Branching Continuous Time and the Semantics of Continuous Action

Ernest Davis*
New York University
251 Mercer St.
New York, New York
DAVISE@CS.NYU.EDU

July 7, 1994

Abstract

It is often useful to model the behavior of an autonomous intelligent creature in terms of continuous control and choice. For example, a robot who moves through space can be idealized as able to execute any continuous motion, subject to constraints on velocity and accelleration; in such a model, the robot can "choose" at any instant to change his accelleration. We show how such models can be described using a continuous branching time structure. We discuss mathematical foundations of continuous branching structures, theories of continuous action in physical worlds, embedding of discrete theories of action in a continuous structure, and physical and epistemic feasibility of plans with continuous action.

1 Continuous Plans

It is often useful to model the behavior of an autonomous intelligent creature in terms of continuous control and choice.

Imagine a creature that can move in the plane up to 1 meter per second but no faster, and that wants to catch unintelligent prey that moves around the plane. The hunter can detect prey up to a distance of 1 meter. A theory of plans should support conclusions like the following:

A. If the prey moves at 0.5 m/sec and is initially within sight, then the hunter can catch it by chasing it.

B. If the prey moves at 0.05 m/sec and is initially within a distance of 2 meters, then the hunter can catch it by encircling it: The hunter first goes in a circle of radius 2 meters, then in a circle of radius 1.13 meters, then in a circle of radius 0.5 meters, until seeing the prey. It then pursues the prey until catching it.[1]

*This research was supported by NSF grants #IRI-9001447 and #IRI-9300446. Thanks to Steve Hanks for helpful discussions.

[1]The plan given in the introduction AIPS-94 version of this paper is, I believe, correct, but I have not worked out a full proof, and I suspect it would be involved. Specifically, a very delicate case analysis would be required to prove that the prey cannot stay out of the range of the hunter sight, the analogue of lemma 20 here. The plan given on the last page of the AIPS paper is certainly incorrect.

C. Suppose that:

   i. The prey is currently at $\langle 0.0, 1.0 \rangle$ and the hunter is at $\langle 0.0, 0.0 \rangle$.
   ii. The prey moves at 2 m/sec in the $\hat{y}$ direction.

   Then the hunter cannot catch the prey.

D. Suppose that:

   i. The prey is currently out of view (more than a meter away).
   ii. The hunter is at the point $\langle 0.0, 0.0 \rangle$.
   iii. The hunter knows that the prey will come to a watering hole at point $\langle 0.0, 0.25 \rangle$.
   iv. The prey moves no faster than 2 m/sec

   Then the hunter can catch the prey by moving to the point $\langle 0.0, 0.25 \rangle$ and waiting there.

E. Let us modify condition (D.iii) above to read "The hunter knows that the prey will come either to $\langle 0.0, 0.25 \rangle$ or to $\langle 0.0, -0.25 \rangle$." Then the hunter cannot be sure of catching the prey.

Note that there is a difference between the impossibility in (C) above and that in (E). (C) is a physical impossibility; the hunter physically cannot catch the prey under the given constraints. (E), however, is an epistemic impossibility. If the hunter could find out which watering hole the prey would go for, it could go there and wait; however, it lacks the necessary information.

A simple model for the hunter is that it can carry out any continuous, piecewise differentiable movement in the plane with a speed that is always less than one m/sec. This is an idealization: actual creatures cannot change velocity discontinuously, they can only execute a planned motion within a certain tolerance, and, having only finite cognitive capacities, they cannot represent an infinitude of different possible behaviors. But the idealization is a useful and natural one, comparable to using a continuous model of matter to approximate the underlying discrete reality. A model that considered these limitations would be more complex and would require more detailed knowledge of the creature.

A model of such an agent can be formulated using a branching, continuous model of time. The time structure is continuous, because we categorize the agent's behavior in continuous terms. It branches, because the agent can choose different ways to go. Indeed, the time structure branches continuously; at every point, the agent has a choice of direction.

The first discussion in the AI literature of branching, continuous time was (McDermott 1982), but though that paper dealt with actions and with continuous processes, it did not study continuous actions in depth. This paper is intended to fill that gap. (Forbus 1989) presented a theory that integrates action with continuous physical change. Forbus used a discrete model of action; actions were sudden discontinuous changes that occur discretely to a continuous world. We will show that such an assumption is not technically necessary and, for many domains, not a natural assumption. (Penberthy 1993) presents a planner that constructs partial plans that include continuous actions that cause parameters to vary as a linear function of time. However, the notion of action used there and the semantics of planning is less general than that developed here. (De Jong 1994) describes explanation-based learning of control strategies. Our paper can be thought of as providing a general ontological framework for programs like those of Penberthy and de Jong.

2 Temporal Theory

The intuition behind our theory is as follows. We characterize a behavior of the agent by a time-varying state that it controls, subject to physical constraints. For example, the state might be

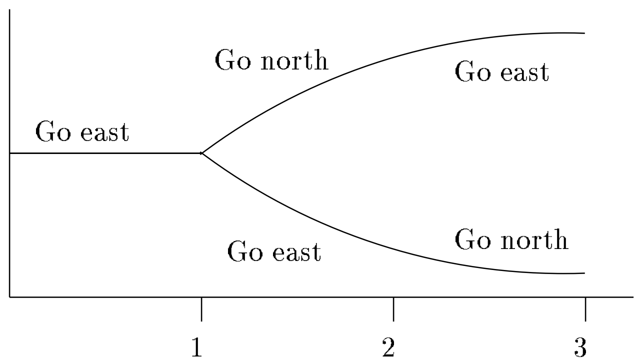

(Note: the only significance of the vertical dimension is to indicate branching.)

Figure 1: Branching Time Structure

position in the plane; or a configuration in some configuration space; or a velocity; or a collection of torques applied at joints; etc. This state can have many components, which may be discrete or continuous. A behavior of the robot is then a function from time to the state space; that is, a fluent that takes values in the state space. We call this the *behavioral fluent.*

The structure of branching time is intended to reflect the robot's choices of behavior. For example, let the state of the robot be his position in space, and suppose that the following two behaviors are possible to the robot: Behavior B1 is first to go east for 2 minutes at 1 meter per second, then to go north for 1 minute; behavior B2 is first to go east for 1 minutes at 1 meter per second, then to go north for 1 minute, then to go east for 1 minute. The time structure for these two behaviors consists of a single path for the first minute, where the two behaviors are the same; a branch point at the one minute mark, corresponding to the choice that must be made; and two separate paths for the remaining two minutes, corresponding to the two possible continuations (Figure 1). Such a time structure has two properties:

- It *exhibits* the two behaviors. That is, each of the behaviors occurs in some branch of the structure.
- It *branches on* the behavior. That is, the two branches coincide in the time structure as long as the behaviors are identical. The branches fork only when they must, because the behaviors have become different.

We will show (Theorem 1) that, given a class $\mathcal{B}$ of possible behaviors, one can construct a time structure that exhibits all the behaviors in $\mathcal{B}$ and branches on the behavioral fluent.

Each point in the time structure is a *situation* (McCarthy & Hayes 1969); that is, a snapshot of the universe. In particular, a situation is more than just a value of the behavioral fluent. Note that branching is only into the future; even though the robot ends up in the same place in $B1$ and $B2$, the final situations are not considered the same.

In this theory, we use branching in the time structure to express only the choice of the single agent. (We do not deal with more than one intelligent agent.) Uncertainty or indeterminacy in the behavior of the external world is treated as uncertainty as to the form of the overall model rather than as branching within the model. Thus, we take the view that the world in fact follows a deterministic path given the agent's behavior, though the nature of that path may be unknown.

For example, consider an agent who controls an output voltage. There is a maximum M on the voltage he can produce, which is known to be between 5 and 7 volts. The agent's output voltage is an input to an amplifier, which multiplies the signal by a factor of k, where k is known to be between 2 and 3. This situation is characterized by positing a branching time structure, in which each possible behavior that stays within M is exhibited in some branch, and in which the amplifier output is always k times the agent's output. The uncertainty in k and M do not generate branches *within* the structure; there is not one branch where the amplification is 2 and another where it is 3. Rather, they generate uncertainty *about* the structure; the amplification within the structure is only partially determined.

3 Mathematical Foundations

Our formal analysis begins with the definition and study of forward branching structures. (Compare (Shoham 1989) (van Benthem 1983).)

Definition 1: Let $\mathcal{S}$ be a set with a binary relation $X < Y$. $\mathcal{S}$ is a *forest* if the following condition holds:

i. (Anti-symmetry) If $X < Y$ then not $Y < X$.

ii. (Transitivity) If $X < Y$ and $Y < Z$ then $X < Z$.

iii. (Forward branching) If $X < Z$ and $Y < Z$ then either $X < Y$, $X = Y$, or $Y < X$. That is, any point Z has a unique, totally ordered past history.

Example: A finite forest is a forest of trees in the usual sense.

A branch in a forest corresponds to an interval in a linear ordering.

Definition 2: Let $\mathcal{S}$ be a forest and let I be a subset of $\mathcal{S}$. I is an *branch* of $\mathcal{S}$ if the following hold:

i. I is totally ordered. That is, for $X, Y \in I$, either $X < Y$, $X = Y$, or $Y < X$.

ii. I is connected. That is, if $X, Y \in I$, $Z \in \mathcal{S}$, and $X < Z < Y$, then $Z \in I$.

Definition 3: Let $\mathcal{S}$ be a forest and let D be a function from $\mathcal{S}$ to the real line $\Re$. D is called a *clock* function. $\mathcal{S}$ (strictly speaking, the triple $\mathcal{S}, <, D$) is a *continuous* forest (CF) if D has the following properties:

i. D is order preserving. That is, if $X, Y \in \mathcal{S}$ and $X < Y$ then $D(X) < D(Y)$.

ii. For any branch $I \subset \mathcal{S}$, $D(I)$ is a real interval.

iii. Let $t \in \Re$ be any clock time, and let s be any element of $\mathcal{S}$. Then there is an $s' \in \mathcal{S}$ such that $s' < s$ and $D(s') < t$. That is, branches go back to $-\infty$.

We now discuss the relation between a branching time structure and a space of possible behaviors of the agent.

Definition 4: An *initial interval* is a real interval which is unbounded below. An *initial behavior* is a function whose domain is an initial interval. Branch I is an *initial branch* if $D(I)$ is an initial interval. A *fluent* is a function whose domain is a CF.

Definition 5: Let $\mathcal{S}$ be a CF with clock function D; let I be an initial branch of $\mathcal{S}$; let F be a fluent over $\mathcal{S}$; and let B be an initial behavior. We say that F *exhibits behavior* B *on* I if the following hold:

i. $D(I)$ is equal to the domain of B;

ii. For all $s \in I$, $F(s) = B(D(s))$.

That is, over branch I, the evolution of F follows B.

Definition 6: Let $\mathcal{S}$ be a CF, and let F be a fluent over $\mathcal{S}$. $\mathcal{S}$ *branches on* $\mathcal{F}$ if the following condition holds: for any behavior B there is at most one initial branch $I \subset \mathcal{S}$ such that F exhibits B on I. That is, any two different initial branches must be distinguished by different behaviors of F.

We now proceed toward the following result: Given a set $\mathcal{B}$ of initial behaviors, we construct a CF called "$\mathrm{CF}_{\mathcal{B}}$" and a fluent called "$\mathrm{FL}_{\mathcal{B}}$" such that $\mathrm{FL}_{\mathcal{B}}$ exhibits all the behaviors in $\mathcal{B}$ and $\mathrm{CF}_{\mathcal{B}}$ branches on $\mathrm{FL}_{\mathcal{B}}$.

Definition 7: Let B_1 and B_2 be initial behaviors, with domains I_1 and I_2 respectively. B_1 is *closed* if I_1 has the form $(-\infty, T]$. B_1 is an *initial segment* of B_2 if $I_1 \subset I_2$ and B_1 is the restriction of B_2 to I_1.

Definition 8: Let $\mathcal{B}$ be a collection of initial behaviors. We construct $\mathrm{CF}_{\mathcal{B}}$ and the fluent $\mathrm{FL}_{\mathcal{B}}$ as follows:

i. X is an element of $\mathrm{CF}_{\mathcal{B}}$ iff X is a closed initial segment of some behavior $B \in \mathcal{B}$.

ii. For $X, Y \in \mathrm{CF}_{\mathcal{B}}$, $X < Y$ if X is an initial segment of Y.

iii. For $X \in \mathrm{CF}_{\mathcal{B}}$ with domain $(-\infty, T]$, $D(X) = T$.

iv. Let B be a behavior in $\mathcal{B}$; let X be a closed initial segment of B; and let $(-\infty, T]$ be the domain of X. Then $\mathrm{FL}_{\mathcal{B}}(X) = B(T)$. It is easily seen that, for fixed X, this determines the same value of $\mathrm{FL}_{\mathcal{B}}(X)$, however B is chosen.

Theorem 1: Let $\mathcal{B}$ be a collection of initial behaviors. Then $\mathrm{CF}_{\mathcal{B}}$ is a CF; $\mathrm{FL}_{\mathcal{B}}$ exhibits all the behaviors in $\mathcal{B}$; and $\mathrm{CF}_{\mathcal{B}}$ branches on $\mathrm{FL}_{\mathcal{B}}$.

Proof: We begin by defining two useful notations:

- If B is an initial behavior, then $\mathrm{Dom}(B)$ is the domain of B. Note that $\mathrm{Dom}(B)$ is an initial interval.
- Let B be an initial behavior, and let T be an element of $\mathrm{Dom}(B)$. Then $B{\downarrow_T}$ is the restriction of B to the interval $(-\infty, T]$. Note that $B{\downarrow_T}$ is an initial segment of T. Note also that, if $T1 < T$, then $(B{\downarrow_T}){\downarrow_{T1}} = B_{T1}$.

The proof is now essentially definition hunting. We make the following observations:

i. The relation "Behavior B_1 is an initial segment of B_2" is a partial ordering on initial behaviors, since it is simply the subset relation on the class of initial behaviors, viewing a behavior as a set of ordered pairs.

ii. Let B_1 and B_2 be initial segments of B_0. Then $B_1 = B{\downarrow_{T_1}}$ and $B_2 = B{\downarrow_{T_2}}$ for some $T_1, T_2 \in$ Dom(B). If $T_1 \leq T_2$ then $B_1 = B_2{\downarrow_{T_1}}$, so B_1 is an initial segment of B_2. Likewise, if $T_2 \leq T_1$ then B_2 is an initial segment of B_1. Thus, the forward branching property (Definition 1 condition iii) holds.

iii. Let $B_1, B_2 \in \mathrm{CF}_{\mathcal{B}}$, let $B_1 < B_2$, and let $\mathrm{Dom}(B_1) = (-\infty, T_1]$ and $\mathrm{Dom}(B_2) = (-\infty, T_2]$. By construction B_1 is a proper initial segment of B_2, so $T_1 < T_2$. But $D(B_1) = T_1$ and $D(B_2) = T_2$. So D is order-preserving, satisfying Definition 3, condition i.

iv. Let J be any branch in $\mathrm{CF}_{\mathcal{B}}$, and let $I = D(J)$. I is thus a set of real numbers. We wish to show that I is an interval; it suffices to show that, if $X, Y \in I$, and $X < Z < Y$ then $Z \in I$. Thus, let $B_X, B_Y \in J$ such that $D(B_X) = X$ and $D(B_Y) = Y$. Thus, $\mathrm{Dom}(B_Y) = (-\infty, Y]$. Let $B_Z = B_Y{\downarrow_Z}$. It is clear that B_X is an initial segment of B_Z and that B_Z is an initial segment of B_Y. Therefore, since J is a branch containing B_X and B_Y, it follows that $B_Z \in J$.

v. Let $T1 \in \Re$ be any time, and let B be an element of $\mathrm{CF}_{\mathcal{B}}$ with domain $(-\infty, T2]$. Choose $T < \min(T1, T2)$. Then $B{\downarrow_T} < B$ and $D(B{\downarrow_T}) < T1$.

vi. Let B be any behavior in $\mathcal{B}$. Let J be the set of all closed initial segments of B; $J = \{B{\downarrow_T} \mid T \in \mathrm{Dom}(B)\}$ It is immediate from the definitions that J is a branch of $\mathrm{CF}_{\mathcal{B}}$ and that $\mathrm{FL}_{\mathcal{B}}$ exhibits behavior B on J.

vii. Let B be any initial behavior and let J, as in (vi), be the set of all closed initial segments of B. We wish to show that J is the only branch on which $\mathrm{FL}_{\mathcal{B}}$ exhibits behavior B. Proof by contradiction: Let $J1$ be another such branch. We note that $D(J1) = D(J) = \mathrm{Dom}(B)$. Therefore, $J1$ is not a proper initial segment of J. So let S be an element of $J1 - J$. Since $S \in J1$, $D(S) \in \mathrm{Dom}(B)$. Since $S \notin J$, S is not an initial segment of B. Thus there is a value $T \leq D(S)$ such that $S(T) \neq B(T)$. Certainly $S{\downarrow_T} \in J1$. But $\mathrm{FL}_{\mathcal{B}}(S{\downarrow_T}) = S{\downarrow_T}(T) = S(T) \neq B(T)$. Thus $\mathrm{FL}_{\mathcal{B}}$ does not exhibit behavior B on $J1$. This completes the contradiction.

Items (i) – (v) establish that $\mathrm{CF}_{\mathcal{B}}$ is an CF. Item (vi) establishes that $\mathrm{FL}_{\mathcal{B}}$ exhibits all the behaviors in $\mathcal{B}$. Item (vii) establishes that $\mathrm{CF}_{\mathcal{B}}$ branches on $\mathrm{FL}_{\mathcal{B}}$. Thus the proof of theorem 1 is complete. $\Box$.

Theorem 2: $\mathrm{CF}_{\mathcal{B}}$ and $\mathrm{FL}_{\mathcal{B}}$ are the unique minimal CF and fluent satisfying the conclusions of Theorem 1, up to isomorphism. More precisely: Let $\mathcal{B}$ be a collection of initial behaviors. Let $<_R$ and D_R be the order relation and the clock function associated with $\mathrm{CF}_{\mathcal{B}}$. Let $\mathcal{S}$ be an CF with order relation $<_S$ and clock function D_S Let $\mathcal{F}$ be a fluent over $\mathcal{S}$ such that $\mathcal{F}$ exhibits all the behaviors in $\mathcal{B}$ and such that $\mathcal{S}$ branches on $\mathcal{F}$. Then there is a subset $\mathcal{S}_{\mathcal{B}}$ of $\mathcal{S}$ and an isomorphism ϕ from $\mathrm{CF}_{\mathcal{B}}$ to $\mathcal{S}_{\mathcal{B}}$ such that

i. ϕ is one-to-one and onto.

ii. For any $B1, B2 \in \mathrm{CF}_{\mathcal{B}}$, $B1 <_R B2$ if and only if $\phi(B1) <_S \phi(B2)$.

iii. For any $B \in \mathrm{CF}_{\mathcal{B}}$, $D_R(B) = D_S(\phi(B))$.

iv. For any $B \in \mathrm{CF}_{\mathcal{B}}$, $\mathrm{FL}_{\mathcal{B}}(B) = \mathcal{F}(\phi(B))$

Proof: Since $\mathcal{S}$ is a forest, it follows that for any element $S \in \mathcal{S}$ there is a unique initial branch whose largest element is S. Let us say that an element $S \in \mathcal{S}$ corresponds to a behavior $B \in \mathrm{CF}_{\mathcal{B}}$ if the fluent $\mathcal{F}$ exhibits behavior B over the initial branch ending in S. Since $\mathcal{F}$ exhibits all behaviors in $\mathcal{B}$, it follows that every behavior $B \in \mathcal{B}$ has a corresponding element $S \in \mathcal{S}$. Since $\mathcal{S}$ branches on $\mathcal{F}$ it follows that this corresponding element is unique. We can therefore define $\phi(B)$ to be the unique element of $\mathcal{S}$ corresponding to B. We define $\mathcal{S}_\mathcal{B}$ to be the image of $\mathrm{CF}_\mathcal{B}$ under ϕ. The above properties then follow immediately.$\Box$

Theorem 3: Let $\mathcal{S}$ be any CF. Then there is a family of behaviors $\mathcal{B}$ such that $\mathcal{S}$ is isomorphic to $\mathrm{CF}_\mathcal{B}$.

Like many such results, the statement seems interesting until you see the proof, which is perfectly circular.

Proof: We construct a family of behaviors whose range is $\mathcal{S}$ itself. For any element $S \in \mathcal{S}$, and for $T \leq D(S)$, let $\theta(S,T)$ to be the unique element of $\mathcal{S}$ such that $D(\theta(S,T)) = T$ and $\theta(S,T) \leq S$. Let $\psi(S)$ be the behavior B such that $\mathrm{Dom}(B) = (-\infty, D(S)]$ and, for any $T \leq D(S)$, $B(T) = \theta(S,T)$. Then it is immediate that $\mathcal{S}$ is isomorphic to $\mathrm{CF}_\mathcal{B}$.

There are two problems with this construction. First, $\mathrm{CF}_\mathcal{B}$ may contain behaviors that are not in $\mathcal{B}$. Second, even if all the behavior in $\mathcal{B}$ have domain $(-\infty,\infty)$, there may be maximal branches in $\mathcal{B}$ of finite length. Intuitively, $\mathcal{B}$ represents constraints on the possible behavior of the agent. $\mathrm{CF}_\mathcal{B}$ then suggests that, by shifting behavior infinitely often, the robot can either violate the constraint, or bring time itself to a sudden end. (Davis 1992a)

Example: Let $\mathcal{B}$ be the space of all continuous bounded functions on $(-\infty,\infty)$. For $k = 2, 3, \ldots$ define the function

$$c_k(t) = \begin{cases} \frac{1}{1-t} & \text{for } t \leq 1 - \frac{1}{k} \\ k & \text{for } t \geq 1 - \frac{1}{k} \end{cases}$$

Then by executing $c_2(t)$ before $t = 1/2$, $c_3(t)$ from $t = 1/2$ to $t = 2/3$, $\ldots$, the robot ends up executing $c_\infty(t) = 1/1 - t$ over the interval $(-\infty, 0)$. This behavior is not an initial segment of any behavior in $\mathcal{B}$ and cannot be extended to $t = 0$ within the given constraint. The corresponding branch in the time structure terminates at time 0. (Figure 2)

In general, a behavior B is exhibited in $\mathrm{CF}_\mathcal{B}$ iff every closed initial segment of B is also an initial segment of some behavior in $\mathcal{B}$. Such a behavior is called an *initial limit* of $\mathcal{B}$.

Two solutions to this problem may be suggested. One is to mark certain branches in the CF as disallowed. Thus a branch like $c_\infty(t)$ that diverges is considered impossible, though every closed segment is possible. The other is to require $\mathcal{B}$ to be closed under the taking of initial limits. For example, these two collections of functions satisfy that condition:

- The class of initial functions bounded by M.
- The class of functions b satisfying the Lipschitz condition $\mid b(x) - b(y) \mid < M \mid x - y \mid$.

4 Causal Theories

One advantage of a discrete theory is that there is a simple basic form for causal theories: a causal theory specifies the result of performing an action in terms of a transition function from one situation to the next. Continuous theories, having no "next" situation, are harder to characterize. Continuous time theories in science are usually posed as differential equations; however, these are often unsuitable to commonsense reasoning (Davis 1988).

The AI literature contains a variety of formal characterizations of continuous physical domains (e.g. (Hayes 1985), (Kuipers 1986), (Sandewall 1989), (Davis 1990)) but no standard framework for such theories has emerged. The incorporation of volitional action into these theories is problematic, in general. Sometimes it is straightforward. For example, in electronic systems, such as those in ENVISION (de Kleer & Brown 1985) an agent can be viewed as generating an exogenous signal.

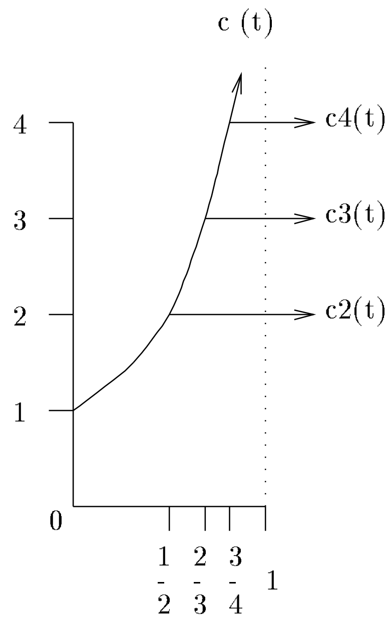

Figure 2: Divergent Behavior

In the microworld of the hunter following prey, the hunter and prey move independently; the only causal rule is that the hunter catches the prey when they are at the same location.

Other domains are harder. Consider the kinematic theory of rigid solid objects. The physical theory is simple (see (Davis 1990), p. 332). There is a collection of objects. An object may be fixed or mobile. Each object has a fixed shape, which is a spatial region satisfying certain regularity conditions. The "place" of an object O is a fluent, whose value in each situation is a region congruent to the shape of O. The "placement" of an object is a fluent whose value in each situation is a rigid mapping from the shape to the place. The theory is characterized by the following axioms:

1. The place of O in situation S is the image of the shape of O under the placement of O in S.
2. The placement of O in S is always a rigid mapping.
3. If $O1 \neq O2$, then the place of $O1$ in S does not overlap the place of $O2$ in S.
4. If O is fixed, then the placement of O is constant.
5. The placement of O is a continuous function of time.

Rules (1) to (4) are domain constraints, which refer only to a single situation S. Rule (5) has non-trivial time-dependence, but it has a simple and standard form.

However, if one of the objects has autonomous choice of motion, then the following rule holds: Any motion of the agent is possible, as long as the other objects can move so as to avoid overlapping it. This rule is much harder to express. (Consider particularly the case where several objects combine to block the agent. (Figure 3)) The simplest formulation that I have found is as follows. We define three new sorts. A *configuration* is a placement of all the objects. A *history* is a function from

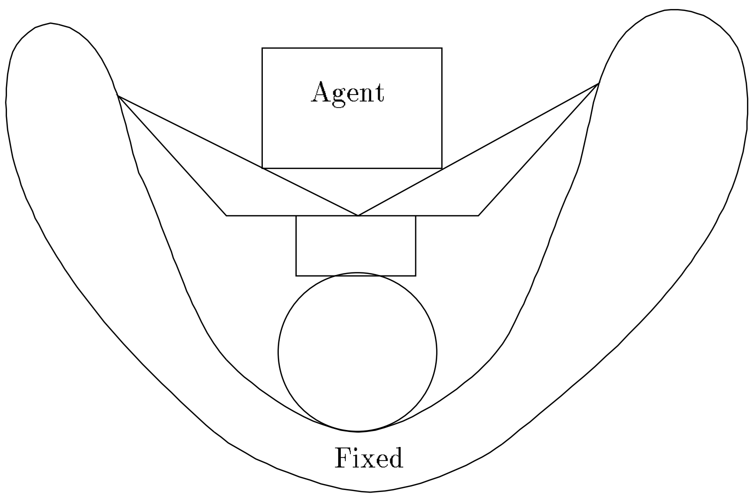

The agent cannot move downward.

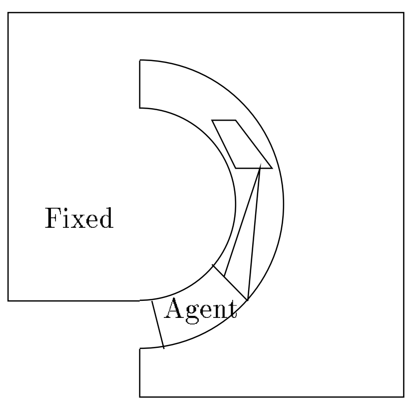

The agent can execute a forward,

upward, twisting motion.

Figure 3: Kinematic Constraints

time to configurations. A history is *feasible* if it is continuous and differentiable, it does not cause objects to overlap, and it leaves every fixed object in a constant position. A *differential motion* is a combination of a translational velocity and an angular velocity; that is, the derivative of a placement. We now state the following rule:

6. Let S be a situation and let C be the configuration of objects in S. Let M be a differential motion of C. If there exists a feasible history H such that the configuration of H at time 0 is equal to C and the derivative of the placement of the agent in H at time 0 is equal to M, then there exists a branch $[S, S1]$ in the time structure such that the derivative at S of the motion of the robot exhibited over $[S, S1]$ is equal to M.

That is, if there is any way for the objects to move around from S so as to permit the agent to move in direction M, then the agent can indeed move in direction M and the objects will move in some permitted manner.

The contrast between the complexity of rule (6) and the simplicity of rules (1-5) is stark. It may possible to derive (6) from a default rule of the form, "A differential motion of the agent is feasible unless it is forbidden by rules (1-5)." However, finding a non-monotonic theory of this form and establishing that it gives all and only suitable answers is hard.

Adding agents to a dynamic theory of objects is easier, at least at the level of metaphysical adequacy (McCarthy & Hayes 1969). Take an agent to be a jointed collection of rigid parts, and characterize its behavior in terms of the torques that it creates at joints. The agent then interacts with the outside world via normal and frictive forces at its surface. The extension is formally simple, but since it operates in terms of joint torques it is nothing like epistemilogically adequate.

5 Embedding Discrete Actions

The standard discrete situation can be easily embedded within a continuous branching time structure as follows. For each situation calculus event E, define a possible value of the behavioral fluent to be a pair $\langle E, T \rangle$ where $0 < T \leq 1$. (We assume that all such events take unit time. This assumption can be modified to have elapsed time depend on the event and the starting situation.) T here represents the fraction of the duration of the execution of E. We can then define the "result" function as follows:

$$
\begin{array}{l}
S1 = \text{result}(E, S0) \Leftrightarrow \\
[S1 > S0 \land \text{clock}(S1) - \text{clock}(S0) = 1 \land \\
\forall_S \; S0 < S \leq S1 \Rightarrow \\
\qquad \text{behavior}(S) = \langle E, \; \text{clock}(S) - \text{clock}(S0) \; \rangle].
\end{array}
$$

That is, during the execution of E, the behavioral fluent always indicates that E is being executed, and the clock for the event advances from 0 to 1.

6 Physical Feasibility of Plans

If the semantics of plan P are defined behaviorally — that is, necessary and sufficient conditions have been given for the assertion "P is executed over interval I" — then the definitions of the physical feasibility and correctness of P is the same as for discrete time. If P is determinate, then the definitions are straightforward:

Definition 9: Determinate plan P is *feasible* in situation S if P is executed over some interval $[S, S1]$.

Definition 10: Determinate plan P *achieves* goal G starting from situation S if, for some $S1$, P is executed over $[S, S1]$ and G holds in $S1$.

To categorize indeterminate plans, we augment the time structure by allowing the behavioral fluent to assume the value "fail". Once the agent enters the failing state, it remains there forever. We then define the semantics of plans so that, if a plan intuitively requires the execution of an action that is currently infeasible or undefined, the agent in fact executes "fail". In such a structure, correctness of plans can be defined thus:

Definition 11: Plan P is *possibly feasible* in situation S if, for some non-failing situation $S1$, P is executed over the interval $[S, S1]$.

Definition 12: Plan P is *necessarily feasible* in situation S if, for every situation $S1$ such that P is executed over the interval $[S, S1]$, $S1$ is not a failing situation.

Definition 13: Plan P *necessarily (possibly) achieves* goal G from situation S if the plan "**begin** P; if G then **no-op** else **fail end**" is necessarily (possibly) feasible in S.

7 Epistemic Feasibility of Plans

The problem of characterizing the conditions under which a discrete plan is epistemically feasible is noted in (McCarthy & Hayes 1969) and was first studied at length by Moore (1985). In (Davis 1994) I propose the following definitions, modified from Moore's. These apply both to determinate and to indeterminate plans. "Executability" corresponds to necessary feasibility; it means that the agent can carry out the plan by executing one step at a time, with no thought except understanding what the plan says to do next. "Epistemic feasibility as a task" corresponds to possible feasibility; it means that, if the agent is assigned to carry out the plan, then, by thinking hard, he can find a way to do it.

Definition 14: Plan P *begins* over interval $[S1, S2]$, if there is an $S3 \geq S2$ such that P executes over $[S1, S3]$.

Definition 15: Plan P' is a *specialization* of plan P in S if every execution of P' starting in S is also an execution of P.

For instance, the plan "begin A;B end" is a specialization of the plan "do both A and B in any order."

Definition 16: Discrete plan P is *executable* for agent A in situation S iff for any $S2$, if P begins over $[S, S2]$ then

i. A knows in $S2$ whether P has completed over $[S, S2]$;

ii. A knows in $S2$ what are all the actions that constitute a next step of P after $[S, S2]$; and

iii. "fail" is not a next step of P after $[S, S2]$.

Definition 17: Discrete plan P is *epistemically feasible as a task* for agent A in situation S if there is a plan P' such that A knows in S that:

i. P' is a specialization of P;

ii. P' is executable in S.

I show (Davis 1994) these definitions are reasonable for a few simple examples and that they have a number of natural properties:

- The more one knows, the more plans are epistemically feasible.
- For an omniscient agent, a plan is executable iff it is necessarily feasible; it is epistemically feasible as task iff it is possibly feasible.
- Moore's (1985) rule[2] for sequences: The plan "sequence($P1, P2$)" is epistemically feasible for A in S if $P1$ is epistemically feasible for A in S and A knows in S that, after executing $P1$, $P2$ will be epistemically feasible.
- Moore's rule for conditionals: The plan "if Q then $P1$ else $P2$" is epistemically feasible for in S if either [A knows in S that Q is true and P1 is epistemically feasible in S] or [A knows in S that Q is not true and $P2$ is epistemically feasible in S].

These definitions have a straightforward generalization to the continuous case. We posit that an agent cannot react instantaneously to perceptions. Rather, that there must be a delay of at least $\Delta > 0$ between perception and reaction.

Admittedly, this requirement of a delay can be a little clumsy. Plans, such as (B) of the introduction, are more naturally stated in a form that idealizes the agent as responding instantaneously ("*As soon as* you see the prey do ..."). However, in looking for a definition of epistemic feasibility that admits instantaneous response, I have always gotten into trouble with unsolvable feedback systems. For example, consider the plan "Move toward the target with speed 1 m/sec greater than the speed of the target". Ordinarily, the feasibility conditions of this plan are easily stated: It is physically feasible if the agent can move faster than the target is moving, and it is epistemically feasible if the agent knows the direction and speed of the target. However, in the special case where the target is attached by a stick to the agent, the plan becomes unachievable, even if both conditions are satisfied. It is analogous to the discrete plan "Put block A on a block which will be clear at the end of the put action." It is not clear to me how this unachievability should be characterized, or how the unachievable case can in general be distinguished from the achievable case.

Under this assumption that there is a delay, we can generalize definitions 16 and 17 to the continuous case. The generalization involves one tricky issue. Let me proceed therefore by presenting first a simple definition that is not quite correct, explaining the gap, and then presenting the more complicated correct definition.

Definition 18.0: (Not quite correct)
A plan P is *executable* for agent A in situation S with delay Δ iff
for any $S1$,
  if A begins to execute P from S to $S1$
    then A will know in $S1$,
      whether P will complete within time Δ
      and how to continue P for time Δ.

The problem is that this definition allows the following scenario: The robot gains some information at time $T_I > T_0 - \Delta$ before it starts execution of the plan at time T_0, and uses that information to guide its execution of the plan. In this scenario, the robot has a reaction time of $T_0 - T_I < \Delta$. We therefore modify the definition to require that the robot knows how to begin the plan for a time Δ before it actually does so.

[2] Contrary to the statement in Moore, these rules are sufficient but not necessary conditions.

Definition 18
A plan P is *executable* for agent A in situation S with delay Δ iff
for any $S1$,
if either A begins to execute P from S to $S1$
or $S1$ precedes S by no more than Δ,
then A will know in $S1$,
whether, starting in S, P will complete within time Δ of $S1$,
and how to continue P after starting in S up to time Δ after $S1$.

Definition 19:
A plan P is *epistemically feasible as a task* for agent A in situation S with delay Δ iff
there is a plan $P1$ such that
A knows in S that
$P1$ is executable for A in S with delay Δ and
$P1$ is a specialization of P.

The above properties of definitions 16 and 17 also apply to definitions 18 and 19. Also, these definitions are monotonic in Δ; if a plan is epistemically feasible with one value of Δ, then it is epistemically feasible with any smaller value of Δ.

Let me emphasize that these definitions do not involve a discretization of time; they introduce a constant delay at every instant of the time line. Thus, at every instant, the agent can make use of all his perceptions up to time Δ before.

In the case of discrete actions, definitions 18 and 19 reduce to definitions 16 and 17, provided that there is always a delay of duration at least Δ between the time that the agent gains information from one step and the time that he applies it in the next step. For example, one could require that the agent always wait for duration Δ between deciding on his next step and executing it.

To fit definition 18, plan (B) must be changed to admit delay between perception and action.

- The hunter must be allowed a delay between seeing the prey and changing from circling to pursuing.
- Thus, the hunter must wait till the prey is well within view, not on the horizon, so that it does not vanish the small interval before he begins pursuit.
- In pursuit, the hunter cannot move toward the current position of the prey. Rather, he moves toward some position the prey has occupied within time Δ.

8 The formal analysis of a plan

We now show how plan (B) can be formally verified as physically and epistemically feasible. Our discussion will be at the ontological level; that is, we will describe a formal model and show that the conclusion holds in the model, but we will not give a full symbolic axiomatization of the model or a full symbolic proof.

8.1 The physical microworld

In this example, the world at an instant consists only of the hunter and its prey, and the only significant characteristics of these are their spatial positions. For these purposes, therefore, we can construct our temporal model as follows: There are two fluents "hunter" and "prey". The value of "prey" in any situation is a point in $\Re^2$. The value of "hunter" in any situation is a pair $\langle \mathbf{X}, F \rangle$ where $\mathbf{X}$ is a point in $\Re^2$ and F (for "failure") is a Boolean value. The set of all values of such a pair will be called "hunter state space." The fluent of the spatial component of "hunter" will be called "h_place" and the fluent of the failure component will be called "h_fail".

We will assume that the positions of the hunter and the prey are continuous and continuously differentiable *from the right* (later times). That is, the limit

$$\lim_{\epsilon \rightarrow 0+} \frac{\mathrm{Q}(t+\epsilon) - \mathrm{Q}(t)}{\epsilon}$$

always exists and is continuous from the right as a function of t. Thus whenever we use the expressions "derivative" we will mean "derivative from the right". Similarly "velocity", "speed" and related expressions are defined in terms of the derivative from the right.

A time structure $\mathcal{T}$ is a CF of situations, branching on the fluent "hunter".

Until the hunter fails, he can execute all and only right-differentiable continuous motions with speeds up to 1 meter per second. The hunter can fail after any situation. Once he has failed, h_fail remains true forever, and by convention we say that he remains in the same place.

We thus have the following definitions:

Definition 20: Let $F(t)$ be a function from the reals to hunter state space. Let I be an initial interval. F is said to be a *realizable hunter behavior* over I if it satisfied the following condition:

There is an initial interval $I1$ such that
(A) Either $I1 = I$ or $I1 = (-\infty, L] \subset I$.
(B) Over $I1$, h_place is right-differentiable, and $\mid$ h_place$'(T) \mid \leq 1$ and
h_fail is always equal to FALSE; and
(C) Over $I - I1$, h_place has the constant value h_place(L) and
"h_fail" is always equal to TRUE.

Definition 21: Let $G(t)$ be a function from $\Re$ to $\Re^2$. G is said to be a *realizable prey behavior* if it is continuous and right-differentiable, and the magnitude of the derivative is no greater than 0.05.

Definition 22: A *realizable time structure* is a CF of situations such that

- Only realizable hunter behaviors are exhibited by hunter.
- Every realizable hunter behavior is exhibited by hunter.
- Only realizable prey behaviors are exhibited by prey.

The simple form of definition 22, where any behavior of "hunter" can be combined with any behavior of "prey", reflects the trivial causal structure of the microworld. Richer worlds, where past behaviors of external objects affect the current choices open to the agent, are harder to characterize. open

8.2 Plan

Plan (B) was originally described in the introduction as follows:

> The hunter first goes around in a circle at radius 2 meters, then at radius 1.13, then at radius 0.5. It then pursues the prey until catching it.

As described in section 7, our theory of epistemic feasibility requires a time delay between the hunter's perception of a fact and his reaction to that fact. This affects the above plan in four respects. First, there must be a time interval between seeing the prey and switching from the encircling part of the plan to the pursuing part of the plan. Second, to make sure that the prey does not escape out of sight during that time interval, we will modify the switchover condition to read that the hunter sees the prey sufficiently far from the limits of his perception. Third, "pursue" must be defined so that the hunter is not required to run precisely in the current direction of the prey, but may run towards the prey's position in the recent past. Finally, the criterion of success cannot be that the hunter reaches the same point as the prey, but that he comes close enough to it.

As a result of this modification, this plan becomes an indeterminate one, because the delay in responding is not fixed, though it is bounded.

Our formal representation of the above plan is as follows:

hunt1 =
sequence(monitor($\lambda(T)$ distance(h_place(T), prey(T)) $\leq$ 0.95,
    sequence(go_circle(0,2), go-circle(2,1.13), go-circle(1.13, 0.5))
    0.01),
  follow(prey, 1, 0.01, 0.01))

go_circle(A, B) =
sequence(go($\lambda(T)$ $\langle 0, A + T\text{sign}(B-A)\rangle$, $\mid B - A \mid$),
  go($\lambda(T)$ $\langle B\cos(T/B), B\sin(T/B)\rangle$, $2\pi B$)).

go($PATH, D$) = for_duration(D,attempt(going($PATH$))).

follow(X, G, L, Δ) =
monitor(dist(h_place,X) $\leq L$,
  forever(attempt(following(X, G, Δ)))).

The meaning of the above primitives is as follows:

- going($PATH$) — $PATH$ is a function of time. Starting at time $_0$, going($PATH$) is the action of moving so that at time $T \geq T_0$, the position of the hunter is PATH($T - T_0$).
- following(X, G, Δ) — X is a fluent whose value in each situation is a position. G is a speed. Δ is a time duration. The action following(X, G, Δ) consists of the hunter moving toward X at speed G with time delay Δ. The precise semantics of the time delay is given in definition 27 below.
- attempt(E) is the plan of attempting to carry out E, or failing if E is not feasible.
- for_duration(D, E) is the plan of carrying out E for time duration D.
- forever(E) is the action of carrying out E forever.
- go($PATH, L$) is thus the action of attempting to go along path P for duration L.
- go_circle($R1, R2$) is thus the following action: Starting at $\langle 0, R1\rangle$, go to $\langle 0, R2\rangle$ then traverse a circle of radius $R2$ around the origin.
- follow(Z, G, L, Δ) is thus the action of attempting to follow Z at speed G with delay Δ for duration L.

- sequence($P_1, P_2 \ldots P_k$) — $P_1 \ldots P_k$ are plans. sequence($P_1 \ldots P_k$) is the plan of doing P_1 through P_k in sequence.
- monitor(Q, P, Δ) — Q is a Boolean fluent. P is a plan. Δ is a time duration. monitor(Q, P, Δ) is the following plan: Execute P, monitoring fluent Q. If Q ever becomes true, then, within time Δ, terminate the "monitor" plan.[3]

In our formal definition, we use "action" and "plan" synonymously A plan is viewed as just a complicated action.

We define the semantics of a plan in terms of the following primitives:

- executes(P, B, Z) — Predicate. Plan P executes completely over branch B of time structure Z.
- succeeds(P, B, Z) — Predicate. Plan P executes successfully over branch B of time structure Z.
- begins(P, B, Z) — Predicate. Plan P begins execution over branch B of time structure Z.

The predicates "succeeds" and "begins" are defined in terms of "executes" and the fluent "h_fail" as follows:

Definition 23: P succeeds over B in T iff P executes over B in T and "h_fail" is false throughout B.

Definition 24: P begins over $B1$ in T iff there is a branch $B2$ of T such that $B1$ is an initial segment of $B2$; P executes over $B2$ in T; and h_fail is false throughout $B1$.

We now proceed with the formal definition of the primitives in the plan above. We first define the semantics of the primitive actions "going" and "following".

Definition 25: Let B be a branch with lower bound B_0 and let $T_0 = \text{clock}(B_0)$. The action "going($PATH$)" is executed over branch B if, for every $S \in B$, $PATH(\text{clock}(S) - T_0) = \text{h\_place}(S)$.

We define the semantics of "following(X, G, Δ)" in terms of a delay function $\Theta(S)$. In any given situation S, the hunter is reacting to a position of G in some situation shortly preceding S. That situation is $\Theta(S)$.

Definition 26: The function $\Theta(T)$ mapping time T to a time is a *delay function bounded by* Δ if the following conditions hold:

i. $T \geq \Theta(T) \geq T - \Delta$

ii. $\Theta(T)$ is monotonically non-decreasing, continuous, and continuously differentiable from the right.

The function $X(\Theta(T))$, the position of X at the delayed time $\Theta(T)$, will be informally called the "ghost" of X.

Definition 27: Action following(X, G, Δ) is executed over interval I if there is a delay function $\Theta(T)$ bounded by Δ such that:

i. If $T \in I$ and $\text{h\_place}(T) \neq X(\Theta(T))$ then the velocity of the hunter is equal to G times the unit direction from $\text{h\_place}(T)$ to $X(\Theta(T))$. That is, if the hunter is not at the ghost of X, then he moves towards the ghost of X with speed G.

[3] In [Davis, 92b], the role of Δ was imbedded in the overall semantics of execution. Incorporating it directly in the control structure leads, I think, both to a more flexible representation and to a cleaner semantics.

ii. If $T \in I$ and h_place$(T) = X(\Theta(T))$ then: Let $V_\Theta(T) = \frac{d}{dT} X(\Theta(T))$. Then the speed of the hunter at time T is equal to $\min(G, V_\Theta(T))$ and the direction of the velocity is parallel to $V_\Theta(T)$. That is, if the hunter is current in the same place as the ghost of X then he tries his best to keep up with the ghost of X, subject to the constraint that the hunter cannot move faster than G. Note that the ghost of X may move faster than X can, if $\Theta(T)$ rapidly catches up to the present.

We next define the semantics of the control structures used in the plan.

The function "attempt(E)" applies only to actions E that are *beginnable* in the sense of the following definition:

Definition 28: Action E is beginnable if it satisfies the following condition: If E is executed over branch B and $B1$ is an initial segment of B, then E is executed over $B1$.

If E is beginnable, then the execution of the attempt to do E is related to the execution of E in the following definition:

Definition 29: Let B be a branch of time structure Z with lower bound S_0; let "clock(S)" be the clock function associated with Z; and let $T_0 = \text{clock}(S_0)$. The action "attempt(E)" is executed over B in Z if either

a. E is executed over B and, for all $S \in B$, h_fail$(S) =$ false; or

b. There is no branch B' with lower bound S_0 such that E is executed over B' and, for all $S \in B$, h_fail$(S) =$ true; or

c. There exists an initial segment $B1$ of B such that

   c.i. E is executed over $B1$ and, for all $S \in B1$, h_fail$(S) =$ false;

   c.ii. There does not exist a $B2$ such that $B1$ is a proper initial segment of $B2$ and such that E is executed over $B2$; and

   c.iii. For all $S \in B - B1$, h_fail$(S) =$ true.

In case (a), the hunter successfully executes E throughout B. In case (b), he has no way to execute E for any length of time so he immediately fails. In case (c), he (i) executes E successfully during the initial segment $B1$; (ii) he then finds that he cannot continue executing E; (iii) he therefore fails for the remainder of the time.

(The planning construct used above, in which the hunter plans to attempt certain actions — that is, he in effect plans to fail if he gets stuck — is certainly counterintuitive. You do not plan to fail; you plan to succeed, you attempt to carry out your plan, and you may fail in your attempt. The object here is to get around the following technical problem: To "attempt" a plan means to begin it and to take it as far as you can. To define this, we need to define "beginning" the execution of a plan that may be impossible to complete successfully. One way to do this, pursued in [Davis, 92b] would be to define the semantics of beginning a plan separately from the semantics of completing it. However, this leads to a clumsy and redundant theory. The approach pursued here is to define "executing" a plan to include all the ways in which the plan can be begun and taken either to the point of successful completion or failure. This gives a more compact and elegant theory, but at the cost of this anomaly, that the plan includes the possibility of failure from the beginning.)

Lemma 4: If E is beginnable, then attempt(E) is beginnable.

Proof: Immediate from the definitions.

Definition 30: The plan "for_duration(D, E)" is executed over branch B if E is executed over B and the duration of B is equal to D.

Definition 31: The plan "forever(E)" is executed over branch B if E is executed over B and B is unbounded above.

Definition 32: Let B be a branch. If B is bounded above, let B_f be the upper bound of B and let T_f = clock(B_f); else, let $T_f = \infty$. The plan "monitor(Q, P, Δ)" is executed over B if the following conditions hold:

A. P is begun over B;

B. If $S \in B$ and clock(S) $< T_f - \Delta$, then Q does not hold in S.

C. One of the following holds:

C.i. B is bounded above, and P is executed over B.

C.ii. B is unbounded above.

C.iii. There is a situation $S_Q \in B$ such that clock(S_Q) $\geq T_f - \Delta$ and such that Q holds in S_Q.

That is, the agent begins the execution of P. If Q ever becomes true, then the execution of P is interrupted within time Δ; otherwise, it continues to completion.

The definition of sequence is standard.

Definition 33: A *k-partition* of branch B is a sequence of branches B_1, $B_2 \ldots B_k$ such that

- $B = B_1 \cup B_2 \cup \ldots \cup B_k$.
- If $i < j$, $S_i \in B_i$, $S_j \in B_j$ then $S_i \leq S_j$.

Definition 34: The plan "sequence($P_1 \ldots P_k$)" is executed over B if there is a k-partition $B_1 \ldots B_k$ of B such that P_i is executed over B_i for $i = 1 \ldots k$.

8.3 Physical Feasibility

The physical correctness of hunt1 is now established in the following theorem:

Postulate 1: In situation s0, h_place(s0) = $\langle 0, 0\rangle$; h_fail(s0) = false; and prey(s0) is a point within distance 2.5 of $\langle 0, 0\rangle$.

Theorem 5: In s0, plan hunt1 is necessarily feasible. Every execution of hunt1 starting in s0 is finite and ends in a situation where the distance from hunter to prey is less than or equal to 0.0005.

Proof: The proof of physical correctness for this plan is actually easier than one might suppose, because it is almost independent of the initial encircling actions. In this microworld a "follow" action always ends with the hunter close to the prey. The purpose of the encircling action is to bring the hunter within view of the prey, in order to establish the epistemic preconditions for the following. The physical preconditions of following are always satisfied.

The proof involves the following steps:

1. Show that the plan step go($\lambda(T)$ $\langle 0, T\rangle$, 2) is determinate and that it is executed successfully over a unique branch [s0,s1] starting in s0. Show that h_place(s1) = $\langle 0, 2\rangle$.
2. Similarly argue that each of the other "go" steps of the plan is determinate and feasible following the previous steps.

3. From (1) and (2), show that the sequence of the three "go_circle" steps is determinate and necessarily feasible in s0.
4. Lemma: If P is necessarily feasible in S, then, for any fluent Q and duration Δ, monitor(Q, P, Δ) is necessarily feasible in S. Thus, the "monitor" subplan of hunt1 is necessarily feasible in s0.
5. For any situation S and for any $\Delta > 0$, there is an unbounded branch B starting in S over which the action "following(prey, 1.0, Δ)" is executed.
6. The action "attempt(following(prey, 1.0, Δ))" never ends in failure, and can be performed over arbitrarily long durations. This follows directly from (5), the definition of "attempt" and the fact that "following" is a liquid event in the sense of [Shoham, 88]; that is, if following is executed over two consecutive intervals, then it is executed over their union.
7. Suppose that "following(X, G, Δ)" is executed over interval I. Let T_0 be the start of I; let L_0 be the distance from the hunter to X in S_0; and let v_m be the maximum velocity attained by X. Then, at any time in I after $T_0 + (L_0/(G - v_m))$, the distance from the hunter to X is at most Δv_m.
8. In any situation S, the action "follow(X, G, L, Δ)" is necessarily feasible. At the end of the execution of "follow(X, G, L, Δ)" the hunter will be at a distance not more than max($L, \Delta v_m$) from the prey. This follows by combining (5) and (7) with the definition of "follow"
9. In any situation S the plan hunt1 is feasible. At the end of the execution of hunt1, the hunter is at a distance not more than $\Delta v_m = 0.005$ from the prey. This follows by combining (4) and (8) with the definition of hunt1.

We will present detailed proofs of steps 5 and 7; the other steps are simple.

Let $\Theta(t)$ be the delay function governing the "following" action. Let $\vec{H}(t)$ be the position of the hunter at time t. Let $\vec{X}(t)$ be the position of the prey. Let $\vec{J}(t) = \vec{X}(\Theta(t))$ be the ghost of the prey. Let v_m be the maximum speed of the prey. For any vector $\vec{V}$ let direction($\vec{V}$) $= \vec{V}/ \mid \vec{V} \mid$ be the unit vector in the direction of V. (direction($\vec{0}$) is an arbitrary unit vector.) We posit that $\Theta(t)$ observes the constraints of definition 25 and that $\vec{X}(t)$ observes the constraints of definition 21.

We begin the proof of 5 by defining abbreviations for the three functions that are used to determine the velocity of the hunter:
Let $\vec{F}1(\vec{v}, t) = G \cdot \text{direction}(\vec{J}(t) - \vec{v})$; let $\vec{F}2(t) = G \cdot \text{direction}(\vec{J}'(t))$; and let

$$\vec{F}(t) = \begin{cases} \vec{F}1(\vec{H}(t), t) & \text{if } \vec{H}(t) \neq \vec{J}(t) \\ \vec{F}2(t) & \text{if } \vec{H}(t) = \vec{J}(t) \text{ and } G <\mid \vec{J}'(t) \mid \\ \vec{J}'(t) & \text{if } \vec{H}(t) = \vec{J}(t) \text{ and } G \geq\mid \vec{J}'(t) \mid \end{cases}$$

Lemma 6: The behavior "following(X, G, Δ)" is executed over I if the equation $\vec{H}'(t) = F(t)$ is satisfied over I.

Proof: Immediate from the definition.

Lemma 7: If the following two conditions hold:

1. $G > v_m$.
2. For all t, $\Theta'(t) = t - \Delta$

then, for any initial value $\vec{H}(0)$, there exists a solution to the equation $\vec{H}'(t) = F(t)$ over the interval $[0, \infty)$. (I believe that condition 2 is actually unnecessary, but it makes the proof much simpler, and it is sufficient for the main theorem to be proven. Also, the solution is in fact unique, but, again, we don't need that property.)

Proof: Note that $| \vec{J}'(t) |= \Theta'(t) | \vec{X}'(\Theta(t)) | \leq v_m < G$. Thus if $\vec{H}(t) = \vec{J}(t)$ over an interval, then during that interval $\vec{F}(t) = \vec{J}'(t) = \vec{H}'(t)$. Therefore, one solution to the initial value problem

$$\vec{H}'(t) = F(t)$$
$$\vec{H}(t_0) = \vec{J}(t_0),$$

is $\vec{H}(t) = \vec{J}(t)$, $t \geq t_0$.

On the other hand, in any region where $\vec{H}(t) \neq \vec{J}(t)$, we have $\vec{H}'(t) = \vec{F}(t) = \vec{F}1(\vec{H}(t), t)$. Now, $\vec{F}1(t)$ is a continuous, bounded function of $\vec{H}$ and t in such a region. Therefore, by the usual existence theorem for ODE's, there exists a solution to the equation $\vec{H}'(t) = \vec{F}1(\vec{H}(t), t)$ for all t, until one leaves the region $\vec{H}(t) \neq \vec{J}(t)$.

So we define $\vec{H}(t)$ as follows:

If $\vec{H}(0) = \vec{J}(0)$
  then $\vec{H}(t) = \vec{J}(t)$ for all $t > 0$
  else let $\vec{Q}(t)$ be the solution to the equation $\vec{Q}'(t) = \vec{F}1(\vec{Q}(t), t)$ as above;
    if there exists a t_1 such that $\lim_{t \rightarrow t_1} \vec{Q}(t) = \vec{J}(t_1)$
      then define $\vec{H}(t) = \vec{Q}(t)$ for $0 \leq t < t_1$, $\vec{H}(t) = \vec{J}(t)$ for $t \geq t_1$;
      else define $\vec{H}(t) = \vec{Q}(t)$ for all $t \geq 0$.

This value of $\vec{H}(t)$ then statisfies the I.V.P.

We shall show below (lemma 9) that the final conditional in the above definition of $\vec{H}(t)$ is, in fact, always satisfied; that is, there always exists a t_1 such that $\lim_{t \rightarrow t_1} \vec{Q}(t) = \vec{J}(t_1)$.

Lemma 8: (Step 5 above). If S is any situation, then there exists an unbounded branch B starting in S such that following(prey,1,Δ) is executed over B.

Proof: Use lemma 7 to construct the solution to the equation $\vec{H}'(t) = F(t)$, $\vec{H}(\text{clock}(S)) =$ h_place(S). It is easily shown that this solution, with h_fail = FALSE, is a hunter-realizable behavior. Therefore, there exists a branch starting in S where it is realized. By lemma 6, following(prey, 1, Δ) is executed on this branch.

We next address step 7 of theorem 5. We begin by defining an objective function $P(t)$ to be the distance from the hunter to the ghost plus a penalty which is the maximal possible distance from the ghost to the current position of the prey (that is, the maximal velocity of the prey times the difference between the true time and the delayed time.)

$$P(t) = | \vec{J}(t) - \vec{H}(t) | + (t - \Theta(t))v_m$$

Lemma 9: If $\vec{H}(t) \neq \vec{J}(t)$ at time t, then $P'(t) < v_m - G$. (If the hunter moves faster than the prey, $v_m - G$ is negative, so P is decreasing.)

Proof:

For any vector function $\vec{V}(t)$,

$$\frac{d}{dt} | \vec{V}(t) |= \frac{d}{dt}(\vec{V}(t) \cdot \vec{V}(t))^{1/2} = \frac{\vec{V}'(t) \cdot \vec{V}(t)}{(\vec{V}(t) \cdot \vec{V}(t))^{1/2}} = \vec{V}'(t) \cdot \text{direction}(\vec{V}(t))$$

In particular

$$\frac{d}{dt} \mid \vec{J}(t) - \vec{H}(t) \mid = (\vec{J}'(t) - \vec{H}'(t)) \cdot \text{direction}(\vec{J}(t) - \vec{H}(t)) =$$

$$\vec{J}'(t) \cdot \text{direction}(\vec{J}(t) - \vec{H}(t)) - \vec{H}'(t) \cdot \text{direction}(\vec{J}(t) - \vec{H}(t)).$$

Now, by definition of "following", $\vec{H}'(t) = G \cdot \text{direction}(\vec{J}(t) - \vec{H}(t))$ so $\vec{H}'(t) \cdot \text{direction}(\vec{J}(t) - \vec{H}(t)) = G$.

Also

$$\vec{J}'(t) = \frac{d}{dt}\vec{X}(\Theta(t)) = \Theta'(t)\vec{X}'(\Theta(t))$$

Since $\mid \vec{X}'(\Theta(t)) \mid \leq v_m$, we have

$$\vec{J}'(t) \cdot \text{direction}(\vec{J}(t) - \vec{H}(t)) \leq \Theta'(t)v_m$$

Thus we have

$$P'(t) = \frac{d}{dt}(\mid \vec{J}(t) - \vec{H}(t) \mid + (t - \Theta(t))v_m \leq \Theta'(t)v_m - G + (1 - \Theta'(t))v_m = v_m - G.$$

Lemma 10: If $\vec{H}(t_0) = \vec{J}(t_0)$ and $G > v_m$ then distance$(\vec{H}(t), \vec{X}(t)) < \Delta v_m$ for all $t \geq t_0$.

Proof: Since the prey never moves faster than v_m and since $J(t) = X(\Theta(t))$, it follows that $\mid \vec{J}(t) - \vec{X}(t) \mid < (t - \Theta(t))v_m$.

Also by the triangle inequality, $\mid \vec{H}(t) - \vec{X}(t) \mid < \mid \vec{H}(t) - \vec{J}(t) \mid + \mid \vec{J}(t) - \vec{X}(t) \mid$. Combinding this with the above inequality and the definition of $P(t)$, we deduce that $\mid \vec{H}(t) - \vec{X}(t) \mid \leq P(t)$.

Therefore, if $\vec{H}(t) = \vec{J}(t)$ then distance$(\vec{H}(t), \vec{X}(t)) < (t - \Theta(t))v_m < \Delta v_m$.

Suppose $\vec{H}(t_0) = \vec{J}(t_0)$, but $\vec{H}(t) \neq \vec{J}(t)$ for some $t > t_0$. Let t_2 be the least upper bound of all times $t_a < t$ for which $\vec{H}(t_a) = \vec{J}(t_a)$. By the continuity of $\vec{H}$ and $\vec{J}$, it follows that $\vec{H}(t_2) = \vec{J}(t_2)$.

We have $P(t_2) = \mid \vec{H}(t_2) - \vec{J}(t_2) \mid + (t_2 - \Theta(t_2))v_m = (t_2 - \Theta(t_2))v_m$. Moreover, since the conditions of lemma 9 hold over the interval $(t_2, t]$, and since $v_m < G$, it follows from lemma 9 that $P(t)$ is a decreasing function of time over $(t_2, t]$. So $P(t) < P(t_2)$. Thus,

$$\mid \vec{H}(t) - \vec{X}(t) \mid < P(t) < P(t_2) = (t_2 - \Theta(t_2))v_m \leq \Delta v_m$$

Lemma 11: (Step 7 above). If "following(X, G, Δ)" is executed over interval $[T_1, T_2]$, $v_m < G$, and

$$T_2 > T_1 + \frac{\mid \vec{H}(T_1) - \vec{X}(T_1) \mid}{G - v_m}$$

then the distance from $\vec{H}(t_2)$ to $\vec{X}(t_2)$ is less than Δv_m.

Proof: Immediate from lemmas 9 and 10.

This completes the proof of theorem 1.

8.4 Epistemic Theory

We use a possible-worlds model for knowledge [Moore, 85], [Halpern and Moses, 85], [Davis, 90]. A possible world is one possible state of the world at an instant. Worlds are connected by an accessibility relation. A statement ϕ is known in world W iff ϕ is true in every world accessible form W. We will assume that the accessibility relation is reflexive and transitive; it is known [Halpern and Moses, 85] that such a model is equivalent to an S4 modal logic.

The epistemic theory is integrated with our temporal theory by identifying possible worlds with situations [Moore, 85]. Thus, we imagine a whole collection of branching time structures, interconnected by knowledge accessibility relation. We posit the following three constraints relating accessibility to the temporal structure.

- The agent knows the clock time.
- "Axiom of memory" [Davis, 90]: The agent remembers everything he once knew.
- The agent knows his own behavior.

We can express these constraints axiomatically as follows:

Axiom K.1: (Veridicality of knowledge) Any situation S is accessible from itself.

Axiom K.2: (Positive introspection). If SB is accessible from SA, and SC is accessible from SB, then SC is accessible from SA.

Axiom K.3: (Knowledge of the time) If SB is accessible from SA, then clock(SB) = clock(SA).

Axiom K.4: (Persistence of memory) If $S1A$ precedes $S2A$ and $S2B$ is accessible from $S2A$, then there exists a situation $S1B$ such that $S1B$ is accessible from $S1A$ and $S1B$ precedes $S2B$.

Axiom K.5: (Memory of behavior) If SB is accessible from SA, then the value of the behavioral fluent is the same in the two worlds.

We now proceed to rephrase definitions 18 and 19 of epistemic feasibility in the language of possible worlds. We begin with a few preliminary definitions.

Definition 35: The pair $\langle S1B, S2B \rangle$ is said to *correspond* to the pair $\langle S1A, S2A \rangle$ if $S1A$ is ordered relative to $S2A$; $S1B$ is ordered relative to $S2B$; $S1B$ is knowledge accessible from $S1A$ and $S2B$ is knowledge accessible from $S2A$.

Definition 36: The interval $[S1, S3]$ is a *continuation* of P within Δ after the pair $\langle S1, S2 \rangle$ if $S1 < S3$ and $S2 \leq S3$; plan P begins over $[S1, S3]$; and clock($S3$) $-$ clock($S2$) $\leq \Delta$.

Definition 37: Intervals $[S1A, S3A]$ and $[S1B, S3B]$ *match in behavior* if clock($S1A$) = clock($S1B$); clock($S3A$) = clock($S3B$) and for every $S2A \in [S1A, S3A]$, $S2B \in [S1B.S3B]$, if clock($S2A$) = clock($S2B$) then the value of the behavioral fluent is the same in $S2A$ and $S2B$.

Definition 38: Intervals $[S1A, S3A]$ and $[S1B, S3B]$ *match in respect to plan P* if they match in behavior; P begins on one iff it begins on the other; and P executes on one iff it executes on the other.

Definition 18.A:
A plan P is *executable* in situation $S1A$ with delay Δ iff
[P is necessarily feasible in $S1A$ and
for any $S2A$, $S1B$, $S2B$,
if $\langle S1B, S2B \rangle$ corresponds to $\langle S1A, S2A \rangle$,

then, for every continuation $[S1A, S3A]$ of P within Δ after $\langle S1A, S2A\rangle$,
there exists a continuation $[S1B, S3B]$ after $\langle S1B, S2B\rangle$
that matches $[S1A, S3A]$ with respect to P; and
for every continuation $[S1B, S3B]$ of P within Δ after $\langle S1B, S2B\rangle$,
there exists a continuation $[S1A, S3A]$ after $\langle S1A, S2A\rangle$
that matches $[S1B, S3B]$ with respect to P.]

The form "For every continuation $[S1A, S3A]$ there exists a matching continuation $[S1B, S3B]$" expresses the statement that each continuation of the plan P is known to be a continuation. The converse, "For each continuation $[S1B, S3B]$ there is a continuation $[S1A, S3A]$" expresses the statement that all the continuations of P are known.

Definition 19.A:
A plan P is *epistemically feasible as a task* in situation SA with delay Δ iff
there is a plan $P1$ such that
for every situation SB accessible from SA,
$P1$ is executable in SB with delay Δ and
$P1$ is a specialization of P starting in SB.

8.5 Proof of Epistemic Feasibility

We now proceed to the proof that plan hunt1 is epistemically feasible as task with delay 0.01 in situation s0.

We begin by noting that the non-deterministic plan "hunt1" is not itself executable with delay 0.01, because some completions of hunt1 — namely, those in which the "monitor" cuts off the circular path before 0.01 seconds have elapsed, or those in which the prey is followed with less than a 0.01 second delay — can only be carried out by an agent that can respond, at least sometimes, in less than 0.01 second. We therefore define a determinate plan "hunt2" in which the hunter always delays to the maximal degree allowed by hunt1. The major part of the proof is showing that hunt2 is known to be executable (lemma 24). Combining this with the fact that hunt2 is known to be a specialization of hunt1 (lemma 25) we verify that hunt1 is epistemically feasible as a task.

Formally, we define hunt2 as follows:

hunt2 =
sequence(monitor2($\lambda(T)$ distance(h_place(T), prey(T)) $\leq$ 0.95,
sequence(go_circle(0,2), go_circle(2,1.13), go_circle(1.13,0.5)),
0.01),
follow2(prey, 1, 0.01, 0.01))

follow2(X, G, L, Δ) =
monitor2(dist(h_place,X) $\leq L$,
forever(attempt(following2(X, G, Δ)))).

This uses the following new predicates:

- following2(X, G, Δ) — Follow X at speed G with delay precisely Δ. That is, in definition 27, choose $\Theta(T)$ to be always equal to $T - \Delta$.

- monitor2(Q, P, Δ) — Execute P, monitoring fluent Q. If Q ever becomes true, then, after exactly time Δ, terminate the "monitor" plan.

To be more precise, monitor2 is defined as follows:

Definition 39: Let Q be a fluent and let B be a branch with lower bound $S1$. Q *becomes true at the start of* B if, for any $SE \in B, SE > S1$, there exists an $SD \in [S1, SE)$ such that Q holds at SD. (Note that SD may be equal to $S1$).

Definition 40: Let B be a branch. If B is bounded above, let B_f be the upper bound of B and let $T_f = \text{clock}(B_f)$; else, let $T_f = \infty$. Let $B1$ be the subbranch of B containing all situation $S \in B$ such that $\text{clock}(S) < T_f - \Delta$ and let $B2 = B - B1$. The plan "monitor2(Q, P, Δ)" is executed over B if the following conditions hold:

A. P is begun over B;

B. Q is false throughout $B1$

C. One of the following holds:

C.i. B is bounded above, and P is executed over B.

C.ii. B is unbounded above.

C.iii. Q becomes true at the start of $B2$.

We also need to augment our characterization of the microworld by specifying the perceptual powers of the robot. For the purposes of the current example, it suffices to postulate that, if the prey is within distance 1 of the hunter, then the hunter knows where it is; if it is more than distance 1, then the hunter knows that it is more than distance 1. (If we wished to prove that a certain plan was epistemically infeasible, due to the limitations of perception, then the more expressive language of [Davis, 88] would be needed.)

Postulate 2: Let SA be a situation, and let SB be knowledge accessible from SA.
If distance(prey(SA), h_place(SA)) ≤ 1, then prey(SB) = prey(SA).
If distance(prey(SA), h_place(SA)) > 1, then distance(prey(SB), h_place(SB)) > 1.

We further postulate that the robot always knows his own position.

Postulate 3: If SB is knowledge accessible from SA, then hunter(SA) = hunter(SB).

We begin by proving sufficient condition for the executability of a sequence. We introduce the following two definitions:

Definition 42: Plan P is *always defined* if, for any situation S, there is a branch B starting in S such that P executes over B. (Keep in mind that "executing" here includes failing.)

Definition 43: Plan P is *always finite* if any branch B such that P executes over B is bounded.

Lemma 12: If plan p is executable in situation s1a, p succeeds over [s1a,s2a], and ⟨s1b,s2b⟩ corresponds to ⟨s1a,s2a⟩, then p succeeds over [s1b,s2b]. (That is, if p is executable in s1a, and succeeds over [s1a,s2a] then the agent knows in s2a that it has succeeded.)

Proof: [s1a,s2a] is a continuation of p after ⟨s1a,s2a⟩. Therefore, by definition of executability, there is a continuation [s1b,sqb] after ⟨s1b,s2b⟩ that matches [s1a,s2a] with respect to p. But now note that clock(sqb) = clock(s2a) (by definition of "matching") = clock(s2b) (by axiom K.3) and that sqb $\geq$ s2b (by definition of continuation). Therefore sqb = s2b, which completes the proof.

Lemma 13: Let p = sequence(pu,pv), where pu and pv are always finite and always defined. Interval [s1,s3] is a continuation of p within Δ after ⟨s1,s2⟩ iff either

1. [s1,s3] is a continuation of pu within Δ after $\langle$s1,s2$\rangle$; or
2. pu executes over [s1,sm], sm $<$ s3 $\leq$ s2 + Δ, and [sm,s3] is a continuation of pv after $\langle$sm,s2$\rangle$.

Proof: There are three parts to the above statement:

i. If [s1,s3] is a continuation of pu then [s1,s3] is a continuation of p. Proof: Let [s1,s3] be a continuation of pu after $\langle$s1,s2$\rangle$ within Δ. Since pu is always finite, there is an s4 $\geq$ s3 such that pu executes over [s1,s4]. Since pv is always defined and finite, there is an s5 $\geq$ s4 such that pv executes over [s2,s5]. By definition of sequence, p executes over [s1,s5]. Therefore, p is begun over [s1,s3], so [s1,s3] is a continuation of p.

ii. If pu executes over [s1,sm] and [sm,s3] is a continuation of pv, then [s1,s3] is a continuation of p. Proof: There is an s4 such that pv executes over [sm,s4]. By definition of sequence, p executes over [s1,s4]. Therefore [s1,s3] is a continuation of p.

iii. If [s1,s3] is a continuation of p, then either [s1, s3] is a continuation of pu, or pu executes over [s1,sm] and [sm,s3] is a continuation of pv. Proof: There is an s4 $\geq$ s3 such that p executes over [s1,s3]. By definition of sequence, there is an sm $\leq$ s4 such that pu executes over [s1,sm], and pv executes over [sm,s4]. If sm $\geq$ s3, then the first disjunct holds. If sm $\leq$ s3 + Δ, then the second disjunct holds.

Lemma 14: Let $S2$ be a situation, and let $T1$ be a time such that $T1 <$ clock($S2$). Then there is a unique situation $S1$ such that $S1 < S2$ and clock($S1$) = $T1$. We will denote this situation $S1$ as pre($S2, T1$).

Proof: Immediate from definition 3.

Lemma 15: Let $S2B$ be knowledge accessible from $S2A$, and let $S1A < S2A$. Then the situation pre($S2B$,clock($S1A$)) is knowledge accessible from $S1A$.

Proof: From lemma 14 and axioms K.3 and K.4.

Lemma 16: Assume that

Plan p = sequence(pu,pv) where pu and pv are plans.
Plan p is necessarily feasible in situation s1a.
Plans pu and pv are always defined and always finite.
Plan pu is executable in situation s1a.
For any sma, if pu is executed over interval [s1a,sma], then pv is executable in situation sma.

Then p is executable in s1a. (Note: these conditions are sufficient for the executability of a sequence but not necessary. [Davis, 94].)

Proof: Let s2a be any situation that is ordered relative to s1a. Let s1b, s2b be situations such that $\langle$s1b,s2b$\rangle$ corresponds to $\langle$s1a,s2a$\rangle$. We must show that (A) for every continuation of p after $\langle$s1a,s2a$\rangle$ there is a matching continuation of p after $\langle$s1b,s2b$\rangle$ and (B) vice versa.

A: Let [s1a,s3a] be a continuation of p after $\langle$s1a,s2a$\rangle$. Using lemma 14, we can distinguish two possibilities:

A.1 [s1a,s3a] is a continuation of pu within Δ after $\langle$s1a,s2a$\rangle$. In this case, since pu is executable in s1a, there must exist a situation s3b such that [s1b,s3b] is a matching continuation of pu after $\langle$s1b,s2b$\rangle$. By lemma 14, [s1b,s3b] is a continuation of p after $\langle$s1b,s2b$\rangle$.

A.2 pu succeeds over [s1a,sma], sma $\leq$ s3a $\leq$ s2a + Δ, and [sma,s3a] is a continuation of pv after $\langle$sma,s2a$\rangle$. Let smb=pre(s2b,clock(sma)) By axiom K.4, smb is accessible from sma and s1b $\leq$ smb $\leq$ s2b. By lemma 12 plan pu succeeds over [s1b,smb]. Therefore, by hypothesis, pv is executable in smb. Clearly $\langle$smb,s2b$\rangle$ corresponds to $\langle$sma,s2a$\rangle$. Since smb is accessible from sma, by definition of executability, there is a continuation [smb,s3b] of pv after $\langle$smb,s2b$\rangle$ that matches [sma,s2a] with respect to pv. It is easily verified that [s1b,s3b] is a continuation after $\langle$s1b,s2b$\rangle$ that matches [s1a,s3a] with respect to p.

B: The proof of B is exactly analogous to that of A.

Lemma 17: For any situation S, path H, and duration D, if h_place(S) = $H(0)$, then the plan "go(H, D)" is executable in S.

Proof: Let P=go(H, D). Let $\langle S1B, S2B \rangle$ correspond to $\langle S1A, S2A \rangle$. Let $[S1A, S3A]$ be a continuation of P after $\langle S1A, S2A \rangle$. Thus P is begun over $[S1A, S3A]$. Using the definitions, we have that clock($S3A$) $-$ clock($S1A$) $\leq D$, and for $SA \in [S1A, S3A]$, h_place(SA) = H(clock(SA) $-$ clock($S1A$)). Since $[S1A, S3A]$ is a continuation of $\langle S1A, S2A \rangle$, it is immediate that, for any SA, if $S1A \leq SA \leq S2A$, then h_place(SA) = H(clock(SA) $-$ clock($S1A$)). (Of course if $S2A < S1A$, then no such SA exists and the statement holds vacuously.)

We know (step 1 of theorem 5) that there is an unbounded branch BB starting in $S1B$ such that going(H) is executed over BB. Since the time structure branches on "hunter" that this branch is unique. Let $S3B$ be the situation in BB such that clock($S3B$) = clock($S3A$). It is immediate that $[S1B, S3B]$ matches $[S1A, S3A]$ with respect to P.

Let SB be any situation such that $S1B \leq SB \leq S2B$. By axiom K.4, SB is accessible from SA. By axiom K.5, the value of the behavioral fluent h_place(SB) = h_place(SA). Note that clock($S2B$) = clock($S2A$) (by axiom K.3) $\leq$ clock($S3A$) = clock($S3B$). Therefore, since the time structure branches on "hunter", it follows that $S2B \leq S3B$, so $[S1B, S3B]$ is a continuation of $\langle S1B, S2B \rangle$.

The second half of the proof (for every continuation $[S1B, S3B]$ there exists a matching continuation $[S1A, S2A]$) is exactly analogous.

Definition 44: A fluent Q is *always known* if the following condition holds: For any situations SA, SB, if SB is accesible from SA, then Q has the same value in SA and SB.

Lemma 18: The fluent "$\lambda(T)$ distance(h_place(T),prey(T)) < 0.8 " is always known.

Proof: From postulates 2 and 3 and axiom K.4.

Lemma 19: Let Q be a fluent that is always known, and let $P1$ be a plan that is executable with delay Δ. Let $\Lambda > \Delta$. The plan P=monitor2($Q, P1, \Lambda$) is executable with delay Δ.

Proof: Let $\langle S1B, S2B \rangle$ correspond to $\langle S1A, S2A \rangle$. Let $[S1A, S3A]$ be a continuation of P within Δ of $\langle S1A, S2A \rangle$. Let SZA = pre($S3A$,clock($S3A$)$-\Lambda$).

Note that, since clock($S2A$) $\geq$ clock($S3A$)$-\Delta >$ clock($S3A$)$-\Lambda$, it follows that $SZA < S2A$.

Define SZB=pre($S2B$,clock(SZA)). Since $S2B$ is knowledge accessible from $S2A$, the axiom of memory implies that SZB is knowledge accessible from SZA, and that any situation in the interval $[S1B, SZB]$ is knowledge accessible from the corresponding situation in the interval $[S1A, SZA]$. Also, any situation in the interval $[SZB, S2B]$ is knowledge accessible from the corresponding situation in $[SZA, S2A]$.

By definition of monitor2, since P begins over $[S1A, S3A]$, Q does not hold in any situation SA such that $S1A \leq SA < SZA$. Therefore, since Q is universally known, Q does not hold in any situation SA such that $S1B \leq SB < SZB$.

Also, by definition of monitor2, $P1$ is begun over $[S1A, S3A]$. Since $P1$ is executable in $S1A$, and $\langle S1B, S2B\rangle$ corresponds to $\langle S1A, S2A\rangle$, there must be a continuation $[S1B, S3B]$ within Δ of $\langle S1B, S2B\rangle$ that matches $[S1A, S3A]$ with regard to $P1$.

We now claim that $[S1B, S3B]$ is a continuation of $\langle S1B, S2B\rangle$ that matches $[S1A, S2A]$ with respect to P. To verify this, we note the following:

- $P1$ is begun over $[S1B, S3B]$. Since $[S1B, S3B]$ matches $[S1A, S3A]$ with respect to $P1$, we know that $P1$ completes execution over $[S1B, S3B]$ iff it completes execution over $[S1A, S3A]$.
- Q is false throughout $[S1B, S3B)$. Therefore, P=monitor2$(Q, P1, \Lambda)$ is begun over $[S1B, S3B]$.
- $[S1A, S3A]$ is a complete execution of P iff either it is a complete execution of $P1$ or if Q become true at the start of $[SZA, S2A]$. The first possibility holds iff $[S1B, S3B]$ is a complete execution of $P1$. The second holds iff Q becomes true at the start of $[SZB, S2B]$. Thus, P completes over $[S1A, S3A]$ iff it completes over $[S1B, S3B]$.

The proof of the converse — for each continuation $[S1B, S3B]$ of $\langle S1B, S2B\rangle$ there is a matching continuation $[S1A, S3A]$ of $\langle S1A, S2A\rangle$ — is exactly analogous. □

The next step is to show that the encircling part of the plan will bring the hunter withing sight of the prey.

Lemma 20: Let ph1 be the plan, "sequence(go_circle(0,2), go_circle(2,1.13), go_circle(1.13, 0.5))."

Let $S1$ be a situation in which the distance from the hunter is at most 2. If ph1 is executed over interval $[S1, S2]$, there will be some situation $SM \in [S1, S2]$ in which the distance from the hunter to the prey is less or equal to than 0.95.

Proof: The hunter is originally at the origin, the prey is originally a distance not more than 2 from the origin, and the prey moves with speed 0.05. The proof is by contradiction: Assume that the prey is always more than 0.95 from the hunter.

We begin with the first step of ph1. Let i1 be the time interval [0,2], during which the hunter goes from $\langle 0, 0\rangle$ to $\langle 0, 2\rangle$; and let i2 be the interval $[2, 2 + 4\pi]$, during which the hunter executes a circle of radius 2. The duration of i1 $\cup$ i2 is $2 + 4\pi < 15$. Within that time, the prey cannot move a distance more than 0.75. Therefore, the distance from the prey to the origin is at most 2.75 throughout i12.

We now establish that at the end of i2, the distance from the prey to the origin is at most 1.68. To do this, we consider three cases:

Case 1: At some instant during i2, the prey is within 1.05 of the origin. If so, then by the end of i2 the prey cannot be $1.05 + 0.05 \cdot 4\pi < 1.68$ from the origin.

Case 2: At some instant during i2, the origin, the prey, and the hunter are all collinear. If so, at such an instant, the prey cannot be more than 1.05 from the origin. Thus, this is a special case of case 1.

Case 3: Throughout i2, the prey is more than 1.05 from the origin, and the prey never lies on the line between the origin and the hunter. Since the prey is always more than 1.05 from the origin, the angle at the origin subtended by the motion of the prey is less than $0.05 \cdot 4\pi/1.05\pi < 0.2$. Let us take the angle of the hunter at the start of i2 to be 0, so that during i2 the hunter goes from 0 to 2π. Let θ be the angle of the prey at the beginning of i5. We will choose θ to have value in $[0, 2\pi)$. Thus, in i2, the prey goes from angle θ to angle $\theta' \in [\theta - 0.2, \theta + 0.2]$ Given the condition that the hunter and the prey are never at the same angle, it follows by continuity that $\theta > 2\pi - 0.2$ and that $\theta' > 2\pi$. That is, the prey begins just behind the hunter, and ends just in front of it.

Thus we have the following constraints: At the start of i2, the prey is clockwise from the hunter,

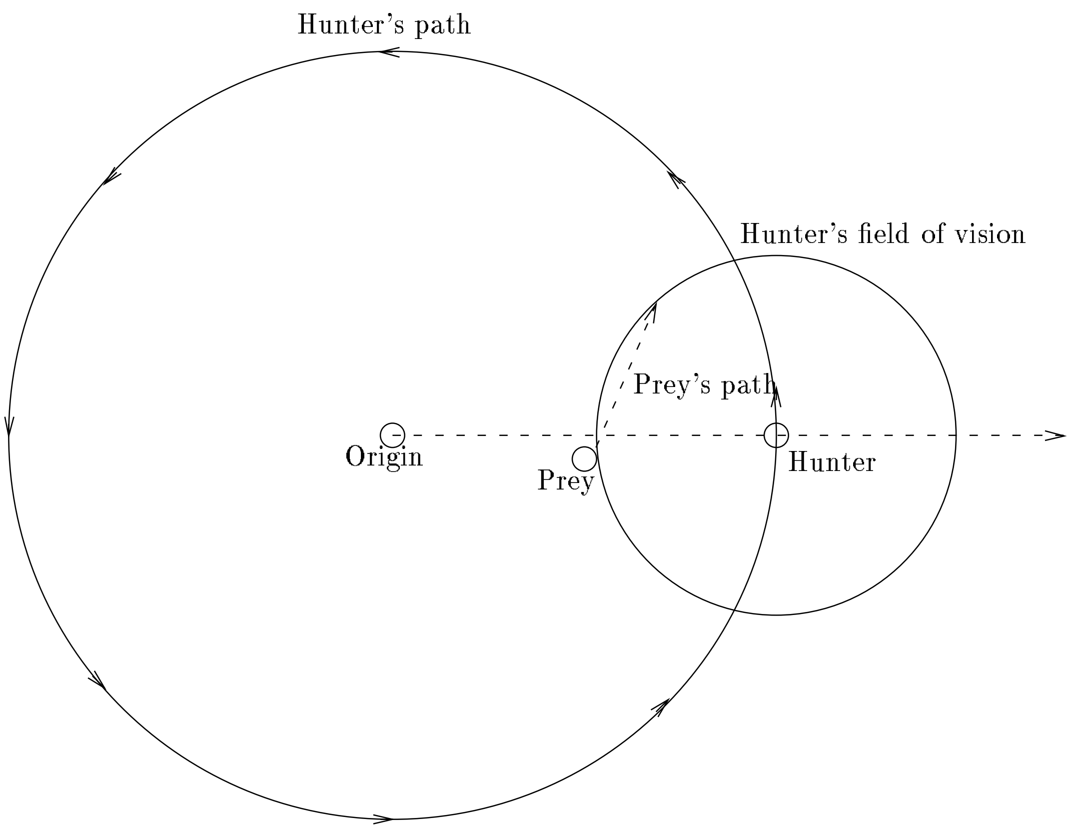

Figure 4: Motion constraint: Case III

is within 2.75 of the origin, and is outside the circle of radius 0.95 of the point $\langle 2, 0 \rangle$. At the end of i2, the prey is counterclockwise from the hunter, and is again outside the circle of radius 0.95 of the point $\langle 2, 0 \rangle$. Moreover the distance from the starting point to the ending point is less than $0.05 \cdot 4\pi < 0.63$ (Figure 4). Thus, we are looking for a chord of length 0.63 across the circle of radius 0.95 with one end in the lower-left quadrant and with the maximal distance from the origin at the other end. It is geometrically obvious, and easy enough to establish formally, that any such chord ends up closer to the origin than the chord that starts at the point $\langle 1.05, 0 \rangle$. But the latter falls within case 1. Thus in case 3 as well, the prey ends up less than 1.68 from the origin.

Similarly, we can show that at the end of completing the circle of radius 1.13 the prey cannot be more than 0.54 from the origin, and that the hunter will therefore encounter the prey during the circle of radius 0.5, completing the contradiction. $\square$.

Lemma 21: Let the plan

monitor2($\lambda(T)$ distance(h_place(T), prey(T)) ≤ 0.95, ph1, 0.01).

be executed over interval $[S1, S2]$. Let SQ=pre($S2$,clock($S2$)$-$0.01). Then

1. At SQ, the distance from hunter to prey is at most 0.95.
2. During interval $[SQ, S2]$, the distance from hunter to prey is never more than 0.9605.

Proof: (1) follows from lemma 20, together with the definition of monitor2. (2) follows from (1) together with the facts that the duration of $[SQ, S2]$ is 0.01, and that the relative velocities of the hunter and the prey is at most 1.05.

Lemma 22: Let $[SQA, S1A]$ be an interval of length Δ throughout which the distance from the hunter to the prey is less than 1. Then the plan following2(prey,1,Δ) is executable with delay Δ in $S1A$.

Proof: Note that following2 is a beginnable action (definition 28) so that beginning over an interval and executing over the interval are equivalent. Let $\langle S1B, S2B \rangle$ correspond to $\langle S1A, S2A \rangle$. Let $[S1A, S3A]$ be a continuation within Δ of $\langle S1A, S2A \rangle$, such that following2(prey,1,Δ) is executed during $[S1A, S3A]$. Let SQB = pre($S1B$,clock(SQA)). Since $S2B$ is knowledge accessible from $S2A$, the position of the hunter is the same in each. Let $S3B$ be the situation such that the behavior of the hunter over $[S2B, S3B]$ is the same as his behavior over $[S2A, S3A]$.

Combining the constraint given here with the consequent of lemma 10, we infer that, throughout the interval $[SQA, S2A]$, the prey is less than 1 from the hunter. Therefore, by postulate 2, the position of the prey is known throughout $[SQA, S2A]$; that is, if $SA \in [SQA, S2A]$, and SB corresponds to SA, then the prey is in the same position in SB as in SA. Now, by definition of "following2", since the duration of $[S2B, S3B]$ is Δ, the behavior that constitutes an execution of following2(prey, 1, Δ) during $[S2B, S3B]$ depends only on the positions of the hunter and prey during $[SQB, S2B]$. Since these positions are the same in $[SQA, S2A]$ as in $[SQB, S2B]$, and since the behavior of the hunter is the same in $[S2B, S3B]$ as in $[S2A, S3A]$, it follows that the hunter continues following2(prey,1,Δ) in $[S2B, S3B]$ iff he continue it in $[S2A, S3A]$. $\square$.

Lemma 23: Let $[SQA, S1A]$ be an interval of length 0.01 throughout which the distance from the hunter to the prey is less than 1. Then the plan follow2(prey, 1, 0.01, 0.01) is executable with delay 0.01 in $S1A$.

Proof: Immediate from lemmas 19 and 22 with postulate 1 and the definition of follow2.

Lemma 24: If the prey is within 2.5 of the hunter and the hunter is at the origin, then hunt2 is executable with delay 0.01.

Proof: Combining lemmas 16, 21, and 23.

Lemma 25: In any situation, hunt2 is a specialization of hunt1.

Proof: Immediate from the definitions.

Theorem 26: If the prey is known to be within 2.5 of the hunter then plan hunt1 is epistemically feasible as task.

Proof: From lemmas 24 and 25 with the definition of "epistemically feasible as task."

8.6 Some remarks about the proof

In the above proof, it is worthwhile between distinguishing three categories of lemmas. Lemmas 13, 16, and 19 characterize the planning language. They can be thus considered built into a planner for that language, and incur only a one-time cost over the lifetime of a domain independent planner. Lemmas 17, 18 and 22 characterize the domain. Lemmas like these must be proven for each new physical domain. Lemmas 20, 21, 23 - 26 characterize the particular plan. Lemmas like these must be proven for each new plan. (The other lemmas are general results about the basic predicates in the model.)

Thus, the most discouraging part of the above analysis is the complexity of lemma 20, which can be directly generalized only to a fairly narrow class of plans in this domain. The complexity of the proof of this lemma reflects the high degree of indeterminacy in the exogenous behavior of the prey that we have to deal with. In effect, we have to prove that no behavior of the prey will enable it to escape. Universally quantified statements of this kind can notoriously involve a complex and careful case analysis.

9 What About Implementation?

If, right now, you had to implement a domain-specific planner for this hunter, what would you do?[4]

I think that one could make a fairly good planner. I would evaluate plans by simulating prey doing random walks, or perhaps executing simple evasive maneuver heuristics. Planning would be done by a combinations of domain-specific heuristics and hill-climbing search in the space of plans. Chiu [in prep.] has implemented such a planner for a similar domain.

If you really had to create a great planner for this domain, then you might be able to come up with high-powered, special purpose algorithms for it.[5] This would have nothing much to do with AI, of course.

What would be the connection between such a planner and all this logic stuff? In other words, how would the logic help you build the planner?

For a simple planner in this fixed domain, it probably wouldn't help at all. One doesn't really need branching time, because planners can, in almost all cases, be viewed as manipulating constraints on a linear time line, rather than dealing with a branching time structure. One can usually avoid the knowledge preconditions problem for primitive actions, by translating knowledge preconditions into physical preconditions. For instance, one can assert that a physical precondition of executing "following" is that the prey be within the range of visibility. One can then avoid the knowledge

[4] Steve Hanks put this cogent question to me in conversation. The other questions in this section are my own.

[5] In view of the obvious military analogues, it seems quite possible that such algorithms exist in the classified literature.

preconditions problem for plans by constructing the control structure primitives in such a way that any physically feasible plan is also epistemically feasible.

One can imagine a more sophisticated planner for the domain which really would have to distinguish between knowledge preconditions and physical preconditions. For example, a hierarchical planner might construct high-level plans where aspects of the plan were described in the language of knowledge. Or you might have a circumstance where you have to reason that "following" is epistemically feasible even though the prey is not within sight, because you somehow know where it is without seeing it. I have no real idea how this could be done, nor how much mileage could be gotten out of it.

Then what's the logic for?

For a broader class of problems; that's what AI is about. For example, if you want to build a domain-independent planner, where you can declaratively specify domain properties, then you will have to construct representations of the domain and of the control strategies, and contemplating definitions 20, 21, 22, 25, and 27 above may at least alert you to some of the kinds of issues that such a representation needs to address. Or if you want to *reason about* plans in ways other than just constructing them — if, for example, you want to automate inference E of the introduction that there is no epistemically feasible plan in that situation — then the logic might be very useful in making explicit the assumptions that underly such an inference and in clarifying the distinctions that have to be made so that these inferences are justified but other, improper, inferences are not.

Well, how would you implement a domain-independent planner for these kinds of domains, or a program that reasons about plans?

I really don't know. I don't even know what the domain representation should look like. If someone were pointing a gun to my head, I would code up the axioms, get myself the best theorem prover available, and tell him that he could probably expect an answer within $10^{10^{100}}$ years.

Penberthy [1993] has implemented a domain-independent planner for continuous actions, but the range of actions and domains is very limited. One thing you might do is to start with that planner and extend it or modify it.

Incidentally, if inference E is a better illustration of the power of the logic, why did you choose B to prove in this paper?

There's nothing wrong with the proof of B as an illustration of the logic. All that the above remarks indicate is that one can probably prove a bastardized version of B with a much weaker logic, which would be much less true of E. As for why I chose B rather than E — mostly chance that I happened to get started on it first; partly that I have the sense that people prefer to see proofs of positive rather than negative results; partly that the proof of E is almost certainly a great deal harder.

The whole thing would be a lot more convincing if you could show a single example of any kind of AI program whose structure had been inspired by this logic, or could be justified in terms of this logic.

Yes, it would.

By the way, the reason I got started writing this particular paper was that I got tired of seeing papers trying to combine reasoning about action with reasoning about physics and starting with the assumptions that actions are discrete and physics is continuous. The point here was originally the small one that action is not necessarily discrete. Then the issue got involved with the issue of knowledge preconditions for continuous plans, a problem I've been working in for five years.

Easier question: At the logical level, what next on this problem?

Actually, I'm not planning more work on this problem at the logical level, at least until I see that some indication that someone in the outside world is seriously interested in this line of thought. As the format of this concluding section indicates, I have gotten to the point of talking to myself, for lack of other people to talk to. This is not a healthy state of mind.

One thing to be done would certainly be to look for a more fundamental definition than the current form of definition 18 and 18.A. These do not command immediate assent, to put it mildly. Another problem would be to look at proofs of epistemic infeasibility, such as example E. My guess is that this theory is essentially sufficient for that problem, but that is a dangerous prediction. Another problem — the most tempting for me personally — would be to apply this in a less trivial physical domain, such as those discussed in (de Jong 1994).

References

Chiu, J.L. (in preparation). Planning in an Imperfect World Using Previous Experiences. NYU Ph.D. Thesis.

Davis, E. 1988. Inferring Ignorance from the Locality of Visual Perception. *Proc. AAAI-88*, pp. 786-790

Davis, E. 1990. *Representations of Commonsense Knowledge.* San Mateo, Cal: Morgan Kaufmann

Davis, E. 1992a. Infinite Loops in Finite Time: Some Observations. *Third International Conference on Knowledge Representation and Reasoning.* San Mateo, CA: Morgan Kaufmann.

Davis, E. 1992b. The Semantics of Tasks that can be Interrupted or Abandoned. *First International Conference on AI Planning Systems.* San Mateo, CA: Morgan Kaufmann

Davis, E. 1994. Knowledge Preconditions for Plans. *Journal of Logic and Computation.* Forthcoming.

DeJong, G. (1994) Learning to Plan in Continuous Domains. *Artificial Intelligence.* vol. 65, pp. 71-141.

de Kleer, J. and Brown. J.S. 1985. A Qualitative Physics Based on Confluences. In D. Bobrow (ed.) *Qualitative Reasoning about Physical Systems.* MIT Press.

Forbus, K. 1989. Introducing Actions into Qualitative Simulation. *Proc. IJCAI-89.* pp. 1273-1278.

Hayes, P. 1985. Ontology for Liquids. In J. Hobbs and R. Moore (eds.) *Formal Theories of the Commonsense World.* Norwood, N.J.: Ablex Pubs.

Kuipers, B. 1986. Qualitative Simulation. *Artificial Intelligence.* 29: 289-338.

McDermott, D. 1982. A Temporal Logic for Reasoning about Processes and Plans. *Cognitive Science,* 6: 101-155.

McCarthy, J. and Hayes, P. 1969. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Michie (eds.) *Machine Intelligence 4.* Edinburgh University Press.

Moore, R. 1985. A Formal Theory of Knowledge and Action. In J. Hobbs and R. Moore (eds.) *Formal Theories of the Commonsense World.* Norwood, N.J.: Ablex Pubs.

Penberthy, J.S. 1993. *Planning with Continuous Change.* University of Washington Department of Computer Science and Engineering Report 93-12-01

Sandewall, E. 1989. Combining Logic and Differential Equations for Describing Real-World Systems.

Proc. First Intl. Conf. on Principles of Knowledge Representation and Reasoning.

Shoham, Y. 1989. Time for Action: On the relation between time, knowledge, and action. *Proc. IJCAI-89.* 954-959, 1173.

van Benthem, J.F.A.K.. 1983. *The Logic of Time.* Dordrecht: Reidel.