The Tree of Knowledge in Action: Towards a Common Perspective

Johan van Benthem and Eric Pacuit

ABSTRACT. We survey a number of decidablity and undecidablity results concerning epistemic temporal logic. The goal is to provide a general picture which will facilitate the 'sharing of ideas' from a number of different areas concerned with modeling agents in interactive social situations.

1 Introduction

When thinking about rational agents facing choices, one appealing mathematical model recurs in the literature. From Borges' story 'The Garden of Forking Paths' to a host of technical paradigms, sometimes at war, sometimes at peace, all invoke the picture of a branching tree of finite sequences of events with epistemic indistinguishability relations for agents between these sequences, reflecting their limited powers of observation. Indeed, tree models for computation, with branches standing for process evolutions over time, have long been studied in computer science, cf. [29, 30, 7, 2, 13]. Philosophers have studied similar models, now enriched with epistemic relations, for the behavior of intelligent human agents facing choices: see Thomason & Gupta [34], Belnap et al. [5] and Horty [18]. Epistemic models of events over time have also been used in computer science by various authors, witness Fagin et al. [8] and Parikh & Ramanujam [26, 27]. Such trees model not only processes, but also games (see Abramsky [1], Halpern [14] and van Benthem [35]). And finally, 'dynamic logics' of communication and information flow in the tradition of Baltag, Moss & Solecki [3] have tree models of events as their natural broader habitat.

Bringing together knowledge and temporal change is a natural move in modeling, but it is also a potentially dangerous one from a complexity perspective, as has been shown forcefully in Halpern & Vardi [15]. The context is clear from the literature cited just now. Rabin's Theorem tells us that the full monadic second-order logic of the tree of events ordered by the relation of 'initial segment', and provided with some finite set of successor functions is decidable [30]. This explains the decidability of purely temporal logics of events such as **CTL**, and others. Likewise, the tree-like nature of models explains the decidability of many modal logics (see [21]). In a slogan, 'Trees are Safe'. But, we also know that the monadic second-order logic, indeed, even the monadic Π^1_1-theory of the grid $\mathbb{N} \times \mathbb{N}$ is undecidable (see [16]). A grid is like a tree, but successors meet, and the resulting confluent structure

Paper submitted to Advances in Modal Logic 2006

is known to cause undecidability in many areas of modal logic ([20]), witness in particular the work of Gabbay et al. on 'product models' [11]. In one more slogan: 'Grids are Dangerous'.

Now, epistemic temporal logics live at a dangerous edge here. On top of Rabin-style tree models, they introduce epistemic indistinguishability relations which generate a 'second dimension', and if the language gets too powerful, enough grid structure can be encoded to cause undecidability. Illustrations for this again come from a wide range of papers. E.g., Thomas [33] points out, following Läuchli, how introducing a relation of 'simultaneity' into the Rabin tree makes the monadic second-order logic undecidable. Likewise, Halpern & Vardi show how epistemic-temporal logics of agents with Perfect Recall and No Learning can become undecidable [15]. But the situation is delicate, as small changes in an epistemic temporal language or class of models can affect the complexity of the logic in drastic ways.

This is the view 'from above', viewing epistemic temporal models as a Grand Stage where events unfold. There is also the view 'from below', found in 'dynamic epistemic logics' which construct successive new event models in definable stages (cf. Baltag, Moss and Solecki [4] and van Benthem, van Eijck and Kooi [39]). The logics tend to be decidable (though cf. [39] and [22]) and this, too, calls for explanation.

In this paper, we position ourselves close to the edge of undecidability in a straightforward system of epistemic temporal logic. We will discuss a number of complexity results, on both sides of the edge, while pointing out how results from all different traditions mentioned here help illuminate the landscape. As a result, we are also able to 'place' dynamic epistemic logics as species of epistemic temporal ones — and find room for comparing ideas from both traditions, e.g., in process algebra and game analysis.

In doing all this, we also have a broader aim. The area that we are describing consists of a number of different frameworks, whose practitioners either do not know about relevant work by others, or are not even on speaking terms. We feel that this is an unfortunate situation, since much is to be gained by seeing the commonality of one area of research here. As we shall see, issues are often the same, and notions and techniques can be borrowed freely. Our paper is one such contribution toward a merge[1].

2 Epistemic Temporal Logic

This section describes the basic models for our study, whose typical interpretations are conversations or games. We are interested in how the agents' knowledge about the situation may change over time. Let Σ be a set of **events**. An event might be a move in some game, or a message sent from one agent to others. Not all agents need be aware of all events. Also, there is a global discrete clock, labelled by natural numbers, which agents may or may not be aware of. Agents do have a finite capacity to remember events, perhaps unbounded.

[1] We emphasize only main lines: cf. [38] for details, here and throughout this paper.

2.1 Epistemic temporal models and structural conditions

We first settle on some notation for the 'playgrounds'. Let Σ be any set of **events**. Given any set X, X^* is the set of finite strings over X and X^ω is the set of infinite strings over X. Elements of $\Sigma^* \cup \Sigma^\omega$ will be called **histories**. Given $H \in \Sigma^* \cup \Sigma^\omega$, $\mathsf{len}(H)$ is the **length** of H, i.e. the number of characters (possibly infinite) in H. Given $H, H' \in \Sigma^* \cup \Sigma^\omega$, we write $H \preceq H'$ if H is a *finite* prefix of H'. If $H \preceq H'$ we call H an **initial segment** of H' and H' an **extension** of H. Given an event $e \in \Sigma$, we write $H \prec_e H'$ if $H' = He$. Finally, let ϵ be the empty string and $\mathsf{FinPre}(\mathcal{H}) = \{H \mid \exists H' \in \mathcal{H} \text{ such that } H \preceq H'\}$ be the set of finite prefixes of the elements of $\mathcal{H}$ and $\mathsf{FinPre}_{-\epsilon}(\mathcal{H}) = \mathsf{FinPre}(\mathcal{H}) - \{\epsilon\}$.

DEFINITION 1. Let Σ be any set of events. A set $\mathcal{H} \subseteq \Sigma^* \cup \Sigma^\omega$ is called a **protocol** provided $\mathsf{FinPre}_{-\epsilon}(\mathcal{H}) \subseteq \mathcal{H}$. A **rooted protocol** is any set $\mathcal{H} \subseteq \Sigma^* \cup \Sigma^\omega$ where $\mathsf{FinPre}(\mathcal{H}) \subseteq \mathcal{H}$.

Intuitively, a protocol is the set of all possible ways an interactive situation may evolve. Given a protocol $\mathcal{H}$ and a finite history $H \in \mathcal{H}$, $\mathsf{Ext}_\mathcal{H}(H) = \{H' \mid H' \in \mathcal{H}, H \preceq H'\}$ is the set of extensions of H from $\mathcal{H}$. If no confusion arises, we write $\mathsf{Ext}(H)$ instead of $\mathsf{Ext}_\mathcal{H}(H)$. Also, $\mathsf{Ext}^{<\omega}(H)$ is the set of **finite extensions** of H and $\mathsf{Ext}^\omega(H)$ the **infinite extensions** of H. Given $t \in \mathbb{N}$ and a history H, H_t is the unique initial segment of H of length t.

Once the underlying temporal structure is in place, we can add the uncertainty of the agents. The most general models we have in mind are 'forests' with epistemic relations between finite branches.

DEFINITION 2. An **ETL frame** is a tuple $\langle \Sigma, \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ where Σ is a (finite or infinite) set of events, $\mathcal{H}$ is a protocol, and for each $i \in \mathcal{A}$, $\sim_i$ is an equivalence relation on the set of finite strings in $\mathcal{H}$.

Making assumptions about the underlying event structure corresponds to "fixing the playground" where the agents will interact. The assumptions of interest are as follows: Let $\mathcal{F} = \langle \Sigma, \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ be an *ETL* frame. If Σ is assumed to be finite, then we say that $\mathcal{F}$ is **finitely branching**. If $\mathcal{H}$ is a rooted protocol, $\mathcal{F}$ is a **tree frame**. We will be interested in **protocol frames** which satisfy both of these conditions. These are finitely branching trees with epistemic relations between the finite branches.

REMARK 3. *Three Equivalent Approaches*: There are at least two further approaches to uncertainty in the literature. The first, discussed in [26], represents agents' "observational" power. That is, each agent i has a set E_i of events it *can* observe[2]. For simplicity, we can assume $E_i \subseteq \Sigma$ but this is not necessary. A **local view** function is a map $\lambda_i : \mathsf{FinPre}(\mathcal{H}) \to E_i^*$. Given a finite history $H \in \mathcal{H}$, the intended interpretation of $\lambda_i(H)$ is "the sequence of events observed by agent i at H". The second approach comes from Fagin et al. [8]. Each agent has a set L_i of **local states** (if necessary, one can also assume a set L_e of environment states). Events e are tuples of local states (one for each agent) $\langle l_1, \ldots, l_n\rangle$ where for each $i = 1, \ldots, n$,

[2] This may be different from what the agent *does* observe in a given situation.

$l_i \in L_i$. Then two finite histories H and H' are i-equivalent provided the local state of the last of event on H and H' is the same for agent i. From a technical point of view, the three approaches to modeling uncertainty are equivalent ([24] provides the relevant intertranslations). However, they may still be different for modeling purposes.

2.2 Agent oriented conditions

Now we turn from the "playground" to the "players". Various types of agents place constraints on the interplay between the epistemic and temporal relations. We survey some conditions from the literature.

DEFINITION 4. Fix an epistemic temporal frame $\langle \Sigma, \mathcal{H}, \{\sim_i\}_{i \in \mathcal{A}} \rangle$. An agent $i \in \mathcal{A}$ satisfies the property **No Miracles** (sometimes called, somewhat misleadingly, **No Learning**) if for all finite histories $H, H' \in \mathcal{H}$ and events $e \in \Sigma$ with $He \in \mathcal{H}$ and $H'e \in \mathcal{H}$, if $H \sim_i H'$ then $He \sim_i H'e$.

Thus, unless a 'miracle' happens, uncertainty of agents cannot be erased by the same event. The next condition is the dual property.

DEFINITION 5. An agent $i \in \mathcal{A}$ satisfies the property **Perfect Recall** provided for all finite histories $H, H' \in \mathcal{H}$ and events $e \in \Sigma$with $He \in \mathcal{H}$ and $H'e \in \mathcal{H}$, if $He \sim_i H'e$ then $H \sim_i H'$.

Perfect Recall means that the histories an agent considers possible can only decrease or remain the same, unless new indistinguishable events occur.

DEFINITION 6. An agent $i \in \mathcal{A}$ is **synchronized** provided for all finite histories $H, H' \in \mathcal{H}$, if $H \sim_i H'$ then $\mathsf{len}(H) = \mathsf{len}(H')$.

Intuitively, if an agent is synchronized, then that agent knows the value of the global clock (this may or may not be expressible in the formal language). For other assumptions that can be made about the interaction between the epistemic relation and time, the reader is referred to [8, 37]. Finally, note that in general we do not assume that all agents have the same reasoning capabilities. When they do, we say, for example, that a frame $\mathcal{F}$ is synchronous if all agents are synchronized.

2.3 Formal languages and truth in a model

Different modal languages can reason about the above structures (see the Handbook chapter [17]), with 'branching' or 'linear' variants. Here we give just the bare necessities.

Let At be a countable set of atomic propositions. We are interested in languages with various combinations of the following modalities: $P\phi$ (ϕ is true *sometime* in the past), $F\phi$ (ϕ is true *sometime* in the future), $Y\phi$ (ϕ is true at *the* previous moment), $N\phi$ (ϕ is true at *the* next moment), $K_i\phi$ (agent i knows ϕ) and $C_B\phi$ (the group $B \subseteq \mathcal{A}$ commonly knows ϕ). Dual operators are written as usual (eg., $\langle i \rangle \phi = \neg K_i \neg \phi$). If X is a sequence of modalities from $\{P, F, Y, N\}$ let $\mathcal{L}_n^X$ be the language with n knowledge modalities $K_1, \ldots, K_n$ together with the modalities from X. For a sequence of modalities X, $\mathcal{L}_C^X$ is the language $\mathcal{L}_n^X$ closed under the common knowledge modality C. Let $\mathcal{L}_{ETL}$ be the full epistemic temporal language, i.e., it contains all of the above temporal and knowledge operators.

Regardless of whether the language has *branching time* or *linear time* temporal operators, formulas express properties about finite histories. The difference lies in the format of the satisfaction relation. In a linear temporal setting, formulas are interpreted at pairs H, t where H is a 'maximal' (possibly infinite) history and t an element of $\mathbb{N}$. The intended interpretation of $H, t \models \phi$ is that *on the branch* H at time t, ϕ is true. In the branching time setting, we only need the moment, and formulas can be interpreted at finite histories H. In the interest of a unified approach we will interpret formulas at branch-time pairs. However, it will sometimes be useful to take the branching time interpretation. This helps draw parallels with results in temporal modal logic and products of modal logics [11].

DEFINITION 7. An **ETL model** based on an ETL frame $\langle \Sigma, \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}} \rangle$ is a tuple $\langle \Sigma, \mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}, V \rangle$ where V is a valuation $V : \mathsf{At} \to 2^{\mathsf{FinPre}(\mathcal{H})}$.

Formulas are interpreted at pairs H, t where $t \in \mathbb{N}$ and $H \in \mathcal{H}$ has length longer than t (finite or infinite). Truth for the languages $\mathcal{L}_n^X$ is defined as usual: see [8] and [17] for details. We only remind the reader of the definition of the knowledge and some temporal operators:

- $H, t \models P\phi$ iff there exists $t' \le t$ such that $H, t' \models \phi$
- $H, t \models F\phi$ iff there exists $t' \ge t$ such that $H, t' \models \phi$
- $H, t \models K_i\phi$ iff for each $H' \in \mathcal{H}$ and $m \ge 0$ if $H_t \sim_i H'_m$ then $H', m \models \phi$

Of course, in addition to our epistemic temporal formulas, there are also the standard logical languages appropriate to these models, such as first-order logic, second-order logic, and other well-known systems.

3 Living at the Edge

Having set up our basic framework, we now want to demonstrate some key facts about the borderline between decidable and undecidable epistemic temporal logics. The previous section did highlight a number of dimensions which may lead undecidability, and even much higher complexity:

1. Expressivity of the formal language. Does the language include a common knowledge operator? A future operator? Both?
2. Structural conditions on the underlying event structure. Do we restrict to protocol frames (finitely branching trees)? Or forrests?
3. Conditions on the reasoning abilities of the agents. Do the agents satisfy Perfect Recall? No Miracles? Synchronization?

Instead of setting up a huge grid of possible model classes and languages, we highlight a few major stages, including (in Section 4) one new highly undecidable epistemic tree logic. The main line of our observations is not all that new by itself, but our presentation and variety of sources is.

3.1 Purely temporal reasoning on protocol models

In this section we fix the underlying event structure and vary other dimensions. The Rabin Tree ([30]) consists of all finite sequences of events from a given finite set, with the binary relation of 'initial subsequence' plus successor functions taking a sequence H to He, for each $e \in \Sigma$.

THEOREM 8 (Rabin [30]). *The monadic second-order logic of the Rabin Tree is decidable.*

This landmark result explains the decidability of many modal and temporal logics, as first pointed out by Gabbay[3] [10]. It applies particularly well to our setting here, since the Rabin Tree has both points and branches, represented as special sets of points. Here is a well-known consequence:

THEOREM 9. *The satisfiability problem for $\mathcal{L}_{TL}$ with respect to TL tree models without epistemic structure is decidable.*

Proof. A formula ϕ involving finitely many events e is true in all protocol models if $\forall A(\text{'}subtree(A)\text{'} \Rightarrow (\phi)_A)$ is true on the corresponding Rabin Tree. Here $(\phi)_A$ is the syntactic relativization of ϕ to the unary predicate A, and '$subtree(A)$' says that A is closed under taking initial segments. ■

A number of authors have noted that seemingly simple extensions to the Rabin tree language leads to undecidability. For example, Läuchli proved that the first-order theory of the Rabin tree expanded with a binary 'equilevel' predicate[4] for nodes is undecidable. Upon first inspection, this appears to be bad news for for the innocent assumption of synchronous communication. However, Thomas [33] provides a more fine-grained perspective: he shows that the monadic second-order theory of the Rabin Tree with an 'equilevel' predicate remains decidable provided that we let the second-order quantification run over *linear chains*, rather than arbitrary subsets. More succinctly: 'Path Logic' over the Rabin Tree with an equilevel predicate is decidable. Path Logic extends our temporal languages, since these talk about initial segments and extensions of the current finite history.

3.2 $\mathcal{L}_{ETL}$ over arbitrary models

First, consider arbitrary ETL tree models ('forests') and the full epistemic-temporal language $\mathcal{L}_{ETL}$. The logic remains simple. Indeed the 'fusion' of epistemic logic (**S5**) with common knowledge pluss a compete temporal logic with past time operators (cf. [11]) will be such an axiomatization. This result is standard, so we only give some relevant details.

THEOREM 10. *The validity problem for arbitrary ETL frames is RE.*

Proof. (Sketch) Any non-theorem of the fusion of an epistemic logic and a temporal logic has a bimodal (Kripke) counter-model $\mathbb{M}$ with one accessibility relation for the temporal modalities and one for the epistemic modalities.

[3] Also relevant here is the emphasis in [41] on the *bounded tree property* as the source of decidability for temporal logics.

[4] That is, the nodes have the same distance from the root.

In order to generate a standard ETL model, we unravel the Kripke model *at each point*. This creates a forest where each tree is rooted by a state from the Kripke structure where we set $H \sim_i H'$ iff $\mathsf{last}(H) \sim_i \mathsf{last}(H')$ in $\mathbb{M}$. The relation from points s in $\mathbb{M}$ to histories with s their last element is a bisimulation. Thus the unraveled $\mathbb{M}$ is an ETL counter model. ■

We do not know if this general logic is decidable, though we suspect that it is, by the general results on transfer of decidability for fusions of modal logics in Gabbay et al. [11], Kurucz [19].

ETL *tree* models will validate some principles not valid in the fusion of epistemic and temporal logic. The first is structural — tree models have a root. It is not hard to find axioms this (see French, van der Meyden and Reynolds [9] for completeness theorems under this assumption). The second principle enforces that each agent knows the underlying protocol. The formula $\langle i \rangle \phi \rightarrow PF\phi$ says that any epistemic alternative is reachable in the tree by going down and moving up again.

THEOREM 11. *The satisfiability problem for the language $\mathcal{L}_{ETL}$ over ETL tree models is RE.*

Proof. (Sketch) The logic of ETL *tree* frames is the fusion of epistemic logic with common knowledge and temporal logic together with the principles discussed above. Starting with a Kripke counter model, we can unravel at the root only, making the above principle true in the model. ■

Of course, behaviour of specific agents will take place in models satisfying additional epistemic-temporal constraints. As we will see in the next sections this can lead to high undecidability results.

3.3 Ideal epistemic agents have a highly undecidable tree logic

Let us now consider the usual idealizations of epistemic logic. For example, Agents have perfect memory, and seeing new events will not confuse them: that is, we have the above Perfect Recall, and No Miracles properties. The resulting interaction of temporal and epistemic structure makes trees look more like grids, and indeed, undecidability strikes. We highlight this result, because it is indicative of the 'danger zone' that we are in. The following result is one of many from a landmark publication:

THEOREM 12 (Halpern & Vardi [15]). *The validity problem for $\mathcal{L}_{ETL}$ on arbitrary ETL frames with No Miracles or Perfect Recall is Π^1_1 complete.*

In fact, these results hold whether or not one assumes that the frames are synchronous (see [15] for details). Essentially, these results show that if we fix the underlying event structure to be an ETL frame (i.e., a forest with arbitrary branching), then any practically any idealization lead to high undecidability as long as we are working in a language with common knowledge and arbitrary future modalities.

One may suspect that the Π^1_1-completeness is due to the underlying event structure and that things are better on event trees instead of forests. However, high complexity still strikes.

THEOREM 13 (Halpern & Vardi [15]). *The validity problem for $\mathcal{L}_{ETL}$ on ETL tree frames with Perfect Recall and No Miracles is Π^1_1-complete.*

On certain playgrounds, these idealizations turn out to be less dangerous.

THEOREM 14 (Halpern & Vardi [15]). *The validity problem for $\mathcal{L}_{ETL}$ with respect to ETL trees that satisfy the no miracles property is co-RE.*

Indeed, under synchronous communication the validity problem even becomes decidable (see [15] for details). These results indicate that when working in a language with both a common knowledge operator and arbitrary future modality there is an interesting interplay between structural assumptions about the underlying event structure and structural about the epistemic capabilities of the agents. Of course, for a full analysis of the situation we need to get our hands dirty and analyze the Π^1_1-completeness results. This will be the topic of Section 4.

This concludes our survey of typical results on decidability and undecidability over epistemic temporal tree models. Not surprisingly, the boundary has to do with the transition from mere trees to grid encoding using the additional epistemic structure. The epistemic setting adds some special flavor, however, in that the small differences which affect complexity represent very concrete assumptions about agents' capabilities, and what we can say about these. Moreover, we have shown how one can learn about relevant results from traditions that look prima facie quite different: epistemic temporal logic, tree languages in the foundations of computation, and (as we shall see in Section 5) current work on products of modal logics.

3.4 Bounded agents have a simple logic

Special agents may also have easier epistemic temporal logics. At the opposite extreme of Perfect Recall, agents with *bounded memory* have some finite bound to the number of preceding events which they can remember. Now, epistemic relations can be defined in terms of temporal ones.

THEOREM 15. *The epistemic temporal logic of memory bounded agents over arbitrary ETL frames is decidable.*

The key observation is that with a finite number of events, the modality $K_i\phi$ is definable. For convenience, we do the case of memory bound one:

$$K_i \leftrightarrow \bigvee_e (P_e\top \wedge U(P_e\top \to \phi))$$

where U is the *universal modality* and $P_e\top$ says the last event was e. The result follows the decidability for the purely temporal language.

4 High Undecidability on Trees

In the previous section we saw that, for a language with a common knowledge and a future operator, varying the underlying event structure and epistemic assumptions about the agents has drastic effects on the decidability of the logic. Now we investigate the tension between the underlying event structure, idealization of the agents and the formal language.

4.1 Tiling arguments

Imagine a finite set of tiles where each side has a different color. Let $\mathcal{T}$ be such a finite set of tile types and for $T \in \mathcal{T}$, let $right(T)$, $left(T)$, $up(T)$ and $down(T)$ be the colors of T. The *tiling problem* (for the first quadrant) asks if there a function $t : \mathbb{N} \times \mathbb{N} \to \mathcal{T}$ such that for each $n, m \in \mathbb{N}$

$$\begin{array}{rcl} right(t(n,m)) & = & left(t(n+1,m)) \\ up(t(n,m)) & = & down(t(n,m+1)) \end{array}$$

That is, can we place the tiles on the $\mathbb{N} \times \mathbb{N}$ plane so that the colors of the edges match. The function t is called a **tiling** of $\mathbb{N} \times \mathbb{N}$. Prima facie this problem looks highly complex (monadic Σ^1_1) as it asserts the existence of a function. However, by appealing to König's Lemma it can be seen to be Π^0_1: it is enough to show the existence of tilings of arbitrarily large *finite* planes. More formally, call any function $t^{(n)} : \{(i,j) \mid 0 \leq i \leq n,\ 0 \leq j \leq n\} \to \mathcal{T}$ that satisfies the above conditions (i.e., tiles match vertically and horizontally) a $(n \times n)$**-tiling** of the plane. Two tilings $t^{(n)}$ and $t^{(m)}$ are **consistent** if one extends the other. Thus each $(n \times n)$-tiling can be thought of as a sequence of partial *consistent* tilings.

LEMMA 16. *Suppose that for each $n > 0$, there is at least one (but only finitely many) partial tilings $t^{(n)}$. Then there is a tiling of the entire plane.*

However, David Harel showed [16] that small changes to the problem greatly increases the complexity. For example, the **recurrent tiling** problem asks, given a set of tiles $\mathcal{T}$ with a distinguished tile $T_1 \in \mathcal{T}$, if there is a tiling t such that T_1 occurs infinitely often in the first row.

THEOREM 17 ([16]). *The recurrent tiling problem is Σ^1_1-complete.*

Thus if there is a formula in the desired language that is satisfiable iff there is a recurrent tiling of the plane, then the satisfiability problem with respect to that language (on the relevant frames) is Σ^1_1-complete. For concreteness, assume that $\mathcal{T} = \{T_1, \ldots, T_k\}$ is a finite set of tiles and $t_1, \ldots, t_k$ is a set of propositional variables.

4.2 A PDL-style tree language

In this section will use a **PDL**-style language which capture features of both linear and branching time languages, an which refers explicity to events. Let $\mathcal{A}$ be a (finite) set of agents and recall that Σ is a (finite) set of events. Define $\mathcal{L}_\Sigma(\mathcal{A})$ inductively as follows:

$$\phi := p \mid \neg\phi \mid \phi \wedge \psi \mid \langle\alpha\rangle\phi$$

$$\alpha := a \mid ?p \mid \alpha;\beta \mid \alpha \cup \beta \mid \alpha^*$$

where $p \in \mathsf{At}$, $a \in \Sigma \cup \mathcal{A}$ and $\sigma \in \Sigma$. Let $\mathcal{L}_\Sigma(\mathcal{A})^-$ be the language $\mathcal{L}_\Sigma(\mathcal{A})$ which allows expressions of the forms $\langle\sigma^-\rangle\phi$.

This language is (strictly) stronger than those described above, as we allow mixing of temporal and epistemic steps under the scope of the $*$-

operator. For example, $\langle (i;e)^*\rangle\phi$ is a well-formed expression of the above grammar, whereas it is not an element of $\mathcal{L}_{ETL}$.

Before defining truth in a model we introduce a relation R_α on the set $\mathsf{FinPre}(\mathcal{H})$, where α is defined by the above grammar. Let H, H' be finite sequences of events and V a valuation (assigning sets of atomic propositions to finite sequences). Suppose $\sigma \in \Sigma$ and $i \in \mathcal{A}$.

- $HR_\sigma H'$ iff $H' = H\sigma$ if $\sigma \in \Sigma$
- $HR_i H'$ iff $H \sim_i H'$
- $HR_{\sigma^-}H'$ iff $\mathsf{len}(H) \geq 1$ and $H = H'\sigma$
- $HR_{?p}H'$ iff $H = H'$ and $p \in V(H)$.

Clauses for the PDL operators are as usual. Truth is also defined as usual, we only give the definition of the modal operator:

- $H, t \models \langle\alpha\rangle\phi$ iff there exists $H' \in \mathcal{H}$ and $m \in \mathbb{N}$ such that $H_t R_\alpha H'_m$ and $H', m \models \phi$

Under the assumption that there are only finitely many events and using a well-known translation of epistemic logic into PDL (with a converse operator), we see that $\mathcal{L}_{ETL}$is a fragment of $\mathcal{L}_\Sigma(\mathcal{A})$. We will write $G\phi$ for $[(\cup_{e\in\Sigma}e)^*]\phi$ and $C\phi$ for[5] $[(\cup_{i\in\mathcal{A}}i)^*]\phi$.

4.3 High complexity over arbitrary ETL frames

We first reprove one of Halpern and Vardi's results from [15] using a tiling argument. [15] use a reduction of the *recurrent Turning machine problem.* They comment that a tiling argument "cannot be straightforwardly applied" in their setting (p. 208). Our argument works thanks to our formulation of the No Miracles and Perfect Recall properties.

THEOREM 18 (Halpern & Vardi [15]). *The validity problem for the $\mathcal{L}_{ETL}$ fragment of $\mathcal{L}_\Sigma(\mathcal{A})$ on finitely branching ETL frames with No Miracles (with at least two agents) is Π^1_1 complete.*

The first step in any tiling argument is to identify a *universal* modality. The combination of the universal temporal and the common knowledge operator (GC) will serve this purpose. The second step is to encode a grid.

The 'x-axis' will be encoded by occurrences of a distinguished event $e \in \Sigma$. The formula $\phi_1 := GC(e\rangle\top)$ says that each accessible finite history has an extension consisting of an infinite sequence of e's. As in [15], the epistemic relations encode the 'y-axis'. Let p be a new propositional variable. Consider the following two formulas: $\phi_2 := GC((p \to Gp) \wedge (\neg p \to G\neg p))$ and $\phi_3 := GC(p \to \langle 1\rangle p) \wedge (\neg p \to \langle 2\rangle p)$. If $H, t \models \phi_2 \wedge \phi_3$, then we can think

[5] Of course this only works if there are finitely many events and finitely many agents. If there are not finitely many events, we assume that F is a primitive operator defined as in Section 2.3. Furthermore, note that we do not need the converse operator here, since we are assuming that the agent's accessibility relations are equivalence relations.

of the histories reachable from H_t as being labeled by p and $\neg p$. Furthermore, there are 1-accessibility relations between from p to $\neg p$ histories and 2-accessibility relations from $\neg p$ to p histories. Thus an 'up-step' is represented by the program $\alpha_u := (?p; 1; ?\neg p; 2)$. Now, the No Miracle property imposes a grid condition on the relations R_e and R_{α_u}.

LEMMA 19. *Suppose that $\mathcal{M}$ is an arbitrary ETL model with no miracles and $H, t \models \phi_1 \wedge \phi_2 \wedge \phi_3$. If H_1, H_2 and H_3 are finite histories reachable from H_t, $H_1 R_{\alpha_u} H_2$ and $H_1 R_e H_3$, then there is an H_4 such that H_4 is reachable from H_t, $H_3 R_{\alpha_u} H_4$ and $H_2 R_e H_4$.*

To complete the proof of Theorem 18, we find a formula that is satisfiable iff there is a recurrent tiling of the plane. The next section sketches how to do this in an analogous case.

4.4 High complexity over ETL protocol frames

Halpern & Vardi mainly consider models where the initial model may be infinite, or there may be infinite branching. In this case, even the 'unmixed' language of Section 2.3 above led to undecidability with No Miracles or Perfect Recall. In this section, we consider finitely branching trees.

Our goal in this section is to sketch a proof of the following theorem.

THEOREM 20. *The satisfiability problem of $\mathcal{L}_\Sigma(\mathcal{A})$ with respect to ETL protocol frames that satisfy No Miracles is Σ^1_1-complete.*

For concreteness, assume $\Sigma = \{l, r\}$ and $\mathcal{A} = \{1, 2\}$. We must find a formula $\phi_\mathcal{T}$ that is satisfiable iff there is a recurrent tiling of $\mathbb{N} \times \mathbb{N}$ using the tiles from $\mathcal{T}$. We begin by describing the formula $\phi_\mathcal{T}$. The formula $\phi_\mathcal{T}$ consists of three parts: 1. a formula which forces the extensions of a finite history to have a particular structure, 2. a formula which forces a grid structure and 3. a formula which places tiles on the grid.

To that end, let ϕ_S be the conjunction of the following formulas: Only $r^* - l^*$ paths: $[r^*; l; l^*]\neg\langle r\rangle\top$; infinite l-paths: $[r^*; l^*]\langle l\rangle\top$; Infinite r-path: $[r^*]\langle r\rangle\top$; Even p paths: $[(r; r)^*][l^*]p$; and Odd $\neg p$ paths: $[r; (r; r)^*][l^*]\neg p$. Then if $H, t \models \phi_S$, the extensions of H_t can be pictured as follows:

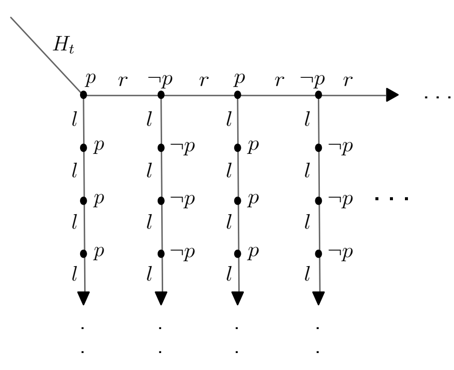

This model represents half of the $\mathbb{N} \times \mathbb{N}$ grid. The idea is to think of the infinite r-path as the y-axis and the first infinite l-path as the x-axis (the fact that the truth value of p alternates between the paths will be used below). We now show how to force the second half of the grid. That is, we need a formula that will be satisfied if there are infinitely many infinite "up" paths. The trick will be to consider the following program:

$$\alpha_u := ?p; l; 1; ?\neg p; l; 2; ?p.$$

Making a step of the above program corresponds to making a 'zig-zag' move through the tree between points which will cross between different branches.

Let ϕ_u be the conjunction of the following two formulas

1. "Up moves" are always possible: $[r^*; l^*]((p \to \langle 1\rangle\neg p) \wedge (\neg p \to \langle 2\rangle p))$
2. There are infinitely many "Up moves": $[l^*][\alpha_u^*]\langle\alpha_u\rangle\top$

Finally, we place the tiles on our tree.

1. One tile at each node: $\phi_1 := [(r;r)^*; l^*](\bigvee_{i=1}^k t_i \wedge \bigwedge_{1\leq i<j\leq k} \neg(t_i \wedge t_j))$
2. Place tiles going across: $\phi_2 := [(r;r)^*; l^*](\bigvee_{right(T_i)=left(T_j)}(t_i \wedge \langle l\rangle t_j))$
3. Place tiles going up: $\phi_3 := [(r;r)^*; l^*](\bigvee_{up(T_i)=down(T_j)}(t_i \wedge \langle\alpha_u\rangle t_j))$
4. Enough tiles: $\phi_4 := [(r;r)^*; l^*](\bigwedge_{up(T_i)\neq down(T_j)}(t_i \to \neg\langle\alpha_u\rangle t_j))$

Let $\phi_{\mathcal{T}} := \phi_S \wedge \phi_u \wedge \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4$.

We first note that any tiling of $\mathbb{N}\times\mathbb{N}$ induces a model for $\phi_{\mathcal{T}}$. The key idea is to remove all vertical lines from the grid and treat the remaining structure as a tree rooted at $(0,0)$. Next, as in the previous section, the No Miracles property imposes a grid structure on the relations H_l and H_{α_u}.

LEMMA 21. *On epistemic temporal frames $\langle\mathcal{H}, \{\sim_i\}_{i\in\mathcal{A}}\rangle$ with No Miracles, if HR_lH' and $HR_{\alpha_u}H''$, there is a H''' with $H''R_lH'''$ and $H'R_{\alpha_u}H'''$.*

Proof of Theorem 20 We need only show that under the assumption of No Miracles, there is a tiling of the plane from a model for $\phi_{\mathcal{T}}$. Indeed, we find a function f from $\mathbb{N}\times\mathbb{N}$ into $\mathsf{Ext}^{<\omega}(H_t)$ such that

- $t(n,m) = T_i$ iff$_{\text{def}}$ $f(n,m) \models t_i$.

is a tiling of $\mathbb{N}\times\mathbb{N}$. We will show that such a function can be extracted from a satisfying model. First of all, since $H, t \models \phi_1$, f is well-defined. Start by defining $f(0,0) = T$ where $H, t \models t$. Once $f(n,m)$ has been defined, we can define $f(n+1,m)$ and $f(n,m+1)$:

LEMMA 22. *Let $H, t \models \phi_{\mathcal{T}}$, $(n,m) \in \mathbb{N}\times\mathbb{N}$, $f(n,m) = H_tH'$ and $H_tH' \models t_i$. Then there are finite extensions of H_t, H^r and H^u, with (1) $H_tH^r \models t_j$, and $right(T_i) = left(T_j)$, and (b) $H_tH^u \models t_k$, and $up(T_i) = down(T_k)$.*

All that remains is to "complete the square". Let $f(n,m)$, $f(n,m+1)$ and $f(n+1,m)$ have finite extensions of H_t as in the Lemma 22: $H_{(n,m)}$,

$H_{(n,m+1)}$ and $H_{(n+1,m)}$ respectively, with mathcing properties for a tiling. Then there is a finite extension $H_{(n+1,m+1)}$ of H_t with the right properties.

LEMMA 23. *Suppose that $H_{(n,m)}, H_{(n,m+1)}$ and $H_{(n+1,m)}$ have been defined as above. Then there is a finite history $H_{(n+1,m+1)} \in \mathsf{Ext}^{<\omega}(H_t)$ such that there is a unique tile proposition with $H_{(n+1,m+1)} \models t$ with $right(T_j) = left(T)$ and $up(T_k) = down(T)$.*

Finally, we must show that there is a formula such that, if satisfiable, implies that a particular tile occurs infinitely often along the x-axis. Let $T_0 \in \mathcal{T}$ be a tile and t_0 the corresponding propositional variable. Consider the formula $\phi_{t_0} := [l^*]\langle l; l^*\rangle t_0$Note that $H, t \models \phi_{t_0}$ implies that t_0 is true infinitely often on the first branch extending H_t. This proves Theorem 20.

Note that we have focused only on the assumption of No Miracles. For systems satisfying Perfect Recall alone epistemic uncertainty is propagated backwards towards the root. Thus, the above arguments will not work. However, by appealing to Lemma 16, we need only show that for each $n \in \mathbb{N}$ there is a tiling of the finite $n \times n$ plane (cf. [38]).

THEOREM 24. *The satisfiability problem with respect to synchronous ETL protocol frames that satisfy perfect recall is Σ^1_1-complete.*

4.5 Synchronized agents on ETL trees

In more special settings we can still get positive results from other areas:

THEOREM 25. *The logic of the language $\mathcal{L}^{P,F}_C$ over synchronous tree models is decidable.*

Proof. Thomas [33] embeds Path Logic into the monadic second-order theory of the Rabin tree, by sending chains to pairs of subsets (A, B) where A encodes the left-most branch on which the chain lies, while B encodes which nodes are on the chain. More precisely, B 'goes left' at levels not represented on the chain, and it 'goes right' at levels where the chain has a node. The equilevel predicate for two nodes is then expressed by saying that they are one-element chains, whose B-sequences go right at the same place. Now, our epistemic relations are subrelations of the 'equilevel' predicate , which latter corresponds roughly to the transitive closure of their union.We can encode this into Path Logic by modifying the chain representation. ■

5 Dynamic Epistemic Logic

Our take on epistemic temporal logic is within the tradition of Fagin et al. [8] and Parikh & Ramanujam [26, 27]. One current paradigm which diverges from these, though addressing similar phenomena, is 'dynamic epistemic logic' ('DEL', [12, 28, 4]). Here, epistemic actions such as announcing a true proposition, or more complex forms of communication and observation, are encoded explicitly in 'action models' $\mathbb{A}$ consisting of the relevant events and the preconditions for their occurrence, plus agents' epistemic relations over these, representing their partial powers of observation. In ur-DEL, preconditions for events are defined by purely epistemic formulas.

What agents learn in such a setting, given some current epistemic model $\mathbb{M}$, is encoded by a new 'product model' $\mathbb{M} \times \mathbb{A}$, where agents are uncertain between worlds (s, e) and (t, f) iff they were uncertain between both the old worlds s, t and the observed events e, f. The language for these models has the usual epistemic operators, plus dynamic modalities $\langle \mathbb{A}, e \rangle \phi$:

$$\mathbb{M}, s \models \langle \mathbb{A}, e \rangle \phi \text{ iff } \mathbb{M} \times \mathbb{A}, (s, e) \models \phi$$

The resulting logic is decidable, and it revolves around 'reduction axioms' for compositional analysis of effects of epistemic events. E.g., a typical reduction axiom analyzes agents' knowledge after public announcement:

$$[!P]K_i\phi \leftrightarrow (P \to K_i[!P]\phi)$$

For more precise definitions and complete dynamic-epistemic logics, cf. van Benthem [35], Baltag & Moss [3], van Benthem [36], van Benthem, van Eijck & Kooi [39], and van Ditmarsch, van der Hoek & Kooi [40].

Prima facie, DEL does not look like ETL: there is no explicit mention of time. And unlike DEL, ETL does not explicitly describe the events that make up its models. But appearances are misleading, and a 'convergence' is easy to find, making mutual borrowing easy and natural.

5.1 Representing DEL models inside ETL models

Product update involves three major ingredients with a logical 'reflection', as was first observed in van Benthem [35]

(a) Product update implies Perfect Recall: $(x, a) \sim_i (y, b)$ implies $x \sim_i y$,

(b) Moreover, 'No Miracles' holds uniformly: actions are either always distinguishable, or never: if $(x, a) \sim_i (y, b)$, then, whenever $u \sim_i v$, also $(u, a) \sim_i (v, b)$ if the latter events can occur at all.

Now, any initial model $\mathbb{M}$ and action model $\mathbb{A}$ induce a natural epistemic model $Tree(\mathbb{M}, \mathbb{A})$. Nodes are finite sequences of events, and the iterated epistemic products with $\mathbb{A}$ are the horizontal levels of the tree. Events only take place when their precondition is satisfied. $Tree(\mathbb{M}, \mathbb{A})$ is a 'forest model', whose initial epistemic model at the root can be arbitrary. (Van Benthem [36], Sadzik [32] explore the special case of trees with finitely many events, looking for epistemic bisimulations between different finite levels.)

The epistemic decoration of the tree models $Tree(\mathbb{M}, \mathbb{A})$ is rather special, since it obeys the above three constraints. Indeed, van Benthem & Liu [37] prove the following representation result:

THEOREM 26. *An epistemic tree model $\mathcal{M}$ is bisimilar to a model of the form $Tree(\mathbb{M}, \mathbb{A})$ if and only if it satisfies (a) Perfect Recall, (b) Uniform No Miracles, and (c) for any event e, the set of nodes where e can take place is closed under epistemic bisimulations inside $\mathbb{M}$.*

Thus, product update corresponds to a special epistemic temporal logic. Indeed, we can unpack the above conditions to the usual axioms for Perfect

Recall ($K_i[a]\phi \to [a]K_i\phi$) and No Miracles, where the uniform version (c) requires the use of universal modalities $E\phi$, $U\phi$ stating that ϕ holds at some world, at all worlds, resp. (cf. Blackburn, de Rijke & Venema [6]):

$$E(\langle a^-\rangle\top \wedge \neg K_i\neg\langle b^-\rangle\top) \to U(\langle a^-\rangle\neg K_i\neg\phi \to K_i[b^-]\phi)$$

Thus, dynamic-epistemic logic describes special idealized agents on epistemic temporal trees, much as discussed in preceding sections. But DEL has further special features. E.g., No Miracles is a much more plausible way of propagating ignorance than the usual No Learning, which seems to say that the passage of time never helps increase knowledge. Here, DEL gives a deeper analysis of the processes that drive information change, instead of merely describing the Grand Stage where all agents live in time. Even so, in the light of Section 3, the very decidability of DEL calls for explanation!

5.2 Why is *basic* DEL decidable? The ETL answer.

In the setting of Theorem 27, DEL-style languages are fragments of our earlier epistemic temporal ones:

THEOREM 27. *Over models $Tree(\mathbb{M}, \mathbb{A})$, DEL is the ETL language of epistemic logic plus one-step future operators $\langle e\rangle$.*

Thus, we do not reach the expressive power needed for undecidability in our earlier arguments. Indeed, for standard DEL, the reduction axioms translate every formula into an equivalent one without action modalities, and one then uses the decidability of the purely epistemic language. But something stronger holds. Let events have preconditions in the epistemic temporal language, as in communication scenarios which refer to the past of the current conversation. In its most blunt form, the precondition for e to occur is then just this: $\langle e\rangle\top$.

THEOREM 28. *The logic of the epistemic temporal language with only operators $\langle e\rangle$ on models with Perfect Recall and No Miracles is decidable.*

Proof. This can be derived from the decidability result for the modal 'product logic' $\mathbf{PDL} \times \mathbf{K}_m$ in Gabbay, Wolter et al. [11], by embedding our language into it. The models for this logic may be viewed as grids with an 'epistemic' **PDL** direction and a 'temporal' one-step **K**-direction. Still, no embedding of tiling problems is possible, because the language does not contain a true universal modality or transitive closure modality accessing all points of the grid (cf. also Marx & Mikulas [20]). ■

Thus DEL can be stronger than it is now, and still stay decidable. To get this positive result, we have appealed to one more tradition in the area of epistemic temporal structure, viz. combining logics through *product constructions*. This model-constructing approach fits quite well with DEL.

5.3 Program structure, true future, and undecidability

DEL style logics do become undecidable when the complete future is added. This would happen, e.g., if one adds sequential structure to action models, modeling, say, conversational processes involving composition and iteration. The landmark paper Miller & Moss [22] proves many relevant results, including this surprising effect of combining two decidable logics:

THEOREM 29 (Miller & Moss). *The dynamic-epistemic logic of public announcement with program iterations is undecidable.*

This shows that the undecidability phenomena already noted in Halpern & Vardi for the language with common knowledge and true future even occur in very restricted settings, where events are just announcements. Indeed, Miller & Moss show that iterated announcement of one single proposition $\Diamond\top$ suffices. However, in the light of our ETL-based Section 3, their analysis also leaves open questions. One of them has to do, again, with our assumption of finite levels. It is unknown whether their undecidability results hold when initial models are assumed to be finite.

5.4 Decidable fragments; tomorrow and yesterday

Our epistemic temporal analysis also suggests various additions to the standard language of DEL which still remain decidable. One typical illustration is the addition of the temporal past of an epistemic process (van Benthem [36]). Yap [42] and Sack [31] analyze such additions, and propose valid axioms. Here is the general setting.

THEOREM 30. *The logic of the epistemic temporal language with one-step future operators $\langle e\rangle$ and one-step past operators $\langle e^-\rangle$ on models with Perfect Recall and No Miracles is decidable.*

Proof. This follows from a simple modification of the above-mentioned decidability proof for $\mathbf{PDL}\times\mathbf{K}$ given in Gabbay, Wolter et al. [11]. ■

We conjecture that decidability still holds when we add many-step Past operators, at least, on our rooted finite-event trees. Such logics can express what agents knew earlier on, but they can also state preconditions for events that reach back in time, such as "say P if you have not already done so".

5.5 Protocols and model constructions

But there are still further features to the comparison of DEL and ETL. First, it is sometimes claimed that DEL lacks an essential resource available in our epistemic temporal models, viz. the choice of a *protocol*, i.e., a set of 'relevant histories'. Now, this is not true, since the above models $Tree(\mathbb{M},\mathbb{A})$ do have explicit restrictions on their available runs, since events can only occur when their precondition is satisfied. Thus, DEL has an explicit calculus for preconditions, as these are encoded in the action models, and these are again available inside the formal language through the modalities $\langle\mathbb{A},e\rangle\phi$. On the other hand, given the special epistemic format for these preconditions, one can only define special protocols, via local restrictions, that must be stated in a purely epistemic language. A more general approach here would merge the two ideas. On the one hand, it seems a good idea to make the protocols explicit in the language, as DEL does. On the other hand, one needs a richer repertoire of definitions for realistic protocols, including temporal operators in their formulation. This is achieved in the following 'Logic of Protocols' (cf. the earlier $\mathcal{L}_\Sigma(\mathcal{A})$):

We first introduce a **PDL** style language with 'protocols' defined by:

$$e \mid \alpha;\beta \mid \alpha \cup \beta \mid \alpha^* \mid \mathsf{skip} \mid \phi?$$

with $e \in \Sigma$ an event and ϕ a formula of $\mathcal{L}_{ETL}$. For example, the protocol $(e \cup \mathsf{skip})^*$ represents the set of histories that contain the event e. A test-free protocol α is a regular expression; and so it represents the set of histories that match α. A more interesting example is a 'Liar Protocol'. Let $\mathsf{send}(i, B, p)$ be the event "agent i sends the message p to the group of agents B", i.e., "i announces p to the set of agents B". Then $((K_i\neg p?; \mathsf{send}(i, \mathcal{A}, p)) \cup (K_i p?; \mathsf{send}(i, \mathcal{A}, \neg p)) \cup \mathsf{skip})^*$ represents a liar protocol. That is, if i knows p then i publically announces $\neg p$, if i knows $\neg p$ then i publically announces p, or i does not say anything.

For each protocol α introduce a modal operator N_α to the language. The intended interpretation of $N_\alpha\phi$ is that ϕ is true at the next moment in *all* extensions of the current history compatible with the protocol α. Thus truth is defined as $H, t \models N_\alpha\phi$ if $H' \in \mathcal{H}, H \preceq_\alpha H'$ and $H', t \models \phi$, where $\preceq_\alpha$ is an extension relation much like the previously defined R_α relations. This addition, though very useful in practice, is arguably a matter of convenience:

THEOREM 31. *ETL with explicit protocols is no more expressive than ETL by an effective translation.*

Introducing explicit protocols is also akin to the use of 'knowledge programs' in Fagin et al. [8]. We forgo the precise connection here.

Our conclusion is that older and newer approaches to dynamic actions and epistemic logic all meet in the same arena of epistemic temporal logic, and that insights can be transferred in illuminating ways.

6 Logics for Model Change

DEL is a calculus of piece-meal model construction, while ETL assumes the 'playground' or Grand Stage has been given as a temporal universe of all possible histories. An update $!P$ is then a minimal move to some available future state where one knows that P- and likewise for belief revision. In this setting, no definable explicit construction takes place for 'the next model' as in DEL - and it is the externally supplied temporal model which decides where things can go. This view is also that of *extensive games*.

As a logic of model construction, DEL, fits well with other modal logics, such as the London-style product logics mentioned before. But note that product update is more like 'direct product', and this shows that there is a large variety of possible constructions. One area where this has been studied in great generality, close to modal bisimulation-based paradigms, is *Process Algebra*, which deals with 'composition', 'choice', 'parallel composition', and 'iteration'. PA gives mainly algebraic calculi for process equivalence, but adding a temporal language would be a very natural step, allowing one to discuss what running a process actually achieves over time. Moreover, adding an epistemic component to current process theories would also make sense, witness the central role of *communication* between processes, as in Milner's recent work ([23]).

Model construction is also found in *game semantics* for logics and programming languages, cf. [25] and [1]. Here operations forming new games reflect the interactions between players. Interestingly, key structures in this case are again branching tree models. Neither Parikh nor Abramsky introduce ETL languages to talk about games, but it makes just as much sense as for the process theory. We can then describe players' activities over time, and define many of their strategies explicitly. As for complexity: the main concern in this paper, basic game logics tend to be decidable - but certain repertoires lead to undecidability: witness the ! of linear logic. We cannot go into details in the context of this paper, but again, we think that the parallels are striking, and worth systematic investigation.

7 Conclusions

This paper shows that epistemic temporal models are a natural meeting place for logicians. In Section 2, we defined basic structures that recur in most major studies of agents' interaction and information. In Section 3, we discussed the decidable/undecidable boundary where many interesting issues live concerning diversity of agent behaviour. This led to a natural merge of insights from different traditions: epistemic temporal logics in computer science, logics of computation, modal logics of products, dynamic-epistemic logics - and eventually also, process algebra and game semantics.

Concerning the relation between all these meeting frameworks, our view is this. Epistemic logics in the style of Fagin et al, and Parikh et al. are largely the same, even up to mutual mathematical representation (cf. [24]). The link between these systems and dynamic-epistemic logics is a bit more complex, but Section 4 has shown some natural merges. Thus, playing up differences between these approaches as different 'paradigms' seems both pointless from a mathematical viewpoint, and harmful from a conceptual or a practical point of view, as it impedes mutual flow of ideas.

Indeed, many further examples of mutual traffic can be found, which we had to leave out here. E.g., we have new results on more exciting scenarios where bounded and ideal agents meet. And, the explicit treatment of model constructions in dynamic-epistemic logic suggests new forms of *process algebra* for agents that can display rational intelligent behavior.

Thus, our approach suggests a turn from competition between frameworks to cooperation. Compare the situation in the 1930s, when many different models were proposed for computation. Instead of creating different churches, logicians started looking for similarities and equivalences (at some appropriate level), and the result was Church's Thesis, usually taken to mean that the field had a stable and mathematically respectable topic. Likewise, convergence, if not downright equivalence, between 'epistemic temporal logics' of agents might signal to a broader world that there is a core notion of genuine interest here concerning 'intelligent interaction', rather than a set of warring religions. Seeing differences may make for short-term gains, seeing analogies leads to a long-term common cause.

BIBLIOGRAPHY

[1] S. Abramsky and R. Jagadeesan. Games and full completeness for multiplicative linear logic. In R. Shyamsundar, editor, *Foundations of Software Technology and Theoretical Computer Science*, pages 291–301. Springer-Verlag, 1992.

[2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 2002.

[3] A. Baltag and L. Moss. Logics for epistemic programs. *Synthese: Knowledge, Rationality, and Action*, 2:165 – 224, 2004.

[4] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *Proceedings of TARK 1998*, 1998.

[5] N. Belnap, M. Perloff, and M. Xu. *Facing the Future*. Oxford University Press, 2001.

[6] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2002.

[7] E. A. Emerson. Temporal and modal logics. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*. 1990.

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

[9] T. French, R. van der Meyden, and M. Reynolds. Axioms for logics of knowledge and past time: Synchrony and unique initial states. In *Proceedings of AiML 2005*. King's College Press, to appear. appeared in the pre-proceedings of the Conference on Advances in Modal Logic, Manchester, Sept 2004.

[10] D. Gabbay. *Investigations into modal and tense logics, with applications to problems in linguistics and philosophy*. Reidel Dordrecth, 1976.

[11] D.M. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier, 2003.

[12] J. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, ILLC, 1998.

[13] V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, to appear 2006.

[14] Joseph Halpern. A computer scientist looks at game theory. *Games and Economic Behavior*, 45(1):114 – 131, 2003.

[15] Joseph Halpern and Moshe Vardi. The complexity of reasoning about knowledge and time. *J. Computer and System Sciences*, 38:195 – 237, 1989.

[16] David Harel. Recurring dominoes: Making the highly undecidable highly understandable. *Annals of Discrete Mathematics*, 24:51 — 72, 1985.

[17] I. Hodkinson and M. Reynolds. Temporal logic. In *Handbook of Modal Logic*. forthcoming.

[18] J. Horty. *Agency and Deontic Logic*. Oxford University Press, 2001.

[19] A. Kurucz. *Handbook of Modal Logic*, chapter Combining Modal Logics, page forthcoming. Elsevier, 2006.

[20] M. Marx and S. Mikulas. Products, or how to create modal logics of high complexity. *Logic Journal of the IGPL*, pages 77 –88, 2001.

[21] M. Marx and Y. Venema. Local variations on a loose theme: modal logic and decidability. In M. Vardi and S. Weinstein, editors, *Finite-Model Theory and Its Applications*, Texts in Theoretical Computer Science. Springer, 2005.

[22] J. Miller and L. Moss. The undecidability of iterated modal relativization. *Studia Logica*, 79(3), 2005.

[23] R. Milner. *Communicating and Mobile Systems: The Pi Calculus*. MIT Press, 1999.

[24] E. Pacuit. Some comments on history based structures. *Journal of Applied Logic*, forthcoming.

[25] R. Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 1985.
[26] Rohit Parikh and R. Ramanujam. Distributed processes and the logic of knowledge. In *Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 256 – 268. Springer, 1985.
[27] Rohit Parikh and R. Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12:453 – 467, 2003.
[28] J. Plaza. Logics of public communications. In *Proceedings, 4th International Symposium on Methodolgies for Intelligent Systems*, 1989.
[29] A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. Foundations of Computer Science*, pages 46 — 57, 1977.
[30] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1 – 35, 1969.
[31] J. Sack. Temporal language for epistemic programs. Unpublished manuscript, Indiana at Bloomington, 2006.
[32] T. Sadzik. Exploring the update universe. Unpublished manuscript, Stanford University, 2005.
[33] W. Thomas. Infinite trees and automaton definable relations over omega-words. *Theoretical Computer Science*, 103(1):143 – 159, 1992.
[34] Richmond Thomason and Anil Gupta. A theory of conditionals in the context of branching time. *The Philosophical Review*, 80:65–90, 1980.
[35] J. van Benthem. Games in dynamic epistemic logic. *Bulletin of Economic Research*, 53:216 – 248, 2001.
[36] J. van Benthem. One is a lonely number. Technical report, ILLC, University of Amsterdam, 2002.
[37] J. van Benthem and F. Liu. Diversity of logical agents in games. *Philosophia Scientiae*, 8(2):163 – 178, 2004.
[38] J. van Benthem and E. Pacuit. The tree of knowledge in action: Towards a common perspective. Technical report, ILLC, University of Amsterdam, 2006.
[39] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. In *Proceedings of TARK 2005*, 2005.
[40] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, forthcoming.
[41] M. Vardi. Why is modal logic so robustly decidable? DIMACS Workshop, 1996.
[42] A. Yap. Product update and looking backward. Unpublished manuscript, Stanford University, 2005.

Johan van Benthem
Plantage Muidergracht 24
Institute of Logic, Language and Computation
University of Amsterdam
`johan@science.uva.nl`

Eric Pacuit
Plantage Muidergracht 24
Institute of Logic, Language and Computation
University of Amsterdam
`epacuit@staff.science.uva.nl`