

Decidability of the interval temporal logic $A\bar{A}B\bar{B}$ over the rationals

Angelo Montanari¹, Gabriele Puppis², and Pietro Sala³

¹ University of Udine `angelo.montanari@uniud.it`
² LaBRI / CNRS `gabriele.puppis@labri.fr`
³ University of Verona `pietro.sala@univr.it`

Abstract. The classification of the fragments of Halpern and Shoham’s logic with respect to decidability/undecidability of the satisfiability problem is now very close to the end. We settle one of the few remaining questions concerning the fragment $A\bar{A}B\bar{B}$, which comprises Allen’s interval relations “meets” and “begins” and their symmetric versions. We already proved that $A\bar{A}B\bar{B}$ is decidable over the class of all finite linear orders and undecidable over ordered domains isomorphic to \mathbb{N} . In this paper, we first show that $A\bar{A}B\bar{B}$ is undecidable over \mathbb{R} and over the class of all Dedekind-complete linear orders. We then prove that the logic is decidable over \mathbb{Q} and over the class of all linear orders.

1 Introduction

Even though it has been authoritatively and repeatedly claimed that interval-based formalisms are the most appropriate ones for a variety of application domains, e.g., [6], until very recently interval temporal logics were a largely unexplored land. There are at least two explanations for such a situation: computational complexity and technical difficulty. On the one hand, the seminal work by Halpern and Shoham on the interval logic of Allen’s interval relations (HS for short) showed that such a logic is highly undecidable over all meaningful classes of linear orders [5], and ten years later Lodaya proved that a restricted fragment of it, denoted BE, featuring only two modalities (those for Allen’s relations *begins* and *ends*), suffices for undecidability [7]. On the other hand, formulas of interval temporal logics express properties of pairs of time points rather than of single time points, and are evaluated as sets of such pairs, that is, binary relations. As a consequence, there is no reduction of the satisfiability/validity in interval logics to monadic second-order logic, and thus Rabin’s theorem (the standard proof machinery) is not applicable here.

In the last decade, a systematic investigation of HS fragments has been carried out. Their classification with respect to the decidability/undecidability of their satisfiability problem is now very close to the end. The outcome of the analysis is that undecidability rules over HS fragments [1, 8], but some meaningful exceptions exist [2, 3, 4, 10, 11]. While setting the status of most and least expressive interval logics is relatively straightforward, e.g., undecidability of full HS can be shown by a reduction from the non-halting problem for Turing machines, decidability of the logic of Allen’s relations *begins* and *begun by* $B\bar{B}$ can be proved by a reduction to the (point-based) linear temporal logic of future

and past, dealing with those fragments that lie on the marginal land between decidability and undecidability is much more difficult. (Un)decidability of HS fragments depends on two factors: their set of interval modalities and the class of linear orders over which they are interpreted. While the first one is fairly obvious, the second one is definitively less immediate. Some HS fragments behave the same over all classes of linear orders. This is the case with the logic of temporal neighbourhood $\mathbf{A}\bar{\mathbf{A}}$, which is NEXPTIME-complete over all relevant classes of linear orders [3]. A real character is, on the contrary, the temporal logic of sub-intervals \mathbf{D} : its satisfiability problem is PSPACE-complete over the class of dense linear orders [2] and undecidable over the classes of finite and discrete linear orders [8] (it is still unknown over the class of all linear orders).

In this paper, we focus our attention on the satisfiability problem for the logic $\mathbf{A}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{B}}$, which pairs the decidable fragments $\mathbf{A}\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}\bar{\mathbf{B}}$. In [11], we proved that the problem is decidable, but not primitive recursive, over finite linear orders, and undecidable over the natural numbers. Here, we first show that undecidability can be lifted to the temporal domain \mathbb{R} , as well as to the class of all Dedekind-complete linear orders. Then, we consider the order \mathbb{Q} . We devise two semi-decision procedures: the first one terminates if and only if the input formula is unsatisfiable over \mathbb{Q} , while the second one terminates if and only if the input formula is satisfiable over \mathbb{Q} . Running the two procedures in parallel gives a decision algorithm for $\mathbf{A}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{B}}$ over \mathbb{Q} . We conclude the paper by showing that decidability over the class of all linear orders follows from that over \mathbb{Q} .

2 The logic

We begin by introducing the logic $\mathbf{A}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{B}}$. Let Σ be a set of proposition letters. The logic $\mathbf{A}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{B}}$ consists of formulas built up from letters in Σ using the Boolean connectives \neg and \vee and the unary modalities $\langle \mathbf{A} \rangle$, $\langle \bar{\mathbf{A}} \rangle$, $\langle \mathbf{B} \rangle$, and $\langle \bar{\mathbf{B}} \rangle$. We will often make use of shorthands like $\varphi_1 \wedge \varphi_2 = \neg(\neg\varphi_1 \vee \neg\varphi_2)$, $[\mathbf{A}]\varphi = \neg\langle \mathbf{A} \rangle\neg\varphi$, $[\bar{\mathbf{B}}]\varphi = \neg\langle \bar{\mathbf{B}} \rangle\neg\varphi$, $\mathbf{true} = a \vee \neg a$, and $\mathbf{false} = a \wedge \neg a$, for $a \in \Sigma$.

To define the semantics of $\mathbf{A}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{B}}$ formulas, we consider a linear order $\mathbb{D} = (D, <)$, called *temporal domain*, and we denote by $\mathbb{I}_{\mathbb{D}}$ the set of all closed intervals $[x, y]$ over \mathbb{D} , with $x \leq y$. We call *interval structure* any Kripke structure of the form $\mathcal{I} = (\mathbb{I}_{\mathbb{D}}, \sigma, A, \bar{A}, B, \bar{B})$, where $\sigma : \mathbb{I}_{\mathbb{D}} \rightarrow \mathcal{P}(\Sigma)$ is a function mapping intervals to sets of proposition letters and $A, \bar{A}, B,$ and \bar{B} are the Allen's relations "meet", "met by", "begun by", and "begins", which are defined as follows: $[x, y] A [x', y']$ iff $y = x'$, $[x, y] \bar{A} [x', y']$ iff $x = y'$, $[x, y] B [x', y']$ iff $x = x' \wedge y' < y$, and $[x, y] \bar{B} [x', y']$ iff $x = x' \wedge y < y'$. Formulas are interpreted over a given interval structure $\mathcal{I} = (\mathbb{I}_{\mathbb{D}}, \sigma, A, \bar{A}, B, \bar{B})$ and a given initial interval $I \in \mathbb{I}_{\mathbb{D}}$ in the natural way, as follows: $\mathcal{I}, I \models a$ iff $a \in \sigma(I)$, $\mathcal{I}, I \models \neg\varphi$ iff $\mathcal{I}, I \not\models \varphi$, $\mathcal{I}, I \models \varphi_1 \vee \varphi_2$ iff $\mathcal{I}, I \models \varphi_1$ or $\mathcal{I}, I \models \varphi_2$, and, most importantly, for all relations $R \in \{A, \bar{A}, B, \bar{B}\}$,

$$\mathcal{I}, I \models \langle R \rangle \varphi \quad \text{iff} \quad \text{there is } J \in \mathbb{I}_{\mathbb{D}} \text{ such that } I R J \text{ and } \mathcal{I}, J \models \varphi.$$

We say that a formula φ is *satisfiable* over a class \mathcal{C} of interval structures if $\mathcal{I}, I \models \varphi$ for some $\mathcal{I} = (\mathbb{I}_{\mathbb{D}}, \sigma, A, \bar{A}, B, \bar{B})$ in \mathcal{C} and some interval $I \in \mathbb{I}_{\mathbb{D}}$.

For example, the formula $[\mathbf{B}]\mathbf{false}$ (hereafter abbreviated π) hold over all and only the *singleton* intervals $[x, x]$. Similarly, the formula $[\mathbf{B}][\mathbf{B}]\mathbf{false}$ (abbreviated

unit) holds over the *unit-length* intervals of a discrete order, e.g. over the intervals of \mathbb{Z} of the form $[x, x + 1]$. The formula $[A][\bar{A}][A][A]\varphi$ ($[G]\varphi$ for short) forces φ to hold universally, that is, over all intervals. The formula $[G](\neg\pi \rightarrow \langle B \rangle \neg\pi)$ (φ_{dense} for short) holds over all and only the interval structures with a dense domain, e.g., the order \mathbb{Q} of the rationals.

Logical types. We now introduce basic terminology and notation that are common in the temporal logic setting. The *closure* of a formula φ is defined as the set $\text{closure}(\varphi)$ of all sub-formulas of φ and all their negations (we identify $\neg\neg\psi$ with ψ , $\neg\langle A \rangle\psi$ with $[A]\neg\psi$, etc.). For a technical reason that will be clear soon, we also introduce the *extended closure* of φ , denoted $\text{closure}^+(\varphi)$, that extends $\text{closure}(\varphi)$ by adding all formulas of the form $\langle R \rangle\psi$ and $[R]\psi$, with $R \in \{A, \bar{A}, B, \bar{B}\}$ and $\psi \in \text{closure}(\varphi)$.

Let $\mathcal{I} = (\mathbb{I}_{\mathbb{D}}, \sigma, A, \bar{A}, B, \bar{B})$ be an interval structure. We associate with each interval $I \in \mathbb{I}_{\mathbb{D}}$ its φ -*type* $\text{type}_{\mathcal{I}}^{\varphi}(I)$, defined as the set of all formulas $\psi \in \text{closure}^+(\varphi)$ such that $\mathcal{I}, I \models \psi$ (when no confusion arises, we omit the parameters \mathcal{I} and φ). A particular role will be played by those types F that contain the subformula $[B]\text{false}$, which are necessarily associated with singleton intervals. When no interval structure is given, we can still try to capture the concept of type by means of a maximal “locally consistent” subset of $\text{closure}^+(\varphi)$. Formally, we call φ -*atom* any set $F \subseteq \text{closure}^+(\varphi)$ such that (i) $\psi \in F$ iff $\neg\psi \notin F$, for all $\psi \in \text{closure}^+(\varphi)$, (ii) $\psi \in F$ iff $\psi_1 \in F$ or $\psi_2 \in F$, for all $\psi = \psi_1 \vee \psi_2 \in \text{closure}^+(\varphi)$, (iii) if $[B]\text{false} \in F$ and $\psi \in F$, then $\langle A \rangle\psi \in F$ and $\langle \bar{A} \rangle\psi \in F$, for all $\psi \in \text{closure}(\varphi)$, (iv) if $[B]\text{false} \in F$ and $\langle A \rangle\psi \in F$, then $\psi \in F$ or $\langle \bar{B} \rangle\psi \in F$, for all $\psi \in \text{closure}(\varphi)$. We call π -*atoms* those atoms that contain the formula $[B]\text{false}$, which are thus candidate types of singleton intervals. We denote by $\text{atoms}(\varphi)$ the set of all φ -atoms.

Given an atom F and a relation $R \in \{A, \bar{A}, B, \bar{B}\}$, we let $\text{req}_R(F)$ be the set of *requests of F along direction R* , namely, the formulas $\psi \in \text{closure}(\varphi)$ such that $\langle R \rangle\psi \in F$. Similarly, we let $\text{obs}(F)$ be the set of *observables of F* , namely, the formulas $\psi \in F \cap \text{closure}(\varphi)$ – intuitively, the observables of F are those formulas $\psi \in F$ that fulfil requests of the form $\langle R \rangle\psi$ from other atoms. Note that, for all π -atoms F , we have $\text{req}_A(F) = \text{obs}(F) \cup \text{req}_{\bar{B}}(F)$ and $\text{req}_{\bar{A}}(F) \supseteq \text{obs}(F)$.

Compass structures. Formulas of interval temporal logics can be equivalently interpreted over the so-called compass structures [14]. These structures can be seen as two-dimensional spaces in which points are labelled with complete logical types (atoms). Such an alternative interpretation exploits the existence of a natural bijection between the intervals $I = [x, y]$ over a temporal domain \mathbb{D} and the points $p = (x, y)$ in the $\mathbb{D} \times \mathbb{D}$ grid such that $x \leq y$. It is convenient to introduce a *dummy atom* \emptyset , distinct from all other atoms, and assume that it labels all and only the points (x, y) such that $x > y$, which do not correspond to intervals. We fix the convention that $\text{obs}(\emptyset) = \emptyset$ and $\text{req}_R(\emptyset) = \emptyset$ for all $R \in \{A, \bar{A}, B, \bar{B}\}$.

Formally, a *compass φ -structure over a linear order \mathbb{D}* is a labelled grid $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$, where the function $\tau : \mathbb{D} \times \mathbb{D} \rightarrow \text{atoms}(\varphi) \uplus \{\emptyset\}$ maps any point (x, y) to either a φ -atom (if $x \leq y$) or the dummy atom \emptyset (if $x > y$).

We observe that Allen’s relations over intervals have analogue relations over points. Figure 1 gives a geometric interpretation of relations A, \bar{A}, B, \bar{B} (by a

slight abuse of notation, we use the same letters for the corresponding relations over the points of a compass structure). Thanks to such an interpretation, any interval structure \mathcal{I} can be converted to a compass one $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$ by simply letting $\tau(x, y) = \text{type}([x, y])$ for all $x \leq y \in \mathbb{D}$. The converse, however, is not true in general, as the atoms associated with points in a compass structure may be inconsistent with respect to the underlying geometrical interpretation of Allen's relations. To ease a correspondence between interval and compass structures, we enforce suitable *consistency conditions* on compass structures. For this, we introduce two relations over atoms F, G :

$$\begin{array}{l}
 F \uparrow G \text{ iff} \\
 \left\{ \begin{array}{l} \text{req}_{\bar{B}}(F) \supseteq \text{obs}(G) \cup \text{req}_{\bar{B}}(G) \\ \text{req}_B(G) \supseteq \text{obs}(F) \cup \text{req}_B(F) \\ \text{req}_{\bar{A}}(F) = \text{req}_{\bar{A}}(G) \end{array} \right. \\
 \\
 F \triangleleft G \text{ iff} \\
 \left\{ \begin{array}{l} \text{req}_A(F) = \text{obs}(G) \cup \text{req}_B(G) \cup \text{req}_{\bar{B}}(G) \\ \text{req}_{\bar{A}}(G) \supseteq \text{obs}(F). \end{array} \right.
 \end{array}$$

Note that the relation \uparrow is transitive, while \triangleleft only satisfies $\triangleleft \circ \uparrow \subseteq \triangleleft$. Observe also that, for all interval structures \mathcal{I} and all intervals I, J in it, if $I \bar{B} J$ (resp., $I A J$), then $\text{type}_{\mathcal{I}}(I) \uparrow \text{type}_{\mathcal{I}}(J)$ (resp., $\text{type}_{\mathcal{I}}(I) \triangleleft \text{type}_{\mathcal{I}}(J)$). Hereafter, we tacitly assume that every compass structure $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$ satisfies analogous consistency properties with respect to its atoms, namely, for all points $p = (x, y)$ and $q = (x', y')$ in $\mathbb{D} \times \mathbb{D}$, with $x \leq y$ and $x' \leq y'$, if $p \bar{B} q$ (resp., $p A q$), then $\tau(p) \uparrow \tau(q)$ (resp., $\tau(p) \triangleleft \tau(q)$). In addition, we say that a request $\psi \in \text{req}_R(\tau(p))$ of a point p in a compass structure $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$ is *fulfilled* if there is another point q such that $p R q$ and $\psi \in \text{obs}(\tau(q))$ – in this case, we say that q is a *witness of fulfilment of ψ from p* . The compass structure \mathcal{G} is said to be *globally fulfilling* if all requests of all its points are fulfilled.

We can now recall the standard correspondence between interval and compass structures (the proof is based on a simple induction on sub-formulas):

Proposition 1 ([11]). Let φ be an $A\bar{A}B\bar{B}$ formula. For every globally fulfilling compass structure $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$, there is an interval structure $\mathcal{I} = (\mathbb{I}_{\mathbb{D}}, \sigma, A, \bar{A}, B, \bar{B})$ such that, for all $x \leq y \in \mathbb{D}$ and all $\psi \in \text{closure}^+(\varphi)$, $\mathcal{I}, [x, y] \models \psi$ iff $\psi \in \tau(x, y)$.

In view of Proposition 1, the satisfiability problem for a given $A\bar{A}B\bar{B}$ formula φ reduces to the problem of deciding the existence of a globally fulfilling compass $\tilde{\varphi}$ -structure $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$, with $\tilde{\varphi} = \langle \mathbf{G} \rangle \varphi$ ($\langle \mathbf{G} \rangle \varphi$ is a shorthand for $\neg[\mathbf{G}]\neg\varphi$, that features the observable $\tilde{\varphi}$ in every point, that is, $\tilde{\varphi} \in \text{obs}(\tau(x, y))$ for all $x \leq y \in \mathbb{D}$).

3 Satisfiability over finite and Dedekind-complete orders

The satisfiability problem for $A\bar{A}B\bar{B}$ was originally addressed in [11]. We first proved that $A\bar{A}B\bar{B}$ is decidable if interpreted over finite linear orders, but not primitive recursive. The decidability result rests on a contraction method that, given a formula φ and a finite compass structure satisfying φ , shows that, under suitable conditions, the compass structure can be reduced in size while pre-




Fig. 1. Geometric interpretation of Allen's relations.

servicing consistency and fulfilment properties. This leads to a non-deterministic procedure that decides whether φ is satisfiable by exhaustively searching all contraction-free compass structures. The proof of termination relies on Dickson’s lemma, while non-primitive recursiveness is proved via a reduction from the reachability problem for lossy counter machines [13]. Then, we showed that the problem becomes undecidable if we interpret $A\bar{A}B\bar{B}$ over a temporal domain isomorphic to \mathbb{N} (in fact, this is already the case with the proper fragment $A\bar{A}B$). The proof is based on a reduction from an undecidable variant of the reachability problem for lossy counter machines, called structural termination [9], which consists of deciding whether a given lossy counter machine admits a halting computation starting from a given location and some arbitrary initial assignment for the counters. Due to an oversight, in [11] we claimed that such an undecidability result can be transferred to any class of linear orders in which \mathbb{N} can be embedded. As a matter of fact, Dedekind completeness is a necessary condition. The following theorem properly states undecidability results for $A\bar{A}B$.

Theorem 1. The satisfiability problem for $A\bar{A}B$ interpreted over \mathbb{N} , \mathbb{R} , and the class of all Dedekind-complete linear orders is undecidable.

In view of the above theorem and the decidability results in [11], the satisfiability problem for $A\bar{A}B\bar{B}$ over \mathbb{Q} , as well as over the class of all interval structures, remains open. In the next section, we will show that, quite surprisingly, both problems are decidable with non-primitive recursive complexity.

4 Satisfiability over the rationals and all linear orders

We begin by describing a fairly simple semi-decision procedure for the *unsatisfiability* of $A\bar{A}B\bar{B}$ formulas over interval structures with a dense temporal domain. The crucial observation is that, whenever a formula φ is unsatisfiable over \mathbb{Q} , this can be witnessed by a *finite* set of intervals with inconsistent requests. Based on this observation, one can enumerate all finite compass structures that witness φ and are distinct up to isomorphism, following the partial order induced by the embedding relation (this relation is defined as an isomorphism between the smaller structure and the restriction of the larger structure to a suitable subset of its temporal domain). The only way the enumeration procedure can terminate is when no refinement is applicable: in this case, one proves that the input formula φ is not satisfiable. Conversely, if the enumeration procedure does not terminate, then the formula φ is satisfied by some compass structure that is obtained from the limit of an infinite series of refinements (suitable fairness conditions for the generated refinements guarantee that the temporal domain of the limit compass structure is isomorphic to \mathbb{Q}). We refer the reader to [12] for the details of the above enumeration procedure of unsatisfiable formulas over \mathbb{Q} .

The rest of the section is devoted to finding a semi-decision procedure that receives an input formula φ and terminates (successfully) iff φ is satisfiable over an interval structure with a dense temporal domain. Differently from the previous procedure, this one is based on enumerating suitable *finite abstractions* of compass structures, which is far from being an easy task.

A first step consists of simplifying the consistency and fulfilment conditions. More precisely, we show how to turn them into more “local” constraints, so as to ease, later, the abstraction task. To this end, recall that the rational line is isomorphic to any countable dense ordering with neither a minimal element nor a maximal one. This means that, for the purpose of studying satisfiability over \mathbb{Q} , it does not matter if we consider interval structures over \mathbb{Q} or over any subset of it that is dense and contains no extremal elements. Similarly, the complexity of the satisfiability problem does not change if we add minimal and maximal elements to the underlying temporal domain – for the sake of brevity, we call the resulting order a *dense order with endpoints*. Now, to turn the consistency and fulfilment conditions into local constraints, we decompose any dense order with endpoints \mathbb{D} into some infinite, finitely-branching tree T whose nodes represent pairs of elements of \mathbb{D} of the form $s = (y_1, y_2)$, with $y_1 < y_2$, and whose edges connect nodes $(y_1, y_2) \in T$ to tuples of nodes $s_1 = (z_1, z_2), \dots, s_n = (z_n, z_{n+1})$, with $n \geq 2$ and $y_1 = z_1 < z_2 < \dots < z_n < z_{n+1} = y_2$ (see Figure 2). Note that the domain \mathbb{D} is not necessarily entirely covered by the time points that appear in the nodes of a decomposition T . Moreover, since all dense orders with endpoints are isomorphic, we will not be concerned with the coordinates of the nodes of T and we will often overlook them in the constructions that follow.




Fig. 2. Decomposition of a compass structure.

Using decompositions of temporal domains we can extract “horizontal slices” of a compass structure. More precisely, given a compass structure $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$ and a node $s = (y_1, y_2)$ of a decomposition T of \mathbb{D} , we define the *slice of \mathcal{G} in s* as the induced sub-structure $\mathcal{G}_s = (\mathbb{D} \times \{y_1, y_2\}, \tau)$. Intuitively, the slice \mathcal{G}_s is obtained from \mathcal{G} by selecting the rows with coordinates y_1 and y_2 and by restricting the labelling function τ to them (to reduce the notational overload, we denote such a restriction of the labelling function by τ).

Below, we introduce suitable abstractions, called profiles, for the labels that can appear in a slice of a compass structure. Intuitively, for each slice $\mathcal{G}_s = (\mathbb{D} \times \{y_1, y_2\}, \tau)$ and each pair of atoms (F, G) , where possibly $F = \emptyset$ or both $F = \emptyset$ and $G = \emptyset$ (dummy atoms), we keep track of the *number* of coordinates $x \in \mathbb{D}$ such that $\tau(x, y_1) = F$ and $\tau(x, y_2) = G$. In particular, in these abstractions, we forget the occurrence order of the pairs of atoms along the x -axis. To this end, we make extensive use of *multisets*. Given a multiset M and an element e in M , we denote by $M(e)$ the *multiplicity*, that is, the number of occurrences, of e in M , and we write $M(e) = \infty$ when M contains infinitely many occurrences of e . We freely use set-theoretic notations with multisets. For example, we denote membership by $e \in M$, containment by $M \subseteq N$, etc. Moreover, given a multiset M of n -tuples and a set $I \subseteq \{1, \dots, n\}$ of component indices, we denote by $M|_I$ the *projection* of M onto I , that is, the multiset that contains exactly $\sum_{e|_I=f} M(e)$ occurrences of each I -tuple f (note that the sum ranges over all n -tuples e that coincide with f on the components indexed in I). Note that, differently

from set projections, projections of multisets are injective, as they send distinct occurrences of tuples to distinct occurrences of tuples. In particular, $|_I$ defines a *bijection* from multiset M to multiset $M|_I$. Finally, we denote by $\text{set}(M)$ the *support* of a multiset M , that is, the set of all elements e such that $M(e) \geq 1$.

We associate with each slice $\mathcal{G}_s = (\mathbb{D} \times \{y_1, y_2\}, \tau)$ of a globally fulfilling compass structure \mathcal{G} , the multiset M defined by $M(F, G) = |\{x \in \mathbb{D} : \tau(x, y_1) = F, \tau(x, y_2) = G\}|$ for all $(F, G) \in (\text{atoms}(\varphi) \uplus \{\emptyset\})^2$. We call this multiset *the profile of the slice* \mathcal{G}_s and we denote it by $\text{profile}(\mathcal{G}_s)$. Note that the projection $\text{profile}(\mathcal{G}_s)|_1$ (resp., $\text{profile}(\mathcal{G}_s)|_2$) onto the first (resp., second) component is a multiset that represents the number of occurrences of each atom along the lower (resp., upper) row of the slice \mathcal{G}_s . Definition 1 below captures a more general notion of profile that does not refer to a particular compass structure. We will then introduce trees labelled with profiles as abstractions of compass structures.

Definition 1. A *profile* is a multiset M of pairs of (possibly dummy) atoms $(F, G) \in (\text{atoms}(\varphi) \uplus \{\emptyset\})^2$ such that: (i) for all $(F, G) \in M$, if $F \neq \emptyset$, then $G \neq \emptyset$ and $F \uparrow G$; (ii) for all $(F, G), (F', G') \in M$, $\text{req}_A(F) = \text{req}_A(F')$ and $\text{req}_A(G) = \text{req}_A(G')$; (iii) M contains infinitely many occurrences of pairs (\emptyset, G) with $G \neq \emptyset$; (iv) M contains *exactly one* occurrence of a pair (F, G) with F π -atom and *exactly one* occurrence of a pair (\emptyset, H) with H π -atom (for short, we denote the two pairs (F, G) and (\emptyset, H) by M_π and M^π , respectively); (v) if $M_\pi = (F, G)$, then $\text{req}_A(F) = \bigcup_{(F', G') \in M} \text{obs}(F')$; similarly, if $M^\pi = (\emptyset, H)$, then $\text{req}_A(H) = \bigcup_{(F', G') \in M} \text{obs}(G')$.

Definition 2. A *profile tree* is an infinite finitely-branching tree $\mathcal{T} = (T, N, E)$, where T is a decomposition of some dense order with endpoints, N is a function mapping nodes of T to profiles, and E is a function mapping nodes of T to multisets of tuples of atoms, such that:

- **(profile-match)** every node $s \in T$ has at least two children, say, s_1, \dots, s_n , with $n \geq 2$, and $E(s)$ is a multiset of $(n+1)$ -tuples such that $E(s)|_{1, n+1} = N(s)$ and $E(s)|_{i, i+1} = N(s_i)$ for all $1 \leq i \leq n$;
- **(profile-finite-req)** for every node $s \in T$ and pair $(F_1, F_{n+1}) \in N(s)$, with $F_1 \neq \emptyset$, if $N(s)(F_1, F_{n+1}) < \infty$, then $E(s)$ contains *exactly* $N(s)(F_1, F_{n+1})$ occurrences of tuples (F_1, \dots, F_{n+1}) such that $\text{req}_{\bar{B}}(F_1) = \bigcup_{2 \leq i \leq n+1} \text{obs}(F_i) \cup \text{req}_{\bar{B}}(F_{n+1})$ and $\text{req}_B(F_{n+1}) = \bigcup_{1 \leq i \leq n} \text{obs}(F_i) \cup \text{req}_B(F_1)$;
- **(profile-infinite-req)** for every node $s \in T$ and pair $(F_1, F_{n+1}) \in N(s)$, with $F_1 \neq \emptyset$, if $N(s)(F_1, F_{n+1}) = \infty$, then $E(s)$ contains *at least one* occurrence of a tuple (F_1, \dots, F_{n+1}) such that $\text{req}_{\bar{B}}(F_1) = \bigcup_{2 \leq i \leq n+1} \text{obs}(F_i) \cup \text{req}_{\bar{B}}(F_{n+1})$ and $\text{req}_B(F_{n+1}) = \bigcup_{1 \leq i \leq n} \text{obs}(F_i) \cup \text{req}_B(F_1)$;
- **(profile-dummy)** for every node $s \in T$ and pair $(\emptyset, G) \in N(s)$, with $G \neq \emptyset$, $E(s)$ contains at least one occurrence of a π -tuple, i.e., a tuple with a π -atom, of the form (F_1, \dots, F_{n+1}) , with $F_1 = \emptyset$ and $F_{n+1} = G$.

In addition, if the profile at the root s_0 of \mathcal{T} contains only the pair $N^\pi(s_0) = (F, G)$, with F π -atom and $\text{req}_{\bar{B}}(G) = \emptyset$, and some pairs of the form (\emptyset, H) , with $H \neq \emptyset$ and $\text{req}_{\bar{B}}(H) = \emptyset$, then \mathcal{T} is said to be a *full* profile tree.

The first item of Definition 2 enforces the matching conditions between the pairs in the profile of a node and the pairs in the profiles of its children. The

second item requires that all requests that appear in a pair (F, G) of the profile of a node s are either “locally fulfilled” by the observables of corresponding pairs in the profiles of the children or transferred to other nodes of the profile tree at the same level as s . This condition, however, concerns only those pairs (F, G) that have finite multiplicity in the profile; for the remaining pairs, we enforce a similar, but weaker condition (third item of the definition). Finally, the fourth item requires that for each atom G , if the profile $N(s)$ contains the pair (\emptyset, G) , then at least one occurrence of this pair is “refined” in the multiset $E(s)$ by an occurrence of a tuple of the form $(\emptyset, \dots, \emptyset, F, \dots, G)$ that contains a π -atom F (such a tuple is called for short π -*tuple*) and that ends with the atom G (possibly $F = G$). We will see later that this condition is necessary for the fulfilment of the requests along the direction \bar{A} .

Below, we show that full profile trees are correct (though not yet finite) abstractions of globally fulfilling compass structures. We present this result with two statements showing, respectively, completeness and a weak form of soundness of profile trees. Note that the two-way correspondence is sufficient for witnessing satisfiability of $\mathbf{A}\bar{\mathbf{A}}\mathbf{B}\bar{\mathbf{B}}$ formulas by means of profile trees.

Proposition 2. For every globally fulfilling compass structure $\mathcal{G} = (\mathbb{D} \times \mathbb{D}, \tau)$ over a dense order with endpoints \mathbb{D} , there is a full profile tree $\mathcal{T} = (T, N, E)$ such that T is a decomposition of \mathbb{D} and, for all nodes $s \in T$, $N(s) = \mathbf{profile}(\mathcal{G}_s)$. Conversely, for every full profile tree $\mathcal{T} = (T, N, E)$, with T decomposition of some dense order with endpoints \mathbb{D} , there is a globally fulfilling compass structure $\mathcal{G} = (\mathbb{D}' \times \mathbb{D}', \tau)$, with $\mathbb{D}' \subseteq \mathbb{D}$ dense order with endpoints, such that $\mathbf{set}(\mathbf{profile}(\mathcal{G}_s)) = \mathbf{set}(N(s))$ for all $s \in T$.

Below, we show how to further restrict ourselves to a complete subset of full profile trees and derive *finite* representations of them. The general idea is to normalise profile trees so as to obtain structures that are sufficiently “regular” to be represented by finite trees. To this end, we introduce a finite variant of the notion of profile tree, called *finite profile tree*, that is obtained by enforcing the conditions of Definition 2 to internal nodes only (accordingly, since the multisets $E(s)$ that are associated with the leaves s in a finite profile tree are not anymore relevant, one can assume that the function E is undefined on the leaves). We also introduce a strengthening of the containment relation on multisets, denoted by \sqsubseteq and defined as follows: $M \sqsubseteq N$ iff $\mathbf{set}(M) = \mathbf{set}(N)$ and $M(\bar{F}) \leq N(\bar{F})$ (resp., $M(\bar{F}) = N(\bar{F})$) for all tuples \bar{F} (resp., π -tuples \bar{F}). The following definition captures precisely the set of profile trees we are interested in.

Definition 3. Let \mathcal{T} be a finite or infinite profile tree. We say that \mathcal{T} is *pseudo-regular* iff for all paths π , there are $s, s' \in \pi$, with s proper ancestor of s' , such that $N(s) \sqsubseteq N(s')$ and $N(s)(\emptyset, G) = N(s')(\emptyset, G)$ for all atoms G .

In the following, we mainly work with profiles that appear at the roots of infinite profile trees (*feasible* profiles for short). We observe that the restriction of the partial order \sqsubseteq to feasible profiles is a *well* partial order: indeed, the definition of profile tree implies that every π -tuple has multiplicity either 0 or 1 in any feasible profile, which in turn means that \sqsubseteq is the conjunction of the well partial order \sqsubseteq and an equivalence of finite index. Hence, by a combination

of Dickson’s and König’s lemmas, every infinite pseudo-regular tree has a finite prefix that is also pseudo-regular (a prefix of a tree is any restriction of it to an upward-closed set of nodes). A converse result also holds:

Proposition 3. For every finite pseudo-regular profile tree \mathcal{T} , there is an infinite profile tree \mathcal{T}' that has the same profile as \mathcal{T} at the root.

The crux of our semi-decision procedure for testing the satisfiability of $A\bar{A}\bar{B}\bar{B}$ formulas is to enumerate all atoms that appear in feasible profiles. Proposition 3 allows us to use finite pseudo-regular profile trees as witnesses of existence of some of these atoms. Unfortunately, this is not yet the end of the story, because not all profile trees are pseudo-regular and hence, a priori, there might exist atoms that appear only in infinite profile trees that are not pseudo-regular. The last piece of the puzzle amounts at showing that this is not the case and that we can indeed safely restrict ourselves to atoms appearing in pseudo-regular profile trees. We will prove this result by normalizing infinite profile trees via a series of operations that “inflate” the profiles as much as possible.

An important aspect that must be taken into account while inflating the profiles in a tree is that there are matching constraints to satisfy. As a matter of fact, these constraints induce dependencies between the multiplicities of pairs (\emptyset, F) in the profile associated with a node s and the multiplicities of corresponding pairs (F, G) in the profile associated with the right sibling of s . As a consequence, there will be differences in the treatment of pairs of the form (\emptyset, F) and pairs of the form (F, G) , with $F \neq \emptyset$. We take a brief interlude to give an example of the type of dependencies that can be enforced.

Example 1. Consider a formula φ that contains, among other conjuncts, the subformula $[G](a \rightarrow [B]_{-a} \wedge [\bar{B}]_{-a})$. Figure 3 describes a slice of a compass structure that may satisfy φ , with some distinguished points annotated with observables and requests. The formula requires that all a -labelled points lie on distinct vertical axes; on the other hand, it allows arbitrarily many a -labelled points to be horizontally aligned. This is a representative example because, in general, forbidding multiple occurrences of an observable along the same horizontal line can be only done using the modal operator $[E]$, which is not available in the logic. As concerns the multiplicities of the example profile, we observe that by inserting multiple a -labelled points along a single horizontal line and by accordingly modifying the upper part of the compass structure, one can get as many pairs of atoms (F, G) , where $\langle \bar{B} \rangle a \in F$ and $[\bar{B}]_{-a} \in G$. On the other hand, to increase the number of pairs (\emptyset, F) , where $\langle \bar{B} \rangle a \in F$, one has to introduce new horizontal lines ending with π -atoms H such that $\langle \bar{A} \rangle a \in H$: this is not always possible as other conjuncts of φ may enforce bounds to the number of π -atoms H .




Fig. 3. Dependencies between multiplicities.

As shown by the above example, the simplest way one can inflate a profile, while preserving its feasibility, is by increasing the multiplicities associated with the pairs (F, G) , where $F \neq \emptyset$. We formalise this in the next lemma.

Lemma 1. If N is a feasible profile and N' is a profile such that $N \sqsubseteq N'$ and $N(\emptyset, G) = N'(\emptyset, G)$ for all atoms G , then N' is feasible too. Moreover, a profile tree with root profile N' can be obtained from a profile tree with root profile N without modifying the underlying decomposition tree.

We describe a second inflation method, which depends on the previous one and can be used to further simplify the reasoning on the matching conditions of a profile tree $\mathcal{T} = (T, N, E)$. In particular, it shows that w.l.o.g. one can assume that the *finiteness* of the multiplicity of any tuple (F_1, \dots, F_{n+1}) in a multiset $E(s)$ depends only on the multiplicity of the first component F_1 in $E(s)|_1$. This property is formalized below by the definition of “pointwise fair” profile tree, followed by a corresponding lemma that shows how to enforce the property.

Definition 4. A multiset E of $(n+1)$ -tuples is *fair* if for all $(n+1)$ -tuples $(F_1, \dots, F_{n+1}) \in E$, with $F_1 \neq \emptyset$, $E|_1(F_1) = \infty$ implies $E(F_1, \dots, F_{n+1}) = \infty$. A profile tree $\mathcal{T} = (T, N, E)$ is *pointwise fair* if all multisets $E(s)$ are fair.

Lemma 2. For every feasible profile N , there is an infinite pointwise fair profile tree that has root profile $N' \supseteq N$. Moreover, one can assume that, for all pairs of atoms (F, G) , if $N|_1(F) < \infty$, then $N(F, G) = N'(F, G)$.

A third inflation method makes use of the fact that the partial order \sqsubseteq restricted to the set of feasible profiles is ω -complete.

Lemma 3. Every sequence of feasible profiles $N_0 \sqsubseteq N_1 \sqsubseteq \dots$ has a supremum $\sup_i N_i$, defined by $(\sup_i N_i)(F, G) = \sup_{i \in \mathbb{N}} (N_i(F, G))$ for all atoms F, G , that is a feasible profile.

We have described three ways of increasing the multiplicities of profiles at the roots of profile trees. In general, these techniques are not applicable to nodes that are strictly below the root. This is why we introduce a new partial order \preceq , incomparable with \sqsubseteq , that is defined only over feasible profiles N, N' as follows:

$$N \preceq N' \quad \text{iff} \quad \begin{cases} N \sqsubseteq N' \\ \text{set}(N|_2) = \text{set}(N'|_2) \\ N(F, G) = N'(F, G) \quad \text{for all atoms } F, G \neq \emptyset. \end{cases}$$

We observe that from any infinite \preceq -chain of feasible profiles, one can extract an infinite sub-sequence that is also a \sqsubseteq -chain. Thus, an immediate consequence of Lemma 3 is that every \preceq -chain has an *upper bound*. In its turn, the existence of upper bounds on \preceq -chains implies the existence of feasible profiles that are *maximal* with respect to \preceq (this can be seen as a consequence of Zorn’s Lemma):

Corollary 1. For all feasible profiles N , there is a \preceq -maximal profile $N' \supseteq N$.

Based on existence of \preceq -maximal profiles, we say that a profile tree is *pointwise \preceq -maximal* if all its profiles are \preceq -maximal. Below, we show that all atoms of feasible profiles appear at the roots of some pseudo-regular profile trees.

Proposition 4. For every infinite pointwise fair profile tree with root profile N , there is an infinite pointwise fair and pointwise \preceq -maximal profile tree with root profile $N' \supseteq N$.

Proposition 5. Every infinite pointwise fair and pointwise \preceq -maximal profile tree is pseudo-regular.

Wrapping up, we can devise a semi-decision procedure that tests the satisfiability of a formula φ over \mathbb{Q} . The procedure works as follows. It first transforms φ into an equi-satisfiable formula $\varphi_{\uparrow\downarrow}$ interpreted over a dense order with endpoints \mathbb{D} . Then, the procedure enumerates all finite full pseudo-regular trees, until a tree is found that contains the formula $\langle G \rangle \varphi_{\uparrow\downarrow}$ as an observable of one of its atoms. The above semi-decision procedure is correct, namely, it terminates successfully iff the input formula φ is satisfiable over \mathbb{Q} . Indeed, if φ is satisfiable over \mathbb{Q} , then $\varphi_{\uparrow\downarrow}$ is satisfiable over a dense order with endpoints \mathbb{D} , and hence there is a globally fulfilling compass structure \mathcal{G} that contains $\langle G \rangle \varphi_{\uparrow\downarrow}$ as an observable of all its atoms. By Propositions 4 and 5, there is also an infinite, pseudo-regular full profile tree \mathcal{T} that witnesses $\langle G \rangle \varphi_{\uparrow\downarrow}$ at the root profile. By the remarks that follow Definition 3, there is also a prefix of \mathcal{T} that is a finite pseudo-regular full profile tree, and eventually this tree must be discovered by the procedure. Conversely, if the procedure terminates with a finite pseudo-regular full profile tree witnessing $\langle G \rangle \varphi_{\uparrow\downarrow}$, then by Proposition 3 there is an infinite full profile tree \mathcal{T} , and hence a compass structure \mathcal{G} , that witness the satisfiability of $\langle G \rangle \varphi_{\uparrow\downarrow}$ over \mathbb{D} . One can then conclude that φ is satisfiable over \mathbb{Q} .

A full decision procedure that solves the satisfiability problem for $A\bar{A}\bar{B}\bar{B}$ over \mathbb{Q} can simply run in parallel the two semi-decision procedures that we described for unsatisfiability and satisfiability of $A\bar{A}\bar{B}\bar{B}$ formulas.

As for the satisfiability problem over the class of all interval structures, one can simply observe the following. The logic $A\bar{A}\bar{B}\bar{B}$, as any other HS fragment, can be viewed as a fragment of first-order logic that uses binary relations to express properties of pairs of elements of the underlying temporal domain. The relation $<$ of the temporal domain can be easily constrained by a first-order formula so as to define a linear order, and Allen's relations can be expressed in first-order logic in term of $<$. From Löwenheim-Skolem theorem, it follows that every interval structure can be assumed to contain only *countably* many intervals. Moreover, since every countable linear order can be embedded inside \mathbb{Q} , satisfiability of formulas of a given HS fragment over the class of all linear orders can be reduced to their satisfiability over \mathbb{Q} , provided that the fragment is powerful enough to express such an embedding. This is the case with $A\bar{A}\bar{B}\bar{B}$: it suffices to introduce a distinguished proposition letter $\#$, to constrain all $\#$ -labelled intervals to be singletons ($[G](\# \rightarrow \pi)$), and to relativize all modalities to intervals with endpoints labelled by $\#$ (intervals that satisfy $\langle B \rangle \# \wedge \langle A \rangle \#$). We conclude by establishing the precise complexity of the satisfiability problem.

Theorem 2. The satisfiability problem for $A\bar{A}\bar{B}\bar{B}$ interpreted over \mathbb{Q} , as well as over the class of all linear orders, is decidable, but not primitive recursive.

5 Conclusions

In this paper we close the open questions concerning the satisfiability problem for the interval temporal logic $A\bar{A}\bar{B}\bar{B}$. First, we generalized the undecidability result from [11] to \mathbb{R} and to the class of all Dedekind-complete linear orders, and then we proved that it is decidable in two interesting cases: \mathbb{Q} and the class of all interval structures. To decide satisfiability of $A\bar{A}\bar{B}\bar{B}$ formulas over \mathbb{Q} we used

a combination of techniques from [4] (tree-shaped decomposition of models) and [11] (encoding of models by systems with counters), plus new key ingredients (separation into two semi-decision procedures, König’s lemma). As concerns the second result, the decidability of $A\bar{A}B\bar{B}$ over the class of all interval structures follows from the decidability over \mathbb{Q} and from Löwenheim-Skolem theorem, which allows us to assume, without loss of generality, that the interval structures are countable and hence embeddable inside \mathbb{Q} . The fact that $A\bar{A}B\bar{B}$ is powerful enough to express the embedding of a countable order inside \mathbb{Q} completes the reduction. It is worth pointing out that the same technique cannot be applied to all HS fragments; for instance, the satisfiability problem for the temporal logic of sub-/super-intervals $D\bar{D}$ is known to be decidable over \mathbb{Q} [2, 10], but it is open for the class of all interval structures.

References

- [1] D. Bresolin, D. Della Monica, V. Goranko, A. Montanari, and G. Sciavicco. The dark side of Interval Temporal Logic: sharpening the undecidability border. In *Proc. of the 18th TIME*, pages 131–138. IEEE, 2011.
- [2] D. Bresolin, V. Goranko, A. Montanari, and P. Sala. Tableaux for logics of subinterval structures over dense orderings. *Journal of Logic and Computation*, 20(1):133–166, 2010.
- [3] D. Bresolin, V. Goranko, A. Montanari, and G. Sciavicco. Propositional interval neighborhood logics: Expressiveness, decidability, and undecidable extensions. *Annals of Pure and Applied Logic*, 161(3):289–304, 2009.
- [4] D. Bresolin, A. Montanari, P. Sala, and G. Sciavicco. What’s decidable about Halpern and Shoham’s interval logic? The maximal fragment $AB\bar{B}L$. In *Proc. of the 26th LICS*, pages 387–396. IEEE, 2011.
- [5] J. Halpern and Y. Shoham. A propositional modal logic of time intervals. *Journal of the ACM*, 38(4):935–962, 1991.
- [6] H. Kamp and U. Reyle. *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory, Volume 42 of Studies in Linguistics and Philosophy*. Springer, 1993.
- [7] K. Lodaya. Sharpening the undecidability of interval temporal logic. In *Proc. of 6th ASIAN*, volume 1961 of *LNCS*, pages 290–298. Springer, 2000.
- [8] J. Marcinkowski and J. Michaliszyn. The ultimate undecidability result for the Halpern-Shoham logic. In *Proc. of the 26th LICS*, pages 377–386. IEEE, 2011.
- [9] R. Mayr. Undecidable problems in unreliable computations. *Theoretical Computer Science*, 297(1–3):337–354, 2003.
- [10] A. Montanari, G. Puppis, and P. Sala. A decidable spatial logic with cone-shaped cardinal directions. In *Proc. of the 18th CSL*, volume 5771 of *LNCS*, pages 394–408. Springer, 2009.
- [11] A. Montanari, G. Puppis, and P. Sala. Maximal decidable fragments of Halpern and Shoham’s modal logic of intervals. In *Proc. of the 37th ICALP - Part II*, volume 6199 of *LNCS*, pages 345–356, 2010.
- [12] A. Montanari, G. Puppis, and P. Sala. Decidability of the interval temporal logic $A\bar{A}B\bar{B}$ over the rationals. Technical Report RR01/2014, Univ. of Udine, 2014.
- [13] P. Schnoebelen. Lossy counter machines decidability cheat sheet. In *Proc. of the 4th RP*, volume 6227 of *LNCS*, pages 51–75. Springer, 2010.
- [14] Y. Venema. Two-dimensional modal logics for relation algebras and temporal logic of intervals. ITLI prepublication series LP-89-03, Univ. of Amsterdam, 1989.