

PDT Logic for Stream Reasoning in Multi-agent Systems

Karsten Martiny and Ralf Möller

Institute for Software Systems
Hamburg University of Technology
{karsten.martiny, moeller}@tuhh.de

Abstract

We present *Probabilistic Doxastic Temporal (PDT) Logic for streams*, a formalism to reason about probabilistic beliefs and their infinite temporal evolution in multi-agent systems. Extending previous work on PDT Logic, this formalism builds on a Markov chain model to represent infinite streams of possible worlds. Within these streams, it enables the quantification of beliefs through probability intervals as well as the representation of temporal relations and epistemic actions. We show how agents can update their beliefs with respect to their observations, provide a model for infinite streams of possible worlds and show how we can map clippings of these streams to finite time windows. Based on these time windows, we introduce an adoption of the semantics of PDT Logic for finite time frames, and show how this provides a means to overcome the limitation of finite time domains.

1 Introduction

Epistemic and doxastic logics are used to reason about agents' knowledge. Formalizing the analysis of knowledge and belief through such logics has been an active topic of research in diverse fields such as philosophy [14], economics [2], game theory [11], and computer science [9]. Numerous extensions to modal epistemic logic have been made to reason about knowledge in multi-agent settings [9], [3], to add probabilistic knowledge [8], [5], and to analyze the dynamic evolution of knowledge [7].

In realistic scenarios an agent usually has only incomplete and inaccurate information about the actual state of the world, and thus considers several different situations as actually being possible. As it receives new information (e.g., it observes some facts), it has to update its beliefs about these possible worlds such that they are consistent with the new information. These updates can for example result in regarding some worlds as impossible, or judging some worlds to be more likely than before. Thus, in addition to analyzing the set of worlds an agent believes to be possible, it is also expedient to quantify these beliefs in terms of probabilities. This provides means to specify fine-grained distinctions between the range of worlds that an agent considers possible but highly unlikely, and worlds that seem to be almost certainly the actual world.

When multiple agents are involved in such a setting, an agent may not only have varying beliefs regarding the facts of the actual world, but also regarding the beliefs of other agents. In many scenarios, the actions of one agent will not only depend on its belief in facts of the actual world, but also on its beliefs in some other agent's beliefs.

Furthermore, an agent is usually also interested in temporal relationships and therefore also maintains beliefs about both the temporal evolution of the world as well as the temporal evolution of other agents' beliefs. To illustrate how reasoning about other agents beliefs can yield significant advantages in practical scenarios, we start with the following informal description of an example from the cyber security domain:

Example 1 (Cyber security). *Suppose that an adversary is trying to break into a computer system. This is usually done by using an attack graph to detect and exploit potential vulnerabilities of the system. An*

attack graph specifies a set of paths (i.e., sequences of actions) to carry out an attack. Several paths of the attack graph might be used in parallel, potentially by different agents (for instance, a number of infected computers controlled by a botnet). Usually, attack patterns specified by one attack graph are used multiple times, which has two important ramifications: the adversary will learn from experience which of the paths yield a high probability of successfully breaking into a system. Defenders in turn will be able to gain knowledge of the attack graph through the repeated observation of certain patterns. Thus, when a system is under attack, the defender will have beliefs about both the chosen attack paths and the adversary's belief regarding the success of the respective path. Thus, the defender can choose countermeasures effectively by reacting only on paths where these nested beliefs are high and which indeed pose a threat according to the system's mission impact model.

To formalize reasoning about such beliefs, we have developed *Probabilistic Doxastic Temporal (PDT) Logic* [17], which enables the representation of and reasoning about dynamically changing quantified temporal multi-agent beliefs through probability intervals. While it facilitates the expression of rich temporal relations, it is limited to finite time frames. The main contribution of this paper is an extension of our previous work to allow for models with an infinite evolution of time while maintaining the temporal expressivity of PDT Logic. Thus, we provide a means to integrate established stream reasoning systems (e.g., the well-known linear road benchmark [1]) and probabilistic multi-agent belief operators into a coherent framework. To illustrate the merit of this extension, consider the above example again: PDT Logic already provides means to formalize the above example, but it would require frequent re-initializations because the formalism is limited to finite time frames. Using the stream extension presented in this paper instead allows for a continuing analysis of current threads.

The remainder of this work is structured as follows: first, we start with a discussion of related work in the next section. Next, after summarizing the syntax in section 3, we define the formal semantics of PDT Logic for streams in Section 4. Then, we discuss an example in Section 5 and, finally, conclude the paper with Section 6.

2 Related Work

Approaches to formalize reasoning about knowledge and belief date back to Hintikka's work on epistemic logic [14]. Classical forms of epistemic logic do not allow for a quantification of an agent's degree of belief in certain facts; it can only be specified whether an agent does or does not know (resp. believe) some fact. To remove this limitation, combinations of logic and probability [6] resulted in numerous logics of knowledge and belief with probabilistic quantifications: [8] and [28] define a belief operator to quantify lower bounds on the probabilities that an agent assigns to a formula. Evaluation of these probabilities is done by evaluating this formula in every state the agent considers possible. [18] introduces Probabilistic Epistemic Logic and uses a similar belief operator to express lower bounds on the probabilities, but evaluates the probabilities by representing models as Bayesian networks.

To reason about dynamically changing knowledge and belief, extensions to epistemic and doxastic logics have been proposed, e.g. [23] (building on the situation calculus from [20]), [12], and [26]. In these works only the single-agent case is considered, and therefore they do not provide for representations of nested beliefs. Multi-agent extensions to these approaches can be found for example in [13], [4], [10], [27], and [7]. A common limitation of these works is that they are only able to reason about step-by-step changes. Explicit representation of and reasoning about time is difficult in these frameworks due to the lack of an explicit notion of time. [21] and [22] alleviate these limitations by combining Dynamic Epistemic Logic [7] with temporal modalities.

[24] and [25] introduce APT Logic, a framework to represent probabilistic temporal evolutions of worlds in finite-length threads. APT Logic assigns prior probabilities to every thread and uses these

probabilities to determine probabilities of events occurring in specific threads. To represent temporal relationships between events, APT Logic introduces the concept of frequency functions. APT Logic by itself does not provide any formalization of knowledge or belief. However, it provides a foundation to integrate probabilistic temporal reasoning into PDT Logic [17], our doxastic multi-agent framework that can explicitly reason about temporal relationships through temporal rules by an adoption of frequency functions.

3 PDT Logic programs: Syntax

In this section, we summarize the syntax of PDT Logic programs.

We assume the existence of a first order logic language with finite sets of constant symbols \mathcal{L}_{cons} , predicate symbols \mathcal{L}_{pred} , and an infinite set of variable symbols \mathcal{L}_{var} . Every predicate symbol $p \in \mathcal{L}_{pred}$ has an *arity*. A *term* is any member of the set $\mathcal{L}_{cons} \cup \mathcal{L}_{var}$. A term is called a *ground term* if it is a member of \mathcal{L}_{cons} . If t_1, \dots, t_k are (ground) terms, and p is a predicate symbol in \mathcal{L}_{pred} with arity n , then $p(t_1, \dots, t_k)$ with $k \in \{0, \dots, n\}$ is a (ground) atom. If a is a (ground) atom, then a and $\neg a$ are (ground) positive resp. negative *literals*. The set of all ground literals is denoted by \mathcal{L}_{lit} . As usual, \mathcal{B} denotes the Herbrand Base of \mathcal{L} , i.e., the set of all ground atoms that can be formed with predicates from \mathcal{L}_{pred} and terms from \mathcal{L}_{cons} .

Time is modeled in discrete steps. The set of time points is denoted by T and ranges over the set of natural numbers. Consequently, deviating from our previous work on PDT Logic [17], the temporal evolution of worlds is not limited to a finite time frame anymore.

The arbitrarily large, but finite set of agents is denoted by \mathcal{A} , and the number of agents ($|\mathcal{A}|$) is denoted by n . To describe what agents observe, we define observation atoms as follows:

Definition 1 (Observation atoms). For a group of agents $G \subseteq \mathcal{A}$ and a ground literal $l \in \mathcal{L}_{lit}$, $Obs_G(l)$ is an *observation atom*. \mathcal{L}_{obs} denotes the set of all observation atoms.

Intuitively, the meaning of a statement of the form $Obs_G(l)$ is that all agents in group G observe that fact l holds. Note that l may be a negative literal and therefore we can explicitly model observations of negative facts (such as “it is not raining”). We assume that the agents in G not only observe that l holds, but that each agent in G is also aware that all other agents in G make the same observation. To describe observations at specific time points, we use *time-stamped observation atoms* of the form $[Obs_G(l) : t]$, specifying that observation $Obs_G(l)$ occurs at time t .

Definition 2 (Formulae). Both atoms and observation atoms are formulae. If F and G are formulae, then $F \wedge G$, $F \vee G$, and $\neg F$ are formulae. A formula is ground if all atoms of the formula are ground.

To express temporal relationships between formulae, we define temporal rules following the approach of APT rules from [24].

Definition 3 (Temporal rules). Let F, G be two formulae, Δt a time interval, and fr a frequency function (describing temporal patterns, as defined in [24]). Then $r_{\Delta t}^{fr}(F, G)$ is called a temporal rule.

The meaning of such an expression is to be understood as “ F is followed by G in Δt time units w.r.t. the frequency function fr ”. For a given set of temporal rules \mathcal{R} , we refer to $\Delta_{t_{max}}$ as the largest time interval used in \mathcal{R} .

Now, we can define the belief operator $B_{it'}^{\ell u}$ to express agents’ beliefs. Intuitively, $B_{it'}^{\ell u}(\cdot)$ means that at time t' , agent i believes that some fact (\cdot) is true with a probability $p \in [\ell, u]$. We call the probability interval $[\ell, u]$ the *quantification* of agent i ’s belief. We use F_t to denote that formula F holds at time t .

Definition 4 (Belief formulae). Let i be an agent, t' a time point, and $[\ell, u] \subseteq [0, 1]$. Then, *belief formulae* (bf) are inductively defined as follows:

- If F is a formula and t is a time point, then $B_{it'}^{\ell u}(F_t)$ is a bf.
- If $r_{\Delta t}^{\text{fr}}(F, G)$ is a temporal rule, then $B_{it'}^{\ell u}(r_{\Delta t}^{\text{fr}}(F, G))$ is a bf.
- If F and G are bf, then so are $B_{it'}^{\ell u}(F)$, $F \wedge G$, $F \vee G$, and $\neg F$.

4 Semantics

In this section, we will provide the formal semantics for stream reasoning in PDT Logic. This semantics is an expansion of the semantics introduced in our previous work [17], such that it enables reasoning about infinite temporal evolutions.

4.1 Possible Worlds

We start with the definition of *worlds* (or *states* in the terminology of [9]). A world ω consists of a set of ground atoms and a set of observation atoms that hold in ω . With a slight abuse of notation, we use $a \in \omega$ and $Obs_G(l) \in \omega$ to denote that an atom a (resp. observation atom $Obs_G(l)$) holds in world ω . Since agents can only observe facts that actually hold in the respective world, we say that a world ω is *consistent* if all observed facts hold (i.e, $l \in \omega$ if l is a positive literal, and $l \notin \omega$ if l is a negative literal).

The set of possible worlds is denoted by $\Omega \subseteq 2^{\mathcal{B}} \times 2^{\mathcal{L}_{obs}}$. The size and configuration of Ω depends on the problem that is to be analyzed. We assume that the specified set of worlds fits the respective problem domain; especially, we assume in the following discussion that Ω does not contain any inconsistent worlds. We define satisfaction of a ground formula F by a world ω in the usual way [16]:

Definition 5 (Satisfaction of ground formulae). A ground formula F is satisfied by a world ω ($\omega \models F$)

- If $F = a$ for some ground atom a , then $a \in \omega$.
- If $F = \neg F'$ for some ground formula F' , then $\omega \not\models F'$.
- If $F = F' \wedge F''$ for formulae F' and F'' , then $\omega \models F'$ and $\omega \models F''$.
- If $F = F' \vee F''$ for formulae F' and F'' , then $\omega \models F'$ or $\omega \models F''$.

4.2 Streams

[24] introduces the concept of threads (equivalent to the concept of *runs* in [9]) as finite-length sequences of possible worlds. We extend this definition to infinite streams by removing the finite-length property:

Definition 6 (Stream). A *stream* is a mapping $St : T \rightarrow \Omega$.

Since T ranges over the set of natural number, a stream is an infinite possible sequence of worlds and $St(t)$ identifies the actual world at time t according to stream St .

In order to provide a bridge between our previously introduced PDT semantics for finite-time models [17] and the notion of infinite streams, we partition streams into fixed-length segments.

Definition 7 (Segment). A *segment* S with fixed length s is a finite sequence of possible worlds : $S : \tau \rightarrow \Omega$, $\tau = \{1, \dots, s\}$.

To identify the temporal position of specific segments, we enumerate segments with a parameter k , such that S_k represents the segment from time points $(k \cdot s) + 1, \dots, (k \cdot s) + s$, i.e., segment S_0 identifies the sequence of possible worlds from time $t = 1$ to $t = s$, segment S_1 identifies the sequence from $t = s + 1$ to $t = 2s$, and so on. This notion gives rise to an alternative representation of streams:

Lemma 1 (Stream segmentation). *Every stream of possible worlds can be equivalently represented as a sequence of segments $(S_k)_0^\infty$.*

To distinguish the notions of time points in segments and streams, we call t an *absolute* time point if $t \in T$ identifies the actual time point in some stream St . We call t a *relative* time point if $t \in \tau$ identifies a time point within a specific segment without specifying the segment's position in the stream.

We assume that all possible streams of the application domain can be modeled with an arbitrarily large, but finite set of possible segments $\mathcal{S} = \{S^1, \dots, S^u\}$. In most scenarios this requirement can be met through an appropriate modeling choice of the segment length s . In order to be able to reason with a set of temporal rules \mathcal{R} , we require that the segment length is set such that $\Delta_{t_{max}}(\mathcal{R}) \leq s$.

To model the transitions from one segment to the next, we model the possible segment sequences as an ergodic Markov chain [19]. An ergodic Markov chain is a transition system such that all states are aperiodic (i.e., returns to them can occur at irregular times) and positive recurrent (i.e., every state has a finite mean recurrence time):

Definition 8 (Markov chain). Let $\mathcal{S} = \{S^1, \dots, S^u\}$ be a set of possible segments. Then, transition probabilities between segments are given through a stochastic matrix $M^{|\mathcal{S}| \times |\mathcal{S}|}$ (i.e., m_{ij} specifies the probability of a transition from segment i to segment j), such that the segment sequence yields an ergodic Markov chain.

An important characteristic of an ergodic Markov chain is its steady state vector:

Definition 9 (Steady state vector). For an ergodic Markov chain, the vector π is called steady state vector, if its elements $\pi_j, (j = 1, \dots, |\mathcal{S}|)$ the following properties:

- $0 \leq \pi_j \leq 1$,
- $\sum_j^{|\mathcal{S}|} \pi_j = 1$,
- $\pi_i = \sum_i^{|\mathcal{S}|} \pi_i m_{ij}$

The last property expresses that the chain converges to the steady state vector in the long run (regardless of the starting state). Thus, π_j specifies the long-term visiting rate of segment S^j .

Due to the ergodicity, we immediately obtain a key property about the steady state vector π that we will exploit later (cf. [19]):

Lemma 2. *The long-term visiting rates π_j of all possible segments $S \in \mathcal{S}$ are nonzero.*

4.3 Time Windows

In order to be able to reason about temporal changes in the near future, we use a time window W_{t, l_w} starting at absolute time t and having a length of l_w time points, i.e., some finite clipping of the infinite stream of possible worlds. Again, the only restriction we place on the size of time windows is that




Figure 1: Schematic depiction of the relations between different temporal concepts.

they must be large enough to include all temporal rules (i.e., $l_w \geq \Delta_{t_{max}}(\mathcal{R})$), and thus they may well span across multiple segments. In a sense, a time window is a myopic representation of the stream. This artificial limitation of time is meaningful to capture the influence of current observations over the probabilities of contemporary events: While the probabilities of all events in the remote future will converge to the long-term visiting rates of the respective segments in the Markov chain, the probabilities of specific events within a certain time window may heavily depend on current observations and therefore significantly deviate from the long-term visiting rates. The different possible temporal evolutions within the time window (i.e., sequences of possible worlds for time $t, \dots, t + l_w$) are denoted by Th^1, \dots, Th^m and are called *threads* below. Equivalently to the notion of segments, we use $Th(1), \dots, Th(l_w)$ to identify the worlds of thread Th at the relative time points $1, \dots, l_w$. For notational convenience, we assume that every thread has an additional prior world $Th(0)$. The possible threads are determined through the set of possible segments \mathcal{S} and the transition matrix P of the Markov chain. Consequently, a thread in the time window W_{t, l_w} can be mapped onto a sequence of segments as

$$Th \rightarrow S_k^{j_1}, S_{k+1}^{j_2}, \dots, S_{k+x}^{j_x} \quad (1)$$

with $k = \lfloor (t/s) \rfloor$, $x = \lfloor (t + l_w - 1)/s \rfloor$, and $S_k^{j_1}, \dots, S_{k+x}^{j_x}$ (not necessarily distinct) segments from \mathcal{S} . Note that time windows may be positioned arbitrarily, i.e., the start of a time window does not necessarily have to coincide with the start of a new segment, nor does the window length l_w have to be an integer multiple of the segment length s .

The set of all possible threads within a time window W_{t, l_w} is denoted by \mathcal{T}_{t, l_w} . If the actual position and length of the time window are not important for the analysis, in the following we simplify notation and write \mathcal{T} to denote the set of threads to be analyzed. A visualization of these temporal concepts is shown in Figure 1.

4.4 Kripke Structures

With the definition of threads, we can use a slightly modified version of Kripke structures [15] to express an agent's possibility relations within the given time window. As usual, we define a Kripke structure as a tuple $\langle \Omega, \mathcal{K}_1, \dots, \mathcal{K}_n \rangle$, with the set of possible worlds Ω and binary relations \mathcal{K}_i on Ω for every agent $i \in \mathcal{A}$. Intuitively, $(\omega, \omega') \in \mathcal{K}_i$ specifies that in world ω , agent i considers ω' as a possible world.

Since we assume that both the set of possible worlds Ω and the set of possible segments \mathcal{S} are defined in a way that they do not yield any blatantly impossible threads, we initialize the Kripke structure such that a priori every agent considers all worlds possible that could occur at relative time $t = 1$ in some thread Th' :

$$\forall Th \in \mathcal{T} : \mathcal{K}_i(Th(0)) := \bigcup_{Th' \in \mathcal{T}} \{Th'(1)\}, \quad i = 1, \dots, n \quad (2)$$

With the evolution of time, each agent can eliminate the worlds that do not comply with its respective observations. Through the elimination of worlds, an agent will also reduce the set of threads it currently considers possible (if - due to some observation - a world ω is considered impossible at a relative time point t , then all threads Th with $Th(t) = \omega$ are considered impossible). We assume that agents have perfect recall and therefore will not consider some thread possible again if it was considered impossible at one point. This leads to the updates of \mathcal{K}_i w.r.t. the agent's respective observations:

$$\begin{aligned} \mathcal{K}_i(Th(t)) &:= \{Th'(t) : (Th'(t-1) \in \mathcal{K}_i(Th(t-1))) \wedge \\ &\{Obs_G(l) \in Th(t) : i \in G\} = \{Obs_G(l) \in Th'(t) : i \in G\}\} \end{aligned} \quad (3)$$

Remark 1. Note that these updates are defined w.r.t. the set of threads induced by a *specific time window* W_{t,l_w} . For other time windows (obtained for example by shifting the original window), the temporal scope changes and therefore, other sets of threads will be considered possible. This is due to the fact that a different temporal scope naturally yields different possible evolutions.

4.5 Subjective Posterior Temporal Probabilistic Interpretations

Each agent has probabilistic beliefs about the expected evolution of a world over time. This is expressed by subjective temporal probabilistic interpretations:

Definition 10 (Subjective posterior probabilistic temporal interpretation). Given a set of possible threads \mathcal{T} , some thread $Th' \in \mathcal{T}$, a relative time point $t \in \{1, \dots, l_w\}$ and an agent i , $\mathcal{I}_{it}^{Th'} : \mathcal{T} \rightarrow [0, 1]$ specifies the *subjective posterior probabilistic temporal interpretation* from agent i 's point of view at time t in thread Th' , i.e., a probability distribution over all possible threads: $\sum_{Th \in \mathcal{T}} \mathcal{I}_{it}^{Th'}(Th) = 1$. We call Th' the *point of view (pov) thread* of interpretation $\mathcal{I}_{it}^{Th'}$.

The prior probability of a thread Th is denoted by $\mathcal{I}_{i0}^{Th'}(Th)$. These priors are induced by the transition matrix M : since we allow for a placement of a time window at arbitrary positions within the stream, the probability of starting the thread in segment S^{j_1} is given through element π_{j_1} of the steady state vector. Then, the probability of segment S^{j_2} occurring next is given through the entry $m_{j_1 j_2}$ in the transition matrix M of the Markov chain, and so on. Thus, prior probabilities of threads can be obtained as:

Lemma 3. Let \mathcal{S} be a set of segments, Th be a thread mapped onto a sequence of x of these segments, M be the transition matrix of the Markov process, and π the corresponding steady state vector. Then, the prior probability of the thread Th can be computed as

$$\mathcal{I}_{i0}^{Th'}(Th) = \pi_{j_1} \prod_{i=1}^{x-1} m_{j_i j_{i+1}} \quad (4)$$

Since this expression builds on global properties of the Markov chain, the prior probabilities over the set of possible threads are identical across all possible pov threads and all agents (i.e., $\mathcal{I}_{i0}^{Th'}(Th) = \mathcal{I}_{j0}^{Th''}(Th) \forall Th, Th', Th'' \in \mathcal{T}$ and $\forall i, j \in \mathcal{A}$).

Using this lemma, we can treat arbitrary clippings from the infinite Markovian stream of possible worlds as finite threads. Thus, we have created a bridge to our previous work on PDT Logic for finite time frames. Consequently, we can simply adopt the subsequent semantics from [17]:

Definition 11 (Interpretation update). Let i be an agent, t a time point, and Th' a pov thread. Then, if the system is actually in thread Th' at time t , agent i 's probabilistic interpretation over the set of possible threads is given by the update rule:

$$\mathcal{I}_{it}^{Th'}(Th) = \begin{cases} \frac{1}{\alpha_{it}^{Th'}} \cdot \mathcal{I}_{it-1}^{Th'}(Th) & \text{if } Th(t) \in \mathcal{K}_i(Th'(t)) \\ 0 & \text{if } Th(t) \notin \mathcal{K}_i(Th'(t)) \end{cases} \quad (5)$$

with $\frac{1}{\alpha_{it}^{Th'}}$ being a normalization factor to ensure that $\sum_{Th \in \mathcal{T}} \mathcal{I}_{it}^{Th'}(Th) = 1$:

$$\alpha_{it}^{Th'} = \sum_{\substack{Th \in \mathcal{T}, \\ Th(t) \in \mathcal{K}_i(Th'(t))}} \mathcal{I}_{it-1}^{Th'}(Th)$$

Essentially, this rule assigns impossible threads a probability of zero and scales the probabilities of the remaining threads such that they are proportional to the probabilities of the previous time point.

Remark 1. We assume that the system is synchronous, i.e., the agents have a global clock. Thus, even if an agent does not observe anything in world $Th(t)$, it is still aware of time passing and can therefore distinguish between worlds $Th(t)$ and $Th(t-1)$.

4.6 Semantics of the Belief Operator

Now, with the definitions of subjective posterior probabilistic temporal interpretations, we can give the formal semantics for the quantified multi-agent belief operators as defined in [17]:

Definition 12 (Belief in ground formulae). Let \mathcal{T} be a set of threads, t and t' be relative time points in the time window specified through \mathcal{T} , and $\mathcal{I}_{it'}^{Th'}$ be agent i 's interpretation at time t' in pov thread Th' . Then, it follows from this interpretation that agent i believes at time t' with a probability in the range $[\ell, u]$ that

1. A formula F holds at time t (denoted by $\mathcal{I}_{it'}^{Th'} \models B_{it'}^{\ell u}(F_t)$) iff:

$$\mathcal{I}_{it'}^{Th'} \models B_{it'}^{\ell u}(F_t) \quad \text{iff} \quad \ell \leq \sum_{Th \in \mathcal{T}, Th(t) \models F} \mathcal{I}_{it'}^{Th'}(Th) \leq u. \quad (6)$$

2. A temporal rule¹ $r_{\Delta t}^{\text{fr}}(F, G)$ holds (denoted by $\mathcal{I}_{it'}^{Th'} \models B_{it'}^{\ell u}(r_{\Delta t}^{\text{fr}})$) iff:

$$\ell \leq \sum_{Th \in \mathcal{T}} \mathcal{I}_{it'}^{Th'}(Th) \cdot \text{fr}(Th, F, G, \Delta t) \leq u. \quad (7)$$

3. A belief $B_{jt}^{\ell_j u_j}(\cdot)$ of some agent (denoted by $\mathcal{I}_{it'}^{Th'} \models B_{jt}^{\ell_j u_j}(B_{jt}^{\ell_j u_j}(\cdot))$) holds iff:

$$\ell \leq \sum_{\substack{Th \in \mathcal{T}, \\ \mathcal{I}_{jt}^{Th} \models B_{jt}^{\ell_j u_j}}} \mathcal{I}_{it'}^{Th'}(Th) \leq u. \quad (8)$$

¹See [17] for a definition of temporal rules $r_{\Delta t}^{\text{fr}}(F, G)$ and frequency functions $\text{fr}(Th, F, G, \Delta t)$. We have included the definition of beliefs in temporal rules only for the sake of completeness, it is not required to follow the remainder of this paper.

Remark 2. *Agent i does not know the actual beliefs of agent j . However, due to the assumption of common and equal priors, agent i is able to reason about agent j 's hypothetical updates given that the system is in a specific thread. Thus, agent i is able to compute (8) without knowing j 's exact beliefs.*

Using these definitions in combination with Lemma 3 allows us to employ the techniques from [17] to analyze the temporal evolution of arbitrary belief expressions within streams of possible worlds.

5 Example

To illustrate the application of PDT Logic for streams, consider the following simplified example: a taxi driver A is positioned at the central station and is occupied with driving two different routes from there, either to the airport (requiring one time step) or to the stadium (requiring two time steps). These routes of the taxi driver are depicted in Figure 2(a). Furthermore, we assume that the driver can take a break at the stadium every once in a while and inform her boss B about it. This means that both A and B know that A is at the stadium (and they are both aware that the other agent knows this), so we can model the effect of this as a shared observation $Obs_{AB}(stadium)$. A possible segmentation of the infinite stream of possible worlds into segments could use segments of length three. This yields the set \mathcal{S} of possible 3-step-sequences shown in Figure 2(b). If we assume that driving to the airport has a probability of, say, 0.7, and accordingly, the probability of driving to the stadium is 0.3, we obtain the Markov chain depicted in Figure 2(c) with the corresponding transition matrix shown in Figure 2(d).

One can easily verify that this yields the longterm visiting rates π_i of Segments S^i :

$$\pi \approx (.172 \ .074 \ .245 \ .150 \ .105 \ .105 \ .032 \ .074 \ .045) \quad (9)$$

To show the evolution of the agents' beliefs, consider a time window of length $l_w = 6$, i.e., a sequence of two segments. Assume that we are interested in the beliefs of, say, A being at the stadium (i.e., $F \equiv stadium$) at relative time $t = 4$, denoted by $W(4) \models F$. Analysis of the set of segments \mathcal{S} shows that $W(4) \models F$ iff the second segment of the considered time window is S^5 . Thus, the probability of being at the stadium at relative time point 4 is equivalent to the property of segment S^5 occurring as the second segment within the time window. Hence, we can obtain the prior belief of both agents in this event from the steady state vector (through $\pi_2 \cdot m_{25} + \pi_4 \cdot m_{49} + \pi_7 \cdot m_{75}$) as

$$B_{i0}^{0.15 \ 0.16}(F_4), i \in \{A, B\}.$$

Now suppose that time evolves for 3 steps and that the actual events are the ones depicted in S^4 . Then, driver A can update her belief in the event F_4 to

$$B_{A3}^{0.3 \ 0.3}(F_4)$$


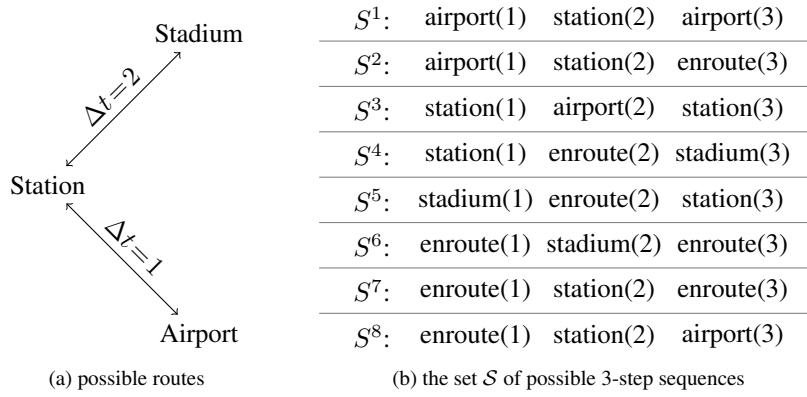
(if she is on her way to the stadium at $t = 3$, she knows that she will be there at $t = 4$ if she takes a break), while B (lacking any new information) maintains the same belief expressed at time 0. Since A is aware that B did not receive new information,

$$B_{A3}^1(B_{B3}^{0.15 \ 0.16}(F_4))$$

also holds.

Finally, consider the situation at $t = 4$: if A decides to take a break, she informs B (represented in S^9), and consequently,

$$B_{i4}^1(F_4), i \in A, B$$



$$M = \begin{pmatrix} 0 & 0 & .7 & .3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ .49 & .21 & 0 & 0 & 0 & .3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & .21 & .49 & .3 \\ .49 & .21 & 0 & 0 & 0 & .3 & 0 & 0 & 0 \\ 0 & 0 & .7 & .3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & .7 & .3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

(d) transition matrix

Figure 2: Taxi driver example: (a) depicts the possible routes, (b) all possible 3-step sequences, (c) the resulting possible segment transitions, and (d) the corresponding transition matrix.

holds. Else, if A does not take a break, her beliefs are updated to

$$B_{A4}^{0,0}(F_4),$$

while B 's beliefs (through elimination of possibility S^9) are updated to

$$B_{B4}^{0,10,0,11}(F_4).$$

This small example shows how a stream of possible worlds can be used to generate an infinite stream of evolving multi-agent beliefs. Furthermore, it shows how the individual beliefs of agents can diverge depending on their respective information. More detailed examples on the application of PDT Logic including temporal rules can be found in [17].

6 Conclusion

In this paper, by extending our previous work on PDT Logic to infinite time domains, we have developed a general framework to reason about the belief change in multi-agent stream scenarios. Although the actual reasoning about temporal beliefs is carried out for finite windows of time, arbitrary placements of these windows within infinite streams of possible worlds enable unlimited temporal reasoning in stream domains. For example, one could employ a sliding window mechanism to provide for continuous reasoning about the evolution of beliefs. As described in the introductory example, using an infinite stream of possible worlds as an input to the formalism can be beneficial because it allows for a continuing analysis of changing situations without requiring frequent re-initializations. While the employed model of infinite streams of possible worlds based on a Markov chain restricts the application domains to scenarios that can be modeled through recurring segments, this still allows for a wide range of applications because the properties of most domains can be captured with a suitable selection of length and number of segments.

PDT Logic for streams as introduced in this work provides the foundation for future work. While we have shown that PDT Logic for streams can be analyzed with the methods introduced previously for the finite-time formalism PDT Logic, we will continue to investigate optimized decision procedures, using both exact and approximate methods. First promising approaches are given through a combination of established Markov chain inference algorithms with the APT Logic decision algorithms described in [25].

References

- [1] A. Arasu, M. Cherniack, E. Galvez, D. Maier, A. S. Maskey, E. Ryzkina, M. Stonebraker, and R. Tibbetts. Linear road: A stream data management benchmark. In *Proceedings of the Thirtieth International Conference on Very Large Data Bases*. VLDB Endowment, 2004.
- [2] R. Aumann. Agreeing to Disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [3] A. Baltag and L. Moss. Logics for Epistemic Programs. *Synthese*, 139(2):165–224, 2004.
- [4] F. Belardinelli and A. Lomuscio. A Quantified Epistemic Logic for Reasoning About Multiagent Systems. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '07, pages 87:1–87:3, New York, NY, USA, 2007. ACM.
- [5] M. W. Cripps, J. C. Ely, G. J. Mailath, and L. Samuelson. Common Learning. *Econometrica*, 76(4):909–933, 2008.
- [6] L. Demey, B. Kooi, and J. Sack. Logic and probability. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Center for the Study of Language and Information, Stanford University, spring 2013 edition, 2013.

- [7] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2007.
- [8] R. Fagin and J. Y. Halpern. Reasoning about Knowledge and Probability. *Journal of the ACM*, 41:340–367, 1994.
- [9] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [10] L. Fang and Y. Liu. Multiagent Knowledge and Belief Change in the Situation Calculus. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, Bellevue, Washington, USA*, 2013.
- [11] J. C. Harsanyi. Games with Incomplete Information Played by ‘Bayesian’ Players. *Management Science*, 14:159–182, 320–324, 486–502, 1967–68.
- [12] A. Herzig. Action Representation and Partially Observable Planning using Epistemic Logic. In *In Proc. IJCAI03*, pages 1067–1072. Morgan Kaufmann, 2003.
- [13] A. Herzig, J. Lang, and P. Marquis. Action Progression and Revision in Multiagent Belief Structures. In *Sixth Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC)*, page 05, 2005.
- [14] J. Hintikka. *Knowledge and Belief*. Ithaca, N.Y., Cornell University Press, 1962.
- [15] S. A. Kripke. Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16(1963):83–94, 1963.
- [16] J. W. Lloyd. *Foundations of Logic Programming, 2nd Edition*. Springer, 1987.
- [17] K. Martiny and R. Möller. A Probabilistic Doxastic Temporal Logic for Reasoning about Beliefs in Multiagent Systems. Technical report, <http://www.sts.tuhh.de/tech-reports/papers.html>, February 2014.
- [18] B. Milch and D. Koller. Probabilistic models for agent’s beliefs and decisions. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence, UAI ’00*, pages 389–396, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.
- [19] J. R. Norris. *Markov chains*. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, 1998.
- [20] R. Reiter. The Frame Problem in Situation the Calculus. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation*, pages 359–380. Academic Press Professional, Inc., San Diego, CA, USA, 1991.
- [21] B. Renne, J. Sack, and A. Yap. Dynamic Epistemic Temporal Logic. In X. He, J. Horty, and E. Pacuit, editors, *Logic, Rationality, and Interaction: Second International Workshop, LORI 2009, Chongqing, China, October 8–11, 2009, Proceedings (LORI-II)*, volume 5834 of *Lecture Notes in Computer Science*, pages 263–277. Springer-Verlag Berlin Heidelberg, 2009.
- [22] J. Sack. Extending probabilistic dynamic epistemic logic. *Synthese*, 169(2):241–257, 2009.
- [23] R. B. Scherl and H. J. Levesque. Knowledge, Action, and the Frame Problem. *Artificial Intelligence*, 144(1-2):1–39, March 2003.
- [24] P. Shakarian, A. Parker, G. Simari, and V. S. Subrahmanian. Annotated Probabilistic Temporal Logic. *ACM Trans. Comput. Logic*, 12(2):14:1–14:44, January 2011.
- [25] P. Shakarian, G. Simari, and V. S. Subrahmanian. Annotated Probabilistic Temporal Logic: Approximate Fixpoint Implementation. *ACM Trans. Comput. Logic*, 13(2):13:1–13:33, April 2012.
- [26] T. C. Son and C. Baral. Formalizing sensing actions A transition function based approach . *Artificial Intelligence*, 125(12):19 – 91, 2001.
- [27] J. van Benthem, J. Gerbrandy, and B. P. Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96, 2009.
- [28] W. van der Hoeck. Some considerations on the logics PFD: A logic combining modality and probability. *Journal of Applied Non-Classical Logics*, 7(3):287–307, 1997.