Elements of a Utilitarian Theory of Knowledge and Action

Piotr J. Gmytrasiewicz* and Edmund H. Durfeet
'Department of Computer Science
Hebrew University, Jerusalem, Israel

'Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, Michigan 48109
piotr@cs.huji.ac.il, durfee@engin.umich.edu

Abstract

According to the utilitarian paradigm, an autonomous intelligent agent's interactions with the environment should be guided by the principle of expected utility maximization. We apply this paradigm to reasoning about an agent's physical actions and exploratory behavior in urgent, time-constrained situations. We model an agent's knowledge with a temporalized version of Kripke structures—as a set of branching time lines described by fluents, with accessibility relations holding among the states comprising the time lines. We describe how to compute utility based on this model which reflects the urgency that the environment imposes on time. Since the physical and exploratory actions that an agent could undertake transform the model of branching time lines in specific ways, the expected utilities of these actions can be computed, dictating rational tradeoffs among them depending on the agent's state of knowledge and the urgency of the situation.

1 Introduction

Reasoning about actions and their effects is a central issue in Artificial Intelligence. Agents that are to be successful in their interactions with the real world have to be able to plan their actions and execute them efficiently in time-constrained situations, without having complete knowledge about the state of their environment, and hence, about its future development.

Our approach to these issues is strongly motivated by the utilitarian paradigm for designing autonomous intelligent agents (see [Doyle, 1992] for convincing arguments supporting this view). According to this paradigm, an agent's rationality is equivalent to its ability to choose among its possible behaviors so as to maximize its expected utility. The possible behaviors of even a mildly sophisticated agent could include: information gathering (informative) actions, like scanning; computational actions, like inference and planning; and physical actions that directly influence the agent's environment. In this paper, we concentrate on physical and informative actions. We present a preliminary framework to model an agent's uncertain knowledge about the temporal evolution of its environment. The model, essentially a temporal Kripke structure, is based on sets of branching time lines [McDermott, 1982]. We augment a set of time lines with accessibility relations to represent the knowledge of the agent evolving in time. The time lines are described by fluents, that is, by attributes with time-varying values.

Our model depicts an agent's knowledge about its actions as how those actions transform the structure of the possible time lines and the accessibility relations among them. We further provide the definitions needed to compute the expected utilities of these actions based on our model, and show how the important notion of urgency, or the value of time, naturally arises from our approach. Urgency, in some cases, makes it rational for an agent to undertake immediate physical action despite uncertainty that could be resolved by an informative action (it can "leap before looking"). We believe that the same basic tradeoffs exist in the problem of whether it is rational to "leap before thinking", that is very much in the middle of the debate of deliberative versus reactive planning.

Our work builds on a large body of previously reported research on utility-based approaches in AI [Doyle, 1992]. The work by Haddawy and Hanks [Haddawy and Hanks, 1990; Hanks and Firby, 1990] is similar to ours, but it differs in its emphasis on symbolic goals as the basis for utility calculations, suggesting that these goals are imposed on the agent during operation. We, on the other hand, see utility as a reflection of general preferences, maybe more appropriately called values, that are inherent in any truly autonomous agent. We also believe that it is possible to represent explicitly the tradeoffs among sensing, acting, and planning, and thus to decide among them in a principled way. Further similarities exist between our model based on branching time lines and the probabilistic temporal projection described in [Hanks, 19881, and although our approach is, at this point, more knowledge-theoretic, the practical methods outlined in [Hanks, 1988] are relevant to our model as well.

Models similar to branching time lines have also been examined as a basis for temporal probabilistic logics for actions and planning [Haddawy, 1991; Pelavin, 1988]. Our contribution in this paper is to show how these models can be used to rationally guide action and perception, in addition to providing semantics of probabilistic temporal logic. Further,

---

°This research was supported, in part, by the Golda Meir Fellowship and the Alfassa Fund administered by the Hebrew University of Jerusalem, and by the National Science Foundation under grant IRI-9015423 and PY1 award IRI-9158473.

while the logical language capturing various forms of action is useful, these logics by themselves are not capable of fully supporting purposeful choice among alternative behaviors [Doyle, 19921.

Other related work on formal, logically oriented approaches [Moore, 1990; Morgenstern, 1987] is based on similar intuitions to ours, but lacks the utilitarian perspective. In particular, we argue that, in realistic settings, an agent will *never* possess enough knowledge about the state of the world to be able to *prove* that executing any particular action, such as dialing a phone, will be successful. Further, from a strictly logical perspective the lack of knowledge is, taken by itself, no obstacle to acting at all. If, for instance, an agent wants to reach someone by phone without actually knowing the telephone number (other than that it contains 7 digits), the agent can simply dial all possible seven digit numbers in sequence, being quite convinced that eventually it will get through. The part that is missing, of course, is the relative value of lime and other resources wasted in the process, but that is precisely what makes this an issue in economics, rather than in logic [Doyle, 1992]. It is in this context that one can understand why looking up a number in the phone book actually *pays off.*

Yoav Shoham also relies on branching time lines to represent an agent's uncertain knowledge, its evolution, and the effects of action [Shoham, 1989]. The major difference between his work and ours is that his does not include the clement of expected utility, and therefore does not address the crucial issue of motivation and choice of action [Doyle, 1992], referring instead to the logically-motivated work of Cohen and Levesque [Cohen and Levesquc, 1987]. Further, Shoham envisions the evolution of the agent's knowledge as a series of time line structures, while we represent it as accessibility relations within one such structure. He also seems to include the uncertainty about the agent's own future actions as branching in his model, while we forbid it.

Finally, our work builds on the enormous body of research done over the years in the fields of mathematical economics, and utility and decision theories. We refer to it in the text of the paper as we go.

In the following sections of the paper, we first introduce our branching time lines model of knowledge. We then describe how the model can be used to reflect results of physical and informative actions, and how expected utility calculations can be performed on it. We then give examples of expected utility calculation for plans involving physical and informative actions, and show how our approach naturally gives rise to the notion of urgency as the value of time. We end with conclusions and further research directions.

2 Modeling Temporal Knowledge by Branching Time Lines

Branching time lines were proposed by McDermott as a model for temporal logic of processes and plans [McDermott, 1982] and they provide a basis for various logics of probability, time and action in other related work [Haddawy, 1991; Pelavin, 1988; Dean and Wellman, 1991]. We have adopted this model here, and augmented it with accessibility relations so that it forms a temporalized version of the Kripke structure (for details on Kripke structures, see [Halpern and Moses, 1990; Halpern and Moses, 1991]). The branching time lines are composed of chronicles consisting of totally temporally ordered states. A chronicle is intended to represent a possible course of events in an agent's environment, and it can be described by a set of fluents, i.e., attributes with values changing over time fGinsberg, 1990; McDermott, 1982].

Loosely speaking, the current uncertainty in an agent's knowledge about the possible past courses of events is represented by the fact that the time lines branch off into the past; the uncertainty of the agent's knowledge about the present state of the world is represented by the fact that there are many branching time lines at the present date, NOW; and the current uncertainty in the agent's knowledge about the future evolution of the world is represented by the fact that the time lines are branching off into the future. In this paper, we are interested in the utilities of the optional future actions of the agent, and therefore we only consider the representations starting at NOW and extending into the future.

Thus, the branching time lines represent the *current* uncertainty in the agent's knowledge. To complement the picture with the knowledge that the agent may have about its future knowledge, we introduce the accessibility relations that can hold among the time lines and the branches in the model. We say that if one of the time lines in the model is accessible from another line, then the agent situated in the world evolving according to the first time line will consider the evolution of the world according to the other time line as possible. If two time lines are accessible from each other then the agent will not be able to tell if the world evolves according to one or the other.[1]

The model of branching time lines introduced above corresponds precisely to a temporal sequence of classical Kripke structures. Let us also note that if the accessibility relations defined above form a partition of possible slates of the world at any given time into a set of exhaustive equivalence classes, then the model of branching time lines becomes identical to what in the economics literature is called the information structure of the agent [Radner, 19821.

In our graphical representation, the accessibility relations, depicted by dotted arrows, will be assumed to hold for the whole section of the time line between the nearest points, which can be either branching points, or temporal marks on the time lines. As an example, let us take the case of an agent sitting by his desk in his office not knowing whether the building is on fire, but having some indication that it might be (burning smell, commotion, etc.). The agent knows that there is a sprinkler system installed in the building, but he does not know whether the system will function properly. Let us assume for the moment that our agent knows that it

---

[1] This is just an intuitive description sufficient for the purpose of this paper. For now let us only note that the accessibility relation connecting two time lines defined above is, in fact, vague, since it does not specify whether each of the states in one line is accessible from each of the states in the other line, or whether it is accessible only from the state with the same date. The first interpretation corresponds to the situation in which the agent does not know what time it is, while the other corresponds to an agent equipped with a clock. In the following discussion we will use the second interpretation of the agent being aware of the exact lime, although we plan to investigate the first alternative in future research.

is positioned so that it would not notice the sprinkler system's operation. The uncertain knowledge about the present state of the world, and about the future developments that will follow, can be represented as two time lines in Figure 1. According to one of the time lines, the building is on fire. This time line has a branching, B, corresponding to the agent's uncertainty as to the automatic activation of the sprinkler system sometime later, in the case of fire. The other time line presents the possibility that the building is not on fire. Both of the time lines are each accessible from the other, and each is accessible from itself, with the accessibility relations depicted in Figure 1 as arrows. The arrows also connect the branches of the time line according to which the building is on fire, since, as assumed, our agent would not be aware of the sprinkler system's operation in the case it is activated. These branches are similarly connected to the other time line to indicate that the sprinkler's operation would not resolve the question of whether the building is on fire for this agent.

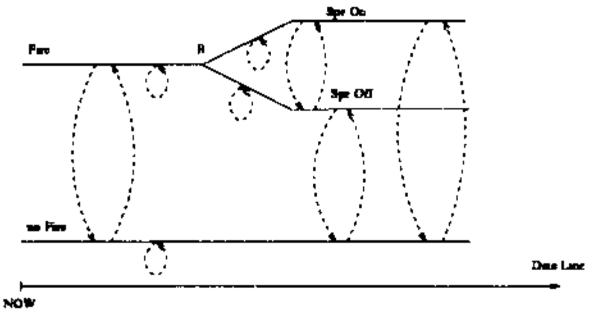

Figure 1: Agent's knowledge when it is uncertain about fire and sprinklers.

Let us now consider a variation of the above situation that indicates how our model can represent somewhat different knowledge of the agent. In this case, we assume that the sprinkler system is also installed in the agent's office, so it will certainly notice its operation. The branching time lines model of this situation is depicted in Figure 2. Note that the two branches of the time line corresponding to the fire are not connected via the accessibility relation. This is because, in this situation, the cases of the sprinkler being on or off are not indistinguishable any more. Further, the upper branch, according to which the sprinkler kicks in, is not connected to the lower time line, while the lower branch is. This is the representation of the fact that the sprinkler's operation would be noticed by the agent and effectively identify for it whether the building is, in fact, on fire. Of course, the agent still would not be able to distinguish the case where the building is not on fire from the case when there is a fire but the sprinkler system is not operational, as indicated in Figure 2.

In what follows, we will call the model of an agent's knowledge that does not include results of any action of this agent (although it may contain effects of actions of other agents, with their corresponding uncertainties) the *base model.*

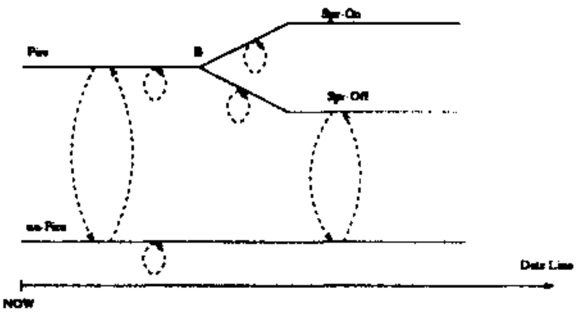

Figure 2: Agent's knowledge if it would notice sprinkler's operation.

3 Acting in the Model of Branching Time Lines

As we have mentioned, an agent can construct the branching time line models that represent its knowledge of the dynamic effects of performing various activities in its environment, and then can compute the expected utility of these activities allowing a rational decision on the proper course of action. Before describing this calculation, we examine how changes in the model can reflect the effects of an agent's possible courses of action.

Let us lake the example of the agent in his office wondering about the building being on fire, as in the previous section. Consider an action our agent might contemplate performing, for instance one of immediately activating the sprinkler system. The result of this action in our representation is a pair of time lines, this time with no branches (see Figure 3). The fact that the sprinkler system has been activated is true in both of them, but, according to one of the possible courses of events, the sprinkler system is extinguishing a real fire, while in the other the building is just getting sprayed. The lack of branching in the time line according to which there is a fire reflects the fact that the agent has some newly acquired certainty; it has not left matters up to an unpredictable sprinkler actuator which was the source of the branching in its model before.

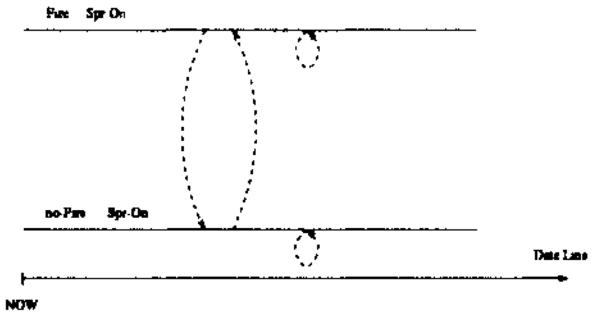

Figure 3: Agent's knowledge of the result of activating the sprinklers.

Let us now take an example of a different type of action - one aimed specifically at gathering information. Let us say that it is possible for our agent to check whether there really is a fire, for instance by performing an action that we will refer to as $A_{FireTest}$, that is expected to take $T_{FireTest}$ time units. The result of this action is sketched in Figure 4, which depicts the two time lines, as before, but now the accessibility relations between the lines have been severed after time $T_{FireTest}$.[2] The severance corresponds to the fact that after the $A_{FireTest}$ has been successfully performed, our agent *will know* whether the building is on fire or not. The time line according to which there is a fire, however, still has a branching in it, corresponding to the agent's uncertainty about the proper functioning of the sprinkler system.

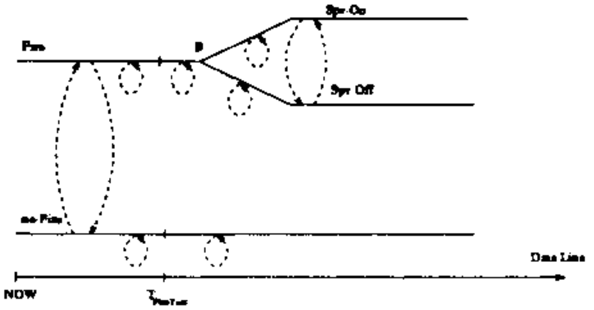

Figure 4: Agent's knowledge of result of $A_{FireTest}$.

The agent can also consider more complex actions, for instance those consisting of a series of simpler actions, those involving conditionals, or those with while loops (see also [Moore, 1990; Morgenstem, 1987]). For example, the agent might consider an action consisting of performing the action $A_{FireTest}$ first, and then, if the fire is burning, on turning the sprinkler system on. The effect of this composite action is depicted in Figure 5. We still have two time lines corresponding to the cases of the fire burning and not, that are accessible from each other for all times before the action $A_{FireTest}$ has been completed, but not accessible afterwards. Also, at the time $T_{FireTest}$, the agent activates the sprinkler system in the upper time line, as indicated in Figure 5.

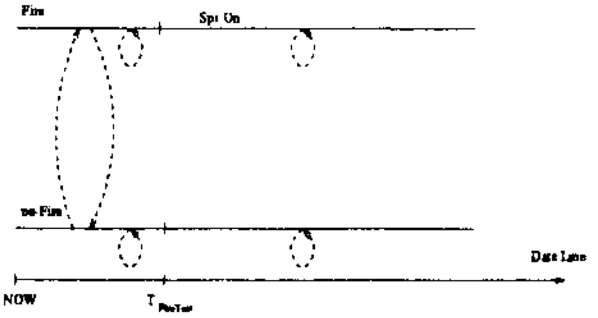

Figure 5: Agent's knowledge of the result of activating the sprinklers if the result of $A_{FireTest}$ is positive.

4 Computing Utilities in the Model of Branching Time Lines

In this section we show how utility associated with the model of branching time lines can be defined. The definition will require that we enrich our model with probabilistic information, supplementing the possibilistic information represented by the model presented so far. Thus, we assume that an agent's knowledge about the likelihoods of time lines, TL,, in its model can be represented by probability distributions, *PTL,*. The states that comprise each time line are in total temporal order, and each of them can be described by the propositions stating the values of the fluents, *F(t),* in this particular stale. The descriptions, *F(t)*, can be then used to assign a numerical value, $U_{TLr}(t)$> to the state, representing the state's desirability to the agent (see [Keeney and Raiffa, 1976] for details on combining the effects of several fluents (or attributes) using a multiattribute utility calculation).

The expected utility, $EU_{TL}(T)$, associated with a set of states belonging to the time lines at time *t,* can be defined as a probabilistic mixture of the utilities, $U_{TL}(t)$, *OF* states on each time line *i* at time *t:*

$$EU_{TL}(t) = \sum_i p_{TL_i} U_{TL_i}(t), \qquad (1)$$

where P*TL,* is the probability of the time line TL_i.

In practice, the part of a set of time lines, 7L, used by the agent to compute the expected utilities of its optional actions lies between the point in time called NOW, corresponding to the current date, and the point that lies in the infinite future, ∞. Consequently, we define the expected utility, *EVTL* of this set of time lines, as:

$$EU_{TL} = \int_{NOW}^{\infty} e^{-\alpha t} EU_{TL}(t)dt,$$

where the constant, α, is the discount rate.[3] The definition above accepts the simplifying assumption of simple additivity of the expected utility over the states (see [Keeney and Raiffa, 1976; McKenzie, 1986] for discussion).

The problem of existence for EU_{TL}, that is, the convergence of the integral in Eq.(2), can be overcome by demanding that the discount rate, α, be large enough. We find this demand counterintuitive at times, and for cases where convergence is a problem, as when $\alpha = 0$, we use the method developed in [Ramsey, 1928], which advocates shifting the utility scale to force the convergence of the integral, or the method introduced in [Weizsaecker, 1965], which uses the so-called overtaking criterion, based on comparing finite partial integrals over time. A somewhat *ad hoc* way of dealing with the nonconvcrgence of the integral in Eq. 2 would be to introduce a finite time horizon, which might be a legitimate simplification at times, but we will not use it here.

5 Utilities of Physical and Informative Actions

Having defined the utilities for the branching time lines model above, it is quite straightforward to define the utilities

---

[2] Recall that the accessibility relations depicted graphically hold only between the nearest branching points or time marks.

[3] The discount rate corresponds, roughly, to the agent's possibly preferring to receive a payoff sooner rather then later. See [Keeney and Raiffa, 1976] for more details.

of physical actions. Since the physical actions that an agent can perform in its environment transform the model of the branching time lines, the utility of an action can be identified as the utility of the branching time lines model that reflects the effects of this action.

To continue the example of the agent wondering about the fire in the office building, the expected utility of the agent's action, A_{SprOn}, of immediately activating the sprinkler system can be computed from the model in Figure 3. To illustrate this example calculation we assume that the fluents, $F(t)$, describing the possible time lines in Figure 3 are the attributes of FireSize and GettingWet. The attribute GettingWet will have a value that we will denote by Well, at time NOW, and monotonically decrease to 0 at time 10, when everything is already wet, in both possible time lines. The attribute FireSize will be equal to zero in the time line according to which the fire is not burning, and will be monotonically decreasing from a value, denoted Sizel, at time NOW, to zero at time 10 (see Figure 6a). Given the above, we will assume that the utility functions are linear over the values of attributes, with the proportionality constant of -1, and additive, so that the utility of a state is the sum of the utilities associated with each of the attributes.[4] Assuming that the discount rate, a, is equal to zero, the total utility of immediately activating the sprinklers can be easily computed, using Equation 2 (see Figure 6a). For the upper time line, the contribution of the fire is the negative of the area under the FireSize fluent line: $-0.5 \times Size1 \times 10$. The sprinkler's contribution is, similarly: $-0.5 \times Wet1 \times 10$. Adding the contribution of the upper time line with probability p_{Fire} and the lower time line with probability $1 - p_{Fire}$, we get:

$$\begin{aligned} U(A_{SprOn}) &= -5p_{Fire}(Size1 + Wet1) \\ &\quad - 5(1 - p_{Fire})Wet1 \end{aligned} \tag{3}$$

utiles, where p_{Fire} is the probability associated with the fire burning (we will compare this result with another shortly).

The case of informative actions is less straightforward.[5] Let us take the action $A_{FireTest}$, and its resulting branching time lines model, depicted in Figure 4. The only change in this model, as compared to one in which the agent is quiescent (Figure 1), is the change in the accessibility relations between the two time lines after the time $T_{FireTest}$. But, according to the equations defining the utility calculations given in the preceding section, the accessibility relations by themselves do not impact the utility of the model in any way.[6] That the possession of knowledge by itself does not impact the agent's welfare is quite understandable; the agent's information only counts as far as it puts it in a better position to make desirable physical changes to the environment. This observation leads us to the conclusion that the desirability of certain actions, for example informative (and computational) ones, has to be always assessed as embedded in other, physical actions, that these non-physical ones facilitate.

As an example, let us take the composite conditional action considered before, consisting of $A_{FireTest}$, and, if the fire is burning, followed by the activation of the sprinkler system, which we will call $A_{FireTest/SprOn}$. The utility of this composite action can be calculated based on the model depicted in Figures 5. To make this calculation concrete we will use the same parameters as in the case of the agent activating the sprinklers without testing, above. The difference boils down to the time delay, $T_{FireTest}$. Thus, in the case that there is the fire, it would go on burning for the time $T_{FireTest}$, at the level Size1, after which the agent would find the positive result of the test and activate the sprinkler system resulting in linear decrease of the FireSize to zero at the time ($T_{FireTest}$ + 10) (see Figure 6b). In the case there is no fire, the sprinkler system is, of course, not activated. With the simplifying assumptions of the utility being additive and inversely proportional to the attributes, the utility of the conditional action can be computed as:

$$\begin{aligned} U(A_{FireTest/SprOn}) &= -p_{Fire}Size1 \times T_{FireTest} \\ &\quad - 5p_{Fire}(Size1 + Wet1) \end{aligned} \tag{4}$$

utiles. The relative value of $A_{FireTest}$, in this particular context of the action activating the sprinklers, could be assessed by comparing the utilities of activating the sprinkler system conditioned on outcome of the $A_{FireTest}$, versus activating the sprinkler system without performing the test. As can be easily seen from Equations 3 and 4, this boils down to comparing the disutility of the office building getting sprayed unnecessarily if there is no fire if the sprinklers were activated without testing: $-5(1 - p_{Fire})Wet1$, and the disutility of the fire, if there is one, burning for additional time $T_{FireTest}$ in the other case: $-p_{Fire}Size1 \times T_{FireTest}$.

Clearly, the result of this comparison is not guaranteed to favor the performance of $A_{FireTest}$. One can easily imagine the situation in which the lengthy time required to perform the test, coupled with the high *a priori* probability that the fire is burning and possibly high losses incurred during the time $T_{FireTest}$, gives priority to activating the sprinklers without the test. For example, if we assume that $Size1 = 10Wet1$, then it is rational to perform the testing operation only if $T_{FireTest} \leq (1 - p_{Fire})/2p_{Fire}$. So, if the probability of a fire were 10%, for example, then running the test if it takes 4 time units would be rational, but if the test takes 5 time units it would be rational to activate the sprinklers NOW instead.

5.1 Urgency: The Value of Time

In the example above, one is tempted to say that the probability of losses due to fire have made the situation urgent, so that wasting time on an action like testing that does not have the immediate physical effects of remedying the situation, may not be worth it. Said more succinctly: *the probability of losses due to inaction have created urgency,*

---

[4] These simplifying assumptions are made only to make the example easier to follow. They are in no way inherent in our approach.

[5] There are, probably, no actions that are purely informative in the sense that no action's sole effect is the change of the agent's knowledge. There will always be some physical resources consumed and possibly other physical side effects. There may be no reason, though, to include them in the branching time lines model if these effects are negligible from the agent's welfare point of view.

[6] That is, unless the agent's possession of knowledge is by itself desirable, and included in the set of fluents, F, in Eq. 1. We do not consider this case here.

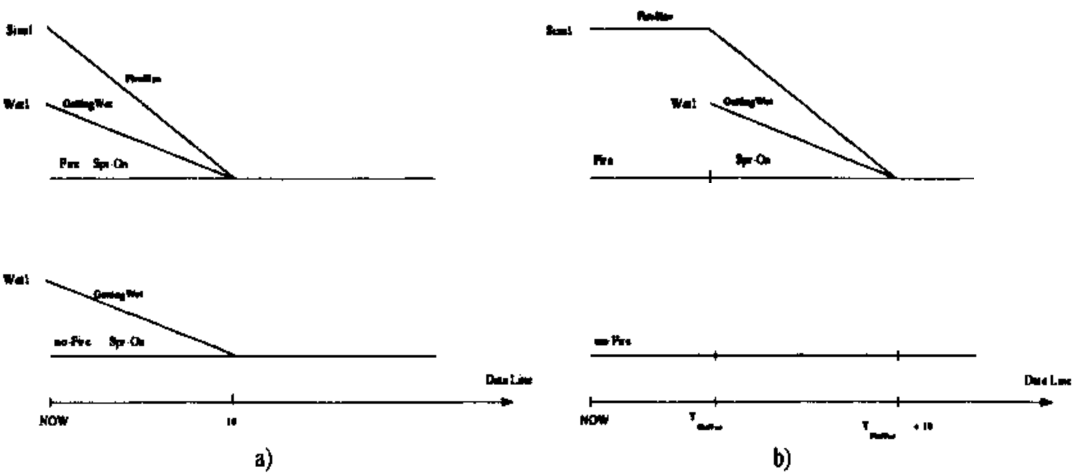

Figure 6: Depiction of the fluents; FireSize and GettingWct, describing the results of a): immediate action of activating the sprinklers, and b): testing for the fire before.

i.e., have imparted value to time. To make the above insight more precise:

Definition: The urgency, *UHG(t),* or the value of time function, is defined as the negative of the expected utility, $EU_{TL}(t)$, of the *base model* of the agent's knowledge, represented by the set of branching time lines, TL, as defined in Equation 1.

Recall that the base model is the model of branching time lines that represents the agent's knowledge about the evolution of the events in the world but does not include any of the agent's own actions. The above definition, therefore, formalizes urgency as the rate of accumulation of disutility to an agent, if it delays physical action.

For instance, for the base model in the example of the agent in the office building that is possibly on fire, the urgency for the points in time NOW and afterwards[7] is: $URG(t) = p_{FireSize} 1$ utiles/sec.

Clearly, the fact that time is valuable forces rational agents to be time-effective in executing physical actions, and crucially impacts on the viability of non-physical action, as shown in the example above of the action $A_{FireTest}$, which becomes ill-advised if it takes too long and the situation is too urgent.

Possibly the most important non-physical action that the urgency of the situation could make ill-advised is, of course, the agent's reasoning, and, in particular, planning. We anticipate that the approach we have taken can prove fruitful in addressing the issue of reactive versus deliberative behavior in rational agents; we will investigate this in our future research.

6 Conclusions and Future Work

In this preliminary report we have presented an application of the utilitarian paradigm (maximization of expected utility) to reasoning about knowledge and action. In our approach, an agent's knowledge is represented as a set of branching time lines, with probability distributions describing the agent's knowledge about the likelihoods of possible courses of action, and with accessibility relations connecting the lines and branches describing the future states of the agent's knowledge. The rational agent can use this model to investigate the expected utilities of the various courses of action it could take, and behave rationally based on this calculation. We have investigated examples of physical action and informative action. The utility of the latter was due to the fact that it would enhance the efficiency of the agent's physical action.

We showed how the important notion of urgency, or the value of time, arises from our approach. Urgency has the intuitive property of favoring immediate physical actions, sometimes making non-physical actions, such as sensing and reasoning, ill advised. Our result here goes along other work previously reported that treats time as a resource fBoddy and Dean, 1989; Dean and Boddy, 1988; Dean, 1990; Horvitz *et al.,* 1989; Russell and Wcfald, 1989], but our contribution is the rigorous derivation of the value of this resource.

In our future work, drawing on our earlier investigation [Gmytrasiewiez and Durfee, 1992], we will address the issue of representing the temporal uncertain knowledge of agents in multiagent situations, in which recursion in models arises due to agents' modeling others modeling them, and so on. We will attempt to apply this framework to the process of coordination and communication in dynamic and uncertain multiagent situations [Gmytrasiewicz *etal.,* 1991a; Gmytrasiewicz *et* a/., 1991b]. Further, we will develop a suitable language of knowledge and action, drawing on previous work in [McDermott, 1982; Moore, 1990; Morgenstern, 1987]. Finally, we plan on investigating the issue of rationality of reasoning under time pressure in an attempt to address the issue of rational choice among sensing, reasoning, and physical actions in time-constrained domains.

References

[Boddy and Dean, 1989] Mark Boddy and Thomas Dean. Solving time-dependent planning problems. In *Proceedings of the Eleventh International Joint Conference on Artificial I nielli gence,* pages 979-984, Detroit, Michigan, August 1989.

---

[7] We neglect here for simplicity the fact that the fire would burn itself out after some long time, and the urgency would then be zero.

[Cohen and Levesque, 1987] Philip R. Cohen and Hector J. Levesque. Intention = choice + commitment. In Proceedings of the National Conference on Artificial Intelligence, pages 410-415, July 1987.

[Dean and Boddy, 1988] Thomas Dean and Mark Boddy. An analysis of time-dependent planning. In Proceedings of the National Conference on Artificial Intelligence, pages 49-54, St. Paul, Minnesota, August 1988.

[Dean and Wellman, 1991] Thomas L. Dean and Michael P. Wellman. Planning and Control. Morgan Kaufamnn, 1991.

[Dean, 1990] Thomas Dean. Planning under uncertainty and time pressure. In Proceedings of the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control, pages 390-395. November 1990.

[Doyle, 1992] Jon Doyle. Rationality and its role in reasoning. Computational Intelligence, 8:376-409, 1992.

[Ginsberg, 1990] Matthew L. Ginsberg. Computational considerations in reasoning about action. In Proceedings of the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control, pages 37-47, November 1990.

[Gmytrasiewicz and Durfee, 1992) Piotr J. Gmytrasiewicz and Edmund H. Durfee. Logic of knowledge and belief for recursive modeling: Preliminary report. In Proceedings of the National Conference on Artificial Intelligence, pages 628-634, July 1992.

[Gmytrasiewicz et al., 1991a] Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David K. Wehe. The utility of communication in coordinating intelligent agents. In Proceedings of the National Conference on Artificial Intelligence, pages 166-172, July 1991.

[Gmytrasiewicz et al., 1991b] Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David K. Wehe. A decision-theoretic approach to coordinating multiagent interactions. In Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, pages 62-68. August 1991.

[Haddawy and Hanks, 1990) Peter Haddawy and Steve Hanks. Issues in decision-theoretic planning: Symbolic goals and numeric utilities. In Proceedings of the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control, pages 48-58, November 1990.

[Haddawy, 1991] Peter Haddawy. A temporal probability logic for representing actions. In Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning, pages 313-324, 1991.

[Halpern and Moses, 1990] Joseph Y. Halpern and Yoram Moses. A guide to the modal logics of knowledge and belief. Technical Report 74007, IBM Corporation. Almaden Research Center, 1990.

[Halpern and Moses, 1991] Joseph Y. Halpern and Yoram Moses. Reasoning about knowledge: a survey circa 1991. Technical Report 50521, IBM Corporation, Almaden Research Center, 1991.

[Hanks and Firby, 1990] Steven Hanks and R. James Firby. Issues in architectures for planning and execution. In Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control, pages 59-70, November 1990.

[Hanks, 1988] Steven Hanks. Practical temporal projection. In Proceedings of the National Conference on Artificial Intelligence, pages 158-163, July 1988.

[Horvitz et al., 1989] E. J. Horvitz, G. F. Cooper, and D. E. Heckerman. Reflection and action under scarce resources: Theoretical principles and theoretical study. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, pages 1121-1127, Detroit, Michigan, August 1989.

[Keeney and Raiffa, 1976] R. L. Keeney and H. Raiffa. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley and Sons, 1976.

[McDermott, 1982] Drew McDermott. A temporal logic for reasoning about processes and plans. Cognitive Science, 6:101-155, 1982.

[McKenzie, 1986] L. W. McKenzie. Optimal economic growth, turnpike theorems and comparative dynamics. In M. D. Intriligator K. J. Arrow, editor, Handbook of Mathematical Economics, Vol. III, pages 1281-1355. North-Holland, 1986.

[Moore, 1990] Robert C. Moore. A formal theory of knowledge and action. In Formalizing Common Sense, pages 319-358. Ablex Publishing Corporation, 1990.

[Morgenstern, 1987] Leora Morgenstern. Knowledge preconditions for actions and plans. In Proceedings of the National Conference on Artificial Intelligence, pages 867-874, Seattle, Washington, August 1987. (Also published in Readings in Distributed Artificial Intelligence, Alan H. Bond and Les Gasser, editors, pages 151-158, Morgan Kaufmann, 1988.).

[Pelavin, 1988] Richard N. Pelavin. A formal approach to planning with concurrent actions and external events. Technical Report TR 254, Computer Science Department, University of Rochester, Rochester, New York, 1988.

iRadner, 1982] R. Radner. Equilibrium under uncertainty. In M. D. Intriiigtor K. J. Arrow, editor, Handbook of Mathematical Economics. North-Holland, 1982.

[Ramsey, 1928] Frank Ramsey. A mathemetical theory of savings. Economic Journal, (38):543-559, 1928.

[Russell and Wefald, 1989] S. Russell and E. Wefald. On optimal game tree search using rational meta-reasoning. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, pages 334-340, Detroit, Michigan, August 1989.

[Shoham, 19891 Yoav Shoham. Time for action: on the relation between time.knowledge and action. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, pages 1168-1172, Detroit, Michigan, August 1989.

[Wcizsaccker, 1965] C. C. Von Weizsaecker. Existence of optimal programs of accumulation for an infinite time horizon. Review of Economic Studies, (32):85-104, 1965.