymbol{x}\}$ and $M = P \setminus \{\boldsymbol{l}, \boldsymbol{r}\}$.

Now suppose only one of $\boldsymbol{l}$, $\boldsymbol{r}$ is a limit of elements of P in the interior of s. Say $\boldsymbol{l}$ is such a limit, but $\boldsymbol{r}$ is not. If $\boldsymbol{r}$ has a direct predecessor $\boldsymbol{r'}$, then $\boldsymbol{r'}$ is in the interior of s. If there are points of P (strictly) due north of $\boldsymbol{r}$, let the northmost one be $\boldsymbol{w}$. Then either $\boldsymbol{w}$ is a limit, and using a join we can reduced to the case of the previous paragraph, or there is a point $\boldsymbol{r'}$ either an immediate predecessor or due west of $\boldsymbol{w}$. In the second case, $\boldsymbol{r'}$ is in the interior of $\boldsymbol{s}$. Similarly if there are points of P due west of $\boldsymbol{r}$. In each case we obtain an $\boldsymbol{r'}$ in the interior of P and the biboundary ∂_{SE} of $s\restriction_{[\boldsymbol{r}\wedge\boldsymbol{r'}, \boldsymbol{r}\vee\boldsymbol{r'}]}$ can be checked in $S(N) + 3\delta$, by Lemma 6.3.

We are going to obtain the biboundary of s as the join of four biboundaries. In the southeast corner is ∂_{SE}. If the biboundary of $s\restriction_{[\boldsymbol{l}\wedge\boldsymbol{r'}, \boldsymbol{l}\vee\boldsymbol{r'}]}$ is of reduced height then we are done. Otherwise, the other three biboundaries we use are modifications of restrictions. For the northwest corner, take the biboundary of $s\restriction_{[\boldsymbol{l}\wedge\boldsymbol{r'}, \boldsymbol{l}\vee\boldsymbol{r'}]}$. Modify its south edge to the bi-trace (c^-, c^-) and its east edge to (c^+, c^+). This is still a biboundary; call it ∂_{NW}. For the northeast corner, take the biboundary of $s\restriction_{[\boldsymbol{r'}, \boldsymbol{t}]}$ (which is ground). Modify its west edge to the bi-trace (c^+, c^+). This is still a biboundary (and still ground). Similarly for the southwest corner.

The biboundary of s is the join of these four biboundaries. It only remains to show ∂_{NW} is fabricated, and that this can be checked efficiently. As ∂_{NW} is of full height, its upper cluster is c^+ and lower c^-. By appealing to the facts that s is a surrectangle and $s{\restriction}_{[\boldsymbol{l}\wedge\boldsymbol{r}',\boldsymbol{l}\vee\boldsymbol{r}']}$ is also a surrectangle, we can see that ∂_{NW} is the shuffle of Δ', using M, where $\Delta' = \{\partial^{s{\restriction}_{[\boldsymbol{x}\wedge\boldsymbol{y},\boldsymbol{x}\vee\boldsymbol{y}]}} \mid \boldsymbol{x},\boldsymbol{y} \in (P\backslash\{\boldsymbol{l},\boldsymbol{r}\})\cap [\boldsymbol{l}\wedge\boldsymbol{r}',\boldsymbol{l}\vee\boldsymbol{r}'],$ and $\boldsymbol{y}$ is the successor of $\boldsymbol{x}\}$, and $M = (P\setminus\{\boldsymbol{l},\boldsymbol{r}\})\cap[\boldsymbol{l}\wedge\boldsymbol{r}',\boldsymbol{l}\vee\boldsymbol{r}']$. The set Δ' is a subset of Δ, so by Lemma 6.4, the biboundary ∂_{NW} can be checked using no more than $S(N)+3\delta+n\delta$ space. Hence ∂ can been checked within $S(N)+\mathcal{O}(n^3)$, as promised.

The case where neither $\boldsymbol{l}$ nor $\boldsymbol{r}$ are limits is similar.

7 Halpern–Shoham logic

In this section we explain how the PSPACE procedure for the temporal validities of $(\mathbb{R}^2,\prec)$ can be modified to procedures performing the same function for certain Halpern–Shoham logics of intervals on the real line.

In Halpern–Shoham logic, intervals are identified with pairs (x,y) of points, with either $x<y$ (**strict interval semantics**) or $x\leq y$ (**non-strict semantics**).[7] There are thirteen different atomic relations that may hold between two strict intervals. Following Allen [1], we call these `equals`, `before`, `after`, `during`, `contains`, `starts`, `started_by`, `finishes`, `finished_by`, `meets`, `met_by`, `overlaps`, and `overlapped_by`.[8] Modalities that may be included in a Halpern–Shoham logic are any corresponding to a relation given by the union of some of these thirteen.

Let $H_<$ be the open half-plane $\{(x,y)\in\mathbb{R}^2 \mid x<y\}$. The frame $(H_<,\prec)$ is precisely the frame of strict intervals of $\mathbb{R}$ with the relation `overlaps`$\cup$`meets`$\cup$`before`. Hence the temporal logic of $(H,\prec)$ is the strict Halpern–Shoham logic of $\mathbb{R}$ with two modalities corresponding to this relation and its converse. As in [5], it can be shown that an arbitrary finitely generated directed partial order is a p-morphic image of the reflexive closure of $(H_<,\prec)$. Hence an arbitrary purely modal formula (not using the past modality) is in **S4.2** if and only if it is valid for the reflexive closure of $(H_<,\prec)$, and hence the validity of temporal formulas over this frame is PSPACE-hard. The PSPACE-hardness of irreflexive $(H_<,\prec)$ itself, follows.

A **surtriangle** is similar to a surrectangle, but the domain of the core map f is either $\{(x,y)\in H_< \mid a\leq x,\ y\leq b\}$, for some $a<b$, or a similar set where either $a\leq x$ is replaced by $a<x$, or $y\leq b$ is replaced by $y<b$, or both. A finite sequence of upper supplementary clusters and transition points are defined along $\{(x,b)\in H_< \mid a<x<b\}$ if contained in the domain of f, and similarly for $\{(a,y)\in H_< \mid a<y<b\}$ with lower supplementary clusters.

---

[7] Thus there is no distinction between open, half-open, and closed intervals, and unbounded intervals are not present.

[8] When point-intervals are present, pairs of intervals of the form $((x,x),(x,y))$ for $x<y$ are considered by Halpern and Shoham to stand in the relation `starts` (and not in `meets`). Similarly for `started_by`, `finishes`, and `finished_by`.

A biboundary is called **triangular** if its domain is disjoint from $\{b, t, r, S, E\}$. A surtriangle determines a triangular biboundary in the obvious way. If τ_1 is a triangular biboundary containing its northern edge, τ_2 a triangular biboundary containing its western edge, and ∂ a (rectangular) biboundary closed on its southern and eastern edges, such that $\tau_1(N) = \partial(S)$ and $\partial(W) = \tau_2(E)$, then the join $J(\tau_1, \partial, \tau_2)$ is the triangular biboundary formed by joining the three parts together. (The inconsequential southeast corner of ∂ is discarded.) A triangular biboundary is **fabricated** if it is either ground, or the join of a ground triangular biboundary, a fabricated rectangular biboundary, and a triangular biboundary of strictly smaller depth. The proofs of the following lemmas and theorem are similar to the proofs of Lemma 5.7, Theorems 5.8, and Theorem 6.1.

Lemma 7.1 *A triangular biboundary is fabricated if and only if it is the triangular biboundary of some surtriangle.*

Lemma 7.2 *There is a* PSPACE *algorithm to determine whether a triangular biboundary is fabricated.*

Theorem 7.3 *Let $H_< = \{(x, y) \in \mathbb{R}^2 \mid x < y\}$, and let R be either the relation given by $(x, y)R(x', y') \iff x < x'$ and $y < y'$ (so $R =$ `overlaps` $\cup$ `meets` $\cup$ `before`) or the reflexive closure of this relation (so $R =$ `equals` $\cup$ `overlaps` $\cup$ `meets` $\cup$ `before`). Then the temporal logic of $(H_<, R)$ is* PSPACE*-complete.*

Appendix

Proof of Lemma 4.2 (ground biboundaries) Let R be the appropriate rectangle (given ∂) with $(0, 1) \times (0, 1) \subseteq R \subseteq [0, 1] \times [0, 1]$. Define $f \colon R \to \mathsf{MCS}$ as filling the interior of R densely with the cluster $\partial(+)$. Define f on any corners of R in the evident way.

If ∂ is defined on N, with $\partial(N) = (c_0^-, c_0^+, b_1, \ldots, c_n^+)$ say, then define $f(\frac{i}{n+1}, 1) = b_i$ for each i, and fill each line segment $((\frac{i}{n+1}, 1), (\frac{i+1}{n+1}, 1))$ densely with the interpolant of c_i^- and c_i^+. Define the supplementary clusters for this edge to be $c_0^+, \ldots, c_n^+$, and the sequence of transition points to be $(\frac{1}{n+1}, 1), \ldots, (\frac{n}{n+1}, 1)$. Similarly if ∂ is defined on S, E, or W.

The conditions on biboundaries in general, together with those for ground fabricated ones in particular, ensure our construction satisfies all conditions necessary to be a surrectangle.

Proof of Lemma 4.3 (joins) We know the bi-trace obtained from the northern edge of s_1 equals that obtained from the southern edge of s_2, and in particular these edges have the same number of transition points. Then by Lemma 4.1, we may assume these two edges and their cluster pairs, transition points, transition values, and values of f at any end-points coincide. As we remarked in the opening of Section 4, this is sufficient to assume s_1 and s_2 agree on the common edge.

We define s with domain the union of the domains of s_1 and s_2 in the obvious way. Then it is straightforward to check that s forms a surrectangle and that $\partial^s = \partial_1 \oplus_- \partial_2$.

The proof for horizontal joins is completely analogous.

Proof of Lemma 4.4 (limits) Since ∂^* is a southeastern limit of ∂_0 using $\partial_1, \partial_2, \partial_3$, the vertical join of ∂_2, ∂_0 is defined, and similarly for other joins. Figure 3(a) illustrates how a surrectangle for ∂^* can be constructed using transformed copies of the surrectangles s_0, s_1, s_2, s_3. The domain of the surrectangle is a rectangle with interior $(0,1) \times (0,1)$. The upper-left quadrant $(0, \frac{1}{2}) \times (\frac{1}{2}, 1)$ together with boundaries dictated by ∂_0 is a copy of s_0. Each rectangle $(0, 1-\frac{1}{2^k}) \times (\frac{1}{2^{k+1}}, \frac{1}{2^k})$, plus appropriate boundaries, is a copy of s_2, for $k \geq 1$. Use Lemma 4.1 to ensure agreement along common boundaries. Transformed copies of s_1 are used dually. The square $[1-\frac{1}{2^k}, 1-\frac{1}{2^{k+1}}] \times [\frac{1}{2^{k+1}}, \frac{1}{2^k}]$ is a copy of s_3, for $k \geq 1$. This covers the interior of the unit square, and northern, western edges, and northwest corner, if included in the domain of ∂^*. Use ∂^* to define the corners b, r, t if included in the domain of ∂^*. Use the interpolant of the lower and upper clusters of $\partial^*(S)$ densely along the southern edge of the unit square, if S is included in the domain of ∂^*, similarly for the eastern edge.

The proof for northwestern limits is completely analogous.

Proof of Lemma 4.5 (shuffles) The closure of the rectangle of s' will be $[0,1] \times [0,1]$. Let $M \subseteq \mathsf{MCS}$ be such that the conditions in Definition 3.4, the definition of a shuffle, are satisfied. Let d be the open diagonal line-segment $((0,1),(1,0))$. We describe how to construct s' in stages, as illustrated in Figure 3(b). First fill the area of $(0,1) \times (0,1)$ below d densely with $\partial'(-)$, and fill the area above this diagonal densely with $\partial'(+)$. Fill d with any element of M.

For any edge that ∂' indicates should be closed in s', assign the appropriate sequence of supplementary clusters, and evenly spaced transition points. Assign the appropriate transition values at the transition points, and between transition points fill the edge densely with the appropriate interpolant. If ∂' indicates any corners are required, assign them the appropriate maximal consistent set.

Next, we successively modify the interior of the construction. Each point will be updated at most once, so this process has a well-defined limit. We maintain a finite set S of disjoint open subsegments of d, initialised to $\{d\}$. At a later stage, pick $d' \in S$ and $b \in \Delta \cup M$. If $b \in \Delta$, reassign the closed rectangle whose diagonal is the central third of d', using the surrectangle s_b. Otherwise, reassign the midpoint $\boldsymbol{m}$ of d' with the maximal consistent set b. In either case, the segment d' is replaced by two in S—in the first case the open initial and final thirds of d', in the second, the two halves of $d' \setminus \{\boldsymbol{m}\}$. Schedule the choices of d' and b so that for every segment d' occurring in this construction, for every $b \in \Delta \cup M$, a choice d^*, b is eventually selected, where d^* is a subsegment of d'.

The limit of this process gives a well-defined map f from the appropriate subset of $[0,1] \times [0,1]$ to MCS, and supplementary clusters where appropriate. We argue that together these form a surrectangle, s'.

The constraints placed on ∂' by the ordering conditions in the definition of a

biboundary, together with the ordering conditions in the definition of a shuffle, ensure that f is preorder preserving. Our construction has the property that every point below d that did not undergo reassignment has an open subset of points in its future that also did not undergo reassignment, *and* has a copy of each s_∂, and each $m \in M$ in its future. Noting this, it is straightforward to check that all future defects are either resolved internally or passed upwards. The other conditions in the definition of a surrectangle are also straightforward to check.

Clearly the biboundary of s' is ∂', as required.

Proof of Lemma 5.1 If Γ^- and Γ^+ are disjoint, then by the boundedness of R, and since they are closed, they are bounded away from each other, with a bound $\varepsilon > 0$ say. Then we may divide R into a *finite* grid of rectangles each with a diagonal shorter than ε. The restriction of s to such a rectangle has height N or less, so yields a fabricated biboundary. Then ∂^s is a join of such biboundaries, and is therefore itself fabricated.

Proof of Lemma 5.2 Using the inductive hypothesis, we may assume that the biboundary of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x},\boldsymbol{y}\vee\boldsymbol{x}]}$ is of height $N+1$, that is, both Γ^- and Γ^+ intersect the interior of $[\boldsymbol{x}\wedge\boldsymbol{y},\boldsymbol{x}\vee\boldsymbol{y}]$. As each $\boldsymbol{x}_i$ is in $R \setminus (\downarrow I^- \cup \uparrow I^+)$, we know $s{\restriction}_{[\boldsymbol{x}_i,\boldsymbol{x}\vee\boldsymbol{y}]}$ (if $[\boldsymbol{x}_i,\boldsymbol{x}\vee\boldsymbol{y}]$ is nondegenerate) has height at most N, so its biboundary is fabricated, similarly for $s{\restriction}_{[\boldsymbol{x}\wedge\boldsymbol{y},\boldsymbol{x}_i]}$. Hence, using joins, it suffices to prove the biboundary of $s{\restriction}_{[\boldsymbol{x}_i\wedge\boldsymbol{x},\boldsymbol{x}_i\vee\boldsymbol{x}]}$ is fabricated, for some i. Then it follows using Lemma 5.1 and the inductive hypothesis that we may assume Γ^- and Γ^+ meet at $\boldsymbol{x}$, that no point due west of $\boldsymbol{x}$ is in Γ^-, and that no point due north of $\boldsymbol{x}$ is in Γ^+. Together, these assumptions imply the upper cluster on the southern edge of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x},\boldsymbol{y}\vee\boldsymbol{x}]}$ is constantly c^-, and the lower cluster on the eastern edge is constantly c^+.

Since there are only finitely many biboundaries, by taking a subsequence, we can assume the biboundary of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x}_i,\boldsymbol{y}\vee\boldsymbol{x}_i]}$ is constant, and by our previous assumptions we may also assume that $\partial^{s{\restriction}_{[\boldsymbol{x}_0,\boldsymbol{y}\vee\boldsymbol{x}_1]}}$ has constant lower cluster $\partial^s(+)$ on its eastern edge, and $\partial^{s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x}_1,\boldsymbol{x}_0]}}$ has constant upper cluster $\partial^s(-)$ on its southern edge. Then the biboundary of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x},\boldsymbol{y}\vee\boldsymbol{x}]}$ is a southeastern limit of the biboundary of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x}_0,\boldsymbol{y}\vee\boldsymbol{x}_0]}$ using the biboundaries of $s{\restriction}_{[\boldsymbol{x}_0,\boldsymbol{y}\vee\boldsymbol{x}_1]}, s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x}_1,\boldsymbol{x}_0]}$, and $s{\restriction}_{[\boldsymbol{x}_0\wedge\boldsymbol{x}_1,\boldsymbol{x}_0\vee\boldsymbol{x}_1]}$. All of these are fabricated, by the hypothesis that $\boldsymbol{y} \equiv \boldsymbol{x}_1$. Hence the biboundary of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x},\boldsymbol{y}\vee\boldsymbol{x}]}$ is fabricated. Furthermore the biboundary of the restriction of s to any nondegenerate $[\boldsymbol{w},\boldsymbol{z}] \subsetneq [\boldsymbol{y}\wedge\boldsymbol{x},\boldsymbol{y}\vee\boldsymbol{x}]$ is clearly fabricated, so $\boldsymbol{y} \equiv \boldsymbol{x}$.

Proof of Lemma 5.3 It is clear that if $[\boldsymbol{p}\wedge\boldsymbol{q},\boldsymbol{p}\vee\boldsymbol{q}]$ is degenerate then $\boldsymbol{p} \equiv \boldsymbol{q}$. If $\boldsymbol{p}$ is an immediate $\lhd$-successor of $\boldsymbol{q}$ pick $\boldsymbol{y} \lhd \boldsymbol{x}_0 \lhd \boldsymbol{x}_1 \lhd \dots$ in the interior of $[\boldsymbol{p}\wedge\boldsymbol{q},\boldsymbol{p}\vee\boldsymbol{q}]$, converging to $\boldsymbol{p}$. Since $\boldsymbol{y},\boldsymbol{x}_i$ are in the interior of the rectangle, the upper and lower clusters of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{x}_i,\boldsymbol{y}\vee\boldsymbol{x}_i]}$ are bounded away, hence the biboundary is fabricated. By Lemma 5.2 the biboundary of $s{\restriction}_{[\boldsymbol{y}\wedge\boldsymbol{p},\boldsymbol{y}\vee\boldsymbol{p}]}$ is fabricated. Similarly, by considering northwestern limits, the biboundary of $s{\restriction}_{[\boldsymbol{p}\wedge\boldsymbol{q},\boldsymbol{p}\vee\boldsymbol{q}]}$ is fabricated and $\boldsymbol{p} \equiv \boldsymbol{q}$. By Lemma 5.2 and its diagonal dual, the $\equiv$-equivalence classes are topologically closed. The lemma follows.

Proof of Lemma 5.4 The upper boundary Γ^- of I^- is homeomorphic to the closed unit interval $[0,1]$, and such a homeomorphism f will map $R(E) \cap \Gamma^-$ to an interval, for any $\approx$-equivalence class E. The set I^- is the disjoint union of the closed sets $R(E) \cap \Gamma^-$ as E ranges over $\approx$-equivalence classes. Hence $\{f(R(E) \cap \Gamma^-) \mid E \in P/\approx\}$ is a partition of $[0,1]$ into closed intervals. It follows that if there is more than one $\approx$-equivalence class, then uncountably many of these closed intervals are singletons; see [7, Lemma 2.9]. For each singleton interval, the corresponding $\approx$-equivalence class must itself be a singleton.

Proof of Lemma 6.2 If Γ^- and Γ^+ are disjoint, they are bounded apart, by ε say. Then either Γ^- is bounded away from north edge of s by $\frac{\varepsilon}{\sqrt{2}}$, or Γ^+ is bounded away from the *west* edge by $\frac{\varepsilon}{\sqrt{2}}$. Without loss of generality, assume the latter. Let s_1 be the restriction of s to the region at a distance no more than $\frac{\varepsilon}{\sqrt{2}}$ from the west edge of s, and s_2 be the restriction of s to the region *at least* $\frac{\varepsilon}{\sqrt{2}}$ from the west edge. Then $\partial = \partial^{s_1} \oplus_| \partial^{s_2}$, and ∂^{s_1} has height no greater than N. Hence it is possible to choose option 1 with ∂^{s_1}, ∂^{s_2}, and check ∂_1 is fabricated using no more than $S(N) + \delta$ space. The tail-recursive call is to the biboundary of s_2, which is smaller than s in one dimension by $\frac{\varepsilon}{\sqrt{2}}$ and—if still of height $N+1$—has upper and lower clusters still bounded apart by ε. Hence iterating the described choice scheme requires space $S(N)+\delta$ each iteration, and must eventually result in a reduction of the height of the second biboundary s_2 in the join, after which the recursive call to s_2 can succeed with $S(N)$ space. The maximum space required is $S(N) + \delta$, as claimed.

Proof of Lemma 6.3 If Γ^- and Γ^+ intersect only at *one* point on the boundary of s, by Lemma 6.2, we may assume (without exceeding our space allowance) it is either $\boldsymbol{l}$ or $\boldsymbol{r}$—without loss of generality, assume $\boldsymbol{r}$. Then for some $\boldsymbol{x}$ and $\boldsymbol{y}$ in the interior of s, the biboundary ∂ is the limit of $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{x}]}$ using $\partial^{s\restriction[\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{y}]}$, $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{y},\boldsymbol{x}]}$, and $\partial^{s\restriction[\boldsymbol{x}\wedge\boldsymbol{y},\boldsymbol{x}\vee\boldsymbol{y}]}$, and both $\partial^{s\restriction[\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{y}]}$ and $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{y},\boldsymbol{x}]}$ are of height no greater than N. Whilst checking $\partial^{s\restriction[\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{y}]}$ and $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{y},\boldsymbol{x}]}$, at most $S(N)$ space is needed, plus 3δ to store the other three biboundaries. Whilst checking $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{y},\boldsymbol{x}]}$, by Lemma 6.2, at most $S(N)+\delta$ is needed, plus δ to store $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{x}]}$. Lastly, $\partial^{s\restriction[\boldsymbol{l}\wedge\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{x}]}$ is checked, and $S(N)+\delta$ is sufficient by Lemma 6.2. The maximum space needed is bounded by $S(N)+3\delta$.

If Γ^- and Γ^+ intersect at *two* points on the boundary of s, then using the same choice scheme, $\Gamma^- \cap [\boldsymbol{l}\wedge\boldsymbol{x}, \boldsymbol{l}\vee\boldsymbol{x}]$ and $\Gamma^+ \cap [\boldsymbol{l}\wedge\boldsymbol{x}, \boldsymbol{l}\vee\boldsymbol{x}]$ will intersect only on the boundary of $s\restriction_{[\boldsymbol{l}\wedge\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{x}]}$, and at at most one point. Then by the previous case, checking the biboundary of $s\restriction_{[\boldsymbol{l}\wedge\boldsymbol{x},\boldsymbol{l}\vee\boldsymbol{x}]}$ can be done with $S(N)+3\delta$ space. As this is done last, the bound $S(N)+3\delta$ remains intact.

Proof of Lemma 6.4 If a biboundary ∂' is a shuffle over Δ' using $M' \subseteq \mathsf{MCS}$ then it is the shuffle of any subset Δ'' of Δ' using any subset M'' of M', so long as any future defect of $\partial'(-)$ is passed up somehow to Δ'' or M'', and any past defect of $\partial'(+)$ is passed down to Δ'' or M''. There are at most n defects, past or future, so any shuffle is the shuffle of at most n biboundaries, using at most n auxiliary maximal consistent sets. Hence by choosing option 3, Algorithm 1 can succeed in checking ∂'. Since each of the up to n subchecks

requires no more than $S(N) + 3\delta$ space, checking ∂' can be accomplished using $S(N) + 3\delta + n\delta$ space.

References

[1] Allen, J. F., *Maintaining knowledge about temporal intervals*, Communications of the ACM **26** (1983), pp. 832–843.

[2] Bresolin, D., D. Della Monica, V. Goranko, A. Montanari and G. Sciavicco, *Decidable and undecidable fragments of Halpern and Shoham's interval temporal logic: Towards a complete classification*, in: I. Cervesato, H. Veith and A. Voronkov, editors, *proceedings of 15th conference on Logic for Programming, Artificial Intelligence, and Reasoning* (2008), pp. 590–604.

[3] Bresolin, D., A. Kurucz, E. Muñoz Velasco, V. Ryzhikov, G. Sciavicco and M. Zakharyaschev, *Horn fragments of the Halpern–Shoham interval temporal logic*, ACM Transactions on Computational Logic **18** (2017), pp. 22:1–22:39.

[4] Bresolin, D., A. Montanari, P. Sala and G. Sciavicco, *What's decidable about Halpern and Shoham's interval logic? The maximal fragment ABBL*, in: *proceedings of 26th IEEE symposium on Logic in Computer Science*, 2011, pp. 387–396.

[5] Goldblatt, R., *Diodorean modality in Minkowski spacetime*, Studia Logica **39** (1980), pp. 219–236.

[6] Halpern, J. Y. and Y. Shoham, *A propositional modal logic of time intervals*, Journal of the ACM **38** (1991), pp. 935–962.

[7] Hirsch, R. and M. Reynolds, *The temporal logic of two dimensional Minkowski spacetime is decidable*, The Journal of Symbolic Logic (2018), (in press), arXiv:1507.04903.

[8] Immerman, N., *Nondeterministic space is closed under complementation*, SIAM Journal on Computing **17**, pp. 935–938.

[9] Marcinkowski, J. and J. Michaliszyn, *The ultimate undecidability result for the Halpern–Shoham logic*, in: *proceedings of 26th IEEE symposium on Logic in Computer Science*, 2011, pp. 377–386.

[10] Marcinkowski, J. and J. Michaliszyn, *The undecidability of the logic of subintervals*, Fundamenta Informaticae **XX** (2013), pp. 1–25.

[11] Montanari, A. and P. Sala, *An optimal tableau system for the logic of temporal neighborhood over the reals*, in: *proceedings of 19th symposium on Temporal Representation and Reasoning*, 2012, pp. 39–46.

[12] Savitch, W. J., *Relationships between nondeterministic and deterministic tape complexities*, Journal of Computer and System Sciences **4** (1970), pp. 177–192.

[13] Shapirovsky, I., *On PSPACE-decidability in transitive modal logics*, in: R. Schmidt, I. Pratt-Hartmann, M. Reynolds and H. Wansing, editors, *Advances in Modal Logic, papers from the 5th conference, held in Manchester (UK), 2004* (2005), pp. 269–287.

[14] Shapirovsky, I., *Simulation of two dimensions in unimodal logics*, in: P. Balbiani, N.-Y. Suzuki, F. Wolter and M. Zakharyaschev, editors, *Advances in Modal Logic, papers from the 8th conference, held in Moscow, 2010* (2010), pp. 371–391.

[15] Shapirovsky, I. and V. B. Shehtman, *Chronological future modality in Minkowski spacetime*, in: F. Wolter, H. Wansing, M. de Rijke and M. Zakharyaschev, editors, *Advances in Modal Logic, papers from the 4th conference, held in Toulouse, 2002* (2003), pp. 437–460.

[16] Shehtman, V. B., *Modal logics of domains on the real plane*, Studia Logica **42** (1983), pp. 63–80.

[17] Spaan, E., *The complexity of propositional tense logics*, in: M. de Rijke, editor, *Diamonds and Defaults: Studies in Pure and Applied Intensional Logic*, Springer, 1993 pp. 287–307.

[18] Szelepcsényi, R., *The method of forced enumeration for nondeterministic automata*, Acta Informatica **26** (1988), pp. 279–284.

[19] Venema, Y., *Expressiveness and completeness of an interval tense logic*, Notre Dame Journal of Formal Logic **31** (1990), pp. 529–547.