ECAI 2014
T. Schaub et al. (Eds.)

doi:10.3233/978-1-61499-419-0-543

Decidability of model checking multi-agent systems against a class of EHS specifications

Alessio R. Lomuscio, Jakub Michaliszyn [1]

Abstract. We define and illustrate the expressiveness of the $A\bar{B}L$ fragment of the Epistemic Halpern–Shoham Logic as a specification language for multi-agent systems. We consider the model checking problem for systems against specifications given in the logic. We show its decidability by means of a novel technique that may be reused in other contexts for showing decidability of other logics based on intervals.

1 Introduction

Multi-agent systems are typically specified by means of formal languages expressing various aspects of their behaviour. One key formalism often used is *epistemic logic*, or logic for knowledge. This is a well-understood modal logic aimed at representing what agents know in the system and how their knowledge evolves over time [6]. A key attractiveness of epistemic logic is that a number of toolkits [7, 13, 11], based on model checking [4], support the verification of systems against temporal-epistemic specifications. In these approaches time is assumed to be discrete, either branching or linear, and formulas are evaluated at states. Other notions of time are however of interest and have recently been thoroughly explored. Notably, in *interval temporal logic* [17, 9] propositions are not evaluated at instants but at intervals of time. By doing so one can express properties of continuous processes; this is useful in several AI areas including planning [8, 19].

It is therefore natural and compelling to investigate extensions of interval temporal logic for the specification of multi-agent systems. An attempt towards this aim was made in [12] where a temporal-epistemic language, called epistemic Halpern-Shoham logic (EHS), based on the interval logic proposed by Halpern and Shoham [9] was introduced. In the paper the authors put forward a notion of knowledge interpreted on intervals, defined the resulting model checking problem and analysed its complexity for some limited fragments. This is shown to be PSPACE-hard for a basic epistemic logic with no temporal operators. It is also shown that model checking interpreted systems against specifications combining epistemic operators with the BDE-fragment of the Halpern and Shoham logic (HS) is PSPACE-complete. The BDE-fragment is defined by considering only the modalities for B ("begins"), D ("during"), and E ("ends").

While this work introduces the model checking problem in the context of multi-agent systems against an epistemic language, only a handful of variants are considered. 2^{12} fragments of HS exist; the majority of them have been studied over the years from a satisfiability point of view [3, 5]. While many of them are undecidable, some very expressive decidable fragments exist.

It is therefore natural to identify fragments of EHS which enjoy a decidable model checking problem. We isolate one such fragment in this paper. We begin in Section 2 by defining the semantics of interpreted systems defined on intervals and the syntax of the $A\bar{B}L$ fragment of EHS (that consists of the modalities "after", "begun by" and "later"), which we call EHS$^{A\bar{B}L}$. We illustrate its expressiveness in Section 3 by discussing an interval-based variant of the well-known bit transmission problem. We turn to the the model checking problem for EHS$^{A\bar{B}L}$ in Section 4 where we show its decidability. The methodology we put forward is novel and includes the introduction of a technique, similar in spirit to the pumping lemma in computability theory, that enables us to check infinitely many intervals by analysing a finite number of them. We discuss the limitations of the technique in Section 5 and provide some remarks as to how these may be overcome.

Related work. The only paper we are aware of investigating the model checking problem for interval temporal logic with or without epistemic operator is [12]. The BDE fragment is shown to have a decidable model checking problem, but in that logic one can only refer to intervals with the same length or shorter, thereby greatly limiting the expressivity of any specification. Since the number of such intervals is finite, decidability is immediate. This is not the case for the fragment that we analyse here which includes the "After" modality A, that can refer to an infinite number of intervals.

2 The epistemic-interval logic EHS$^{A\bar{B}L}$

In this section we define a variant of interpreted systems based on intervals and introduce the model checking problem for an expressive fragment of the epistemic-interval logic EHS. We follow the presentation given in [12], although we simplify it by removing the notion of "generalised Kripke structure" discussed there.

Definition 1. *Given a set of agents* $A = \{0, 1, \dots, m\}$, *an* interpreted system *is a tuple* $IS = (\{L_i\}_{i \in A}, \{l_i^0\}_{i \in A}, \{ACT_i\}_{i \in A}, \{P_i\}_{i \in A}, \{t_i\}_{i \in A}, L)$, *where:*

- L_i *is a finite set of local states for agent i;*
- $l_i^0 \in L_i$ *is the initial state for agent i;*
- ACT_i *is a finite set of local actions available to agent i,*
- $P_i : L_i \to 2^{ACT_i}$ *is a local protocol function for agent i, returning the set of possible local actions in a given local state;*
- $t_i \subseteq L_i \times ACT_0 \times \cdots \times ACT_m \times L_i$ *is a local transition relation returning the next local state when a joint action is performed by all agents and the environment on a given local states;*
- $L : S^2 \to 2^{Var}$ *is a labelling function, where* $S = L_0 \times L_1 \times \cdots \times L_m$ *is the set of possible global states for the system and* Var *is a set of propositional variables.*

[1] Department of Computing, Imperial College London, UK

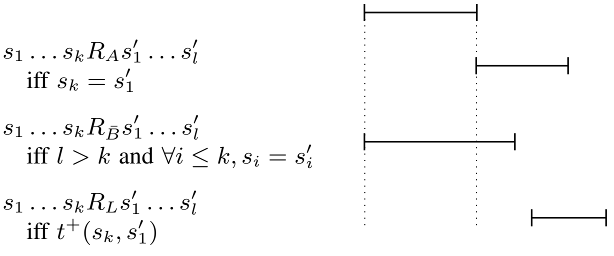

$s_1 \ldots s_k R_A s'_1 \ldots s'_l$
iff $s_k = s'_1$

$s_1 \ldots s_k R_{\bar{B}} s'_1 \ldots s'_l$
iff $l > k$ and $\forall i \leq k, s_i = s'_i$

$s_1 \ldots s_k R_L s'_1 \ldots s'_l$
iff $t^+(s_k, s'_1)$

Figure 1. Three Allen's relations. t^+ denotes the transitive closure of t.

Sometimes we refer to agent 0 as the environment e in the system.

By composing t_i for all agents and the environment we obtain the global transition relation t. We now define models of an IS on sets of paths from is initial state.

Definition 2. *Given a set of agents $A = \{0, 1, \ldots, m\}$ and an interpreted system $IS = (\{L_i\}_{i \in A}, \{l_i^0\}_{i \in A}, \{ACT_i\}_{i \in A}, \{P_i\}_{i \in A}, \{t_i\}_{i \in A}, L)$, an* interval-based interpreted system *(IBIS), or simply* the model of the *IS, is a tuple $M = (S, s_0, t, L)$, where*

- *The set $S = L_0 \times L_1 \times \cdots \times L_m$ is the set of possible global states;*
- *The state $s_0 = (l_0^0, \ldots, l_m^0)$ is the initial state of the system;*
- *$t \subseteq S^2$ is the global transition relation;*
- *L is the labelling function.*

Given an IBIS M, an *interval* in M is a finite path on M, i.e., a sequence of states $I = s_1 s_2 \ldots s_n$ such that $t(s_i, s_{i+1}), 1 \leq i \leq (n-1)$. A *point interval* is an interval that consists of exactly one state. Given an interval $I = s_1 s_2 \ldots s_n$, by *first*(I) we denote the first state of I, namely s_1, by $last(I)$ we denote the last state of I, namely s_n, and by $pi(I)$ we denote whether I is a point interval.

Notice that the above definition is different than the one in [12], where the set of states of an IBIS is the result of applying the standard unravelling procedure to the set of the global states and the global transition relation. Since here we only consider forward modalities (i.e., formulas can only refer to the future), we obtain exactly the same semantics of the $A\bar{B}L$ fragment of the epistemic Halpern-Shoham logic defined in [12].

For a global state $s = (l_0, l_1, \ldots, l_m)$ we denote by $l_i(s)$ the local state $l_i \in L_i$ of agent $i \in A$ in s.

We now define the syntax of the specification language we focus on in this paper. The temporal operators we consider represent some of the relations between intervals as originally defined by Allen [1]. These are depicted in Figure 1: R_A represents "**A**fter" or "meets"; $R_{\bar{B}}$ stands for "**B**egun by" or "started by"; and R_L encodes "**L**ater".

Definition 3. *The syntax of logic EHS$^{A\bar{B}L}$ is defined by the following BNF.*

$$\varphi ::= p \mid pi \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid C_G\varphi \mid \langle A \rangle\varphi \mid \langle \bar{B} \rangle\varphi \mid \langle L \rangle\varphi$$

where $p \in$ Var is a propositional variable, $i \in A$ is an agent, and $G \subseteq A$ is a set of agents.

The logic EHS$^{A\bar{B}L}$ is a fragment of the epistemic-interval logic EHS introduced and studied in [12] as the proposition pi (**p**oint **i**nterval) can be seen as an abbreviation of $\neg\langle B \rangle\top$. We write $[X]\varphi$ for $\neg\langle X \rangle\neg\varphi$ and we use Boolean connectives $\vee$, $\Rightarrow$, $\Leftrightarrow$ and constants $\top, \bot$ in the standard way.

We say that two global states g, g' are such that $g \sim_i g'$ iff $l_i(g) = l_i(g')$, i.e., two global states are epistemically equivalent for agent i if its local states are the same in the two global states [6]. Following [12] we say that two intervals $I = s_1, \ldots s_k, I' = s'_1, \ldots s'_l$ are such that $I \sim_i I'$ iff $k = l$ and for all $j \leq k$, $l_i(s_j) = l_i(s'{}_j)$. In other words agent i cannot distinguish between the corresponding states in the intervals I, I'. We extend this definition to the common knowledge case by considering for any group of agents G, $\sim_G = (\bigcup_{i \in G} \sim_i)^*$, where * denotes the transitive closure.

We can now define when a formula is satisfied in an interval.

Definition 4 (Satisfaction)**.** *Given an EHS$^{A\bar{B}L}$ formula φ, an IBIS M, and an interval I, we inductively define whether φ holds in the interval I, denoted $M, I \models \varphi$, as follows:*

1. *For all $p \in$ Var, we have $M, I \models p$ iff $p \in L($first$(I), last(I))$.*
2. *$M, I \models pi$ iff I is a point interval.*
3. *$M, I \models \neg\varphi$ iff it is not the case that $M, I \models \varphi$.*
4. *$M, I \models \varphi_1 \wedge \varphi_2$ iff $M, I \models \varphi_1$ and $M, I \models \varphi_2$.*
5. *$M, I \models K_i\varphi$, where $i \in A$, iff for all $I' \sim_i I$ we have $M, I' \models \varphi$.*
6. *$M, I \models C_G\varphi$, where $G \subseteq A$, iff for all $I' \sim_G I$ we have $M, I' \models \varphi$.*
7. *$M, I \models \langle X \rangle\varphi$ iff there exists an interval I' such that IR_XI' and $M, I' \models \varphi$, where R_X is an Allen's relation as above.*

In this paper we are interested in analysing the model checking problem for the logic above.

Definition 5. *Given an EHS$^{A\bar{B}L}$ formula φ, an interpreted system IS defining the model M, and an interval I,* the model checking problem for *$\mathcal{L}$ amounts to checking whether or not $M, I \models \varphi$.*

It is instructive to identify expressive fragments for which verification is decidable. As we see later the logic EHS$^{A\bar{B}L}$ is one of such fragments. Before showing this, we turn to analyse the expressiveness of the logic EHS$^{A\bar{B}L}$ by means of a well-known scenario in AI and epistemic logic. It is worth mentioning that the knowledge-free fragment of EHS$^{A\bar{B}L}$ is known to have a satisfiability problem in EXPTIME over the naturals [16].

3 An EHS$^{A\bar{B}L}$-based analysis of the bit transmission protocol

The bit transmission protocol (BTP) is a well-known communication scenario that has been analysed by means of temporal-epistemic specifications [6]. In the BTP two agents, a "Sender" S and a "Receiver" R, communicate over a faulty channel, which may drop messages but may not flip them. We here present a revised version of the protocol where the sender needs to compute what message to send before initiating communication; we refer to the existing literature for more details [6]. As in the original protocol we here consider only one bit of information, either 0 or 1; the protocol can be generalised with no difficulty. As usual we assume that S keeps sending the bit until he gets an acknowledgement from R who, in turn, remains silent until he gets the bit; from then on R keeps sending an acknowledgement back to S. A CTLK specification often considered when analysing the BTP is $AG(reckack \rightarrow K_S(K_R(bit = 0) \vee K_R(bit = 1)))$; in other words, when an ack has been received by the sender, the sender knows that the receiver knows the value of the bit. While the specification has been shown to be useful, discrete notions of time do not enable us to describe sequences of contiguous or overlapping epistemic states of affairs in the runs. Intuitively, in the absence of any fairness constraint, a property we would like to ensure is that runs of the protocol consist of potentially unbounded intervals in which S is first computing the value to send, then S is waiting for the acknowledgement, and finally enters an unbounded interval in which S knows that R knows the value of the bit. Differently from the CTL-based specification the emphasis here is on specifying what holds at *sequences of intervals* which may be related among them following the Allen relations. In what follows we show that the EHS$^{A\bar{B}L}$ logic can provide an expressive specification for the variant of the BTP here described.

To do this we first model the revised BTP in the formalism of the previous section. The sender is modelled by considering locals states of the form $(status, bit) \in L_S$, where $status \in \{computing, sending, acked\}$ and $bit \in \{0, 1, \lambda\}$. We take S's initial state to be $l_s = (computing, \lambda)$. The actions for S consist of $ACT_S = \{compute, send_0, send_1, \epsilon\}$, where $compute$ represents the action of computing the bit to be sent and ϵ encodes a null action.

The receiver agent R is modelled by taking $L_R = \{\lambda, 0, 1\}$. R's initial state is $l_r = \lambda$, when R is waiting for the bit to be received. R's actions are $ACT_R = \{\epsilon, sendack\}$ where ϵ is the null action.

We take the environment's local states to consist of a single state $l_e = \lambda$ from which it may non-deterministically perform the actions $\rightarrow, \leftarrow, \leftrightarrow$ and ϵ, representing, respectively, messages being delivered from S to R, from R to S, in both directions, and in no direction.

The protocols mapping states to possible actions can be formalised by following the description above. The transition relation t_S for S is such that a loop may be formed on the local state $(computing, \lambda)$ by means of any joint action that includes the local action $compute$. Under the same conditions the relation t_S also includes a non-deterministic transition to the states $(sending, 0)$, $(sending, 1)$, from which S starts sending the bit. S remains in one of these states until he receives an acknowledgement from R, triggered by either the joint actions $(sendbit, sendack, \leftarrow)$ or $(sendbit, sendack, \leftrightarrow)$. From that point onward S moves either to the local state $(acked, 0)$, or to $(acked, 1)$ depending on the value of the bit and loops on that state for the rest of the run.

The transitions for R can similarly be formalised. The relation t_R includes a loop on the initial state λ where R performs the action ϵ. From there R makes a transition either to the state 0 or 1 following the joint actions $(sendbit, \epsilon, \rightarrow)$ and $(sendbit, \epsilon, \leftrightarrow)$. From that state R can only loop in combination with the local action $sendack$.

From the description of the IS for the BTP above we can generate the IBIS M. We consider a labelling function L for M such that $p \in L(s, s')$, where $s = (\lambda, (status_S, bit_S), bit_R)$, $s' = (\lambda, (status'_S, bit'_S), bit'_R)$ iff:

- $p = sending$ and $status'_S \neq acked$,
- $p = computing_{bit'_S}$, $status_S = status'_S = computing$ and $bit'_S \neq \lambda$, or
- $p = b^R_{bit_R}$ and $bit_R \neq \lambda$.

We are interested in verifying the following property: In any interval beginning with an interval in which S is computing the bit, if S stops sending the bit, having started at some point after its computation began, then in all intervals from that point onwards S knows that R knows the value of the bit. This represents the natural flow of intervals for the protocol culminating in an interval where an epistemic postcondition holds.

Let $[G]\varphi = \varphi \wedge [\bar{B}]\varphi \wedge [A]\varphi \wedge [L]\varphi$ be an operator that $[G]\varphi$ holds if φ holds in all the reachable intervals. The specification above can be expressed by means of the following EHS$^{A\bar{B}L}$ formula.

$$\bigwedge_{b \in \{0,1\}} [G](computing_b \rightarrow [\bar{B}](\neg sending \Rightarrow [A] K_S K_R bit^R_b))$$

It can be checked that the property holds in M. Note that this specification is not expressible in any other fragment of EHS for which the model checking problem is known to be decidable; in particular, it is not expressible in the BDE fragment analysed in [12].

In the next section we will show that the model checking problem against EHS$^{A\bar{B}L}$ specifications is decidable.

4 Decidability of the model checking problem

To begin, observe that the modality $\langle L \rangle$ can be expressed by using $\langle A \rangle$; indeed, for any φ, $\langle L \rangle \varphi \equiv \langle A \rangle (\neg pi \wedge \langle A \rangle \varphi)$. Given this, in what follows we assume that the formulas do not contain $\langle L \rangle$ operators.

Let KM be the set of the epistemic modalities, i.e., $KM = \{K_i \mid 1 \leq i \leq m\} \cup \{C_G \mid G \subseteq \{1, \ldots, m\}\}$, and $SM = KM \cup \{\langle A \rangle, \langle \bar{B} \rangle\}$ be the set of all the operators in EHS$^{A\bar{B}L}$. For convenience, for each $X \in SM$ we define a relation R_X as follows: $R_{\langle A \rangle} = R_A$, $R_{\langle \bar{B} \rangle} = R_{\bar{B}}$, $R_{K_i} = \sim_i$ and $R_{C_G} = \sim_G$.

Given a formula φ, a *top-level subformula* of φ is a modal subformula of φ which is not in the scope of any modality. For example, the top level subformulas of $\langle A \rangle K_1 p \wedge C_{\{1\}} \langle \bar{B} \rangle q$ are $\langle A \rangle K_1 p$ and $C_{\{1\}} \langle \bar{B} \rangle q$. Assume an IBIS M such that $|S_{IS}| = n$ states. Let $f^M(\varphi)$ be defined recursively as follows: $f^M(\varphi) = 2n^2 2^{f^M(\varphi_1)} \ldots 2^{f^M(\varphi_k)}$, where $X_1 \varphi_1 \ldots X_k \varphi_k$ are the top-level subformulas of φ with $X_i \in SM, i = 1, \ldots, k$. If φ contains no modalities, then $f(\varphi) = n^2$. Clearly f is non-elementary in the size of φ.

A key consideration in our decidability proof for the model checking problem for EHS$^{A\bar{B}L}$ is that, as we will see later, the problem can be solved by considering only a *bounded* number of intervals. To show this, we give a bounded satisfaction definition and show that this is equivalent to (unbounded) satisfaction of Definition 4.

Definition 6 (Bounded satisfaction)**.** *Given an EHS$^{A\bar{B}L}$ formula φ, an IBIS M, and an interval I, we inductively define whether $M, I \models_B \varphi$, as follows:*

1. *For all $p \in Var$, we have $M, I \models_B p$ iff $p \in L(I)$.*
2. *$M, I \models pi$ iff I is a point interval.*
3. *$M, I \models_B \neg\varphi$ iff it is not the case that $M, I \models_B \varphi$.*
4. *$M, I \models_B \varphi_1 \wedge \varphi_2$ iff $M, I \models_B \varphi_1$ and $M, I \models_B \varphi_2$.*
5. *$M, I \models_B K_i\varphi$, where $i \in A$, iff for all $I' \sim_i I$ we have $M, I' \models_B \varphi$.*
6. *$M, I \models_B C_G\varphi$, where $G \subseteq A$, iff for all $I' \sim_G I$ we have $M, I' \models_B \varphi$.*
7. *$M, I \models_B \langle X \rangle \varphi$ iff there exists an interval I' such that $|I'| \leq |I| + f^M(\varphi)$, $I R_X I'$ and $M, I' \models_B \varphi$, where X is A or $\bar{B}$.*

It follows from the above that to determine the truth value of a formula in a given interval of a system w.r.t. the bounded semantics, one only needs to consider a bounded number of intervals. This is because there are only finitely many intervals of the same size as I (cases 5, and 6) and finitely many intervals whose size is less than or equal to a given bound which depends on the formula to be checked (case 7). This leads to the following.

Theorem 7. *The model checking problem for EHS$^{A\bar{B}L}$ on bounded semantics is decidable.*

Algorithm 1 The model checking procedure for the EHS$^{A\bar{B}L}$ logic.

```
procedure VERIFY($M$, $I$, $\varphi$)
    if $\varphi = p$ then return $p \in L(first(I), last(I))$
    if $\varphi = pi$ then return $pi(I)$
    if $\varphi = \neg\varphi'$ then return NOT(VERIFY($M$, $I$, $\varphi'$))
    if $\varphi = \varphi_1 \wedge \varphi_2$ then return AND(VERIFY($M$, $I$, $\varphi_1$), VERIFY($M$, $I$, $\varphi_2$))
    if $\varphi = K_i\varphi'$ where $i \in A$ then
        for all $J$ s.t. $I \sim_i J$ do
            if NOT(VERIFY($M$, $J$, $\varphi'$)) then return false
        return true
    if $\varphi = C_G\varphi'$ where $G \subseteq A$ then
        for all $J$ s.t. $I \sim_G J$ do
            if NOT(VERIFY($M$, $J$, $\varphi'$)) then return false
        return true
    if $\varphi = X\varphi'$ where $X \in \{\langle A\rangle, \langle\bar{B}\rangle\}$ then
        for all $J$ s.t. $IR_X J$ and $|J| \leq f(\varphi) + |I|$ do
            if VERIFY($M$, $J$, $\varphi'$) then return true
        return false
```

Proof of Theorem 7. The procedure VERIFY() (Algorithm 1) solves the model checking problem. Firstly, note that the procedure always stops. In case of the knowledge modalities, VERIFY(M, I, φ) calls itself recursively at most $|S|^{|I|}$ times on the relevant subformula. The bound $|S|^{|I|}$ corresponds to the number of intervals of length $|I|$. In case of the temporal modalities, the algorithm calls itself at most $|S|^{|I|+f^M(\varphi)}$ times on the nested subformula. Since f cannot be bounded elementarily in the size of φ, the whole procedure is non-elementary.

To see that the procedure solves the model checking problem, observe that the exit calls correspond to the definition of the bounded semantics. □

The main technical result of this section is the equivalence between the bounded and the unbounded semantics. To achieve this we introduce the notion of modal context tree and some results pertaining to those.

Definition 8 (Modal Context Tree)**.** *Given an IBIS M, the* modal context tree *of an interval I w.r.t. an EHS$^{A\bar{B}L}$ formula φ, denoted by MCT_I^φ, is the unranked tree with labelled nodes and edges defined recursively as follows.*

- *The root of the tree is labelled by I.*
- *For each top-level subformula $X\psi$ of φ and each interval I' such that $IR_X I'$, the root of MCT_I^φ has an X-successor $MCT_{I'}^\psi$ (X indicates the labelling of an edge).*

In other words, MCT_I^φ contains all the intervals that need to be considered to determine the value of φ in I. Modal context trees are usually infinite. Below we present their finite counterparts.

Definition 9 (Restricted Modal Context Tree)**.** *Given an IBIS M, the* restricted modal context tree *of an interval I w.r.t. an EHS$^{A\bar{B}L}$ formula φ, denoted by $RMCT_I^\varphi$, is the unranked tree obtained from MCT_I^φ first by changing each node label from I to $first(I), last(I), pi(I)$ and then by applying recursively the following operation in the bottom-up manner:*

- *If w is a node connected to a subtree T by an edge labelled by some X, then remove all the other subtrees T' that are identical to T and such that w is connected to T' by an edge labelled by X.*

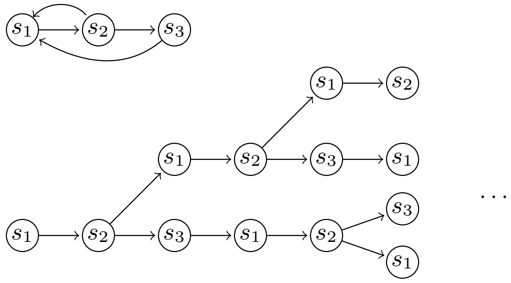

Figure 2. The agent 1 from Example 10 (top left) and its unraveling.

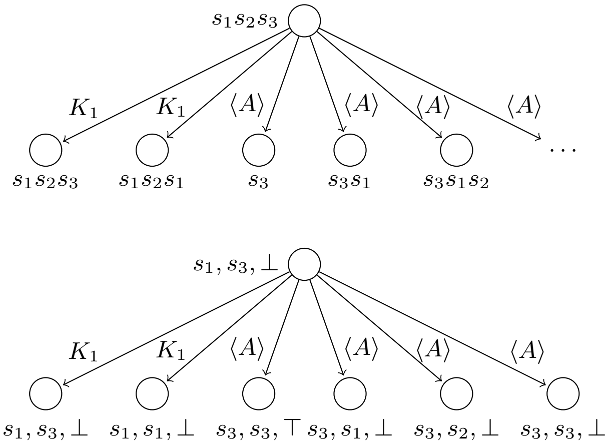

Figure 3. MCT_I^φ from Example 10 (top) and the corresponding $RMCT_I^\varphi$ (bottom).

So $RMCT_I^\varphi$ is obtained from MCT_I^φ by replacing intervals in the labels by their endpoints and removing identical subtrees.

Example 10. *Consider an agent 1 with local states $L_1 = \{s_1, s_2, s_3\}$ and one actions $ACT_1 = \{\epsilon\}$ such that $t_1 = \{(s_1, \epsilon, s_2), (s_2, \epsilon, s_3), (s_3, \epsilon, s_1), (s_2, \epsilon, s_1)\}$ and the environment such that $L_e = \{s_e\}$ and $ACT_e = \{\epsilon\}$ (see Figure 2). Assume that $\sim_1 = \{(s_1, s_1), (s_2, s_2), (s_3, s_3), (s_1, s_3), (s_3, s_1)\}$.*

Consider a formula $\varphi = K_1 p \wedge \neg\langle A\rangle p$ and an interval $I = s_1s_2s_3$ (to simplify the notion, we ignore the environment states here, writing s_i for the IBIS states instead of (s_e, s_i)). The root of MCT_I^φ (Figure 3, top) is labelled by I. The top level subformulas of φ are $K_1 p$ and $\langle A\rangle p$. The root of MCT_I^φ has exactly two K_1-successors: $s_1s_2s_3$, $s_1s_2s_1$ (since $s_3 \sim_1 s_1$), and infinitely many $\langle A\rangle$-successors: s_3, s_3s_1, $s_3s_1s_2$, $s_3s_1s_2s_1 \ldots$.

While the tree MCT_I^φ is infinite, the tree $RMCT_I^\varphi$ (Figure 3, bottom) is finite. For example, the nodes s_3s_1 and $s_3s_1s_2s_1$ are represented by the same node $s_3, s_1, \bot$ that represents all the $\langle A\rangle$-successors of the root starting in s_3 and ending in s_1.

Lemma 11. *Given an IBIS M and a formula φ, the following facts hold.*

1. *$|\{RMCT_I^\varphi \mid I \text{ is an interval in } M\}| < f^M(\varphi)$.*
2. *If I, I' are intervals such that $RMCT_I^\varphi = RMCT_{I'}^\varphi$, then $M, I \models \varphi$ if and only if $M, I' \models \varphi$.*
3. *If I, I' are intervals such that $RMCT_I^\varphi = RMCT_{I'}^\varphi$ and an*

interval J is such that $last(J)$ is a predecessor of $\mathit{first}(I)$, then $RMCT^{\varphi}_{JI} = RMCT^{\varphi}_{JI'}$.

Proof. We show Part 1. by induction on φ. Clearly, if a formula has no modalities, then $\{RMCT^{\varphi}_{I} \mid I \text{ is an interval in } M\}$ contains trees with only one node. For $n > 1$, the number of such trees can bounded by the number of different labelling of a node, i.e., $n^2 + n$, which is greater that $2n^2$.

Consider a formula φ with the top-level subformulas $X_1\varphi_1, \ldots, X_k\varphi_k$. Each tree for φ consists of one of $n^2 + n$ possible roots and, for each i, any subset of subtrees for φ_i. Therefore, $|\{RMCT^{\varphi}_{I} \mid I \text{ is an interval in } M\}| < 2n^2 2^{f^M(\varphi_1)} \ldots 2^{f^M(\varphi_k)} = f^M(\varphi)$.

Part 2 can also be shown by induction on φ. Assume that $\varphi = p$ for some variable p. The root of the $RMCT^{\varphi}_{I}$ is labelled by the endpoints of I, and the root of the $RMCT^{\varphi}_{I'}$ is labelled by the endpoints of I'. Since the two trees are equal, the endpoints are the same and since the labelling depends only on the endpoints of an interval, it follows that $M, I \models p$ iff $M, I' \models p$.

Assume that $\varphi = pi$. The root of the $RMCT^{\varphi}_{I}$ is labelled by $pi(I)$, and so is the root of $RMCT^{\varphi}_{I'}$, and therefore $pi(I) = pi(I')$.

Assume that $\varphi = \langle A\rangle\varphi'$ for some φ'. As above, we know that the last point of I and I' is the same point s. Therefore, $M, I \models \langle A\rangle\varphi'$ iff there is a path starting from s satisfying φ' which is iff $M, I' \models \langle A\rangle\varphi'$.

Assume that $\varphi = \neg\varphi'$ for some φ'. By the inductive assumptions, $M, I \models \varphi'$ iff $M, I' \models \varphi'$, so $M, I \models \varphi$ iff $M, I' \models \varphi$.

Assume that $\varphi = \varphi_1 \wedge \varphi_2$ for some φ_1, φ_2. By the induction assumption, $M, I \models \varphi_1$ iff $M, I' \models \varphi_1$ and $M, I \models \varphi_2$ iff $M, I' \models \varphi_2$, so $M, I \models \varphi$ iff $M, I' \models \varphi$.

Assume that $\varphi = K_i\varphi'$ for some φ' and i. Assume that $M, I \models \varphi$. Consider any interval J' such that $I' \sim_i J'$. By the definition, in the tree $MCT^{\varphi}_{I'}$ the subtree $MCT^{\varphi'}_{J'}$ is an K_i-successor of the root. It follows that in the tree $RMCT^{\varphi}_{I'}$, $RMCT^{\varphi'}_{J'}$ is an K_i-successor of the root. Let J be such that $I \sim_i J$ and $RMCT^{\varphi'}_{J'} = RMCT^{\varphi'}_{J}$. Such a J exists because $RMCT^{\varphi}_{I'} = RMCT^{\varphi}_{I}$. Clearly, since $M, I \models \varphi$, $M, J \models \varphi'$. By the inductive assumptions, $M, J' \models \varphi'$. Therefore $M, I' \models \varphi$.

The proof for the cases of $\varphi = C_G\varphi'$ and $\varphi = \langle\bar{B}\rangle\varphi'$ are similar and omitted.

Part 3. Given a formula φ, an IBIS M, an interval I and a state s such that $t(s, \mathit{first}(I))$, $RMCT^{\varphi}_{sI}$ can be computed on the basis of M and $RMCT^{\varphi}_{I}$. Therefore, if we consider two intervals I, I' of the same model such that $RMCT^{\varphi}_{I} = RMCT^{\varphi}_{I'}$, then $RMCT^{\varphi}_{sI} = RMCT^{\varphi}_{sI'}$. The consideration above can be used to prove Part 3.

To do this, consider the procedure PREPEND in Algorithm 2. We show that the result of PREPEND(s, T), where s is a state and $T = RMCT^{\varphi}_{I}$ for some interval I, is $RMCT^{\varphi}_{sI}$. In the algorithm we use the function $singe\_node\_tree(l)$ to define a new tree containing only the root labelled with l; $label(T)$ returns the label of a node; $root(T)$ returns the root of T; $subtree(T, w)$ returns a subtree of T rooted in w; and $add\_a\_subtree(t, l, T)$ adds T as an l-successor of t. Recall that t is the transition function of the interpreted system.

We now show by induction that for any φ and any I, s, PREPEND$(s, RMCT^{\varphi}_{I}) = RMCT^{\varphi}_{sI}$.

The roots of PREPEND$(s, RMCT^{\varphi}_{I})$ and $RMCT^{\varphi}_{sI}$ are labelled by $(s, last(I), \perp)$, so they are equal.

Assume that $X_1\varphi_1, \ldots, X_k\varphi_k$ are the top-level subformulas of φ and $i \in \{1, \ldots, k\}$ (if there are no such formulas, then the result follows).

Assume $X_i = \langle A\rangle$. Observe that for any interval J, IR_AJ iff sIR_AJ. Therefore the $\langle A\rangle$-successors of the root in $RMCT^{\varphi}_{I}$ are the same as $\langle A\rangle$-successors of the root in $RMCT^{\varphi}_{sI}$, and therefore they are the same in PREPEND$(s, RMCT^{\varphi}_{I})$ and $RMCT^{\varphi}_{sI}$.

Algorithm 2 The procedure for Part 3 of Lemma 11

```
procedure PREPEND($s$, $T$)
    $(f, l, pi) \leftarrow label(root(T))$
    $T' \leftarrow single\_node\_tree((s, l, \perp))$
    for all $\langle A\rangle$-successor $w$ of $root(T)$ do
        $add\_a\_subtree(root(T'), \langle A\rangle, subtree(T, w))$
    for all $\langle\bar{B}\rangle$-successor $w$ of $root(T)$ do
        $T'' \leftarrow$ PREPEND$(s, subtree(T, w))$
        $add\_a\_subtree(root(T'), \langle\bar{B}\rangle, T'')$
    for all $X \in KM$ and $X$-successor $w$ of $root(T)$ do
        $(f', l', pi') \leftarrow label(w)$
        for all $s' R_X s$ such that $t(s', f')$ do
            $T'' \leftarrow$ PREPEND$(s', subtree(T, w))$
            $add\_a\_subtree(root(T'), X, T'')$
    return $T'$
```

Assume $X_i = \langle\bar{B}\rangle$. Observe that for any interval J, IR_BJ iff sIR_BsJ. Therefore, $RMCT^{\varphi_i}_{sJ}$ is an $\langle\bar{B}\rangle$-successors of the root in $RMCT^{\varphi_i}_{sI}$ iff $RMCT^{\varphi_i}_{J}$ is an $\langle\bar{B}\rangle$-successors of the root in $RMCT^{\varphi_i}_{I}$. By the inductive hypothesis, for any J the trees PREPEND$(s, RMCT^{\varphi_i}_{J})$ and $RMCT^{\varphi_i}_{sJ}$ are the same; therefore, the set of the $\langle\bar{B}\rangle$-successors of the root is the same in PREPEND$(s, RMCT^{\varphi}_{I})$ and $RMCT^{\varphi}_{sI}$.

Assume $X_i \in KM$. Observe that for an interval J, and state s' s.t. $s' R_{X_i} s$ and $t(s', \mathit{first}(J))$, $IR_{X_i}J$ iff $sIR_{X_i}s'J$. As in the previous case, we have that $RMCT^{\varphi_i}_{s'J}$ is an X_i-successor of the root in $RMCT^{\varphi}_{sI}$ iff $RMCT^{\varphi_i}_{J}$ is an X_i-successor of the root in $RMCT^{\varphi}_{I}$ and $s' R_{X_i} s$; so we conclude that the sets of X_i-successors of the root is the same in PREPEND$(s, RMCT^{\varphi}_{I})$ and $RMCT^{\varphi}_{sI}$. □

Having established the Lemma above, we can now give the main result of this section.

Theorem 12. *Given an EHS$^{A\bar{B}L}$ formula φ, an IBIS M and an interval I, $M, I \models \varphi$ if and only if $M, I \models_B \varphi$.*

Proof. The proof is by induction on the structure of φ.

The cases for $\varphi = p$, $\varphi = pi$, $\varphi = \neg\varphi'$, $\varphi = \varphi_1 \wedge \varphi_2$, $\varphi = K_i\varphi'$, or $\varphi = C_G\varphi'$ for some subformulas $\varphi', \varphi_1, \varphi_2$, follow given the fact that the semantic rules are the same in both semantics.

Assume that $\varphi = X\varphi'$ for some φ' and $X \in \langle A\rangle, \langle\bar{B}\rangle$. If $M, I \models_B \varphi$, then there is an interval I' of bounded size such that $M, I' \models_B \varphi'$ and IR_XI'. By the induction hypothesis, $M, I' \models \varphi'$ and therefore $M, I \models \varphi$. If $M, I \models \varphi$, then there is an interval I' such that $M, I' \models \varphi'$ and IR_XI'. Let I' be the shortest possible interval with this property. We show that $|I'| \leq |I| + f^M(\varphi)$.

Let $I' = s_1 \ldots s_t$ and I'_k denote the suffix of I' starting at s_k, i.e., $s_k \ldots s_t$. Assume that $|I'| > |I| + f^M(\varphi')$. By Lemma 11.1 we have that among $I'_{|I|+1} \ldots I'_{|I|+f^M(\varphi')}$ there are two suffixes I'_k, I'_l such that $|I| < k < l$ and $RMCT^{\varphi'}_{I'_k} = RMCT^{\varphi'}_{I'_l}$.

Let $J = s_1 \ldots s_{k-1}s_l \ldots s_t$. By Part 3 of Lemma 11, we have that $RMCT^{\varphi'}_{J} = RMCT^{\varphi'}_{I'}$, and so by Part 2 of Lemma 11 we have that J is an interval such that $M, J \models \varphi'$ and IR_XJ. Clearly, $|J| < |I'|$; this is a contradiction given our assumption that $|I'| > |I| + f^M(\varphi')$.

Notice that the requirement that $k > |I|$ is only needed in the case of $\langle\bar{B}\rangle$ since J has to contain I as a prefix. □

We can now derive the main technical result of the paper.

Theorem 13. *The model checking problem for EHS$^{A\bar{B}L}$ is decidable.*

The proof follows immediately by considering Theorem 7 and Theorem 12

Consider a relation R_N such that $s_1 \dots s_k R_N s'_1 \dots s'_l$ iff $t(s_k, s'_1)$ and a corresponding modality $\langle N \rangle$. This modality is a counterpart of the X operator of CTL. In EHS, this can be defined by assigning $\langle N \rangle \varphi = \langle A \rangle (\neg pi \wedge \langle B \rangle \langle B \rangle \bot \wedge \langle A \rangle \varphi)$. One significant limitation of EHS$^{A\bar{B}L}$ is that it cannot define $\langle N \rangle$. However, the technique above can be extended to the case of this operator in a straightforward way.

Proposition 14. *The model checking problem for EHS$^{A\bar{B}L}$ extended with the modality $\langle N \rangle$ is decidable.*

5 Conclusions and Future Work

Since the early proposals, dating back to Prior [18], to use logic to express temporal concepts, there has always been an active interest in exploring different models of time. The usual dichotomy between linear and branching models is one example of this, but others exist, including discrete and continuous models. Interval temporal logic was put forward in the 90s [17, 9] as a powerful mechanism to represent and reason about continuous processes. These often occur, for instance, in planning where one needs to express facts that occur at intervals and not at time instances. This was preceded by investigations dating back at least to the 70s [10].

The current literature on HS logic focuses on the study of subtly different logics expressing intervals which can be defined by using subsets of the operators corresponding to the Allen's relations. While the number of possible fragments is 10^{12}, most of them are known to have undecidable satisfiability problems [5, 14]. A key avenue of research has so far involved the identification of fragments for which satisfiability and validity are decidable [2, 3, 15].

The logic EHS combining the interval temporal logic HS and epistemic logic has recently been introduced [12]. Since EHS is a proper extension of HS, its satisfiability problem is also undecidable. However, it was shown that the model checking problem for its BDE fragment, as well as a number of other weak logics, is decidable. While these results are positive, the BDE fragment is not particularly expressive; for example, all the intervals it may refer to are of a bounded length. In this paper we showed the decidability of the model checking problem for the $A\bar{B}L$ fragment of the logic. As we discussed, specifications written in EHS$^{A\bar{B}L}$ enable us to refer to intervals of arbitrary length. The BTP example that we discussed in Section 3 demonstrates this.

One possible future direction of study is to characterise the expressive power of the logics here discussed (under the locality assumption, see [12]) to that of more popular formalisms such as CTLK.

While a key result of the paper is the decidability result, the reduction technique put forward in its proof, enabling us to reduce the model checking problem for infinitely many intervals to one on intervals on bounded length, seems significant on its own. It is possible that the decidability for other fragments may be obtained by adapting the scheme of the proof here introduced.

We conclude by remarking that it is currently not known whether the model checking problem for the full EHS logic is undecidable, although we suspect that is the case. Further research on identifying the precise border of decidability is therefore required.

ACKNOWLEDGEMENTS

This research was funded by the EPSRC under grant EP/I00520X.

REFERENCES

[1] J. F. Allen, 'Maintaining knowledge about temporal intervals', *Communications of the ACM*, **26**(11), 832–843, (1983).

[2] D. Bresolin, D. Della Monica, V. Goranko, A. Montanari, and G. Sciavicco, 'Metric propositional neighborhood logics: Expressiveness, decidability, and undecidability.', in *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI10)*, pp. 695–700, (2010).

[3] D. Bresolin, D. Della Monica, A. Montanari, P. Sala, and G. Sciavicco, 'Interval temporal logics over finite linear orders: the complete picture', in *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI12)*, pp. 199–204, (2012).

[4] E. M. Clarke, O. Grumberg, and D. A. Peled, *Model Checking*, The MIT Press, Cambridge, Massachusetts, 1999.

[5] D. Della Monica, *Expressiveness, decidability, and undecidability of interval temporal logic*, Ph.D. dissertation, University of Salerno, 2011.

[6] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning about Knowledge*, MIT Press, Cambridge, 1995.

[7] P. Gammie and R. van der Meyden, 'MCK: Model checking the logic of knowledge', in *Proceedings of 16th International Conference on Computer Aided Verification (CAV04)*, volume 3114 of *Lecture Notes in Computer Science*, pp. 479–483. Springer, (2004).

[8] V. Goranko, A. Montanari, and G. Sciavicco, 'A road map of interval temporal logics and duration calculi', *Journal of Applied Non-Classical Logics*, **14**(1-2), 9–54, (2004).

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

[10] C. L. Hamblin, 'Instants and intervals', *Studium Generale*, **27**, 127–134, (1971).

[11] M. Kacprzak, W. Nabialek, A. Niewiadomski, W. Penczek, A. Pólrola, M. Szreter, B. Wozna, and A. Zbrzezny, 'Verics 2007 - a model checker for knowledge and real-time', *Fundamenta Informaticae*, **85**(1-4), 313–328, (2008).

[12] A. Lomuscio and J. Michaliszyn, 'An epistemic Halpern-Shoham logic', in *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI13)*, pp. 1010–1016. AAAI Press, (2013).

[13] A. Lomuscio, H. Qu, and F. Raimondi, 'MCMAS: A model checker for the verification of multi-agent systems', in *Proceedings of the 21th International Conference on Computer Aided Verification (CAV09)*, volume 5643 of *Lecture Notes in Computer Science*, pp. 682–688. Springer, (2009).

[14] J. Marcinkowski and J. Michaliszyn, 'The ultimate undecidability result for the Halpern-Shoham logic', in *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science (LICS11)*, pp. 377–386. IEEE Computer Society, (2011).

[15] A. Montanari, G. Puppis, and P. Sala, 'Maximal decidable fragments of Halpern and Shoham's modal logic of intervals.', in *Proceedings of 37th International Colloquium on Automata, Languages and Programming (ICALP10)*, volume 6199 of *Lecture Notes in Computer Science*, pp. 345–356, (2010).

[16] A. Montanari, G. Puppis, P. Sala, and G. Sciavicco, 'Decidability of the interval temporal logic ABBbar over the natural numbers', in *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (STACS10)*, volume 5, pp. 597–608, (2010).

[17] B. C. Moszkowski, *Reasoning about digital circuits*, Ph.D. dissertation, Stanford University, Stanford, CA, USA, 1983.

[18] A. N. Prior, 'Possible worlds', *Philosophical Quarterly*, **12**, 36–43, (1962).

[19] B. Richards, Y. Jiang, and H. Choi, 'On interval-based temporal planning: An iq strategy', in *Methodologies for Intelligent Systems*, volume 542 of *Lecture Notes in Computer Science*, 226–235, Springer, (1991).