

Model Checking Multi-Agent Systems against Epistemic HS Specifications with Regular Expressions

Alessio Lomuscio
Imperial College London, UK

Jakub Michaliszyn
University of Wrocław, Poland

Abstract

We introduce EHS^+ , a novel temporal-epistemic logic defined on temporal intervals characterised by regular expressions. We investigate the complexity of verifying multi-agent systems against EHS^+ specifications for a number of fragments of EHS^+ with results ranging from PSPACE-completeness to non-elementary time. The findings show that, at least for the fragments under analysis, the increase in expressiveness obtained by using regular expressions rather than end-points as standard, can be achieved without increasing the complexity of the problem. We show that the expressiveness of regular expressions can also be adopted at the level of specifications without severe computational cost. To do so we introduce a further temporal-epistemic logic, called EHS^{RE} , in which regular expressions are used within propositions, and give a polynomial time reduction of the model checking problem from EHS^{RE} to EHS^+ .

1 Introduction

A relatively recent area of interest in epistemic logic (Ditmarsch et al. 2015), or logics for knowledge, has been the study of its model checking problem (Clarke, Grumberg, and Peled 1999; Lomuscio and Penczek 2015). This is of relevance when *reasoning about multi-agent systems* (MAS), in which autonomous, self-interested agents interact with one another and the environment based on their knowledge. Given a MAS, it is of interest to predict what epistemic states the agents will achieve in their interaction. This may involve the knowledge they have about the changing environment, how their own knowledge evolves over time, the knowledge they have about the knowledge of other agents, or more complex notions such as common or distributed knowledge, as is the case, for example, in security (Boureau, Cohen, and Lomuscio 2009), distributed diagnosis (Ezekiel and Lomuscio 2009; Ezekiel et al. 2011), and services (Lomuscio, Qu, and Solanki 2008; Belardinelli, Lomuscio, and Patrizi 2012). Traditionally, the mainstream modal logic $S5_n$, augmented with group modalities, has been adopted as the underlying formalism to model the knowledge of agents in a MAS (Fagin et al. 1995).

The assumptions made about the nature of time are a key factor when reasoning about evolving knowledge or

knowledge about dynamic environments. Time is normally assumed to be discrete and formulas are typically interpreted at states. Model checking approaches against epistemic properties under linear time (Meyden and Shilov 1999; Gammie and van der Meyden 2004) and branching time (Penczek and Lomuscio 2003; Raimondi and Lomuscio 2005) have been put forward and open-source model checkers have been released (Gammie and van der Meyden 2004; Lomuscio, Qu, and Raimondi 2015; Kacprzak et al. 2008). However, other approaches are possible: for example in (Lomuscio, Penczek, and Woźna 2007) a bounded model checking technique for the verification of MAS against an epistemic language enriched with real time was put forward. The choice of the appropriate notion of time is often influenced by the particular application under study. For example, while discrete time is adequate for reasoning about knowledge and privacy in the context of security applications, protocols for real-time networks normally require timed-automata and continuous time.

Interval temporal logic (Moszkowski 1983; Halpern and Shoham 1991) has been proposed in the past as a formalism to reason about time when formulas are naturally interpreted on *intervals* rather than states. An instance of these problems is planning where it is at times useful to reason about properties that hold continuously between two states. In recent work (Lomuscio and Michaliszyn 2013) an epistemic variant of the Halpern-Shoham logic (HS) was introduced by considering the notion of epistemic indistinguishability of intervals for the agents in the system. The resulting Epistemic Halpern-Shoham logic (or EHS) consists of a family of 2^{12} fragments of varying expressivity and complexity depending on which set of operators for intervals is adopted from the 12 modalities available (A for “after”, B for “begins”, D for “during”, E for “ends”, L for “later”, O for “overlaps” and their respective inverses \bar{A} , \bar{B} , \bar{D} , \bar{E} , \bar{L} , \bar{O}). Some fragments of EHS admit a model checking problem of complexity ranging from PTIME to PSPACE-hard. Further expressive fragments of EHS have been identified and studied. For example, the $\bar{A}\bar{B}L$ fragment of EHS is known to have a decidable model checking problem (Lomuscio and Michaliszyn 2014). The decidability of other fragments is currently open.

A fundamental feature of EHS is that the labelling function is defined on the endpoints of the intervals. This is a

widely adopted setup in the literature, and it corresponds to the intuitive representation of intervals as pairs. However, other choices are possible. For example, (Montanari et al. 2014) considers the labelling for an interval as the intersection of the labellings of all its elements. This increases the expressiveness of the resulting formalism. In this paper, we introduce considerably more expressive labellings in the context of epistemic specifications in that we allow the labelling function to be given by *any regular expression on the states of the interval*. We argue that this results in a considerable increase in the expressiveness of the specifications at no computational cost in terms of the corresponding model checking problem.

Related Work. Work on interval temporal logic has until recently concerned the satisfiability problem only. This is known to be undecidable in general (Halpern and Shoham 1991; Bresolin et al. 2014b; 2014a; Della Monica 2011), even when HS is restricted to some unimodal fragments (Bresolin et al. 2011a). Notable decidable fragments are the $A\bar{A}$ fragment with length constraints (Bresolin et al. 2013; 2010), the $AB\bar{B}L$ fragment (Bresolin et al. 2011b), and the recently identified Horn fragment (Artale et al. 2015). Some fragments are decidable only over some particular classes of orderings. For example, the $B\bar{B}D\bar{D}LL$ fragment was shown to be decidable over the class of all dense orders (Montanari, Puppis, and Sala 2009), while the D fragment is undecidable over discrete orders (Marcinkowski and Michaliszyn 2011; 2014). The same logic is decidable under the assumption that an interval is its own infix (Montanari, Pratt-Hartmann, and Sala 2010). While a wealth of results have been put forward, open questions remain. For example, the decidability of the D fragments over the class of all orders is currently open. None of this work considers epistemic modalities.

EHS, the first epistemic logic based on intervals, was defined and studied in (Lomuscio and Michaliszyn 2013), which also introduced the model checking problem for interval temporal logic. While the present work follows the definition of the epistemic modalities, including common knowledge, introduced in (Lomuscio and Michaliszyn 2013), it differs from it in the novel labelling function here introduced. This leads to an altogether different notion of models, as we describe below. An open question on the decidability of the model checking problem for a particularly expressive fragment of EHS was solved in (Lomuscio and Michaliszyn 2014). But the semantics adopted there is the same as that in (Lomuscio and Michaliszyn 2013); in contrast, the present work considerably extends the expressiveness of the formalism.

The model checking problem for an interval temporal logic was also studied in (Montanari et al. 2014); however, no knowledge specifications were considered in this work, and the labelling is less expressive than the one here studied even when limited to the temporal modalities.

The present contribution is also related to the very first approach to Interval Temporal Logic (Moszkowski 1983), where regular expressions can be used in the context of any subformula of the language. In contrast, in EHS^{RE} , the second logic we here propose, we permit regular expressions

to appear as propositions only. However, Interval Temporal Logic expresses properties of a single interval, whereas we here study the properties of different branches and focus on the low complexity of some of the fragments, such as the BDE one.

Two further formalisms related to the work here pursued are PDL (Harel, Tiuryn, and Kozen 2000) and its linear counterpart LDL (De Giacomo and Vardi 2013). An epistemic version of PDL, E-PDL, was proposed in (van Benthem, van Eijck, and Kooi 2006). However, epistemic modalities in E-PDL are interpreted on points, not intervals as in our work. This is largely the reason why the logic we study here is more expressive than E-PDL and the model checking problem for E-PDL is decidable in polynomial time (Lange 2006), whereas the model checking problem for EIT, a simple fragment of EHS, is already PSPACE-hard. Two further differences are that E-PDL cannot be used to reason about the past, but can be used to reason about actions explicitly.

The correspondence between regular expressions and HS was studied in (Montanari and Sala 2013), where it was shown that each ω -regular language can be encoded in the $AB\bar{B}$ fragment of HS, interpreted over the naturals. This result, however, is limited to the satisfiability problem, and cannot be used for the model checking problem.

The rest of this paper is organised as follows. We begin by defining a novel class of interpreted systems, called interpreted systems with regular labellings, that we use to model multi-agent systems. We then define the logic EHS^+ , whose syntax is the same as EHS, but whose semantics is interpreted over the proposed models over regular expressions. We then analyse the model checking problem of EHS^+ and show that it shares with EHS all the positive results known for it. We continue by investigating the expressive power EHS^+ . To do so, in order to be able to express properties of standard point-based models, we define and study the logic EHS^{RE} . In EHS^{RE} regular expressions can appear within the atomic propositions rather than just in the labelling function. We show polynomial time reductions between the model checking problems for EHS^{RE} and EHS^+ and characterise the expressive power of the former.

2 Interpreted Systems with Regular Labelling

Given a finite set X , the set of regular expressions over X , denoted by RE_X , is defined by the following BNF expression: $e ::= \emptyset \mid \epsilon \mid s \mid e \cdot e \mid e + e \mid e^*$, where $s \in X$. We allow parentheses for grouping and often omit the concatenation symbol “.”. Let $\text{Lang}(e)$ stand for the language denoted by a regular expression e , defined in the usual way.

We first introduce the semantics of interpreted systems with labellings on regular expressions by generalising the interval-based interpreted systems from (Lomuscio and Michaliszyn 2013).

Definition 1. *Given a set of agents $A = \{0, 1, \dots, m\}$, an interpreted system with labelling on regular expressions, ISRL for short, is a tuple $IS = (\{L_i, l_i^0, ACT_i, P_i, t_i\}_{i \in A}, \lambda)$, where for each $i \in A$:*

- 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 ,
- $P_i : L_i \rightarrow 2^{ACT_i}$ is a local protocol function for agent i , returning the set of possible actions in a given local state,
- $t_i \subseteq L_i \times ACT \times L_i$, where $ACT = ACT_0 \times \dots \times ACT_m$, is a local transition relation for agent i returning the next local state when a joint action is performed by all agents on a given local state.

Furthermore, $\lambda : Var \rightarrow REG$ is a labelling function, where $G = L_0 \times L_1 \times \dots \times L_m$ is the set of global configurations and Var is a finite set of propositional variables.

We assume that agent 0 is the *environment* for the system.

We now define models of an ISRL on sets of paths from its initial configuration. Let $t^G \subseteq G^2$ be a relation such that $t^G((l_0, \dots, l_m), (l'_0, \dots, l'_m))$ iff there exists a joint action $(a_0, \dots, a_m) \in ACT$ such that for all i we have $a_i \in P_i(l_i)$ and $t_i(l_i, (a_0, \dots, a_m), l'_i)$.

Definition 2. Given an ISRL $IS = (\{L_i, l_i^0, ACT_i, P_i, t_i\}_{i \in A}, \lambda)$ over a set of agents $A = \{0, \dots, m\}$, the model of IS is a tuple $M = (S, s_0, t, \{\sim_i\}_{i \in A}, \lambda)$, where

- $S \subseteq G^+$ is the set of global states, i.e., non-empty sequences $g_0 \dots g_k$ such that $g_0 = (l_0^0, \dots, l_m^0)$ and for each $i < k$ we have $t^G(g_i, g_{i+1})$,
- $s_0 = g_0 = (l_0^0, \dots, l_m^0)$ is the initial state of the system,
- $t \subseteq S^2$ is the global transition relation such that $t(g_0 \dots g_k, g'_0 \dots g'_l)$ iff $l = k + 1$ and for all $i \leq k$ we have $g_i = g'_i$,
- $\sim_i \subseteq S^2$ is the epistemic equivalence relation for agent i such that $g_0 \dots g_k \sim_i g'_0 \dots g'_l$ iff $g_k = (l_0, \dots, l_m)$, $g'_l = (l'_0, \dots, l'_m)$ and $l_i = l'_i$, and
- λ is the labelling function.

Intuitively, S denotes the set of global configurations of the ISRL equipped with information about all their predecessors. This is the standard construction used for defining unravelling in temporal logic (see, e.g., Definition 4.51 in (Blackburn, de Rijke, and Venema 2001)). Information about the predecessors is kept to interpret backward modalities. The epistemic relations for states are defined on the basis of local equality of the corresponding local states; we extend this notion in the next section to deal with intervals.

Given a model M , an *interval* in M is a finite path on M , i.e., a sequence of states $I = s_1, s_2, \dots, s_n$ such that $t(s_i, s_{i+1})$, for $1 \leq i \leq n - 1$. A *point interval* is an interval that consists of exactly one state. We assume $pi(I) = \top$ if I is a point interval and $pi(I) = \perp$ otherwise. By $fst(I)$ and $lst(I)$ we denote the first and the last state of I .

For a state of $s = g_0, \dots, g_k \in S$, we assume $G(s) = g_k$. So $G(s)$ denotes the current global configuration of s , not its history. We extend G to intervals by assuming $G(I) = G(s_0) \dots G(s_k)$ for every interval $I = s_0, \dots, s_k$. For $g = (l_0, l_1, \dots, l_m)$, by $l_i(g)$ we denote the local state $l_i \in L_i$ of agent $i \in A$ in g . For a global state $s = g_0, \dots, g_k$, we assume $l_i(s) = l_i(g_k)$.

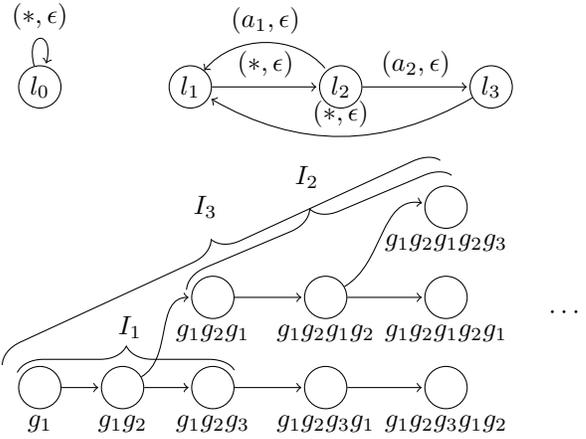


Figure 1: The agents from Example 3 (top; * stands for any action) and a fragment of the model of IS_{ex} (bottom). I_1, I_2 and I_3 are labelled by p , as $G(I_1) = G(I_2) = g_1g_2g_3$ and $G(I_3) = g_1g_2g_1g_2g_3$ belong to $\text{Lang}(\lambda(p))$.

Now we give an example of an interpreted system and of its model. We will use this example in the following sections to illustrate other constructions.

Example 3. Consider a set of agents $A = \{0, 1\}$, an ISRL $IS_{ex} = (\{L_i, l_i^0, ACT_i, P_i, t_i\}_{i \in A}, \lambda)$ and a propositional variable p , where

$$\begin{aligned} L_0 &= \{l_0\}, L_1 = \{l_1, l_2, l_3\}, l_0^0 = l_0, l_1^0 = l_1, \\ ACT_0 &= \{a_1, a_2\}, ACT_1 = \{\epsilon\}, \\ P_0(l_0) &= ACT_0, P_1(l_1) = P_1(l_2) = P_1(l_3) = ACT_1, \\ t_0 &= \{(l_0, (a_1, \epsilon), l_0), (l_0, (a_2, \epsilon), l_0)\}, \\ t_1 &= \{(l_1, (a_1, \epsilon), l_2), (l_1, (a_2, \epsilon), l_2), (l_2, (a_2, \epsilon), l_3), \\ &\quad (l_2, (a_1, \epsilon), l_1), (l_3, (a_1, \epsilon), l_1), (l_3, (a_2, \epsilon), l_1)\}, \\ \lambda(p) &= g_1(g_1 + g_2)^*g_3 \text{ where } g_i = (l_0, l_i). \end{aligned}$$

Figure 1 depicts the agents of IS . We have $G = \{g_1, g_2, g_3\}$ and $t^G = \{((l_0, l_1), (l_0, l_2)), ((l_0, l_2), (l_0, l_3)), ((l_0, l_2), (l_0, l_1)), ((l_0, l_3), (l_0, l_1))\}$. The model M_{ex} of IS_{ex} is infinite. A fragment is depicted in Figure 1.

3 The Logic EHS⁺

We now define the syntax of the specification language we focus on in this paper. We use temporal operators to represent relations between intervals as defined in (Halpern and Shoham 1991). Six of these relations are presented in Figure 2: R_A (“after” or “meets”), R_B (“begins” or “starts”), R_D (“during”), R_E (“ends”), R_L (“later”), and R_O (“overlaps”). Six additional operators can be defined corresponding to the six inverse relations. Formally, for each $X \in \{A, B, D, E, L, O\}$, we also consider the relation R_X , corresponding to R_X^{-1} .

For convenience, we also consider the “next” relation R_N such that $IR_N I'$ iff $t(lst(I), fst(I'))$ (Lomuscio and Michaliszyn 2014). Let $\mathbb{HS} = \{A, \bar{A}, B, \bar{B}, D, \bar{D}, E, \bar{E}, L, \bar{L}, N, \bar{N}, O, \bar{O}\}$.

Definition 4. The syntax of the Epistemic Halpern–Shoham

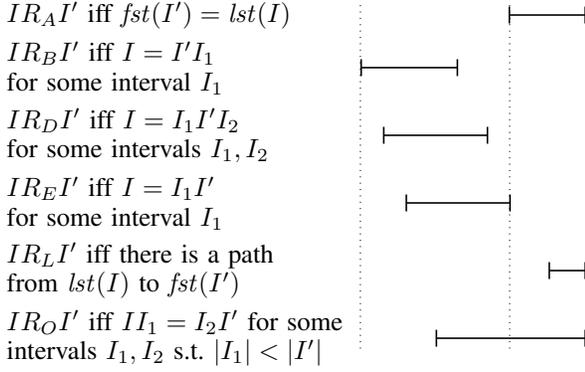


Figure 2: Basic Allen relations.

Logic (EHS^+), \mathcal{L}_{EHS^+} is defined by the following BNF.

$$\varphi ::= pi \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid C_\Gamma\varphi \mid \langle X \rangle\varphi$$

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

We define that $s_1, \dots, s_k \sim_i s'_1, \dots, s'_l$, read as the two intervals are *epistemically indistinguishable for i* , if $k = l$ and for all $j \leq k$ we have $s_j \sim_i s'_j$. In other words, for two intervals to be indistinguishable to agent i the two intervals need to be of the same length and the agent cannot distinguish any corresponding point in the interval. This is a generalisation to intervals of the point-based knowledge modalities traditionally used in epistemic logic (Fagin et al. 1995). For example, in the model presented in Example 3, we have $I \sim_0 I'$ if and only if $|I| = |I'|$ and $I \sim_1 I'$ if and only if $G(I) = G(I')$; in general these relations may be more complicated. We extend this definition to the common knowledge case by considering $\sim_\Gamma = (\bigcup_{i \in \Gamma} \sim_i)^*$, for any group of agents $\Gamma \subseteq A$, where $*$ denotes the transitive closure. For further explanations we refer to (Lomuscio and Michaliszyn 2013).

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

Definition 5 (Satisfaction). *Given an EHS^+ formula φ , an ISRL IS , its model $M = (S, s_0, t, \{\sim_i\}_{i \in A}, \lambda)$, and an interval I , we inductively define whether φ holds in the interval I , denoted $M, I \models \varphi$, as follows:*

- (i) $M, I \models pi$ iff I is a point interval,
- (ii) $M, I \models p$ iff $G(I) \in \text{Lang}(\lambda(p))$,
- (iii) $M, I \models \neg\varphi$ iff it is not the case that $M, I \models \varphi$,
- (iv) $M, I \models \varphi_1 \wedge \varphi_2$ iff $M, I \models \varphi_1$ and $M, I \models \varphi_2$,
- (v) $M, I \models K_i\varphi$, where $i \in A$, iff for all $I' \sim_i I$ we have $M, I' \models \varphi$,
- (vi) $M, I \models C_\Gamma\varphi$, where $\Gamma \subseteq A$, iff for all $I' \sim_\Gamma I$ we have $M, I' \models \varphi$,
- (vii) $M, I \models \langle X \rangle\varphi$ iff there is an interval I' such that $IR_X I'$ and $M, I' \models \varphi$, where R_X is an Allen relation as above.

We write $IS, I \models \varphi$ if $M, I \models \varphi$, where M is the model of IS , and $IS \models \varphi$ if $IS, s_0 \models \varphi$.

We use standard abbreviations, including $[X]\varphi$ for $\neg\langle X \rangle\neg\varphi$ and the usual Boolean connectives $\vee, \Rightarrow, \Leftrightarrow$ as well as the constants \top, \perp in the standard way.

Note that the modality $\langle N \rangle$ is a counterpart of the EX operator of CTL. While $\langle N \rangle$ is redundant in EHS^+ since $\langle N \rangle\varphi = \langle A \rangle(\neg pi \wedge [B][B]\perp \wedge \langle A \rangle\varphi)$, it is useful in fragments of EHS^+ that do not contain B and E .

4 An Interval-based Analysis of the Bit Transmission Protocol

We exemplify the expressive power of EHS^+ by extending the bit transmission protocol (BTP), a well-known communication scenario that is often analysed by means of temporal-epistemic specifications (Fagin et al. 1995). The BTP models a scenario where an agent S (“Sender”) attempts to communicate with an agent R (“Receiver”) over a faulty channel, which may drop messages but may not flip them. A version of the BTP was discussed in context of interval temporal logic in (Lomuscio and Michaliszyn 2014). In that variant the sender computes for some time the message to send before initiating communication. The modelling presented in (Lomuscio and Michaliszyn 2014) is suited for sending one bit and can be adapted for sending any fixed number of bits. We here generalise the scenario to allow for an *unbounded number of bits* to be sent.

We assume that the string of bits sent by S includes an error detecting code (EDC). We assume that the EDC can be computed in constant memory, and that it consists of two bits at the end of the message, representing whether the number of 0s sent, respectively 1s, is odd.

We stipulate that S sends bits one by one, and keeps sending a bit until he gets an acknowledgement; when he does, he either ends the communication or sends the next bit. To distinguish consecutive bits, S adds a parity bit to the message. R remains silent until he receives a bit, then he keeps acknowledging the bit until he receives the other one (which he distinguishes by the parity bit).

We model the revised BTP as an ISRL IS as follows. S 's local states are of the form $(status, bit, \mathcal{P})$, where $status \in \{cmp, snd, acked\}$, $bit \in \{0, 1, -\}$, and $\mathcal{P} \in \{0, 1\}$. We take S 's initial local state to be $(cmp, -, 0)$. The actions for S are $ACT_S = \{send_0^0, send_0^1, send_1^0, send_1^1, \epsilon\}$. The set of local states of R is $L_R = \{-, bit_0^{\mathcal{P}}, bit_1^{\mathcal{P}} \mid \mathcal{P} \in \{0, 1\}\}$, where we assume that R only remembers the last bit received. R 's actions are $ACT_R = \{\epsilon, sendack^0, sendack^1\}$ where ϵ is the null action. The environment E has a single local state and four actions $ACT_E = \{\rightarrow, \leftarrow, \leftrightarrow, \epsilon\}$, representing, respectively, messages being delivered from S to R , from R to S , in both directions, and in no direction.

The transition relation t_S for S is such that S may either loop in the state $(cmp, -, \mathcal{P})$, where $\mathcal{P} \in \{0, 1\}$, or move to (snd, b, \mathcal{P}) for some $b \in \{0, 1\}$. From this state S starts sending the bit b by means of the action $send_b^{\mathcal{P}}$. S remains in one of these states until he receives an acknowledgement from R , triggered by either the joint actions $(send_b, sendack^{\mathcal{P}}, \leftarrow)$ or $(send_b, sendack^{\mathcal{P}}, \leftrightarrow)$. From that point onward S moves to the local state $(acked, b, \mathcal{P})$. S may loop on this state for the rest of the run or non-

deterministically jump to $(cmp, -, 1 - \mathcal{P})$, which encodes the computation and the sending of another bit. The transitions for R can similarly be formalised. The relation t_R includes a loop on the initial state $-$, where R performs the action ϵ . From this state R makes a transition to the state bit_b^0 following the joint actions $(send_b^0, \epsilon, \rightarrow)$ and $(send_b^0, \epsilon, \leftrightarrow)$. In a state $bit_b^{\mathcal{P}}$, R uses the action $sendack^{\mathcal{P}}$ and remains in this state unless the action is $(send_b^{1-\mathcal{P}}, \epsilon, \rightarrow)$ or $(send_b^{1-\mathcal{P}}, \epsilon, \leftrightarrow)$, in which case R moves to $bit_b^{1-\mathcal{P}}$. The protocols are defined accordingly.

We consider a labelling function λ such that $\lambda(snd)$ defines intervals not containing any acknowledgements, starting with cmp and ending with snd ; $\lambda(cmp_b)$ defines intervals in which S computes the same bit of the message, and for each $b \in \{0, 1\}$, $\lambda(bit_b^R)$ defines intervals in which all the local states of R are bit_b^0 or bit_b^1 ; finally, $\lambda(correctEDC)$ defines intervals whose starting point is S 's initial local state, and whose ending point is of the form $(acked, b, \mathcal{P})$, and in which the message sent by the sender has the correct EDC.

We are interested in evaluating the following specification: 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 there is a successive interval where S knows that R knows the value of the bit. This is expressed by the EHS^+ formula:

$$\bigwedge_{b \in \{0,1\}} [G](cmp_b \Rightarrow \langle \bar{B} \rangle (snd \wedge \langle A \rangle K_S K_R bit_b^R))$$

where $[G]$ is an operator such that $[G]\varphi$ holds if φ holds in all the reachable intervals (this can be easily defined in EHS^+). It can be checked that, following our intuition, the property holds in the model of IS . Note that this specification is not expressible in EHS , in which the labelling depends only on the endpoints of the intervals.

A further specification of interest is whether over any interval starting in the initial state and ending in a state when the status of the sender is $acked$, if the EDC sent over this interval for this message was correct, then the sender knows that the receiver knows this. This is captured by the formula:

$$\langle \bar{B} \rangle (correctEDC \Leftrightarrow K_S K_R correctEDC)$$

It can be checked that the formula holds in the point interval consisting of the initial state.

These specifications cannot be expressed in EHS , $CTLK$, or even E -PDL (see discussion below). The property does hold in the model here studied as, intuitively, over the interval in question the sender knows all the bits of the message received by the receiver. Notice that this does not imply that the receiver really computes the EDC, but rather that the receiver has enough information to compute it.

5 Expressive Power

To investigate the expressive power of EHS^+ , we now introduce EHS^{RE} , a variant of EHS^+ defined over point-based interpreted systems, defined as follows.

Definition 6. An ISRL is point-based if λ only labels the point intervals, i.e., for each $v \in Var$ we have $\lambda(v) =$

$\sum_{g \in G'} g$ for some $G' \subseteq G$. An ISRL is endpoint-based if λ is defined on the endpoints of the intervals, i.e., for each $v \in Var$ we have $\lambda(v) = \sum_{g \in G'} (g + gG^*g) + \sum_{(g,g') \in P} gG^*g'$ for some $G' \subseteq G$, $P \subseteq G^2 \setminus \{(g, g) \mid g \in G\}$.

In the above, $g + gG^*g$ is a regular expression that denotes all the intervals whose current global configuration of both endpoints is the global state g , whereas gG^*g' denotes intervals starting at g and ending in g' . The models of the point-based ISRL can be seen as standard Kripke structures; the models of the endpoint-based ISRL are the generalised Kripke structures of (Lomuscio and Michaliszyn 2013).

We show that the model checking problems for EHS^{RE} and EHS^+ admit a polynomial time reduction to one another on the corresponding semantics. We also observe that EHS^{RE} can represent properties not expressible by $CTLK^*$, the epistemic version of CTL^* (and therefore $LTLK$ and $CTLK$).

For a labelling function λ and a regular expression r , let $\lambda \circ r$ be the regular expression obtained from r by replacing each propositional variable p by $\sum_{g \in \text{Lang}(\lambda(p))} g$. Notice that in point-based systems $\text{Lang}(\lambda(p))$ consists of single global configurations only.

Definition 7. The language of EHS^{RE} , $\mathcal{L}_{EHS^{RE}}$, is defined as follows:

$$\varphi ::= pi \mid r \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid C_\Gamma\varphi \mid \langle X \rangle\varphi$$

where $r \in RE_{2var}$, $i \in A$, $\Gamma \subseteq A$, and $X \in \text{HS}$.

The semantics of EHS^{RE} results from replacing the second rule in Definition 5 by (ii') $M, I \models r$ iff $G(I) \in \text{Lang}(\lambda \circ r)$.

Notice that in the above we have $r \in RE_{2var}$ rather than $r \in RE_{var}$. This is because we may want to express properties such that each point of interval is labelled by p and not q , which would not be possible with the latter definition as we could only state one variable at a time. For convenience, we allow to use p and $\neg p$ in the regular expressions, by defining $p = \sum_{X \subseteq Var, p \in X} X$ and $\neg p = \sum_{X \subseteq Var, p \notin X} X$.

Intuitively, EHS^{RE} is the result of adapting EHS^+ by moving the regular expressions from the labelling function into the language.

Let \mathbb{L}_{var} be the set of all the possible labellings of interpreted systems with variables of Var , and $\mathbb{L}_{var}^{pi} \subseteq \mathbb{L}_{var}$ be the set of all such labellings for point-based interpreted systems.

Theorem 8. There exist polynomial time computable functions $f : \mathbb{L}_{var} \times \mathcal{L}_{EHS^+} \rightarrow \mathbb{L}_{var}^{pi} \times \mathcal{L}_{EHS^{RE}}$ and $f' : \mathbb{L}_{var}^{pi} \times \mathcal{L}_{EHS^{RE}} \rightarrow \mathbb{L}_{var} \times \mathcal{L}_{EHS^+}$ such that:

1. If $IS, I \models \varphi$, then $IS', I \models \varphi'$ for any ISRL $IS = (\{Ag_i\}_{i \in A}, L)$, $IS' = (\{Ag_i\}_{i \in A}, L')$ and φ, φ' such that $f(L, \varphi) = (L', \varphi')$.
2. If $IS, I \models \varphi$, then $(\{L_i, l_i^0, ACT_i, P_i, t_i\}_{i \in A}, L'), I \models \varphi'$ for any point based ISRL $IS = (\{Ag_i\}_{i \in A}, L)$, $IS' = (\{Ag_i\}_{i \in A}, L')$ and φ, φ' such that $f'(L, \varphi) = (L', \varphi')$.

Proof sketch. Intuitively, the functions f and f' replace the regular expressions from the labelling to the formula and the other way round. The function f is such that $f(\lambda, \varphi) =$

(λ', φ') , where $\lambda'(g) = g$ for all the states s and φ' is the result of replacing each propositional variable q in φ by $\sum_{g \in \lambda(q)} g$. The function f' is such that $f'(\lambda', \varphi') = (\lambda, \varphi)$, where for each regular expression r in φ' , we replace r with a unique propositional variable q^r and we let $\lambda(q^r) = \lambda' \circ r$. It can easily be checked that both functions are as required. \square

Given Theorem 8, we can say that EHS^+ and EHS^{RE} can describe the same properties of corresponding interpreted systems. Since EHS^{RE} expresses properties of point-based interpreted systems, whose models are standard Kripke structures, we can formally compare the expressive power of EHS^{RE} to that of some more widely known formalisms.

Definition 9. *Given two logics $\mathcal{L}_1, \mathcal{L}_2$ interpreted over point-based ISRL, we write $\mathcal{L}_1 \subseteq \mathcal{L}_2$ if for each formula φ_1 of \mathcal{L}_1 there is a formula φ_2 of \mathcal{L}_2 such that for every point-based ISRL IS we have $IS \models \varphi_1$ iff $IS \models \varphi_2$.*

One can easily show that $\text{EHS}^{\text{RE}} \not\subseteq \text{CTLK}^*$. Consider the temporal property “all the paths starting in the initial state satisfy $(p; \text{True})^\omega$ ”. This property cannot be expressed in CTLK^* (Wolper 1983). However, the property can be verified by evaluating the EHS^{RE} formula $p \wedge [A]((p; \top)^* \Rightarrow [N](p; \top^*))$.

Also observe that the property above cannot be expressed in the logic EHS considered over point-based ISRL either. So over point-based ISRL we have that $\text{EHS}^{\text{RE}} \not\subseteq \text{EHS}$

In terms of limitations, EHS^{RE} can only express properties of finite intervals. For example, the CTL property AFp expressing the fact that each infinite path satisfies p at some point cannot be encoded by any EHS^{RE} formula. Therefore $\text{CTLK} \not\subseteq \text{EHS}^{\text{RE}}$; similarly we have $\text{LTLK} \not\subseteq \text{EHS}^{\text{RE}}$.

Since EHS^{RE} does not allow us to name actions explicitly, we have that E-PDL $\not\subseteq \text{EHS}^{\text{RE}}$. It can also be shown that $\text{EHS}^{\text{RE}} \not\subseteq \text{E-PDL}$, since E-PDL cannot express the property $\langle A \rangle (K_1(pq^*r))$ as the epistemic modalities in E-PDL is based on states rather than time-intervals.

6 The Model Checking Problem

We now investigate the complexity of the model checking problem for fragments of the logics explored so far.

Definition 10. *Given a formula φ of a logic L , an ISRL IS and an interval I , the model checking problem for L amounts to checking whether or not $IS, I \models \varphi$.*

By establishing the above, we say we have model checked the model M against the specification φ at an interval I . Notice that the formula is verified only at the given interval; however, one can easily check whether *all* the initial intervals satisfy a formula φ by checking whether $M, s_0 \models [A]\varphi$.

The $\overline{ABL}N$ fragment of EHS^+ , denoted as $\text{EHS}_{\overline{ABL}N}^+$, is the subset of EHS^+ where the BNF is restricted to the modalities $K_i, C_\Gamma, \langle A \rangle, \langle \overline{B} \rangle, \langle L \rangle$, and $\langle N \rangle$ only. Similarly, the BDE fragment of EHS^+ , denoted as EHS_{BDE}^+ , is the restriction of EHS^+ to $K_i, C_\Gamma, \langle B \rangle, \langle D \rangle$ and $\langle E \rangle$.

Theorem 11. *Model checking ISRLs against EHS_{BDE}^+ specifications is decidable and PSPACE-complete.*

Proof. The lower bound follows from the lower bound for the endpoint-based variant of ISRL that was shown in (Lomuscio and Michaliszyn 2013) for the same syntax. For the upper bound, we consider a polynomial time alternating algorithm that works recursively as follows.

For a given model M , interval I and a formula φ , if φ is a propositional variable, return whether $G(I) \in \text{Lang}(\lambda(p))$. If $\varphi = pi$, return $pi(I)$. If φ is a Boolean operator, then compute recursively the values of the subformulas and return the result of applying the operator to the computed values. If φ is an epistemic formula $K_i\varphi'$ where $i \in A$ (resp. $C_\Gamma\varphi'$ where $\Gamma \subseteq A$), then universally select J such that $J \sim_i I$ (resp. $J \sim_\Gamma I$) and return the value of a recursive call for M, J, φ' . If φ is of the form $\langle X \rangle \varphi'$ where $X \in \{\langle B \rangle, \langle D \rangle, \langle E \rangle\}$, then existentially select J such that $IR_X J$ and return the value of a recursive call for M, J, φ' .

The complexity follows from the fact that each existentially or universally selected interval has the size bounded by the size of the initial interval. Since $\text{APTIME} = \text{PSPACE}$, the theorem follows. \square

Theorem 12. *Model checking ISRLs against $\text{EHS}_{\overline{ABL}N}^+$ specifications is decidable in non-elementary time.*

We prove this by generalising the proof of Theorem 13 given in (Lomuscio and Michaliszyn 2014). To do so, below we introduce a bounded semantics and link it to the unbounded one. This will enable us to give the proof of the theorem at the end of this section.

A *top-level* sub-formula of a formula φ is a sub-formula of φ of the form $X\varphi'$, for some modality X of $\text{EHS}_{\overline{ABL}N}^+$, that is not in the scope of any modality. Assume an ISRL IS . Let $f^{IS}(\varphi)$ be defined recursively as

$$f^{IS}(\varphi) = (2|G|^2 \prod_{q \in \text{Var}} 2^{|\lambda(q)|}) \cdot 2^{f^{IS}(\varphi_1)} \dots \cdot 2^{f^{IS}(\varphi_k)}$$

where $X_1\varphi_1, \dots, X_k\varphi_k$ are the top-level sub-formulas of φ . The idea is that $f^{IS}(\varphi)$ is an upper bound on the number of different *interval types* w.r.t. φ ; an interval type specifies whether an interval is a point interval or not (hence 2), what are its endpoints (hence $|G|^2$), what are the states of the automata corresponding to the regular expressions after reading the interval (hence the product) and the types of intervals related to the interval w.r.t. the top level sub-formulas of φ (hence the recursive part).

We define a bounded satisfaction relation \models_B for $\text{EHS}_{\overline{ABL}N}^+$, for which the decidability of the model checking is straightforward. The rules (i'-vi') of the definition of \models_B are the same as the rules (i-vi) from Definition 5 except that \models is replaced with \models_B . The last rule, however, is different:

- (vii') $M, I \models_B \langle X \rangle \varphi$ if and only if there exists an interval I' such that $|I'| \leq |I| + f^{IS}(\varphi)$, $IR_X I'$ and $M, I' \models_B \varphi$, where X is A, \overline{B}, L , or N .

It is not hard to see that model checking is decidable for the bounded semantics. It turns out that, in the $\text{EHS}_{\overline{ABL}N}^+$ case, the relations \models and \models_B are the same, and therefore the model checking procedure for the bounded semantics solves

Algorithm 1 The model checking procedure for $EHS_{ABL N}^+$.

```

1: procedure VERIFY( $M, I, \varphi$ )
2:   if  $\varphi = p$  then return  $I \in \text{Lang}(\lambda(p))$ 
3:   if  $\varphi = pi$  then return  $pi(I)$ 
4:   if  $\varphi = \neg\varphi'$  then return  $\neg\text{VERIFY}(M, I, \varphi')$ 
5:   if  $\varphi = \varphi_1 \wedge \varphi_2$  then
6:     return  $\text{VERIFY}(M, I, \varphi_1) \wedge \text{VERIFY}(M, I, \varphi_2)$ 
7:   if  $\varphi = E\varphi'$  where  $E$  is  $K_i$  or  $C_\Gamma$  then
8:     for all  $J$  s.t.  $IR_E J$  do
9:       if  $\neg\text{VERIFY}(M, J, \varphi')$  then return false
10:    return true
11:   if  $\varphi = X\varphi'$  where  $X \in \{\langle A \rangle, \langle \bar{B} \rangle\}$  then
12:     for all  $J$  s.t.  $IR_X J$  and  $|J| \leq f(\varphi) + |I|$  do
13:       if  $\text{VERIFY}(M, J, \varphi')$  then return true
14:     return false

```

the model checking problem for the unbounded semantics. The details follow.

Observe that $\langle L \rangle$ can be defined in terms of $\langle A \rangle$: 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 $\langle L \rangle$. We now define some auxiliary notions.

For convenience, for each modality X of $EHS_{ABL N}^+$, 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 $R_{C_\Gamma} = \sim_\Gamma$.

Theorem 13. *Model checking ISRL under bounded semantics against $EHS_{ABL N}^+$ specifications is decidable.*

Proof. The procedure VERIFY given in Algorithm 1 solves the model checking problem. Clearly, it always terminates and its computation time is non-elementary. \square

The key result below links bounded to unbounded semantics.

Theorem 14. *Given an $EHS_{ABL N}^+$ formula φ , a model M , and an interval I , $M, I \models \varphi$ if and only if $M, I \models_B \varphi$.*

Proof. Consider a model $M = (S, s_0, t, \{\sim_i\}_{i \in A}, \lambda)$. For each $p \in \text{Var}$ we denote by \mathcal{A}^p the minimal deterministic finite state automaton (Hopcroft and Ullman 1979) recognising the language $\text{Lang}(\lambda(p))$. By $\mathcal{A}_w(p)$, where $p \in \text{Var}$, we denote the state of \mathcal{A}^p after reading a word w ; in the following, we treat \mathcal{A}_w as a function from Var to automata states.

Definition 15 (Modal Context Tree). *Given a model M , the modal context tree of an interval I w.r.t. an $EHS_{ABL N}^+$ formula φ , denoted by MCT_I^φ , is the minimal unranked tree with labelled nodes and edges defined recursively as follows.*

- The root of the tree is labelled by the tuple $G(\text{fst}(I))$, $G(\text{lst}(I))$, $pi(I)$, A_I .
- For each top-level sub-formula $X\psi$ of φ and each interval I' such that $IR_X I'$, the root of $\text{MCT}_{I'}^\psi$ has an $X\psi$ -successor $\text{MCT}_{I'}^\psi$ (X indicates the labelling of an edge).

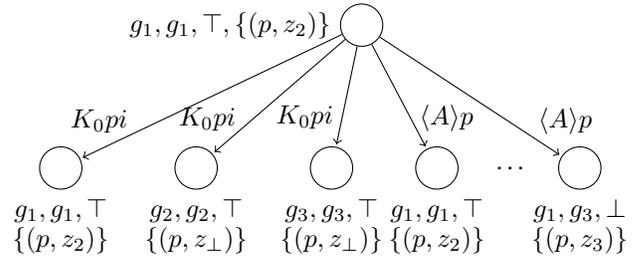
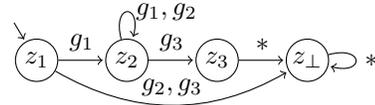


Figure 3: MCT_I^φ from Example 16. The omitted $\langle A \rangle p$ successors are labelled by: $g_1, g_2, \perp, \{(p, z_2)\}$; $g_1, g_1, \perp, \{(p, z_2)\}$; $g_1, g_1, \perp, \{(p, z_\perp)\}$; $g_1, g_2, \perp, \{(p, z_\perp)\}$; $g_1, g_2, \perp, \{(p, z_\perp)\}$.

In other words MCT_I^φ contains sufficient information about all the intervals that need to be considered to determine the value of φ in I as well as the states of the automata after reading I .

Example 16. *Consider the ISRL IS_{ex} from Example 3, the formula $\varphi = K_0 pi \wedge \neg \langle A \rangle p$, and an interval $I = g_1$. To build the modal context tree, we use the following automaton for $\lambda(p) = g_1(g_1 + g_2)^*g_3$. The only accepting state is z_3 .*



The top level sub-formulas of φ are $K_1 pi$ and $\langle A \rangle p$. MCT_I^φ (Figure 3) represents I . Notice that there are infinitely many R_A successors of I , but MCT_I^φ needs only 7 $\langle A \rangle p$ -successors. For example, the successor labelled by $g_1, g_2, \perp, \{(p, z_2)\}$ represents all the intervals I such that $G(I)$ is of the form $g_1(g_1 + g_2)^*$.

We now show that the number of modal context trees for a given formula is bounded. We use this later as a kind of pumping argument to show that if an interval is long enough, then some of its prefixes have the same modal context tree.

Lemma 17. *Given a model M and a formula φ , $|\{\text{MCT}_I^\varphi \mid I \text{ is an interval in } M\}| < f^{IS}(\varphi)$.*

Proof. We show the lemma by induction on φ . If a formula has no modalities, then $\{\text{MCT}_I^\varphi \mid I \text{ is an interval in } M\}$ contains trees with only one node, that can be labelled with $2|G|^2 \prod_{q \in \text{Var}} 2^{|\lambda(q)|}$ different labels.

Consider a formula φ with the top-level sub-formulas $X_1\varphi_1, \dots, X_k\varphi_k$. Each tree for φ consists of one of $2|G|^2 \prod_{q \in \text{Var}} 2^{|\lambda(q)|}$ possible roots and, for each i , any subset of subtrees for φ_i . Therefore, $|\{\text{MCT}_I^\varphi \mid I \text{ is an interval in } M\}| < f^{IS}(\varphi) = 2|G|^2 \prod_{q \in \text{Var}} 2^{|\lambda(q)|} 2^{f^{IS}(\varphi_1)} \dots 2^{f^{IS}(\varphi_k)}$. \square

We show that the modal context tree does not depend on the histories.

Lemma 18. Consider a model $M = (S, s_0, t, \{\sim_i\}_{i \in A}, \lambda)$ and a formula φ . If I and I' are intervals such that $G(I) = G(I')$, then $\text{MCT}_I^\varphi = \text{MCT}_{I'}^\varphi$.

Proof. We show this by induction.

The roots of MCT_I^φ and $\text{MCT}_{I'}^\varphi$ have the same labels, since $G(\text{fst}(I)) = G(\text{fst}(I'))$, $G(\text{lst}(I)) = G(\text{lst}(I'))$, $\text{pi}(I) = \text{pi}(I')$ and the labelling is defined on $G(I)$.

Consider a $\langle X \rangle \varphi'$ -successor T of the root of MCT_I^φ , where $\langle X \rangle \varphi'$ is a top-level sub-formula of φ and $X \in \{A, \bar{B}, N\}$.

There is an interval J such that $IR_X J$ and $\text{MCT}_I^\varphi = T$. So there exists a J' such that $I'R_X J'$ and $G(J) = G(J')$, because X is a “forward modality” so the R_X successors of I' do not depend on the history. By the inductive hypothesis, $\text{MCT}_J^{\varphi'} = \text{MCT}_{J'}^{\varphi'}$, and therefore the roots of MCT_I^φ and $\text{MCT}_{I'}^\varphi$ have the same $\langle X \rangle \varphi'$ successors.

As for the $X\varphi'$ successors where X is an epistemic modality, it is enough to observe that $IR_X I'$, and therefore I and I' are related to the same intervals by the equivalence relation R_X . The lemma follows. \square

We argue that if two intervals have the same modal context tree w.r.t. φ , then either both satisfy φ or none of them.

Lemma 19. Consider a model $M = (S, s_0, t, \{\sim_i\}_{i \in A}, \lambda)$ and a formula φ . If I and I' are intervals such that $\text{MCT}_I^\varphi = \text{MCT}_{I'}^\varphi$, then $M, I \models \varphi$ if and only if $M, I' \models \varphi$.

Proof. We show it by induction on φ .

- $\varphi = p$ for some variable p . The root of the MCT_I^φ is labelled by the state of an automaton corresponding to $\lambda(p)$ after reading I , and the root of the $\text{MCT}_{I'}^\varphi$ is labelled by the state of an automaton corresponding to $\lambda(p)$ after reading I' . Since the two trees are equal, the automaton is in the same state in both cases, either accepting or rejecting, and therefore $M, I \models p$ if and only if $M, I' \models p$.

- $\varphi = \text{pi}$. The root of the MCT_I^φ is labelled by $\text{pi}(I)$, and so is the root of $\text{MCT}_{I'}^\varphi$, and therefore $\text{pi}(I) = \text{pi}(I')$.

- $\varphi = \neg \varphi'$. By the inductive assumptions, $M, I \models \varphi'$ if and only if $M, I' \models \varphi'$, so $M, I \models \varphi$ if and only if $M, I' \models \varphi$.

- $\varphi = \varphi_1 \wedge \varphi_2$. By the induction assumption, $M, I \models \varphi_1$ if and only if $M, I' \models \varphi_1$ and $M, I \models \varphi_2$ if and only if $M, I' \models \varphi_2$, so $M, I \models \varphi$ if and only if $M, I' \models \varphi$.

- $\varphi = K_i \varphi'$. Assume that $M, I \models \varphi$. Consider any interval J' such that $I' \sim_i J'$. By definition, in the tree MCT_I^φ

the subtree $\text{MCT}_{J'}^{\varphi'}$ is a $K_i \varphi'$ -successor of the root. It follows that in the tree $\text{MCT}_I^\varphi (= \text{MCT}_{I'}^\varphi)$, $\text{MCT}_{J'}^{\varphi'}$ is a $K_i \varphi'$ -successor of the root. Let J be such that $I \sim_i J$ and $\text{MCT}_{J'}^{\varphi'} = \text{MCT}_J^{\varphi'}$. Clearly, since $M, I \models \varphi$, $M, J \models \varphi'$. By the inductive assumptions, $M, J' \models \varphi'$. Therefore $M, I' \models \varphi$.

- $\varphi = C_\Gamma \varphi'$. Assume that $M, I \models \varphi$ and J' is such that $I' \sim_\Gamma J'$. Again, in MCT_I^φ , the subtree $\text{MCT}_{J'}^{\varphi'}$ is a $C_\Gamma \varphi'$ -successor of the root. It follows that in the tree MCT_I^φ , $\text{MCT}_{J'}^{\varphi'}$ is a $C_\Gamma \varphi'$ -successor of the root. Let J be such that $I \sim_\Gamma J$ and $\text{MCT}_{J'}^{\varphi'} = \text{MCT}_J^{\varphi'}$, then $M, J \models \varphi'$, and by the inductive assumptions, $M, J' \models \varphi'$. Therefore $M, I' \models \varphi$.

- $\varphi = \langle A \rangle \varphi'$. We have $M, I \models \langle A \rangle \varphi'$ if and only if there is an interval J starting in $\text{lst}(I)$ satisfying φ' . Since

$G(\text{lst}(I)) = G(\text{lst}(I'))$, the intervals starting from $\text{lst}(I)$ and $\text{lst}(I')$ are the same (modulo histories), and therefore there exists an interval J' starting in $\text{lst}(I')$ such that $G(J) = G(J')$. By Lemma 18 we have $\text{MCT}_J^{\varphi'} = \text{MCT}_{J'}^{\varphi'}$.

- $\varphi = \langle \bar{B} \rangle \varphi'$. Assume that there is an interval J such that $IR_{\bar{B}} J$ and $M, J \models \varphi'$. Then, $\text{MCT}_J^{\varphi'}$ is an $\langle \bar{B} \rangle \varphi'$ successor of the root in MCT_I^φ , and so in $\text{MCT}_{I'}^\varphi$. So there is an interval J' such that $I'R_{\bar{B}} J'$ and $\text{MCT}_J^{\varphi'} = \text{MCT}_{J'}^{\varphi'}$. By the inductive hypothesis, $M, J' \models \varphi'$ and therefore $M, I' \models \varphi$.

- $\varphi = \langle N \rangle \varphi'$. Similarly to the case for $\langle A \rangle \varphi'$. \square

As we remarked earlier, if an interval I is long enough, then I has two prefixes with the same modal context tree w.r.t. a formula φ . Intuitively speaking, we would like to replace the longer prefix by the shorter one, thereby obtaining an interval I' , and show that the modal context trees of I and I' are the same. By the above lemma, it would follow that they both satisfy the given formula. What remains to be proved is that if we have two prefixes with the same modal context tree, and we append the same interval to both, the results will also have the same modal context tree.

We use the following terminology. A *partial state* is a sequence of states $g_1 \dots g_k$ such that for all $i < k$, we have $t^G(g_i, g_{i+1})$. Each state of the model is a partial state; but partial states are not required to start at g_0 . A *partial interval* is a sequence $s_1 \dots s_k$ of partial states such that for each $i < k$ we have that $s_{i+1} = s_i g_i$ for some partial state g_i . A partial interval $I = s_1 \dots s_k$ is *clear* if $s_1 = g$ for some partial state g . We extend the functions fst , lst , and g and the other notions to partial intervals in the obvious way.

We define the operation of adding context to partial intervals as follows. Given a partial interval I and a clear partial interval $I' = s_1 \dots s_k$ where $t^G(G(\text{lst}(I)), G(\text{fst}(I')))$, by $I \oplus I'$ we denote the partial interval $I \bar{s}_1 \dots \bar{s}_k$ such that for each i we have that $\bar{s}_i = \text{lst}(I) s_i$. So \oplus joins two intervals in a way that accounts for the history of the partial states. Clearly, $I \oplus I'$ is an interval if and only if I is an interval. We also define the operation \circ such that $I \circ I' = \bar{s}_1 \dots \bar{s}_k$, i.e., it only returns the adjusted partial states of I' .

Lemma 20. Consider a model M , a formula φ , two intervals I, I' , and a partial interval J . If $\text{MCT}_I^\varphi = \text{MCT}_{I'}^\varphi$, and $t^G(G(\text{lst}(I)), G(\text{fst}(J)))$, then $\text{MCT}_{I \oplus J}^\varphi = \text{MCT}_{I' \oplus J}^\varphi$.

Proof. Consider a formula φ , a model M , two intervals I, I' and a partial state $s = g$ such that $t^G(G(\text{lst}(I)), g)$. We show that $\text{MCT}_I^\varphi = \text{MCT}_{I'}^\varphi$ implies $\text{MCT}_{I \circ s}^\varphi = \text{MCT}_{I' \circ s}^\varphi$. This can be used to prove the lemma by induction.

Assume that the root of MCT_I^φ is labelled by $f, l, \text{pi}, \mathcal{A}_I$. Then the roots of both $\text{MCT}_{I \circ s}^\varphi$ and $\text{MCT}_{I' \circ s}^\varphi$ are labelled by f, g, \perp, \mathcal{A} , where for each $p \in \text{Var}$ we put $\mathcal{A}(p)$ equal to the state that the automaton for p reaches from $\mathcal{A}_I(p)$ after reading g . Assume that $X_1 \varphi_1, \dots, X_k \varphi_k$ are the top-level sub-formulas of φ and $i \in \{1, \dots, k\}$ (if there are no such formulas, then the result follows directly). We show that for each i , the roots of $\text{MCT}_{I \circ s}^\varphi$ and $\text{MCT}_{I' \circ s}^\varphi$ have the same $X_i \varphi_i$ -successors.

- X_i is an epistemic modality. Consider any interval J such that $I \oplus s R_{X_i} J$. Let $J = J' \oplus s'$. By the definition, $J' R_{X_i} I$

and $sR_{X_i} s'$. By the former, we have that $\text{MCT}_{J'}^{\varphi_i}$ is an $X_i \varphi_i$ -successor of the root in $\text{MCT}_{I \oplus s}^{\varphi}$, and so $\text{MCT}_{J'}^{\varphi_i}$ is an $X_i \varphi_i$ -successor of the root in $\text{MCT}_{I'}^{\varphi}$. So there is $J'' R_{X_i} I'$ such that $\text{MCT}_{J'}^{\varphi_i} = \text{MCT}_{J''}^{\varphi_i}$. Therefore, $J'' \oplus s' R_{X_i} I' \oplus s$, and thus $\text{MCT}_{J'}^{\varphi_i}$ is the $X_i \varphi_i$ -successors of the root of $\text{MCT}_{I \oplus s}^{\varphi}$.

- $X_i = \langle A \rangle$. Consider any interval J such that $I \oplus s R_A J$. Then there is a clear partial interval \bar{J} such that $J = I \circ \bar{J}$. Let $J' = I' \circ \bar{J}$. It holds that $I' \circ s R_A J'$. By Lemma 18, we have $\text{MCT}_J^{\varphi_i} = \text{MCT}_{J'}^{\varphi_i}$. Therefore, the $\langle A \rangle \varphi_i$ -successors of the root in $\text{MCT}_{I \oplus s}^{\varphi}$ are also $\langle A \rangle \varphi_i$ -successors of the root in $\text{MCT}_{I' \oplus s}^{\varphi}$. The other direction is similar.

- $X_i = \langle \bar{B} \rangle$. Consider any interval J such that $I \oplus s R_{\bar{B}} J$. Then, there is a clear partial interval \bar{J} such that $J = (I \oplus s) \oplus \bar{J}$. Let $J' = (I' \oplus s) \oplus \bar{J}$. It holds that $I' \oplus s R_{\bar{B}} J'$. By Lemma 18, we have $\text{MCT}_J^{\varphi_i} = \text{MCT}_{J'}^{\varphi_i}$. We conclude that the $\langle \bar{B} \rangle \varphi_i$ -successors of the root in $\text{MCT}_{I \oplus s}^{\varphi}$ are the same as $\langle \bar{B} \rangle \varphi_i$ -successors of the root in $\text{MCT}_{I' \oplus s}^{\varphi}$.

- $X_i = \langle N \rangle$. The proof is similar to the case of $\langle A \rangle$. \square

By exploiting the Lemma above, we can now give the proof of Theorem 14 by induction on the structure of φ . The cases for φ equal to $p, pi, \neg \varphi', \varphi_1 \wedge \varphi_2, K_i \varphi',$ and $C_{\Gamma} \varphi'$ for some sub-formulas $\varphi', \varphi_1, \varphi_2$, follow from the fact that the semantic rules are the same in both semantics.

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

Let $I' = s_1 \dots s_t$ and I'_k denote the prefix $s_1 \dots s_k$ of I' . Assume that $|I'| > |I| + f^{IS}(\varphi')$. By Lemma 17 there are two prefixes I'_k, I'_l such that $|I| < k < l$ and $\text{MCT}_{I'_k}^{\varphi'} = \text{MCT}_{I'_l}^{\varphi'}$. Let J be a clear partial interval such that $I' = I'_l \oplus J$. By Lemma 20, we have that $\text{MCT}_{I'_k \oplus J}^{\varphi'} = \text{MCT}_{I'_l \oplus J}^{\varphi'}$. Clearly, $|I'_k \oplus J| < |I'|$ and, by Lemma 19, $M, I'_k \oplus J \models \varphi'$. Since $k > |I|$, it follows that $IR_X I'_k \oplus J$ (the condition $k > |I|$ is only required for $\langle \bar{B} \rangle$ since J has to contain I as a prefix). But we assumed that I' was the shortest interval; so this is a contradiction. It follows that $|I'| \leq |I| + f^{IS}(\varphi)$. \square

We prove Theorem 12 as follows. By Theorem 14, the bounded semantics and the unbounded semantics are equivalent. By Theorem 13, model checking the $A\bar{B}LN$ fragment of EHS^+ with bounded semantics is decidable. Therefore, model checking the $A\bar{B}LN$ fragment of EHS^+ with unbounded semantics can be solved by Algorithm 1.

By employing the polynomial time reductions of Theorem 8, we can show that model checking point-based ISRL against BDE fragment of EHS^{RE} specifications is PSPACE-complete and that model checking point-based ISRL against $A\bar{B}LN$ fragment of EHS^{RE} specifications is decidable.

7 Conclusions and Future Work

Epistemic logic has traditionally been developed on underlying notions of time that are state-based. In this paper we have extended previous work (Lomuscio and Michaliszyn 2013; 2014) on an epistemic logic whose underlying temporal aspects are based on intervals. Specifically, we have put forward the logic EHS^+ which can express epistemic properties in the context of labellings possibly describing several, possibly overlapping stages.

We focused on the model checking aspects of these logics. We showed that the model checking for the BDE fragment of EHS^+ is decidable and PSPACE-complete, and that the model checking problem for the $A\bar{B}LN$ fragment of the logic is decidable. So, while the complexity of the problem for EHS^+ and EHS is the same, EHS^+ is more expressive.

Further ahead we intend to study more expressive fragments of EHS^+ . We believe that the technique presented here can be extended to backward modalities, such as $\langle \bar{A} \rangle, \langle \bar{D} \rangle, \langle \bar{E} \rangle, \langle \bar{L} \rangle$ and $\langle \bar{N} \rangle$. However, a deeper investigations are required, since in the case of backward modalities one cannot simply disregard the histories.

Finally, we are interested in implementing an efficient model checking toolkit for EHS^{RE} specifications. We intend to develop efficient algorithms on symbolic representations and a suitable predicate abstraction technique for EHS^{RE} .

Acknowledgments The authors would like to thank Angelo Montanari whose comments on (Lomuscio and Michaliszyn 2014) led to the present investigation.

This research was funded by the EPSRC under grant EP/I00520X. The second author acknowledges support from the Polish National Science Center, grant 2014/15/D/ST6/00719.

References

- Artale, A.; Kontchakov, R.; Ryzhikov, V.; and Zakharyashev, M. 2015. Tractable interval temporal propositional and description logics. In *Proc. of AAAI'2015*, 1417–1423.
- Belardinelli, F.; Lomuscio, A.; and Patrizi, F. 2012. An abstraction technique for the verification of artifact-centric systems. In *Proc. of KR'2012*, 319–328. AAAI Press.
- Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.
- Boureau, I.; Cohen, M.; and Lomuscio, A. 2009. A compilation method for the verification of temporal-epistemic properties of cryptographic protocols. *Journal of Applied Non-Classical Logics* 19(4):463–487.
- Bresolin, D.; Della Monica, D.; Goranko, V.; Montanari, A.; and Sciavicco, G. 2010. Metric propositional neighborhood logics: Expressiveness, decidability, and undecidability. In *Proc. of ECAI'2010*, 695–700.
- Bresolin, D.; Monica, D.; Goranko, V.; Montanari, A.; and Sciavicco, G. 2011a. The dark side of interval temporal logic: Sharpening the undecidability border. In *Proc. of TIME'2011*, 131–138. IEEE.

- Bresolin, D.; Montanari, A.; Sala, P.; and Sciavicco, G. 2011b. What's decidable about Halpern and Shoham's interval logic? The maximal fragment ABBL. In *Proc. of LICS'2011*, 387–396. IEEE.
- Bresolin, D.; Montanari, A.; Sala, P.; and Sciavicco, G. 2013. Optimal decision procedures for MPNL over finite structures, the natural numbers, and the integers. *Theoretical Computer Science* 493:98–115.
- Bresolin, D.; Della Monica, D.; Goranko, V.; Montanari, A.; and Sciavicco, G. 2014a. The dark side of interval temporal logic: marking the undecidability border. *Annals of Mathematics and Artificial Intelligence* 71(1-3):41–83.
- Bresolin, D.; Della Monica, D.; Montanari, A.; Sala, P.; and Sciavicco, G. 2014b. Interval temporal logics over strongly discrete linear orders: expressiveness and complexity. *Theoretical Computer Science* 560:269–291.
- Clarke, E. M.; Grumberg, O.; and Peled, D. A. 1999. *Model Checking*. Cambridge, Massachusetts: The MIT Press.
- De Giacomo, G., and Vardi, M. Y. 2013. Linear temporal logic and linear dynamic logic on finite traces. In *Proc. IJCAI'2013*, 854–860. AAAI Press.
- Della Monica, D. 2011. *Expressiveness, decidability, and undecidability of interval temporal logic*. Ph.D. Dissertation, University of Udine.
- Ditmarsch, H. v.; Halpern, J. Y.; Hoek, W. v.; and Kooi, B., eds. 2015. *Handbook of Epistemic Logic*. College Publications.
- Ezekiel, J., and Lomuscio, A. 2009. Combining fault injection and model checking to verify fault tolerance in multi-agent systems. In *Proc. of AAMAS'2009*, 113–120. IFAAMAS Press.
- Ezekiel, J.; Lomuscio, A.; Molnar, L.; and Veres, S. 2011. Verifying fault tolerance and self-diagnosability of an autonomous underwater vehicle. In *Proc. IJCAI'2011*, 1659–1664. AAAI Press.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge*. Cambridge: MIT Press.
- Gammie, P., and van der Meyden, R. 2004. MCK: Model checking the logic of knowledge. In *Proc. CAV'2004*, volume 3114 of *LNCS*, 479–483. Springer-Verlag.
- Halpern, J., and Shoham, Y. 1991. A propositional modal logic of time intervals. *Journal of The ACM* 38:935–962.
- Harel, D.; Tiuryn, J.; and Kozen, D. 2000. *Dynamic Logic*. Cambridge, MA, USA: MIT Press.
- Hopcroft, J., and Ullman, J. D. 1979. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Publishing Company.
- Kacprzak, M.; Nabialek, W.; Niewiadomski, A.; Penczek, W.; Pólrola, A.; Szreter, M.; Woźna, B.; and Zbrzezny, A. 2008. Verics 2007 - a model checker for knowledge and real-time. *Fundamenta Informaticae* 85(1):313–328.
- Lange, M. 2006. Model checking propositional dynamic logic with all extras. *Journal of Applied Logic* 4(1):39–49.
- Lomuscio, A., and Michaliszyn, J. 2013. An epistemic Halpern-Shoham logic. In *Proc. of IJCAI'2013*, 1010–1016. AAAI Press.
- Lomuscio, A., and Michaliszyn, J. 2014. Decidability of model checking multi-agent systems against a class of ehs specifications. In *Proc. of ECAI'2014*, 543–548.
- Lomuscio, A., and Penczek, W. 2015. *Handbook of Epistemic Logic*. College Publications. chapter Model Checking Temporal Epistemic Logic.
- Lomuscio, A.; Penczek, W.; and Woźna, B. 2007. Bounded model checking knowledge and real time. *Artificial Intelligence* 171(16-17):1011–1038.
- Lomuscio, A.; Qu, H.; and Raimondi, F. 2015. MC-MAS: A model checker for the verification of multi-agent systems. *Software Tools for Technology Transfer*. <http://dx.doi.org/10.1007/s10009-015-0378-x>.
- Lomuscio, A.; Qu, H.; and Solanki, M. 2008. Towards verifying compliance in agent-based web service compositions. In *Proc. of AAMAS'2008*, 265–272. IFAAMAS Press.
- Marcinkowski, J., and Michaliszyn, J. 2011. The ultimate undecidability result for the Halpern-Shoham logic. In *Proc. of LICS'2011*, 377–386. IEEE Computer Society.
- Marcinkowski, J., and Michaliszyn, J. 2014. The undecidability of the logic of subintervals. *Fundamenta Informaticae* 131(2):217–240.
- Meyden, R. v., and Shilov, H. 1999. Model checking knowledge and time in systems with perfect recall. In *Proc. of FST&TCS'1999*, volume 1738 of *LNCS*, 432–445. Springer.
- Montanari, A., and Sala, P. 2013. Interval logics and ω B-regular languages. In *Language and Automata Theory and Applications*, volume 7810 of *LNCS*. Springer. 431–443.
- Montanari, A.; Murano, A.; Perelli, G.; and Peron, A. 2014. Checking interval properties of computations. In *Proc. of TIME'2014*, 59–68. IEEE.
- Montanari, A.; Pratt-Hartmann, I.; and Sala, P. 2010. Decidability of the logics of the reflexive sub-interval and super-interval relations over finite linear orders. *Proc. of TIME'2010* 27–34.
- Montanari, A.; Puppis, G.; and Sala, P. 2009. A decidable spatial logic with cone-shaped cardinal directions. In *Proc. of CSL'2009*, 394–408.
- Moszkowski, B. C. 1983. *Reasoning about digital circuits*. Ph.D. Dissertation, Stanford University, Stanford, CA, USA.
- Penczek, W., and Lomuscio, A. 2003. Verifying epistemic properties of multi-agent systems via bounded model checking. In *Proc. of AAMAS'2003*, 209–216. IFAAMAS.
- Raimondi, F., and Lomuscio, A. 2005. Automatic verification of multi-agent systems by model checking via OBDDs. *Journal of Applied Logic* 5(2):235–251.
- van Benthem, J.; van Eijck, J.; and Kooi, B. 2006. Logics of communication and change. *Information and Computation* 204(11):1620–1662.
- Wolper, P. 1983. Temporal logic can be more expressive. *Information and control* 56(1):72–99.