Electronic News Journal on Reasoning about Actions and Change 117
Vol. 2: *117-130*, 31.3.1998
http://www.ida.liu.se/ext/etai/actions/nj/9803-2/

Protocol of the discussion during March, 1998 on

Ontologies of Time

Edited by: Erik Sandewall
Linköping University, Sweden

This is the first segment of the ENRAC discussion on the **ontologies of time**. It started in the context of the discussion about one article at the Commonsense workshop. The ENRAC Newsletter of 12.3.1998 carried the following introduction:

> *The discussion that started around Peter Grünwald's article continues in fairly broad circles. Today, Murray Shanahan responds to a question earlier in the discussion re the use of causal approaches to ramification in the event calculus. Pat Hayes, in a response to an earlier question by Vladimir Lifschitz, argues for the need to characterize instantaneous changes. (Compare also answer by Jixin Ma to a question re the paper by Knight, Peng and Ma at the Commonsense workshop,...)*

It soon became evident that the discussion was so broad that it required a session of its own. The following is how it evolved:

From: Jixin Ma on 13.3.1998

Dear Erik,

Since, as you mentioned in the Newsletter ENRAC 12.3 (98026), Pat Hayes' opinion about instantaneous changes has a close relation to our previous work (In fact, Pat did raise the similar question at the Commonsense workshop), I would like to make the following claims/arguments:

(1) For general treatment, both intervals and points are needed.

(2) To overcome the so-called Dividing Instant Problem, that is the problem in specifying whether intervals is "open" or "closed" at their ending-points, both intervals and points should be treated as primitive on the same footing. Neither intervals are constructed out of points, nor points are defined as the "meeting place" of intervals. Points have zero duration and are non-decomposable, while interval have positive duration and are either decomposable or non-decomposable (moments). Intervals Meets/Met-by points or other intervals. Therefore, although conceptually there is no definition of the ending-points for intervals, one may still say if an interval is open or closed at a point when the corresponding knowledge is available. E.g., if we know interval I Meets point P1, we may say I is (right) open at P1; if we know that interval I1 Meets Interval I2 and point P2 Meets I2, then we may say I1 is (right) closed at P2. In fact, this interpretation is consistent with the conventional definition about the closed and open nature of

intervals that are constructed out of points such as reals or rationals. (For full details of the axiomatization of such a time structure based on both intervals and points as primitive, see [j-cj-37-114]).

(3) Now consider the classical example of switching on a light. The arguments really depend on what knowledge is given/available for such a case. First of all, one can image there is an interval I immediately before the switching point P, and another interval J which is immediately after P. That is: $Meets(I, P) \wedge Meets(P, J)$, where $HOLDS(\neg LightOn, I) \wedge HOLDS(LightOn, J)$. Now, what about the switching point P?

By Commonsense, at any time, the light is either on or off, and cannot be both on and off. In other words, one should be able to express the example in terms of two adjacent intervals, I1 and I2, where over I1 the light is off and over I2 the light is on, that is $Meets(I1, I2) \wedge HOLDS(\neg LightOn, I1) \wedge HOLDS(LightOn, I2)$. This is in fact the intention of Allen's approach, which, by excluding the concept of points, overcomes the Dividing Instant Problem, successfully. However, with Allen's logic, one cannot talk about anything about the switching point P, which is intuitively there anyway.

The question now is that, by taking both intervals and points as primitive temporal objects, on the one hand, we can talk about time points such as the switching point P. However, on the other hand, can we still successfully express the Commonsense knowledge for the above example, without bearing the DIP? The answer is YES, since it really depends on what knowledge is given/available. In fact, there are three possible cases:

Case a) We have no knowledge about the state of the Light at the switching point P, though we may insist that there is a switching point, but we don't know if the LightOff interval I1, or the LightOn interval I2 is open or closed at the switching point P. What we know is just that the light changes from state "Off" to state "On". Hence, such a case can be simply expressed as (1):

$$Meets(I1, I2) \wedge HOLDS(\neg LightOn, I1) \wedge HOLDS(LightOn, I2)$$

Case b) We do know, or we impose (by some reason for the specified application) that the Light is on at the switching point P, that is, $HOLDS(LightOn, P)$. In this case, we still can express it as (1), but with the additional knowledge that $I1 = I$, $I2 = P + J$. Therefore, we may say that the LightOff interval I1 is right-open at the switching point P, and the LightOn interval I2 is left-closed at P.

Case c) As an alternative to b), we may know, or we may impose that the Light is still off at the switching point P, that is, $HOLDS(\neg LightOn, P)$. In this case, the additional knowledge becomes $I1 = I + P$, $I2 = J$. Therefore, we may say that the LightOff interval I1 is right-closed at the switching point P, and the LightOn interval I2 is left-open at P.

Jixin

References:

[j-cj-37-114] Jixin Ma and Brian Knight. *A General Temporal Theory.* Computer Journal, vol. 37 (1994), pp. 114-123.

From: Pat Hayes on 14.3.1998

Im largely in agreement with Jixin about points and intervals, although I dont think its got anything to do with knowledge, and I think there's a simpler way to say it all.

First, just forget about whether intervals are open or closed. This issue arises only if we insist (as the standard mathematical account of the continuum does) that an interval is a set of points. But if we take points and intervals as basic, there is no need to do this. Points, as Allen suggested long ago, can be thought of as places where intervals meet each other, not as the substance out of which intervals are constructed. It took me a long time to see how powerful this idea is. The question of which interval 'contains' the meeting point is meaningless. This gives a very simple, elegant formulation in which points are totally ordered, intervals are uniquely defined by their endpoints (which are also the points they fit between) and two intervals meet just when the endpoint of the first is the startpoint of the second. That's all the structure one needs. Truths hold during intervals. One can allow instantaneous intervals, whose endpoints are identical and which have no duration. One can, if one wishes, identify the interval $\langle t, t\rangle$ with the point t, since such an interval makes no 'space' when interposed between two others, ie if $\langle a, b\rangle$ meets $\langle b, b\rangle$ meets $\langle b, c\rangle$, then $\langle a, b\rangle$ also meets $\langle b, c\rangle$; but it is also consistent, if one wishes, to distinguish t from $\langle t, t\rangle$, or even to forbid instantaneous intervals completely. It is also quite consistent to have arbitrary amounts of density, discreteness, etc.; for example, one can say that time is continuous except in a certain class of 'momentary' intervals whose ends are distinct but have no interior points. (Vladimir might find these more congenial that points as the intervals which things like flashes of lightning must occupy.)

(Instantaneous intervals have the odd property of meeting themselves, by the way; in fact this is a way to characterise them without mentioning points explicitly. It is also perfectly consistent to have 'backward' intervals whose end is earlier than their beginning, and which have negative durations. Axiomatic details can be found in a rather long document available as two postscript files

```
http://www.coginst.uwf.edu/~phayes/TimeCatalog1.ps
http://www.coginst.uwf.edu/~phayes/TimeCatalog2.ps
```

One can cast the whole theory in terms of a single three-place relation MEETS-AT between two intervals and a point, much as Allen's original theory can be cast in terms of MEETS.)

In this theory, to talk of the set of points 'in' an interval requires one to specify what it means for a point to be 'in' an interval. If a point is later than the beginning and earlier than the end, its clearly in the interval, but we have some freedom with the endpoints. One could insist that interval endpoints are 'in' the interval. But this is now OK, since truths hold not at points but during intervals, so the apparent contradiction of the light being both on and off at the splitting point simply doesnt arise. The light isn't either on or off at a single point: if you want to know whether the light was on or not, you have to say which interval you are talking about. P may be true during $\langle a, b\rangle$ but false during $\langle b, c\rangle$, even if b is considered to be 'in' both the intervals.

Jixin says that "one cannot talk about anything about the switching point P, which is intuitively there anyway." Well, the point is certainly there, and we can talk about it (for example, its relation to other points and intervals) but the question is whether it makes sense to say that something is true *at* it. Some truths may be instantaneous, ie true only at points; others make sense only when asserted to hold during noninstantaneous intervals. Lights being on or off, for example, might be enduring, while changes in illumination, or isolated flashes, can be instantaneous. So for example suppose it is dark during interval $\langle s,t\rangle$ and the switch is hit at t. If the light stays on, we have two meeting intervals. If the light flashes and immediately burns out (put a 120 V bulb in a 230 V socket), one could say that there is a flash at t, surrounded on both sides by extended intervals of darkness. Both stories are perfectly consistent. It follows, for example, that a random timepoint during a period of extended illumination is not a flash, in spite of its being a timepoint at which the light is on.

Pat Hayes

In the Newsletter issue where this contribution by Pat Hayes was published, the introductory sentence read:

> *Today, Pat Hayes answers Jixin Ma and argues in favor of an ontology for time where intervals are the only elementary concept and timepoints play a secondary role.*

Some of the following contributions refer to that introduction.

From: Jixin Ma on 17.3.1998

After reading Pat's answers to our claims/arguments about the ontology for time, we would like to raise the following questions/arguments:

1. First of all, it is not clear what's the exact role that time points play in Pat's formulation, although, according to Erik's understanding, Pat Hayes "argues in favour of an ontology for time where intervals are the only elementary concept and timepoints play a secondary role". As Pat points out in his answers (in agreement with our opinion as stated in our claims), "if we take points and intervals as basic, there is no need to do this", i.e., deal with the question of whether intervals are open or closed. However, it is not clear what's the exact meaning of "taking points and intervals as basic". Are they both taken as primitive temporal objects, or, as Allen suggests, points are thought as places where intervals meet each other?

2. Pat argues that "the question is whether it makes sense to say that something is true at points". However, his argument is quite confused: in the first place, he claims "truths hold not at points but during intervals" (as for the case when one insists that interval endpoints are "in" the interval). Later, he states "Some truths may be instantaneous, ie true only at points; others make sense only when asserted to hold during noninstantaneous intervals". So, what's the answer to the question "whether it makes sense to say that something is true at points"?

3. Pat's claims that one may identify interval $\langle t,t\rangle$ with point t, or distinguish $\langle t,t\rangle$ from t, or even forbid instantaneous intervals completely. However, what's the choice? Do we need points (instantaneous intervals) or not? Let's consider the case that we do (in fact, for general treatments, we **do** need them). For this case, Pat's states that if $meets(\langle a,b\rangle,\langle b,b\rangle) \wedge meets(\langle b,b\rangle,\langle b,c\rangle)$ then $meets(\langle a,b\rangle,\langle b,c\rangle)$. (In fact, it seems in Pat's formulation, we always have $meets(\langle a,b\rangle,\langle b,c\rangle) \wedge meets(\langle a,b\rangle,\langle b,b\rangle)$ etc., since the

intervals are uniquely defined by their endpoints). Below are some problems with this formulation:

I. As noted by Pat himself, "an instantaneous interval meets itself", though the "basic" points are totally ordered. How to characterise the relation between them? Pat's gives a suggestion: to characterise them without mentioning points explicitly. Then, what's the relationship between points and intervals?

II. How to define other relationships between intervals like those introduced by Allen? For instance, it is intuitive to say that $meets(\langle a, b\rangle, \langle b, c\rangle) \wedge meets(\langle b, c\rangle, \langle c, d\rangle) \rightarrow before(\langle a, b\rangle, \langle c, d\rangle)$. However, in this case, one would have both $meets(\langle a, b\rangle, \langle b, b\rangle)$ and $before(\langle a, b\rangle, \langle b, b\rangle)$, and hence "meets" and "before" would not be exclusive to each other.

III. By saying $meets(\langle a, b\rangle, \langle b, b\rangle)$, $meets(\langle b, b\rangle, \langle b, c\rangle)$, and $meets(\langle a, b\rangle, \langle b, c\rangle)$, one can only express the first case, that is case a), but not the other two cases, that is case b) and case c), as we demonstrated in our former arguments.

4. Pat argues that "I'm largely in agreement with Jixin about points and intervals, although I dont think its got anything to do with knowledge". But it does. In fact, as pointed out by Pat himself, "if you want to know whether the light was on or not, you have to say which interval you are talking about". In other words, if the (additional) knowledge of "which interval you are talking about" is given (e.g., in terms of which interval is open/closed at the switching point, or in terms of the corresponding meets relations - "knowledge"?), we can say whether the light was on or not.

5. Pat also argues that his formulation is simpler (and elegant). In what aspects, compared with which formulation? It seems that it still needs a lot of axioms to characterise the formal structure, especially when issues such as density, linearity, boundness, etc, are to be addressed.

Jixin & Brian

From: Pat Hayes on 27.3.1998

Sorry, I wasnt sufficiently clear, and my carelessness in using intuitive phrasing led to misunderstanding.

First, in my view there is no single answer to many of the issues that Jixin raises. One can make various choices, each internally consistent but not consistent with the others. (That is why I called the cited paper a 'catalog' of time theories, rather than a single theory of time.) This freedom means that one must be clear which alternative one is using, as confusion follows when one tries to put together bits and pieces of incompatible views. (For example, the critique of Allen's account by Galton in [j-aij-42-159] in 1990 (wrongly) assumes that Allen's intervals are sets of points on the real line.) Having said this, however, there does seem to be a simple, basic, account which can be extended in various ways to produce all the other alternatives, and this core theory is the one I was referring to.

Second, I dont agree with Erik's introduction of my note (14.3) as putting intervals before points. As Allen and I showed some time ago, the choice is arbitrary, since points can be transparently defined in an interval theory and vice versa, so the choice of either one as somehow more basic is, er, pointless; and one gets a more useful account simply by allowing them both as primitive. (Actually, if anything, the simple theory I outlined seems more to rely on points as basic, since an interval there is completely defined by its two endpoints and has no other structure, and all the temporal relations between intervals can be inferred from the total ordering of points.)

Jixin asks:

> *However, it is not clear what's the exact meaning of "taking points and intervals as basic". Are they both taken as primitive temporal objects, or, as Allen suggests, points are thought as places where intervals meet each other?*

Both. These arent incompatible alternatives. The basic idea in the 'simple' theory is essentially Allen's, that points are meeting-places. Still, there's no harm in being able to mention these meeting-places as real objects, and doing so makes it easier to say quite a lot of things, such as 'when' some change happens. Clock times seem to be associated more naturally with points than intervals, for example.

> *2. Pat argues that "the question is whether it makes sense to say that something is true at points". However, his argument is quite confused: in the first place, he claims "truths hold not at points but during intervals" (as for the case when one insists that interval endpoints are "in" the interval). Later, he states "Some truths may be instantaneous, ie true only* **at points***; others make sense only when asserted to hold during noninstantaneous intervals".*

(In the above I was careless at the place marked by boldface, sorry. I should have said 'pointlike interval'. It gets hard to speak about this stuff clearly in English, since I need to distinguish our intuitive notion of 'point' from the way that a particular theory encodes this intuition, and different theories do it differently. I will use scare-quotes to refer to the intuitive concept.)

> *So, what's the answer to the question "whether it makes sense to say that something is true at points"?*

There is no (single) answer: one can craft the theory to suit various different intuitions on matters like these. The way I prefer, myself, is to say that propositions hold only during intervals, so that it is simply ill-formed to assert a proposition of a single point; but to allow the possibility of pointlike intervals, of the form $\langle t, t\rangle$, to be the temporal durations of propositions which are (intuitively) thought of as happening at a single 'point'. Or, put another way, some intervals may consist of just a single point, and some points may completely fill an interval. These pointlike intervals are the way that (this version of) the theory encodes the times when instantaneous truths hold.

This doesnt require us to say that *every* point fills an interval, notice: since 'interval' is a basic predicate, it is perfectly consistent to say $\neg interval(\langle t, t\rangle)$; this would entail, for example, that nothing changed at that particular time. But it allows us to consider the proposition that a tossed ball's vertical velocity is zero, and assert that it is true at a single 'point', ie formally, that its interval of truth was pointlike. And since it is easy to characterise pointlike in the theory:

$$pointlike(i) \leftrightarrow begin(i) = end(i)$$

one can, for example, say something like

$$illuminated(i) \vee dark(i) \rightarrow \neg pointlike(i)$$

so that the light is neither on nor off AT the switching point. In this theory, every proposition has a 'reference interval' during which it is true, and a proposition might not be true of subintervals of its reference interval. (Though some propositions might be. This kind of distinction has often been made in the linguistic literature. Note however that this intuition is basically incompatible with the idea that an interval is identical to the set of the points it contains.)

3. Pat's claims that one may identify interval $\langle t,t\rangle$ with point t, or distinguish $\langle t,t\rangle$ from t, or even forbid instantaneous intervals completely. However, what's the choice? Do we need points (instantaneous intervals) or not?

We certainly need something corresponding to 'points', I agree. I meant only that the formal theory can be crafted in the way Ive outlined above, or alternatively by identifying the pointlike intervals with their endpoints, and allowing a proposition to hold at a single point. This is in many ways more intuitively transparent but it is formally a bit more awkward, since pointlike isnt definable any more, and one has to put in special axioms forbidding points to meet each other. The 'reference interval' of a proposition could now be a single point in the theory. This is essentially the theory that Allen and I described in our 1985 IJCAI paper [c-ijcai-85-528], though it takes a little work to see it.

Let's consider the case that we do (in fact, for general treatments, we **do** *need them). For this case, Pat's states that if $meets(\langle a,b\rangle,\langle b,b\rangle) \wedge meets(\langle b,b\rangle,\langle b,c\rangle)$ then $meets(\langle a,b\rangle,\langle b,c\rangle)$. (In fact, it seems in Pat's formulation, we always have $meets(\langle a,b\rangle,\langle b,c\rangle) \wedge meets(\langle a,b\rangle,\langle b,b\rangle)$ etc., since the intervals are uniquely defined by their endpoints).*

Yes, exactly. Interval relations are comletely determined by endpoint orderings,and Allen's huge transitivity table can be painstakingly derived from the assumption of total ordering. That's all it amounts to, in fact.

Below are some problems with this formulation:

I. As noted by Pat himself, "an instantaneous interval meets itself", though the "basic" points are totally ordered. How to characterise the relation between them? Pat's gives a suggestion: to characterise them without mentioning points explicitly. Then, what's the relationship between points and intervals?

The relations are quite simple and transparent: intervals lie between endpoints, and points have intervals extending between them. Self-meeting is the interval-interval relation corresponding to equality in the point ordering. Again, if one has an intuitive objection to self-meeting intervals, then one can take the second alternative mentioned earlier. (All these alternatives are got by extending the basic theory.)

II. How to define other relationships between intervals like those introduced by Allen? For instance, it is intuitive to say that $meets(\langle a,b\rangle,\langle b,c\rangle) \wedge$

$meets(\langle b,c\rangle,\langle c,d\rangle) \rightarrow before(\langle a,b\rangle,\langle c,d\rangle)$. However, in this case, one would have both $meets(\langle a,b\rangle,\langle b,b\rangle)$ and $before(\langle a,b\rangle,\langle b,b\rangle)$, and hence "meets" and "before" would not be exclusive to each other.

True, and indeed the Allen relations only have their usual transitivity properties when applied to intervals which are nonpointlike and forward-oriented. Of course both these are properties expressible in the theory, so that the Allen transitivity relationships can be stated there, suitably qualified. (When the alternative extension axioms are added, the qualifications become tautologous.)

BTW, the claim that "meets" and "before" being exclusive is "intuitive" depends on how one's intuition is formed. Part of what I learned by having to construct alternative formalisations is that intuition is very malleable. Having gotten used to pointlike intervals, I dont find this exclusivity condition at all intuitive.

III. By saying $meets(\langle a,b\rangle,\langle b,b\rangle)$, $meets(\langle b,b\rangle,\langle b,c\rangle)$, and $meets(\langle a,b\rangle,\langle b,c\rangle)$, one can only express the first case, that is case a), but not the other two cases, that is case b) and case c), as we demonstrated in our former arguments.

But these cases only make sense if one thinks of interval and points in the usual mathematical way, which is exactly what Im suggesting we don't need to do. We can get almost everything we need just from the ordering structure: we don't need to get all tied up in distinguishing cases which can only be formally stated by using all the machinery of real analysis.

4. Pat argues that "I'm largely in agreement with Jixin about points and intervals, although I dont think its got anything to do with knowledge". But it does. In fact, as pointed out by Pat himself, "if you want to know whether the light was on or not, you have to say which interval you are talking about". In other words, if the (additional) knowledge of "which interval you are talking about" is given (e.g., in terms of which interval is open/closed at the switching point, or in terms of the corresponding meets relations - "knowledge"?), we can say whether the light was on or not.

Again I was careless in using the word "knowledge", sorry. I should have said: in order to answer the question whether the light is on or off, one has to specify the interval with respect to which this question is posed. On this view, the truth or otherwise of a proposition is only meaningful with respect to certain intervals. I dont mean that the facts are determined by knowing more about the details of the interval, but that the question is a different question when asked about one interval than when asked about another, and for some intervals in may be simply meaningless. Is the light on or off at (exactly) 3.00 pm? The only way to answer this is to find a suitable non-pointlike interval of light or darkness completely surrounding 3.00 pm, because 'being on' is the kind of proposition that requires a nonpointlike reference interval.

This has nothing to do with whether an interval is open or closed: in fact, there is no such distinction in this theory. It only arises in a much more complicated extension which includes set theory and an extensionality axiom for intervals.

5. Pat also argues that his formulation is simpler (and elegant). In what aspects, compared with which formulation?

Perhaps I should have said, of all the various formalized temporal theories I have ever examined in detail, which amounts now to maybe 25 or so, this seems to distill out the essence. The others can all be described as extensions of this one (some a little artifically, but mos tof them quite naturally.) The conventional picture of intervals as sets of points carries with it a lot of excess conceptual baggage, and removing this gives a theory which is simple and intuitive (once you get used to it :-), and is a sound formal 'core' which can be extended to give many other theories.

It seems that it still needs a lot of axioms to characterise the formal structure, especially when issues such as density, linearity, boundness, etc, are to be addressed.

Yes; any theory needs to be extended, of course, to deal with density, boundedness, etc., but again a merit of this very simple framework is that it can be transparently extended in these different ways more or less orthogonally to each other. One can establish unboundedness with one very obvious axiom (there's always a future and past to any timepoint) and density is also very easy. Lack of density, ie discrete time, is harder; in fact, theres a sense in which no first-order theory can describe this, since it assumes the integers. But again, this is a matter of adding one (rather complicated) induction axiom, in a way that is mathematically very ordinary. Or, alternatively, one can just assume that the integers are defined elsewhere, and declare that every point has an integer 'date', which gives the theory implicitly used by most temporal databases. It can even be extended into the standard real line, if you wish, by distinguishing 'open' and 'closed' intervals as triplets of the form $\langle point, interval, point \rangle$.

The theory is basically linear in its nature, since it assumes timepoints are totally ordered. One can easily weaken it to allow partial orders, but then the extensions involving density, etc.,, get rather tricker. I think the universe is deterministic in any case, so linearity doesnt bother me :-)

Pat Hayes

References:

[c-ijcai-85-528] James Allen and Pat Hayes. *A Common-Sense Theory of Time.* Proc. International Joint Conference on Artificial Intelligence, 1985, pp. 528-531.

[j-aij-42-159] Anthony Galton. *A critical examination of Allen's theory of action and time.* Artificial Intelligence, vol. 42 (1990), pp. 159-188.

From: Ray Reiter on 29.3.1998

During all the years that the debate has raged about time points vs intervals, we devotees of the sitcalc have never seen it as an issue. Here's why I think this is so.

In the sitcalc, a fluent ($LightOn$) has a truth value only with respect to a situation (= sequence of action occurrences). So, we might have

$$LightOn(do(switchOn, do(switchOff, S0)))$$

and

$$\neg LightOn(do(switchOff, do(switchOn, S0)))$$

In the sitcalc with explicit time, the first might become

$$LightOn(do(switchOn(3.14), do(switchOff(1.41), S0)))$$

meaning that as a result of the action history consisting of first switching off the light at time 1.41, then switching on the light at time 3.14, the light will be on. Notice that there is no way of expressing the claim that the light is, or is not on at time 3.14 (or 3.5), independently of the situation leading up to this time. On the other hand, time based formalisms do allow one to write $LightOn(3.14)$, without expicitly referencing, in their notation, the history leading up to the time 3.14 at which the fluent's truth value is to be determined. This seems to be the source of all the problems about open vs closed vs semi-open intervals and predicate truth values over these, and also why these seem to be non-issues for the sitcalc.

Now, one could rightly object to the above account because it provides only for fluent truth values at discrete time points, namely at the action occurrence times. So we are tempted to understand

$$LightOn(do(switchOn(3.14), do(switchOff(1.41), S0)))$$

to mean that the light is on at time 3.14, but it tells us nothing about time 3.5 say. This is particularly bad for (functional) fluents that vary continuously with time, for example, the location of a falling object. To handle this, introduce a time argument for fluents, in addition to their situation argument. For the light, one can write:

$$LightOnT(t, s) \leftrightarrow LightOn(s) \wedge t \geq start(s).$$

Here, $start$ is defined by $start(do(a, s)) = time(a)$, where $time(a)$ is the time at which the action a occurs in the history $do(a, s)$.

An instance of this would be

$$LightOnT(t, do(switchOn(3.14), do(switchOff(1.41), S0))) \leftrightarrow t \geq 3.14$$

Here we have committed to the light being on at exactly the time of the switchOn action, and forever thereafter, relative to the history

$$do(switchOn(3.14), do(switchOff(1.41), S0))$$

In other words, provided $switchOff(1.41)$ and $switchOn(3.14)$ are the **only** actions to have occurred, then the light will come on at time 3.14, and remain on forever. Notice especially that we would have **both**

$$LightOnT(3.14, do(switchOn(3.14), do(switchOff(1.41), S0)))$$

and

$$\neg LightOnT(3.14, do(switchOff(1.41), S0))$$

without contradiction. This seems to be precisely the point at which purely time-based formalisms run into difficulties, and the sitcalc version of this problem illustrates the role that explicit situation arguments play in resolving these difficulties.

Now, we can axiomatize falling objects:

$$\begin{aligned} &positionT(t, s) = position(s) + \\ &\qquad velocity(s) * (t - start(s)) + 0.5 * g * (t - start(s))^2 \end{aligned}$$

From: John McCarthy on 29.3.1998

When my car accelerates, there is a time point at which it passes 65 miles per hour. It is awkward to describe this point in a language not providing for time points.

From: Graham White on 31.3.1998

John McCarthy wrote:

> *When my car accelerates, there is a time point at which it passes 65 miles per hour. It is awkward to describe this point in a language not providing for time points.*

Indeed so; but all this shows is that it's awkward to combine an interval-based approach to time with a point-based approach to other continua (such as, for example, velocity).

Graham White

From: John McCarthy on 31.3.1998

[1][McCarthy and Hayes (1969)] used time as a fluent on situations, i.e. time(s). One motivation was that people, and perhaps future robots, often do not know the time with sufficient resolution to compare two situations, e.g. Ray Reiter's recent message with times 1.41 and 3.14. A second motivation for making situations primary was to make it correspond to human common sense. Many people who can reason about the consequences of actions in situations perfectly well do not know about real numbers, and some don't know about numbers at all. The falling body example was also in that paper with time as a fluent. Galileo did know about real numbers.

It's not clear that either of these considerations is of basic importance for AI.

My previous message gave a reason for including time points in a theory of events and actions. The theory could be founded so as to regard them as degenerate intervals, but I don't see any advantage in that, although I suppose the idea stems from the fact that people and robots can't measure time precisely.

From: Pat Hayes on 31.3.1998

Responses to Ray Reiter and John McCarthy. Ray wrote:

> *In the sitcalc, a fluent (LightOn) has a truth value only with respect to a situation (= sequence of action occurrences). So, we might have*
>
> $$LightOn(do(switchOn, do(switchOff, S0)))$$
>
> *and*
>
> $$\neg LightOn(do(switchOff, do(switchOn, S0)))$$
>
> *In the sitcalc with explicit time, the first might become*
>
> $$LightOn(do(switchOn(3.14), do(switchOff(1.41), S0)))$$

[1]Ref: http://www-formal.stanford.edu/jmc/mcchay69.html

meaning that as a result of the action history consisting of first switching off the light at time 1.41, then switching on the light at time 3.14, the light will be on. Notice that there is no way of expressing the claim that the light is, or is not on at time 3.14 (or 3.5), independently of the situation leading up to this time.

I think of Ray's 'sequences of actions' as alternative ways the temporal universe might be, ie possible timelines (or histories, as Ray sometimes calls them.) The point/interval controversy is about reasoning within, or with respect to, one of these possible timelines; sitcalc gets this muddled up with reasoning about alternative futures for the partial timeline up to the present. (Think of the tree of accessible situations in a state's future: the distinction is between reasoning about a single branch, and comparing two different branches.)

On the other hand, time based formalisms do allow one to write $LightOn(3.14)$, *without expicitly referencing, in their notation, the history leading up to the time 3.14 at which the fluent's truth value is to be determined. This seems to be the source of all the problems about open vs closed vs semi-open intervals and predicate truth values over these, and also why these seem to be non-issues for the sitcalc.*

This isn't where the difficulties lie. Even if there is only one possible future and only one thing that could happen at each situation, these conceptual problems about points and intervals would still arise and some solution for them would need to be found.

Now, one could rightly object to the above account because it provides only for fluent truth values at discrete time points, namely at the action occurrence times. So we are tempted to understand

$$LightOn(do(switchOn(3.14), do(switchOff(1.41), S0)))$$

to mean that the light is on at time 3.14, but it tells us nothing about time 3.5 say. This is particularly bad for (functional) fluents that vary continuously with time, for example, the location of a falling object. To handle this, introduce a time argument for fluents, in addition to their situation argument.

Hold on! What kinds of things are these 'times' supposed to be? They seem to be something like clock-times, ie temporal coordinates (maybe understood with respect to a global clock of some kind.) OK, but notice that this isn't what I mean by a 'timepoint'. There are at least six distinct notions of 'time' (physical dimension, time-plenum, time-interval, time-point, time-coordinate and duration.) I think the nearest thing in Reiter's ontology to what I call a time-point is something like the pairing of a clock-time with a situation ('3.14 in situation s').

For the light, one can write:

$$LightOnT(t, s) \leftrightarrow LightOn(s) \wedge t \geq start(s).$$

Here, start is defined by $start(do(a, s)) = time(a)$, *where* $time(a)$ *is the time at which the action* a *occurs in the history* $do(a, s)$.

> *An instance of this would be*
>
> $$LightOnT(t, do(switchOn(3.14), do(switchOff(1.41), S0))) \leftrightarrow t \geq 3.14$$
>
> *Here we have committed to the light being on at exactly the time of the switchOn action, and forever thereafter, relative to the history* $do(switchOn(3.14), do(switchOff(1.41), S0))$. *In other words, provided* $switchOff(1.41)$ *and* $switchOn(3.14)$ *are the* **only** *actions to have occurred, then the light will come on at time 3.14, and remain on forever. Notice especially that we would have* **both**
>
> $$LightOnT(3.14, do(switchOn(3.14), do(switchOff(1.41), S0)))$$
>
> *and*
>
> $$\neg LightOnT(3.14, do(switchOff(1.41), S0))$$
>
> *without contradiction...*

This seems to be the half-open-interval solution, where intervals contain their endpoints but not their starting points. This makes sense for the sitcalc, which focusses on the *results* of actions, but seems ad-hoc and unintuitive in a broader context. (Also, BTW, the idea that one can ever say that some finite list of actions is **all** the actions that have occurred seems quite unrealistic. After all, people's fingers probably pushed the switch and something somewhere was generating electricity. Surely one should be able to actually infer this from a reasonably accurate common-sense description of light-switching.)

> *...This seems to be precisely the point at which purely time-based formalisms run into difficulties, and the sitcalc version of this problem illustrates the role that explicit situation arguments play in resolving these difficulties.*

This isnt where the difficulties lie. These alternatives are obviously incompatible if they are asserted of the same timeline (the light can't have been both switched on and switched off at the same timepoint) and there is no contradiction is saying that p is true at time t in one possible timeline but not in another, as these are different timepoints.

The problem is that even if we stick to talking about a single timeline (eg the unique past, or one alternative future) there still seems to be an intuitive difficulty about timepoints like the time when a light came on. The solution I suggested - that is, truth at a point has to be defined relative to a reference interval containing the point (which is not my idea, let me add) - is similar in many ways to Reiter's , except it applies not just across timelines but also within a single one.

John wrote:

> *When my car accelerates, there is a time point at which it passes 65 miles per hour. It is awkward to describe this point in a language not providing for time points.*

Yes, I agree. Examples like this are what motivate the inclusion of both points and intervals as first-class objects. It is pretty awkward to do without points in any case if one wants to refer to the places (...that is, the times...) where (...that is, when...) intervals meet. However, points **can** be defined in terms of intervals , in principle, so having them around is essentially a matter of convenience more than a point of basic ontology. There are there, in a sense, whether one wants them or not.

The 65mph example is logically similar to the point at the top of a trajectory when the vertical velocity is zero. Examples like this appeal to a basic intuition about continuous change, that it has no 'jumps', so if it is $< x$ at $t1$ and $> x$ at $t2$, then it must $= x$ somewhere between $t1$ and $t2$. One can state this quite directly in the basic theory (For strict first-order syntax, replace (X...) by (value X ...)):

```
(continuous X i) =df
   (forall (y)(implies (between (X (begin i)) y (X (end y)))
                       (exists t) (and (in t i) (= y (X t))))

(strictlycontinuous X i) =df
   (forall (j) (implies (subint j i)(continuous X j)))
```

where `subint` is the Allen union `begins, inside, ends`. (This assumes that the timeline itself is dense; if not, strictlycontinuous is trivially true everywhere.) Other conditions like monotonicity and so forth also transcribe directly from their usual mathematical formulations.

Pat Hayes