From: AIPS 1994 Proceedings. Copyright © 1994, AAAI (www.aaai.org). All rights reserved.

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

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.

Continuous Plans

It is often useful to model the behavior of an autonomous intelligent creature, biological or robotic, 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 2.5 meters, then the hunter can catch it by encircling it. The hunter first goes in a circle at radius 3 meters, then at radius 2, then at radius 1, until it sees the prey. It then pursues the prey until catching it.

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 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.

*This research was supported by NSF grants #IRI-9001447 and #IRI-9300446

From: AIPS 1994 Proceedings. Copyright © 1994, AAAI (www.aaai.org). All rights reserved.

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.

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 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. 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 a 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.

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*

From: AIPS 1994 Proceedings. Copyright © 1994, AAAI (www.aaai.org). All rights reserved.
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 "$\text{CF}_{\mathcal{B}}$" and a fluent called "$\text{FL}_{\mathcal{B}}$" such that $\text{FL}_{\mathcal{B}}$ exhibits all the behaviors in $\mathcal{B}$ and $\text{CF}_{\mathcal{B}}$ branches on $\text{FL}_{\mathcal{B}}$.

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 $\text{CF}_{\mathcal{B}}$ and the fluent $\text{FL}_{\mathcal{B}}$ as follows:

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

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

iii. For $X \in \text{CF}_{\mathcal{B}}$ 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 $\text{FL}_{\mathcal{B}}(X) = B(T)$. It is easily seen that, for fixed X, this determines the same value of $\text{FL}_{\mathcal{B}}(X)$, however B is chosen.

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

Theorem 2: $\text{CF}_{\mathcal{B}}$ and $\text{FL}_{\mathcal{B}}$ are the unique minimal CF and fluent satisfying the conclusions of Theorem 1, up to isomorphism.

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

The proofs are given in the full paper.

There are two problems with this construction. First, $\text{CF}_{\mathcal{B}}$ 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. $\text{CF}_{\mathcal{B}}$ 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

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

Then by executing $c_2(t)$ before $t = -1/2$, $c_3(t)$ from $t = -1/2$ to $t = -1/3$, ..., the robot ends up executing $c_\infty(t) = -1/t$, 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.

In general, a behavior B is exhibited in $\text{CF}_{\mathcal{B}}$ 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 $\mid b(x) - b(y) \mid < M \mid x - y \mid$.

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), (Davis 1988), (Sandewall 1989), (Davis 1990)) but no standard framework for such theories has emerged. The incorporation of volitional

From: AIPS 1994 Proceedings. Copyright © 1994, AAAI (www.aaai.org). All rights reserved.
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. 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.) 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 function from time to configurations. A history is ***feasible*** if it is continuous and differentiable, it does not cause objects to overlap, and it leaves 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.

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:

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

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.

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

From: AIPS 1994 Proceedings. Copyright © 1994, AAAI (www.aaai.org). All rights reserved.
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.

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[1] 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. Under this assumption, we may adopt the following definitions.

Definition 18:
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.

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.

In many cases, it seems natural to idealize an agent as being able to react instantaneously to stimuli. For example, plan (B) of the introduction supposes that the hunter can start following the prey *as soon as* it is seen. For any real hunter, this is, of course, an idealization, but a natural one. I do not know whether there is a reasonable definition of epistemic feasibility that admits instantaneous reactions.

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.

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

From: AIPS 1994 Proceedings. Copyright © 1994, AAAI (www.aaai.org). All rights reserved.

- 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.

Taking $\Delta = 0.01$, we modify plan (B) as follows:

sequence(
  monitor($\lambda(T)$ distance(h_place(T), prey(T)) ≤ 0.9,
    sequence(go($\lambda(T) \prec 0, T \succ$, 3),
      go($\lambda(T) \prec 3\cos(T/3), 3\sin(T/3) \succ$, 6π)
      go($\lambda(T) \prec 0, 3 - T \succ$, 1),
      go($\lambda(T) \prec 2\cos(T/2), 2\sin(T/2) \succ$, 4π)
      go($\lambda(T) \prec 0, 2 - T \succ$, 1),
      go($\lambda(T) \prec \cos(T), \sin(T) \succ$, 2π))
    0.01),
  follow(prey, 1, 0.01, 0.01))

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($PATH$) — Starting from time T_0, move so that at time $T \geq T_0$ your position is PATH($T - T_0$).
- following(X, G, Δ) — Follow fluent F at speed S with time delay Δ.
- attempt(E) — Attempting to carry out E. If impossible, execute "fail."
- for_duration(D, E) — Carry out E for duration D.
- forever(E) — Carrying out E forever.
- sequence($P_1, P_2 \dots P_k$) — Do plans 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.[2]

The full version of this paper contains a proof that, given suitable assumptions, this plan is necessarily physically feasible and executable with $\Delta = 0.01$, and ends with the hunter within distance 0.01 of the prey.

References

Davis, E. 1988. A Logical Framework for Commonsense Predictions of Solid Object Behavior. *AI in Engineering*, 3: 125-140.

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. 1994b. Knowledge Preconditions for Plans. *Journal of Logic and Computation*. Forthcoming.

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.

[2]In (Davis 1992b), Δ was part of the overall semantics of execution. Incorporating it explicitly in the control structure gives both a more flexible representation and a cleaner semantics.