CS-TR-3249 | March 30, 1994
UMIACS-TR-94-39

Thinking Takes Time: A Modal Active-logic for Reasoning IN Time

Madhura Nirkhe
Sarit Kraus
Don Perlis

Institute for Advanced Computer Studies
Computer Science Department
AV Williams Bldg
University of Maryland
College Park, MD 20742

madhura@cs.umd.edu
sarit@umiacs.umd.edu
perlis@cs.umd.edu

Abstract

Most common sense reasoning formalisms do not account for the passage of time as the reasoning occurs, and hence are inadequate from the point of view of modeling an agent's *ongoing* process of reasoning. We present a modal active-logic that treats time as a valuable resource that is consumed in each step of the agent's reasoning. We provide a sound and complete characterization for this logic and examine how it addresses the problem of logical omniscience.

Work done with partial support from NSF grant IRI9311988.

Thinking takes time: A modal active-logic for reasoning in time*

Madhura Nirkhe[1] Sarit Kraus[4,2] and Donald Perlis[2,3]

[1]Department of Electrical Engineering
University of Maryland, College Park, MD 20742
[2]Institute for Advanced Computer Studies
University of Maryland, College Park, MD 20742
[3]Department of Computer Science
University of Maryland, College Park, MD 20742
[4]Department of Mathematics and Computer Science
Bar Ilan University Ramat Gan, 52900 Israel

madhura@cs.umd.edu, sarit@bimacs.cs.biu.ac.il, perlis@cs.umd.edu

Abstract

Most common sense reasoning formalisms do not account for the passage of time as the reasoning occurs, and hence are inadequate from the point of view of modeling an agent's *ongoing* process of reasoning. We present a modal active-logic that treats time as a valuable resource that is consumed in each step of the agent's reasoning. We provide a sound and complete characterization for this logic and examine how it addresses the problem of logical omniscience.

*This work was supported in part by NSF grant No. IRI-9123460. The first author wishes to thank Carolina Ruiz for improving the mathematical rigor of the paper.

1 Introduction

All agents, whether human or automated, that function in the real-world are subject to the fact that time is spent as their reasoning progresses. Most common sense reasoning formalisms do not account for the passage of time as the reasoning occurs, and hence are inadequate from the point of view of modeling an agent's *ongoing* process of reasoning. There are numerous problems in AI-planning and common sense reasoning where the capacity to reason and act *in* time is of paramount importance. Below is a list of few sample problems in which the passage of time (as the agent reasons) is crucial:
1. *Nell & Dudley and the railroad tracks* : Nell is tied to the railroad tracks and the agent Dudley must figure out and enact a plan to save her in time before an oncoming train approaches.
2. *Examination problem* : A student who is taking an examination must figure out a strategy to decide which problems to work on, how much time to allocate to each, etc. And yet, every second spent in such decision making is a second less to actually solve the problems. Deliberation time is a significant chunk of the total time available which the agent must factor in the reasoning.
3. *The three wise men problem* : The well known puzzle [McC78, ED88, KL88] where three wise men, each wearing a cap from a pool of one white and two black caps, are lined up and each is asked to announce the color of the cap on his own head as soon as he knows what it is. The agents must be able to keep track of the time spent in reasoning to solve this problem realistically.

Most formal approaches do not have an appropriate representational framework to tackle time-situated reasoning problems such as the above. They assume that an agent is able to reason forever in a timeless present as if the world had stopped for the agent's benefit. Resource limitations have been of some concern in formal work. In particular, the problem of *logical omniscience* has received attention in the epistemic logic literature. It concerns the difficulty with the classical Hintikka possible world semantics [Hin62] that the agent always knows the logical consequences of her beliefs. However, no existing works provide a semantics addressing the issue of how the reasoning progresses vis a vis the passage of time. Although work in temporal logic involves reasoning *about* time (e.g., [All84, McD82, ER87]), time is not treated as a crucial resource that must be carefully rationed by the agent, as it is spent in every step of reasoning.

Step-logics [EDP90, PEDM] were introduced as a formal apparatus to model an agent's ongoing process of reasoning. They have since been extended and renamed as *active logics*. In [KNP90, NKP91, NKP93] the step-logic framework is used to create an active logic based *fully* deadline-coupled planning and reasoning mechanism which is a combination of declarative and procedural approaches that is capable of solving the above mentioned problems[1]. Although active logics have been characterized and implemented, only limited attempts have been made to give a formal semantics for the step-like reasoning process.

This paper is intended to bridge the gap between previous modal approaches to knowledge and belief and time-situated frameworks such as step-logics which have a means for attributing time to the reasoning process. We discuss the various aspects of logical omniscience and their treatment in section 2. We briefly describe the step-logic approach to reasoning in section 3. In section 4 we present a modal active-logic that is step-like in spirit and motivated by the work on step-logics, but for which, unlike the former, we can provide a sound and complete modal semantics in section 4.2. In section 5 we examine how our approach addresses the logical omniscience problem and summarize our contribution.

2 The various aspects of omniscience and its treatment

Fagin and Halpern [FH88] have analyzed what is meant by the notion of *logical omniscience*. They define an agent to be logically omniscient if whenever he believes formulas in a set Σ, and Σ logically implies the formula ϕ, then the agent also believes ϕ. They further identify three cases of special interest: (a) *closure under implication*, namely, whenever both ϕ and $\phi \rightarrow \psi$ are believed then ψ is believed, (b) *closure under valid implication*, namely, if $\phi \rightarrow \psi$ is valid and ϕ is believed then ψ is believed and (c) *belief of valid*

[1]An active logic for fully deadline-coupled planning has several inference rules for planning with deadlines that are domain independent. These include inference rules for temporal projection and book-keeping, checking deadline feasibility, and plan formulation and execution. Domain specific axioms describe the particular instance of the planning problem. A limited-memory model that addresses two other resources of value: space and parallelism has been integrated into the deadline-coupled reasoning.

formulas, namely, if ϕ is valid, then ϕ is believed.

The agent in the classical model of knowledge [Hin62] has all the undesirable properties (a), (b) and (c) above. Several improvements have been suggested, and they have been broadly classified as "syntactic" and "semantic" approaches. In the syntactic approach e.g. [Ebe74, MH79], what the agent knows is represented by a set of formulas and hence is not constrained under consequence. But such approaches are difficult to analyze, since they are not guided by knowledge-based principles. A commendable syntactic approach is presented by Konolige in his *deduction model* [Kon83] which gives a formal characterization of explicit beliefs and captures how agents syntactically derive new beliefs, possibly with an incomplete set of inference rules.

In contrast, semantic approaches attempt to give semantics similar in most cases to the possible world semantics, but with "fixes". Levesque [Lev84] gives a semantic account of *implicit* and *explicit* belief where implicit beliefs are the logical consequences of explicit belief. A solution to (a) and the possibility of having contradictory beliefs is achieved by introducing an artificial notion of *incoherent or impossible worlds*. Levesque's approach was subject to the criticism that an agent in the logic is a perfect reasoner in relevance logic. Levesque's ideas have been extended in [PS85] and [Lak86]. Montague has given a possible world semantics that gets around problem (a) of logical consequence. We use the main idea in this model, namely, to define knowledge as a relation between a world and a set of sets of possible worlds. However, we provide the distinction of incorporating time-situatedness. Vardi [Var86] provides a co-relation between restrictions on models in the Montague semantics and the corresponding agent properties that they characterize.

Fagin and Halpern [FH88] have presented a series of interesting approaches to limited reasoning that marry the syntactic and semantic approaches. They provide an extension to Levesque's approach for the multi-agent case, and introduce a notion of *awareness*. They also provide an approach to local reasoning that they call a *society of minds approach*. Fagin and Halpern's awareness notion, in their logic of general awareness acts like a filter on semantic formulations. It has been evaluated and criticized in [Kon86]. One of the criticisms is that the model is unintuitive, since it is unlikely that an agent can compute all logical consequences, discarding the one's that it is not aware of, say, because of memory limitations, because in fact, agents are also affected by time limitations. There are a number of works that have consid-

$$
\begin{array}{llcl}
0: & & \emptyset & \\
\vdots & & & \\
i: & \ldots & \alpha & \ldots \\
i+1: & \ldots & \alpha \rightarrow \beta, \beta \rightarrow \gamma & \ldots \\
i+2: & \ldots & \beta & \ldots \\
i+3: & \ldots & \gamma & \ldots \\
\vdots & & &
\end{array}
$$

Figure 1: Step-logic studies

ered logics of knowledge and time e.g. [Sat77, Leh84, KL88, LR86, Ash88]. Fagin and Halpern discuss the possibility of capturing bounded and situated reasoning by letting the awareness set vary over time. However, no attempt has been made to systematically study and model situations where the passage of time is a critical issue.

3 The step-logic approach to reasoning

Step-logics [EDP90] were introduced to model a common sense agent's ongoing process of reasoning in a changing world. A step-logic is characterized by a language, observations and inference rules. A *step* is defined as a fundamental unit of inference time. Beliefs are parameterized by the time taken for their inference, and these time parameters can themselves play a role in the specification of the inference rules and axioms. The most obvious way time parameters can enter is via the expression $Now(i)$, indicating the time is now i. Observations are inputs from the external world, and may arise at any step i. When an observation appears, it is considered a belief in the same time-step. Each step of reasoning advances i by 1. At each new step i, the only information available to the agent upon which to base his further reasoning is a snap-shot of his deduction process completed up to and including step $i-1$. Figure 1, adapted from [ED88] illustrates three steps in a step-logic with Modus Ponens as one of its inference rules.

Elgot-Drapkin also characterized an array of eight step-logics in increasing order of sophistication with respect to three mechanisms : self-knowledge

$$\frac{i : \ldots \quad \alpha,\beta \quad \ldots}{i{+}1 : \ldots \quad \alpha\wedge\beta \quad \ldots} \qquad \text{Conjunction}$$

$$\frac{i : \ldots \quad \alpha\wedge\beta \quad \ldots}{i{+}1 : \ldots \quad \alpha \quad \ldots} \qquad \text{Detachment}$$

$$\frac{i : \ldots \quad \alpha \quad \ldots}{i{+}1 : \ldots \quad \alpha \quad \ldots} \qquad \text{Inheritance}$$

$$\frac{i : \ldots}{i{+}1 : \ldots \quad \alpha_{obs} \quad \ldots} \qquad \text{Observation}$$

Figure 2: Inference rules for an SL_5 logic

(S), time (T) and retraction (R)[2]. According to this classification, SL_5 is the simplest dynamic deductive logic with time and self-knowledge capability, but no retraction mechanism (no ability to handle contradictions). An SL_5 logic is a triple $\langle \mathcal{L}, OBS, INF \rangle$ where $\mathcal{L}$ consists of propositions (with the addition of time), OBS is an observation function describing inputs from the world at each step, and INF is a set of inference rules. We describe an SL_5 step-logic to which we provide a modal active-logic analog. The set INF for it is shown in in figure 2.

We have chosen the simplest SL_5 since our main interest is in the treatment of time and in modeling agents with nested beliefs. We will impose an additional constraint on models that does not allow for contradictions in the agent's beliefs[3].

4 A modal active-logic for reasoning *in* time

With SL_5 as the motivation, we provide a time-situated modal logic. This modal logic is based on Montague's intensional logic of belief [Mon70], that uses structures referred to in the literature as *neighborhood structures* or *minimal structures* [Che80]. They were first used in [Mon68] and in [Sco70].

[2] SL_0:none; SL_1:S; SL_2:T; SL_3:R; SL_4:S,R; SL_5:S,T; SL_6: R,T; and SL_7:S,T,R.

[3] However, this condition may be relaxed if for example, we desire to model an agent with default reasoning capability. Step-logics are inherently nonmonotonic and allow for implicit and explicit contradictions in the agent's reasoning. The modal logic approach which is motivated by the step-logic work is powerful enough to deal with contradictions.

Montague gives a possible world semantics to epistemic logic where, unlike in the classical model[4], knowledge is defined as a relation between a world and a set of sets of worlds. An *intension* of a formula ϕ denoted by $||\phi||$ is the set of worlds w such that $w \models \phi$.

We prefer to use *timelines* instead of possible worlds, since this gives us a way to naturally incorporate time into our framework. L denotes the set of timelines [TSSK91]. We consider time lines that are restricted to be finite from one side and infinite from the other (i.e., are rays). At every time point in each timeline some propositions are true and the rest are false. In particular, there is one timeline of most interest, that captures the *real* history of occurrences in the world. We call this line $l_h \in L$ the *history timeline*.

4.1 Syntax and semantics

In the logic proposed, the agent reasons in a propositional language with time. The interest is in sentences such as:
p: Nell is tied to the railroad tracks at *3 pm*.
q: Dudley is at home at *3:30 pm*.

Formally, we assume that there is a set P of propositions and a set TC of time point constants. We define $PT = P \times TC$ as the set of propositions extended to include time arguments. The formulas in PT are the basic elements of our language, and we will denote them as $p(\tau)$ where $p \in P$ and $\tau \in TC$. The language $\mathcal{G}$ is the smallest set that contains PT, and is closed under the $\neg, \wedge, \vee, \rightarrow$ connectives, and contains $B_\tau \phi$ whenever ϕ is in the language and $\tau \in TC$[5]. This language can easily be extended to include multiple agents, by the use of an additional parameter i, so that $B^i_\tau \alpha$ denotes "at time τ agent i believes in α", where α may include beliefs of other agents.

[4]The classical possible-worlds model is based on the idea that besides the true world, there are other possible worlds, some of which may be indistinguishable to the agent from the true world. An agent is said to *believe* a fact ϕ if ϕ is true in all the worlds that she thinks possible. A semantics based on Kripke structures for this classical model suffers from the well known drawback from the point of view of logical omniscience that $K\phi \wedge K(\phi \rightarrow \psi) \rightarrow K\psi$ is an inherent axiom.

[5]In this language one can express formulas such as p and q above, belief formulas such as $B_{\tau_1} p(\tau_2)$ to mean "at time τ_1 the agent believes that p is true at time τ_2", or nested beliefs formulas such as $B_{\tau_1}(B_{\tau_1+2} p(\tau_2) \vee B_{\tau_1+2} q(\tau_3))$ to mean "at time τ_1 the agent believes that two time points later she will believe $p(\tau_2)$ or she will believe $q(\tau_3)$".

Time is a pair $\langle T, \prec \rangle$ where $\prec$ is a total order on T. A structure in the proposed time-embedded active logic is: $M = \langle L, T, v, \prec, \pi, \mathcal{B}, \mathrm{OBS} \rangle$ where

- L is the set of timelines, $\langle T, \prec \rangle$ is a time structure.
- $v : TC \rightarrow T$ is the interpretation function for time point constants[6].
- $\pi : P \times T \rightarrow 2^L$ is a truth assignment to the formula $p \in P$ for each timeline $l, l \in L$ and time point $t \in T$. Thus π defines the intensions of the base formulas of our language.
- $\mathcal{B} : L \times T \rightarrow 2^{2^L}$ is a belief accessibility relation, defined for each timeline, time point pair $(l, t), l \in L, t \in T$.
- $\mathrm{OBS} : L \times T \rightarrow \mathcal{G}$ is the observation function.

We will use $\mathcal{B}_t(l)$ to denote $\mathcal{B}(l, t)$, which is the set of sets of time lines related to l at time t through the $\mathcal{B}$ relation. Note the use of the pair (l, t). We are interested in epistemic behavior over time, and this is depicted by the evolution of beliefs (and the corresponding accessibility relations) from (l_h, t) to $(l_h, t+1)$ in the real timeline.

Analogous to the Montague intensional logic, we define $B_\tau \phi$ to denote that an agent "believes a formula ϕ at time τ" and define a satisfiability relation for timelines based on intensions. An intension of a formula ϕ in a structure M denoted by $||\phi||$ is $\{l \mid l \in L,\ M, l \models \phi\}$.

Figure 3 illustrates the neighborhood structures for our modal logic[7].

We impose restrictions on models to reflect the step-like reasoning behavior between successive time instances. These restrictions make certain axioms sound in our system. We further characterize the modal active-logic

[6] For all $\tau, \tau' \in TC$ if $v(\tau) \prec v(\tau')$ and there is no $\bar{t} \in T$ such that $v(\tau) \prec \bar{t} \prec v(\tau')$ we will use $\tau + 1$ to denote τ'. Similarly, for $t, t' \in T$ where $t \prec t'$ and there is no such $\bar{t}$, we use $t + 1$ to denote t'.

[7] We comment here that is is possible to extend the modal active-logics to multiple agents reasoning in time. A structure for an active logic with multiple agents is $M = \langle L, T, v, \pi, \mathcal{A}, \mathcal{B}^1, \ldots, \mathcal{B}^n \rangle$. L is the set of time lines and π is the truth assignments to base formulas as before. $\mathcal{A}$ is the set of agents $\{1, \ldots, n\}$, and each of $\mathcal{B}^i, i = 1, \ldots, n$ associates with a timeline, time point pair (l, t), a set of set of timelines that are belief-accessible from l at time t from the perspective of agent $i \in \mathcal{A}$. In problems such as the three wise men problem mentioned in the introduction, a multi-agent logic where the time of all agents increments synchronously can provide an elegant solution to the problem.

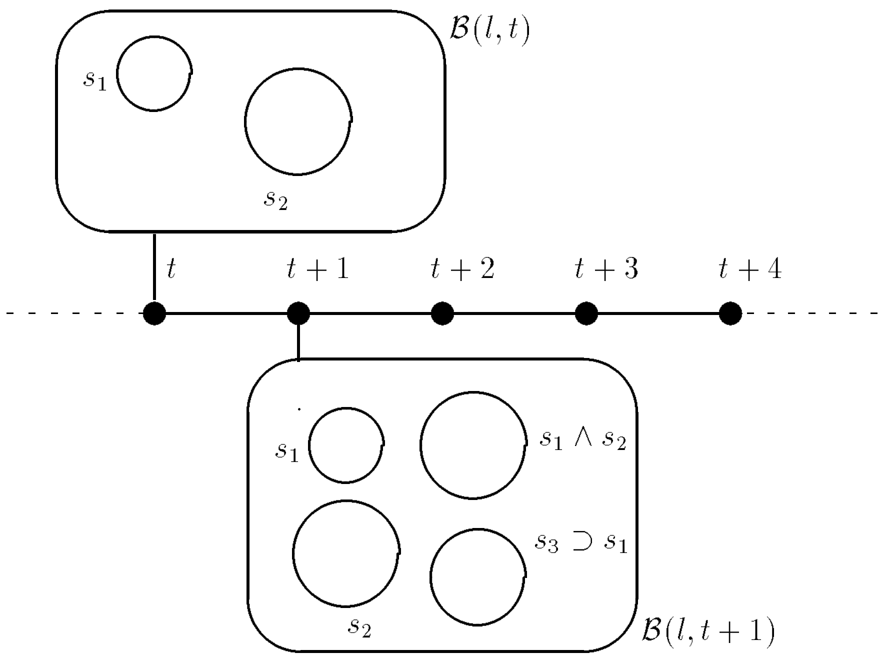

Figure 3: Neighborhood structures for the belief operator

by a sound and complete set of axioms and inference rules. Time is an essential resource in this framework and is consumed in the reasoning process. This logic captures the reasoning process of a non-omniscient resource-limited agent.

We formally define $\models$ for the structure $M = \langle L, T, v, \prec, \pi, \mathcal{B} \rangle$ described above as follows:

1. This defines satisfiability of the base formulas of our language.
   $M, l \models p(\tau)$ if $l \in \pi(p, v(\tau))$
2. This defines satisfiability of negated formulas
   $M, l \models \neg\phi$ iff $M, l \not\models \phi$
3. This defines satisfiability of formulas formed with the $\wedge$ connective.
   $M, l \models (\phi \wedge \psi)$ iff $M, l \models \phi$ and $M, l \models \psi$

4. This defines the satisfiability of the belief formulas.
   $M, l \models B_\tau \phi$ iff $||\phi|| \in \mathcal{B}_{v(\tau)}(l)$

The satisfability of $\vee$ and $\rightarrow$ is defined accordingly. We impose the following restrictions on our models to describe an agent who reasons in a step-like fashion like its motivating step-logic agent described by SL_5.

1. $\forall l \in L, \forall t \in T \quad \{\} \notin \mathcal{B}_t(l)$ and if $s_1, s_2 \in \mathcal{B}_t(l)$ then $s_1 \cap s_2 \neq \{\}$.

   This says that the agent's belief set will be consistent at every time point. As explained before, we introduce this restriction to model a simple agent without contradictory beliefs and without any mechanisms for retraction[8].

2. $\forall l \in L, \forall t \in T$ if $s_1, s_2 \in \mathcal{B}_t(l)$, then $s_1 \cap s_2 \in \mathcal{B}_{t+1}(l)$

   This restriction constrains models at successive time points to be one step richer than their predecessors, in the sense that the agent has added all possible pairwise conjunctions of previous beliefs to the current step, but each pair participates just once[9].

3. $\forall l \in L, \forall t \in T$ if $s_1 \in \mathcal{B}_t(l)$, and $s_2 \supseteq s_1$ then $s_2 \in \mathcal{B}_{t+1}(l)$

   This restriction says that detachments of beliefs from a time step ago are added to the current set.

4. If $\phi \in \mathrm{OBS}(l, t)$ then $||\phi|| \in \mathcal{B}_t(l)$.

   An agent situated in the real world must have the ability to acquire new information through observation. This restriction allows for that capability.

5. $\forall l \quad ||true|| = L \in \mathcal{B}_{t_0}(l)$.

   Since the set T is ordered under $\prec$, and timelines are defined as rays, we can define a start point t_0. This restriction says that the agent believes in $true$ at the beginning of time.

[8] This restriction can be relaxed if the intent is to model a fallible agent who does default reasoning and may be permitted to have contradictory beliefs at any given time. Without the above restriction the neighborhood structures possibly allow for both $M, l \models B_\tau \phi$ and $M, l \models B_\tau \neg\phi$, since both $||\phi||$ and $||\neg\phi||$ could belong to $\mathcal{B}_{v(\tau)}(l)$.

[9] For example, if $M, l \models B_\tau \alpha$, $M, l \models B_\tau \beta$ and $M, l \models B_\tau \gamma$ then $M, l \models B_{\tau+1}(\alpha \wedge \beta)$, $M, l \models B_{\tau+1}(\alpha \wedge \gamma)$ and $M, l \models B_{\tau+1}(\beta \wedge \gamma)$ but $M, l \models B_{\tau+1}(\alpha \wedge \beta \wedge \gamma)$ does not follow from this restriction, however $M, l \models B_{\tau+2}(\alpha \wedge \beta \wedge \gamma)$ does.

4.2 Soundness and completeness

Theorem : The set of axioms (A1–A4) and inference rules (R1–R5) in Figure 4 provide a sound and complete axiomatization of the modal active-logic for reasoning in time.

Proof (sketch) : Soundness follows in a straightforward fashion from the interpretation of $\wedge$ and $\neg$ in the definition of $\models$ and from the restrictions on the models described in section 4. The proof of completeness hinges on the definition of a *canonical* model M^c in which every consistent[10] formula is satisfiable. In M^c we have a timeline corresponding to every *maximal consistent set* V. For definition and properties of maximal consistent sets we refer to [HY92]. We let $M^c = \langle L, T, v, \prec, \pi, \mathcal{B}\rangle$ where
$L = \{\ l_V\ :\ V$ is a maximal consistent set $\}$,
$l_V \in \pi(p, v(\tau))$ iff $p(\tau) \in V$, and
$\mathcal{B}_t(l_V) = \{\ I_{\psi,V} \mid B_t\psi \in V\ \} \cup \{\ S \mid S \supseteq S', S' \in \mathcal{B}_{t-1}(l_V)\ \}$
where $t \succ 0$ and $I_{\psi,V} = \{\ l_W \mid \psi \in W\ \}$, and $\mathcal{B}_0(l_V) = \{\ I_{\psi,V} \mid B_t\psi \in V\ \}$.

We then prove using induction that $M^c, l \models \phi$ iff $\phi \in V$, which proves that all consistent formulas are satisfiable in this structure. For space conservation, we do not provide the details of the proof here.

5 Conclusion

Active logics capture the process of reasoning of a resource-limited agent as it goes on in time. As time progresses, the agent draws more inferences (new beliefs) at each time step. Thus, an agent does not draw all the consequences of its current set of beliefs Σ all at once, but continues to add conclusions to this set in accordance with a set of inference rules. This is reflected by the increasing size of $\mathcal{B}(l_h, t)$, where l_h denotes the *real* history of occurrences in the world, and $\mathcal{B}(l_h, t)$ reflects what the agent believes in time t. The agent is certainly not guilty of omniscience under (a) logical consequence[11] since it is trivial to provide a counter-model to $B_\tau\alpha \quad B_\tau(\alpha \rightarrow \beta), \quad \neg B_\tau\beta$. By virtue of a description that is based on *intensions* of formulas, it is difficult

[10] A formula ϕ is *provable* if ϕ is one of the axioms or follows from provable formulas by application of one or more inference rules. A formula ϕ is *consistent* if $\neg\phi$ is not provable.

[11] The agent may eventually compute all logical consequences of its belief set if it has a set of complete agent inference rules.

Axioms :
(**A1**) All tautologies of propositional logic.
(**A2**) $\neg B_\tau false.$ Consistency
(**A3**) $B_\tau\phi \wedge B_\tau\psi \rightarrow B_{\tau+1}(\phi \wedge \psi).$ Conjunction
(**A4**) $B_\tau(\phi \wedge \psi) \rightarrow B_{\tau+1}\phi.$ Detachment[a]

Inference Rules :
(**R1**) From $\vdash \phi$ and $\vdash \phi \rightarrow \psi$ infer $\vdash \psi$ Modus Ponens
(**R2**) From $\vdash \phi \rightarrow \psi$ infer $\vdash B_\tau\phi \rightarrow B_{\tau+1}\psi$ Weak closure under valid consequence
(**R3**) From $\vdash \phi \leftrightarrow \psi$ infer $\vdash B_\tau\phi \leftrightarrow B_\tau\psi$ Belief in equivalent formulas
(**R4**) From $\vdash \phi$ infer $\vdash B_\tau\phi$ Belief in tautologies
(**R5**) If ϕ *is observed at time* τ infer $\vdash B_\tau\phi$ Observation

[a]Inheritance follows from either (A3) or (A4) when $\phi = \psi$.

Figure 4: Characterization of the modal active-logic

to distinguish between semantically equivalent beliefs. As such, (c) belief of valid formulas and a weak version of (b) closure under valid consequence[12] follows.

However, it is possible to modify our logic by providing a syntactic way to curtail the size of the belief set by introducing an additional element G to the structure M. $G \subseteq \mathcal{G}$ is defined as the agent's language and is closed under subformulas. An agent believes in ψ (i.e., $B_\tau\psi$) only if $\psi \in G$. For this new structure, the set of axioms and inference rules are suitably modified to capture this change (e.g., in (A3) $\phi \wedge \psi \in G$ and in (R4) $\phi \in G$ is added) and appropriate restrictions are placed on M. In essence, $\mathcal{B}_t(l)$ sets are filtered by G for all t and l. It can be proven that the modified set of axioms and inferences are sound and complete with respect to the modified structure. If the model includes more than one agent, each of them may have a different language G. This restricts an agent who believes in ϕ, to only that subset of $[\phi]$ (the equivalence class of ϕ) which is in the agent's language. The agent also believes only those tautologies that are in G. Hence the scope of (b) and (c) is reduced in the modified structure. The agent's language G has similarities to the awareness set concept of [FH88]. If one

[12]In our logic, by (R2), it takes an agent one period of time to deduce the consequence.

considers multiagent belief operators $\mathcal{B}^i$ without a time parameter then a modified version of Axioms (A2), (A4) and (R4) from figure 4 are true in the model of local reasoning of [FH88], (without modalities for implicit belief). Note, that we have only *explicit* beliefs, and there is no notion of *implicit* beliefs. In [FH88] the models are still static, in that even though they suggest incorporating reasoning *about* time, and *changing* awareness functions, there is no way to account for inference time in their models.

Parallels to more advanced step-logics for planning and default reasoning have also been developed. These logics have two additional modal operators $\mathcal{D}$ and $\mathcal{P}$ for default beliefs and temporal projections respectively. Both operators are defined using the Montague style semantics similar to that used to define $\mathcal{B}$ but with a different set of restrictions to allow for contradictions arising due to defaults and to reflect the notion of persistence. A future paper will describe these logics.

References

[All84] J. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154, 1984.

[Ash88] N. Asher. Reasoning about belief and knowledge with self-reference and time. In *Proceedings of the second conferennce on Theoretical Aspects of Reasoning About Knowledge*, pages 61–81. Morgan Kaufmann, 1988.

[Che80] B. Chellas. *Modal Logic: An Introduction.* Cambridge University Press, 1980.

[Ebe74] R. A. Eberle. A logic of believing, knowing and inferring. *Synthese*, 26:356–382, 1974.

[ED88] J. Elgot-Drapkin. *Step-logic: Reasoning Situated in Time.* PhD thesis, Department of Computer Science, University of Maryland, College Park, Maryland, 1988.

[EDP90] J. Elgot-Drapkin and D. Perlis. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2(1):75–98, 1990.

[ER87] McKenzie E. and Snodgrass R. Extending the relational algebra to support transaction time. In *Proceedings of the SIGMOD conference.* ACM, 1987.

[FH88] R. Fagin and J. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39–76, 1988.

[Hin62] J. Hintikka. *Knowledge and Belief.* Cornell University Press, 1962.

[HY92] J. Halpern and Moses Y. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[KL88] S. Kraus and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174, 1988.

[KNP90] S. Kraus, M. Nirkhe, and D. Perlis. Deadline-coupled real-time planning. In *Proceedings of 1990 DARPA workshop on Innovative Approaches to Planning, Scheduling and Control*, pages 100–108, San Diego, CA, 1990.

[Kon83] K. Konolige. A deductive model of belief. In *Proceedings of the 8th Int'l Joint Conference on Artificial Intelligence*, pages 377–381, Karlsruhe, West Germany, 1983.

[Kon86] K. Konolige. What awareness isn't. In J. Halpern, editor, *Proceedings of the first conferennce on Theoretical Aspects of Reasoning About Knowledge*, Monterey, CA, 1986. Morgan Kaufmann.

[Lak86] G. Lakemeyer. Steps towards a first order logic of explicit and implicit belief. In J. Halpern, editor, *Proceedings of the first conferennce on Theoretical Aspects of Reasoning About Knowledge*, Monterey, CA, 1986. Morgan Kaufmann.

[Leh84] D. J. Lehmann. Knowledge, common knowledge and related puzzles. In *Proceedings of the third annual ACM conference on principles of distributed computing*, pages 62–67, 1984.

[Lev84] H. Levesque. A logic of implicit and explicit belief. In *Proceedings of the National Conference on Artificial Intelligence*, pages 198–202, Austin, TX, 1984. AAAI.

[LR86] R. Ladner and J. H. Reif. The logic of distributed protocols. In J. Halpern, editor, *Proceedings of the first conferennce on Theoretical Aspects of Reasoning About Knowledge*, Monterey, CA, 1986. Morgan Kaufmann.

[McC78] J. McCarthy. Formalization of two puzzles involving knowledge. 1978. Unpublished note, Stanford University.

[McD82] D. McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6, 1982.

[MH79] R. C. Moore and G. Hendrix. Computational models of beliefs and the semantics of belief sentences. Technical Report Technical note 187, SRI International, Menlo Park, 1979.

[Mon68] R. Montague. Pragmatics. *Contemporary Philosophy*, pages 101–121, 1968. La Nuova Italia Editrice.

[Mon70] R. Montague. Universal grammar. *Theoria*, 36:373–398, 1970.

[NKP91] M. Nirkhe, S. Kraus, and D. Perlis. Fully deadline-coupled planning: One step at a time. In *Proceedings of the sixth international symposium on methodologies for intelligent systems*, Charlotte, NC, 1991.

[NKP93] M. Nirkhe, S. Kraus, and D. Perlis. Situated reasoning within tight deadlines and realistic space and computation bounds. In *Common Sense 93*, Austin, Texas, 1993.

[PEDM] D. Perlis, J. Elgot-Drapkin, and M. Miller. Stop the world! – I want to think! Invited paper, *International J. of Intelligent Systems*, special issue on Temporal Reasoning, K. Ford and F. Anger (eds.), vol. 6, 1991, pp. 443–456.

[PS85] P.F. Patel-Schneider. A decidable first order logic for knowledge representation. In *Proceedings of the 9th Int'l Joint Conference on Artificial Intelligence*, pages 455–458, Los Angeles, CA, 1985.

[Sat77] M. Sato. A study of kripke-style methods of some modal logics by gentzen's sequential method. *Publications of the REsearch Institute for Mathematical Sciences*, 13(2), 1977.

[Sco70] D. Scott. Advice on modal logic. *Philosophical Problems in Logic*, pages 143–173, 1970.

[TSSK91] Becky Thomas, Yoav Shoham, Anton Schwartz, and Sarit Kraus. Preliminary thoughts on an agent description language. *International Journal of Intelligent Systems*, 6(5):497–508, August 1991.

[Var86] M. Vardi. On epistemic logic and logical omniscience. In J. Halpern, editor, *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 293–305, Monterey, CA, 1986. Morgan Kaufmann.