A Catalog of Temporal Theories

Pat Hayes

The following document was written during 1995-6, while I was a Visiting Scholar at the Beckman Institute at the University of Illinois. It was originally issued, and may be cited, as:

Technical Report UIUC-BI-AI-96-01, University of Illinois.

It is here reproduced exactly as it appeared there, without any corrections or additions. There are a few typos, and the appendix is missing (I no longer have a copy of it.)

A CATALOG OF TEMPORAL THEORIES

Patrick J. Hayes
Beckman Institute and Departments of Philosophy and Computer Science
University of Illinois[1]

Abstract

This document surveys several structures that time can be taken to have, discusses the different intuitions which justify them, and gives organised collections of axioms to describe them.

Introduction

Many ontologies or axiomatic or formal descriptions somehow involve or assume a formalisation of time. The actual structure of time itself is often taken for granted when constructing these formalisations. For example, most temporal database work simply assumes that time is a discrete series of countable clock-ticks; some discussions assume that the timeline is the real line R, while other authors use axioms which seem to be in conflict with Dedekind continuity. There does not seem to be a single account ot the structure of time which is accepted by everyone. Hardly any claim about temporal models is uncontroversial. Some philopshers have even wondered if times are partially ordered. This 'catalog' tries to give a coherent overview of several of these ideas and synthesise them as far as possible.

This is probably not a complete survey[2], and many temporal issues are deliberately excluded. The aim here is only to look at ways of describing the actual structure of time, but not ways in which language and beliefs are related to it (which would require a much larger document.) Mixing epistemic and temporal languages raises several difficult problems in reasoning about what will happen when new knowledge becomes available, for example. Much research has been devoted to reasoning about how facts persist through time; the famous frame problem arises squarely in this area. As time passes, objects are created and destroyed, and people gain, and forget, information. Reasoning involving quantifiers and descriptions of states of knowledge therefore needs to be sensitive to the changes that time can produce. There is nothing intrinsically temporal about these issues, in fact – exactly similar kinds of complexity can be produced with spatial variations – but they seem to be particularly acute in a temporal setting, probably because people are so familar with the need to reinterpret past assertions with the wisdom of hindsight. But in any case, these issues are beyond the scope of this document, which is only concerned with the actual structure of the ‘time-line’[3].

Some of the concepts in this document are new, including the 'vector continuum' described in section 5 and the formalization of clocks and calendars in section 6, but many are taken from previously published work, especially (Van Benthem 1984) , (Freska 1992) and papers by James Allen and myself. Van Benthem’s book is an especially insightful and thorough survey. Particular citations are given in the text. The axioms in this particular form, and the overall organisation, are new.

---

[1]This work was also supported by the University of Southern California on grant NASA NAG2-864
[2] Suggestions for material which should be covered in later editions are welcomed.
[3]Although it need not be a line.

1. What "time" means

The English word "time"has at least six different senses. The first, perhaps most fundamental, refers to one of the basic physical dimensions, on a par with *length*, *mass* and *voltage*. We will have little to say about this other than to note its status as a physical dimension, called the time-dimension. Time-dimension is a physical-dimension in the sense of Gruber & Olsen (1994).

The next idea is that of the universe of time, or temporal continuum; a large temporal "space"[1] within which all events are located, perhaps one dimension of the whole history of the universe. We will call this the *time-plenum*. It is often called the time-line, but it need not always be regarded as linear. For example, relativistic time has a partial ordering abstracted from space-time, and planners often treat the time-plenum as branching into the future. Sometimes the plenum is thought of as a sequence of 'worlds', each of which is therefore considered to be timeless; in modal semantics these are often called *temporally possible worlds*.

The third concept is of pieces of time; physical entities whose sole dimension is time-dimension. These are variously called time-periods or time-intervals, or simply intervals[2]. Examples include the *duringthe 1994 winter Olympics*, *the sixteenth century* and *10:50 to 11.00 a.m. on 30 May 1993*. These are particular pieces of time located in (or perhaps, parts of) the time-plenum. Intervals are in many ways the most central concept for temporal reasoning since they are the temporal extents of things. Events typically are thought of as occupying them, propositions are true during them, and they are the lifetimes of objects. We restrict attention to the simplest and most widely used notion of a contiguous interval (containing all its subintervals, having no gaps) but the idea of an intermittent interval, such as *every Wednesday afternoon during August 1973*, is often useful.

A fourth notion is that of a timepoint. Exactly what counts as a point, and the relationships between points and intervals, seem to be particularly controversial and sensitive questions, and many of the formalisations in use in computer science have taken one or another stance on the answers to these questions. (In particular: whether or not a point can be thought of as an infinitesimal interval; whether or not intervals are sets of points; and whether or not propositions can be true at single points.) We will pay detailed attention to these issues later.

A fifth notion is that of an amount of time. Such things as *a century*, *25 minutes* and *as long as it takes for the kettle to boil*, are amounts of time, which I will call durations. It is natural (although not necessary) to assume that every interval has a duration, but the concepts are distinct. The relationship between an interval and a duration is like that between a particular piece of real estate - say, the site of Sherwood Forest - and its area - say, 45.63 square miles. English is often ambiguous between these meanings. For example 'century' might refer to a duration of 100 years, or to a particular interval such as the 19th century. A duration is a *constant-quantity* of time in the sense of Gruber & Olsen (1994). In their ontology, constant quantities of time are amounts that can be compared, added, and scaled; dividing a time quantity by a standard reference time quantity (a unit of measure) produces a real number. In our development here, more oriented towards clock-ticking, we use integers.

The last notion of 'time' is of a position in a temporal coordinate system. Examples include dates such as *14 March 1994*, day-times such as *3.45 pm*, or such things as stopwatch lap timings in a race. These are usually the appropriate answer to a query concerning when something happened. Temporal positions are often thought of as points. If a temporal position has no

---

[1] The pun in unavoidable. Sorry.

[2] Although a timeinterval need not be an interval in the sense of real analysis, so some care is needed in terminology.

duration, in contrast to an interval, this seems appropriate. (It does not seem to make sense to ask how long 3.45 pm lasts, for example.) On the other hand, it is quite consistent to have positions in a calendar which are themselves intervals, with a finer coordinate system defining 'inner' positions of hours and minutes, and one might claim that such refinement of the temporal coordinates can always be achieved. So timepositions might be modelled in either way.

Although axioms do not always make these distinctions, this basic categorisation and the terminology of time(dimension), (time)plenum, (time)interval, (time)duration, (time)point and (time)position will be used throughout this document.

These ideas are clearly related to one another, but the exact relationships can be defined in a variety of ways. A temporal theory may take points as basic and define intervals as a pair of endpoints, or allow only intervals and find the notion of 'point' incoherent. Some theories identify an interval with the set of points it contains, while others are incompatible with this interpretation. A duration can be defined in terms of a 'standard' interval – such as a clock tick or a day – or it can be given an independent mathematical description. A point can be regarded as an infinitesimal interval, or as an 'atomic' interval, or as a 'quiescent' interval during which no change takes place; or an theory may strictly separate the two categories of interval and point.

Two relationships in particular deserve longer discussion.

1.1 Subinterval inheritance

In some theories, asserting that a proposition is true in an interval entails that it is true at all points, or in all subintervals, of the interval. Other theories explicitly deny the necessity of this *subinterval inheritance*, allowing something to be true during an interval without being true in all subintervals. This difference seems to reflect a fundamental split between two rival intuitions, which can be illustrated by considering a bend in a road.

On a four-day drive from the east to the west coast, a bend in the winding road can mean that one is driving in an easterly direction for ten minutes, say; and yet it seems still true, in some sense, to say that one is driving westward.

One way to describe this says that 'driving west' is true in the four-day interval *I*, but false in the ten-minute subinterval *J*. Examples like this, then, seem to show that a proposition can be true in an interval without being true in all its subintervals. If asked, "Are you driving west now?", the appropriate answer, on this view, would be to inquire what sense of "now" was intended, since driving west is false if "now" is taken to be a short surrounding interval, but true if it is taken to refer to the longer interval which contains the whole journey.

This view fits with the idea that propositions are true only during intervals: the interval is necessary to establish the appropriate context. It hardly makes sense, on this view, to ask for the truthvalue of a proposition at an isolated timepoint.

Another way to describe the bent road, however, distinguishes two senses of *driving west*. One means bound towards a western destination, the other means driving with one's vehicle aimed at the western point of the compass. The correct way to describe the anomalous situation, according to this view, is that *driving west-1* is true throughout *I* (including during *J)*, but *driving west-2* is false throughout *J*. If asked, during *J*, whether one were driving west, the appropriate reply would be to ask which sense of "driving west" was meant, because one is true and the other is false. This view regards a claim of truth during an interval as always implying truth during all subintervals, and insists that apparent counterexamples always involve an ambiguity of meaning.

The second intuition fits very naturally with the view of an interval as a set of points. Given this set-theoretic vision of an interval, it is not easy to see what it could mean for something to be true throughout the interval without it being true at every point in the interval, and hence through every subinterval.

For example, Allen (1984) argues that someone can be writing a novel during a period of, say, a year, without this meaning that they never eat or sleep; hence the process of writing a novel is not inherited by subintervals. Galton (1990) however responds by distinguishing two senses of "writing a novel" (one means having that as one's current professional goal, the other means actually hitting the keyboard) and claims that both of these are inherited by subintervals.

While the second position often seems philosophically convincing, its practical effect can be to create many predicates with subtle nuances of meaning solely for the purpose of maintaining temporal consistency. Fortunately the second view can be modelled in the first one in a way that avoids this conceptual promiscuity. If one accepts the first view, in which truth is always relative to an interval, the second view can be modelled by claiming that the two different senses of the proposition which the second view requires, are just the proposition relativised to different intervals. On this view, we can allow that *P-in-interval-I* is true in all subintervals of *I*, ie that a proposition relativized to an interval is inherited in all its subintervals. One way to express this distinguishes truth **on** an interval from truth **in** an interval. The idea here is that a proposition being true **on** an interval identifies that interval as an appropriate reference interval for the proposition. Truth **in** an interval means that there is a containing interval **on** which the proposition is true. Subinterval inheritance of truth **in** an interval then follows simply from the transitivity of the subinterval relation. Truth **in** an interval need not entail truth **on** that interval; which is exactly the case when the road has a bend in it.

It is quite consistent to accept both positions, and allow some propositions to be inherited by subintervals and some not. That a car motor is running, for example, seems to be a plausible candidate for subinterval inheritance. This can be described by saying that for such predicates, every subinterval of a reference interval is itself a reference interval.

1.2 . Intervals and points

Intervals can be thought of as related to points in several different ways.

The first view is that points *are* intervals, but intervals which are as short as possible: single clock-ticks, as it were. If time is thought of as being discrete this is quite coherent, and it means that there is no real need for the concept of timepoint, since all times are intervals. However, since other theories maintain the distinction, I will use the term *moment* to refer to such a basic interval. A moment is an interval with no subintervals, or an interval with no seperable timepoints inside it. Moments cannot overlap or be contained in one another. They have no internal structure; they are*point-like* intervals. Another way of describing this is that a moment is an interval during which nothing can happen; a quiescent interval. This can be misleading, however, as any beginning student of calculus will testify, since something can be happening at a point – something might be moving, for example – even though it is too small to be when a complete event takes place.

Some theories identify timepoints as being the moments; others interpret them as lying between the moments, since moments are intervals andpoints are places where intervals meet.

These interpretations seem incompatible, but they can be made consistent if one proceeds carefully, as we show in section 5.

Some views insist that time is continuous, so that there can be no moment-intervals. Several different ideas are still possible, however. The intuitions behind these alternatives can be illustrated by the following ancient puzzle about the continuum. Consider dividing an interval into two equal halves. The division must happen at a point; but which half contains that point of division? Whichever half contains it must be not exactly equal to the other: but by hypothesis the halves are equal. Exactly the same intuitive problem arises when we consider intervals which meet, without there seeming to be any reason to put the meeting-point in either one.(Allen 1984)

This puzzle does not have the logical status of Russell's paradox, but it is of interest here because it serves to identify two rival intuitions about the continuum. According to one, now standard in mathematics, points are first-class objects and an interval is identified with the set of points that it contains. According to the other, points serve to locate positions within or between intervals, which are first-class objects with extents which can be compared. If we stick to either intuition carefully the puzzle vanishes. On the first, the conclusion is simply that it is impossible to divide an interval exactly symmetrically in half, and we are led to distinguish open and closed intervals. The second intuition insists that such equal splitting must be possible, and even happen at a point, but rejects the conclusion that this point must be *contained* in either half. If one thinks of an interval as like a piece of glass filament, something which can be snapped neatly in two, then to ask which half 'contains' the point of division would seem perverse, since points are not themselves to be thought of as *parts* of the physical continuum.

Contemporary mathematical theory has firmly adopted the first idea, which we will call the *point continuum*. However, the second is also a coherent position, which we will call the *glass continuum*. In the glass continuum, intervals are not definable as sets of points, but are things in their own right, intervals *sui generis*. Several of the temporal theories given later describe the glass continuum. We demonstrate that the idea is coherent by describing models for those theories. In the point continuum intervals are either open or closed. In the glass continuum, endpoints are not contained in either interval, and there is no open/closed distinction. In the point continuum it is possible to have a closed interval consisting of a single point (which is also both the endpoints of that interval), but this is impossible in the glass continuum (although it is possible in the vector continuum: see section 5)

There are examples which argue for the intuitive plausibility of both ideas. Consider a light going out. The intervals of its being on and then off seem to meet one another, and intuition suggests that the question of whether the light is on or not at the point of extinguishing is meaningless. On the other hand, consider a ball tossed into the air. Qualitative reasoning suggests that the intervals of the ball's upward velocity being positive and negative also meet at a point, but here the meeting-point seems clearly distinguished by a predicate – that the ball is motionless – which is true there and nowhere else. This is impossible in the glass continuum, but very natural in the point continuum. Later we will consider theories which allow both kinds of meeting.

It is sometimes claimed that any physical truth must hold during an interval, although perhaps a very short one, and that points are mere mathematical abstractions. What is coventionally called a timepoint must therefore be understood to really be a very short interval, shorter than the current 'grain size' of the theory. While this idea is physically plausible, I do not think the mathematical consequences have been fully explored. This is compatible with the glass continuum, where it is used to explain the tossed-ball example by insisting that the ball is motionless for a very short period. However, it is also compatible with the point continuum, where it can be used to argue that changes of truthvalue are never instantaneous, so such intervals do not exactly meet but are separated by short linking moments during which the transition happens. In the point continuum these intervals can be single points, so that the intervals of the lights being on and being off would meet at a point at which the light was neither on nor off, but was going out.

Finally, a quite different idea of the relationship between points and intervals is based on the idea of information. On this view an interval is an expression of doubt about the exact position of a point. Decreasing the size of an interval increases the amount of information about the location of the point, i.e. the degree of precision. This 'information' view of intervals is often seen as incompatible with the first idea of identifying an interval by its endpoints, because this would assume infinite precision. It leads to axioms which focus on different relationships. For example, several of Allen's thirteen relations now become meaningless, since it is impossible to distinguish meeting from overlapping.

2. Styles of temporal language

So far we have discussed only the nature of time itself. But the language used to describe time can also vary. Formalisms have been adapted in various ways to refer to temporal relationships, to propositions whose truth may vary with time, and things whose properties may change with time. In this section we will briefly survey the main options: a full survey would require a book-length document.

Time can be involved with knowledge representations in at least three distinct ways. First, we might represent knowledge of time in the same way we might represent knowledge of anything else, simply by writing descriptions of time. Second, the expressions of the knowledge itself might be thought of as being temporally relativised in some way, so that for example the language might involve tenses, or the assertions be time-indexed in some way. And third, the inference machinery which manipulates the knowledge might itself be thought of as embedded in time, so that there is a notion of 'now' and issues of temporal truth maintenance become relevant.

2.1 Timeless quantification

The most direct way to describe time simply treats times as objects and describes them by axioms which relate times to other things. For example, one common way to acknowledge the temporal sensitivity of some relations (and functions) is simply by allowing times as an extra argument, often with some convention such that it must be the last argument. To say that Joe and Anne were married during 1993 one might then simply write

```
(married Joe Anne 1993)
```

Alternative styles include thinking of the world as consisting of things that have duration, so that one might write

```
(contains (time-of (marriage Joe Anne)) 1993)
```

where time-of is a function from things to the timeperiods they occupy.

Rather different axioms result if one thinks of these times as points or as intervals; but they share the property that the language itself is not temporally embedded, so that all quantifiers are timeless. Any temporal restrictions on quantification must be made explicit. It is therefore usually necessary to have some way of asserting that something exists at a time. This can be done by a relation existswhen between things and times, for example.

A familiar use of the temporal-argument technique is seen in the 'situation calculus', which describes actions in this way:

```
(forall (?t ?action)(=> (and (<action-conditions> ?action)
                              (<antecedent-conditions> ?t)    )
                         (<consequent-conditions> (do ?action ?t))
))
```

where `do` is taken to be a function from a time and an action, to the time at the end of the action. Here the time-ordering is defined simply by the syntactic nesting of '`do`' terms, so that `t` is always earlier than `(do a t)`. Since this is a partial order, the time-plenum is considered to be branching, with alternative futures corresponding to alternative courses of action.

The situation calculus assumes that the universe is stable unless actions happen to alter it. Thus, although a situation seems to be an interval, since nothing happens during these intervals they are point-like in the sense just described. Notice this is not here meant to imply that they are of short duration, only that they have no internal temporal structure which can be described in the language. The only temporal structure in the plenum assumed by the situation calculus is a partial ordering of such intervals; they cannot overlap or be contained in one another.

Variations on this style of axiomatic description have been much used in planning, and have become almost a standard in parts of AI. It has obvious connections with state-space descriptions of computation. For our purposes it is sufficient to note that while this style of description treats times as point-like and does not utilise the more complex interval relationships, the use of temporal arguments does not necessarily restrict one to this limitation.

2.2 Holding true

The second style asserts that sentences 'hold true' at times, so that one would write something like the following.

```
(holds (married Joe Anne) 1993)
```

(The use of a different font for this holds is deliberate, as `holds` already has a distinct meaning in KIF. The above expresion is not legal KIF!) Notice that the "inner sentence" is timeless. The symbols holds and `married` cannot here both be understood to denote conventional extensional relations, for then the first argument to holds would be a truthvalue. If the "inner" sentences are indeed sentences, then holds is essentially a modal operator. An alternative view is to regard all the inner expressions as terms denoting propositions, so that relation symbols such as "married" become functions from individuals to propositions. This has the awkward consequence of needing almost the entire syntax of first-order logic to be mirrored in an "inner" language of terms.

It is not exactly clear what holds means, however. If times are regarded as points (or quiescent intervals) and if the 'inner' language does not have quantifiers, then this can be straightforwardly translated back into plain logical syntax by applying the following recursive rules:

```
(holds (and φ,ψ)  ?t)  --->   (and (holds φ ?t) (holds ψ ?t))
(holds (or φ,ψ)  ?t)   --->   (or (holds φ ?t) (holds ψ ?t))
(holds (not φ)  ?t)    --->   (not (holds φ ?t))
(holds (R a1 ... an) ?t)  --->  (R a1 ... an ?t)
```

(Sometimes it may be appropriate to just forget the last rule. It depends on whether or not the relation is thought to be temporally sensitive. For example, the natural translation of

```
(holds (=> (and (man ?x)(married ?x Julia))(happy ?x)) T)
```

would be

```
(=> (and (man ?x)(married ?x Julia T))(happy ?x T))
```

since men may change their states of marriage and happiness, but rarely their sex.)

However, if the times are understood to be nonpointlike intervals, and holds means holding true throughout the interval, then negation and disjunction need to be treated more carefully. If it is possible for both ϕ and (not ϕ) to be true during parts of the interval, then

```
        (not (holds φ ?t))
```

is a weaker claim than

```
        (holds (not φ) ?t)
```

and the third rule is inadequate (although not incorrect).

Following Allen, we might replace the negation rule with

```
(holds (not φ)  ?t)  --->
                              (forall (?s)(=>(in ?s ?t)(not (holds φ ?s))))
```

where `in` is the relation between an interval and a containing interval. This only makes sense, of course, if appropriate axioms are provided for `in`. If `in` is taken to be irreflexive then this assumes that `?t` always has a proper subinterval, and this is not true in discrete time where `?t` might be a single clock-tick. The reflexive interpretation (where every interval is `in` itself) therefore seems more natural here; or we could assume that time is dense.

The disjunction rule *cannot* be similarly replaced with

```
** (holds (or φ,ψ)  ?t) --->
                              (forall (?s)
                              (=>(in ?s ?t)(or (holds φ ?s)(holds ψ ?s))))
```

To see why, consider an oscillator and let ϕ and ψ be respectively the propositions that it is in its two states. Then `(or φ ψ)` is always true, but there may be many subintervals during which the oscillator is not constantly in one state or the other, especially if we turn up the frequency. DeMorgan's laws provide the following translation for disjunction, which is regrettably complicated:

```
(holds (or φ,ψ)  ?t) --->  (forall (?s)(=>(in ?s ?t)
                                (exists (?u)(and(in ?u ?s)
                                       (or (holds φ ?u)(holds ψ ?u))))))
```

Quantification is also a little more complicated. If we allow quantified sentences to `hold` at a time, allowing expressions like

```
        (forall (?x)(=> (holds (exists (?y ?z)
                                (and (not (= ?y ?z))
                                       (married ?x ?y)
                                       (married ?x ?z)))
                        ?t)
                (exists (?t1) (and (later ?t ?t1)
                        (holds (troubled ?x) ?t1)))
        ))
```

then it is natural to think of quantifiers inside holds as ranging over only those individuals that exist at the time mentioned, so that one might write

```
        (holds (not (exists (?x)(and (heavierthanair ?x)
                                       (flyingmachine ?x))))
               1800)
```

With this interpretation, holds is referentially opaque.

The simplest way to interpret this language (still thinking of times as points or point-like intervals) can then be defined by extending the recursive translation with the following two rules:

```
(holds (forall (?x) φ)  ?t)   --->
            (forall (?x)(=> (existswhen ?x ?t)(holds φ ?t)))

(holds (exist (?x) φ)  ?t)   --->
            (exist (?x)(and (existswhen ?x ?t)(holds φ ?t)))
```

where existswhen is the relation mentioned earlier which relates something to the time when it is exists. The quantifiers after this translation are syntactically identical to those before, but have a rather different meaning from: they are timeless, while those before are temporally restricted. Again, however, if propositions hold during (nonquiescent) timeintervals, then the existential translation has to be more baroque:

```
(holds (exist(?x) φ)  ?t) --->
                  (forall (?s)(=>(in ?s ?t)
                        (exist (?x)(and (existswhen ?x ?s)
                              (exist (?u)(and (in ?u ?s)(holds φ ?u)))
                               ))
                              ))
)
```

(To see the need for all this fuss consider a relay race, let φ mean that ?x is carrying the baton, and suppose that the first runner dies of heart failure after his lap, but before the race is over. Somebody is carrying the baton even when one of the earlier carriers has ceased to exist. Or consider a belt delivering pieces of coal to a furnace, and the claim that it is never empty.)

One claimed advantage of this holds notation is that it allows a variety of different ways that something might be true at a time. Allen (1984) for example distinguishes between propositions (he calls them properties), which hold and satisfy subinterval inheritance, events which occur and are never inherited by subintervals, and processes which also occur but which are inherited by some subintervals. Galton (1990) similarly distinguishes three kinds of holding; holds-at, meaning true at a point, holds-in and holds-on. I have not yet found a need for this apparent advantage, however. The distinctions seem to always be describable as differences in the kind of interval or kind of proposition, obviating the need for the notational and axiomatic complexity introduced by such intricate distinctions of vocabulary.

2.3 Tenses[1]

Tense logics extend conventional logics by modal operators corresponding to the English past and future tenses, so that one would write

```
(F (married joe anne) )
```

to mean that Joe will be married to Anne at some time in the future, without mentioning times explicitly. P similarly refers to the past. Tense logics do not usually provide any way to refer to dates directly, and one would need to write something like

[1] Tense logics are thoroughly described in the logical literature and their properties well understood. I mention them here only for completeness. This brief section is not an adequate introduction to the subject.

```
(F (and (married joe anne) (date-is 1993)) )
```

to convey such information (and provide axioms for date-is, of course). Tense operators are part of the sentential syntax in just the same way as the propositional connectives, so that one can have such assertions as

```
(F (and (not (married joe anne)) (P (married joe anne))))
```

which says that Joe and Anne will be divorced at some time in the future.

Notice that there is an implicit "now" in such an assertion. These languages are temporal in a rather deeper sense than the previous ones, since here the very act of assertion is itself understood to have a temporal location. A tense-logical sentence only makes sense relative to a time when it is understood to be asserted. The usual model theory for such languages interprets a sentence to mean a function from temporally possible worlds to truth-values. Different axioms for F and P correspond to different assumptions about the relationships between these possible worlds, which amount to different assumptions about the structure of the time-plenum.

Tense logics can usually be translated into a theory written using `holds`. The usual method of translation introduces a binary ordering relation between times. It amounts to a recursive application of the following transformation function Σ, which takes a sentence ϕ in tense logic and a variable `?t` denoting a time, to a sentence in the holds language:

If ϕ contains no modal operators then

```
Σ[ϕ,?t]          -->      (holds ϕ ?t)
```

and otherwise

```
Σ[F(ϕ),?t]  ->  (exists (r)(and(before ?t, ?r)(holds Σ[ϕ, ?r] ?r)) )
Σ[P(ϕ),?t]  ->  (exists (r)(and(before ?r, ?t)(holds Σ[ϕ, ?r] ?r)) )
```

where `?r` is a variable different from any other free variables. This captures the standard semantics of the modal operators, which contain an implicit quantification over temporally possible world, here regarded as times. Different modal logics correspond to different axioms describing the relation `before`.

For example, the earlier assertion concerning the future divorce of Joe and Anne translates into:

```
(exists (?t1,?t2) (and (later ?t1,?t)
                       (later ?t1,?t2)
                       (holds (married Joe Anne), ?t2)
                       (holds (not (married Joe Anne)),?t1)  ))
```

Notice that the variable `?t` is free in this expression . The translation process always leaves a single free temporal variable corresponding to the implicit ‘now’.

This rather simple translation from modal tense logic to a nonmodal language is not always completely adequate to capture some of the more complex tense-logics, in particular those involving a “continuous present” in which there is a modal operator corresponding to the assertion that something is happening, as opposed to merely true. Nevertheless, similar translations into a nonmodal language are often possible, and many of the special temporal logics which have been developed for AI use are quite unnecessary.

2.4 Temporal knowledge-bases

Finally, the language need have no explicit or implicit temporal reference, but simply be understood to be asserted (or believed) in a temporal framework, with some other mechanism keeping track of when it is supped to be true. Typically, the context is a database or knowledge-base of some kind which is keeping track of a changing world, and the problem is to maintain consistency with the changing state of the environment. Examples include dynamic control of a robot or an industrial system, knowledge fusion in military command situations, 'temporal databases' and truth-maintenance methods. While such a wide variety may not be solvable by a single technique, they all involve temporally sensitive representations in which the representation language itself temporally 'located' and often therefore need not make explicit reference to times.

In these cases there is often a sense that processes which manipulate sentences need to be made sensitive to the passage of time, in contrast to the typical planner which uses the first kind of notation to reason about times but has no way to express the concept of 'now'.

If a temporally located language is itself temporally expressive, then the situation becomes more complex. Under these circumstances the meaning of 'now' is constantly changing in a way which can be expressed in the language, so that future is constantly becoming the present. If the language is suitably expressive this introduces many complex representational problems which are beyond the scope of this document. For example, I know of no really satisfactory way of formalising the meaning of a statement of urgency such as “we must act soon or the bridge will collapse”, one that would enable us to *infer* the need to stop wasting further time doing more *inference*. It seems to be necessary to assume something like an interrupt architecture in the structure of the reasoner, for example.

Issues like this arise in manipulating temporal databases, especially when data entries which refer to the past are liable to correction in the future. For example, temporal database terminology (Snodgrass *et al* 1993) distinguishes the transaction time of a record - the time when it was entered into the database - from the valid time, which is the time that the event or proposition it represents is asserted to be true.

3. Temporal Theories

The theories we consider differ in several respects, both in their axiomatic style and in the intuitions they support. Most theories are consistent with the following ideas:

1. A timeinterval is a connected piece of the time-plenum. Things that have a temporal extent, or which occupy time in some way, must have a timeinterval which is their life-span.
2. Every interval has a unique temporal magnitude.
3. An interval has two endpoints, and is uniquely determined by those endpoints.
4. A point can be uniquely determined by the magnitude of the interval between it and some special 'zero' temporal position (such as midnight, or January 1 of the year zero)

All of this is compatible with various ideas about the structure of the continuum. Some theories describe the point continuum, others the glass continuum, and others assume discrete atoms of time. All the axioms use nontemporal quantification, referring to times explicitly as objects.

Notation

Theory boundaries are indicated in the text and given exactly in Appendix 3. NOTE: **simply conjoining all the axioms listed in a section may not always form a coherent theory**, since alternatives are often listed together. All the axioms and definitions are given in KIF, (although KIF-style notation is also used to discuss other axioms which are not correct KIF, notably second-order axioms.) Sequence quantification is sometimes used, and variadic relations are used when convenient, but most of the axioms are equivalent to simple first-order logic. KIF text is written in `Courier` font.

Every axiom has been given a name, following some simple naming conventions. A theory name is all uppercase. An axiom name like LP-foo means an axiom called foo in the theory LP (or sometimes, to save space, a contraction of the full name is used), but a name like foo-LP means the *definition* of the name foo in the theory LP, and a name like foo-LP-baz means *part of* the definition of foo in the theory LP.

Whether or not a theory constitutes an 'ontology' is still an open question at the time of writing. The chief problem seems to be the question of definitions. In an ontology every name must have a 'definition', but in a nontrivial logical theory it is impossible for every relation and function name to have a defintion, since defined terms are eliminable. In Appendix 3, definitions are indicated as far as possible using the distinctions in KIF, so that ' := ' indicates a logical definition while the weaker sense of 'definition' really only means something like 'relevant to', and indicates a collection of axioms which constitute a theory of the name in question. Some names can't have a definition, notably the basic sortal predicates such as `timepoint`; while in other cases, it is impossible to separate the axioms into disjoint sets each of which defines one name. The name index in Appendix 2 gives, for each relation and function name, a list of the theories which can be taken to establish its meaning in one way or another.

For many of the theories we discuss several possible models or classes of models, especially nonstandard models which differ in some respect from the intended intuition. This is a deliberate attempt to show the limitations of expressive power that many axiom sets have. That axioms have nonstandard models is not necessarily a condemnation of them, however, and is not meant to be so taken. Any first-order theory of arithmetic has nonstandard models, for example.

The chief categories are indicated by the relations `timepoint`, `timeinterval` and `timeduration`. These relations are all *predicative*, i.e. they are variadic but simply assert that a unary predicate holds of all their arguments. A relation Ψ is predicative when:

```
(forall (?x @l) (<=> (Ψ ?x @l)
                     (and (Ψ ?x)(Ψ @l))))
```

Another use of variadic notation throughout this document will be in propositions involving chained relations. A relation is *chained* when it asserts a binary relation between each of its arguments in succession, ie when:

```
(forall (?x ?y @l) (<=> (Ψ ?x ?y @l)
                     (and (Ψ ?x ?y)(Ψ ?y @l))))
```

Examples of chained relations include equality, the temporal ordering of timepoints and the `meets` relation between intervals. A chained relation may or may not be transitive.

These axioms can be transcribed into KIF by using the 'holds' notation by the following three axioms, which constitute the theory BASIC-SYNTACTIC-TOOLS used throughout the rest of the catalog:

```
;;predicative-BASIC
(defrelation predicative (?r) :=
     (forall (?x @l)
             (<=>
              (holds ?r ?x @l)
              (and (holds ?r ?x) (holds ?r @l)))))

;;chained-BASIC
(defrelation chained (?r) :=
     (forall (?x ?y @l) (<=> (holds ?r ?x ?y @l)
                             (and (holds ?r ?x ?y) (holds ?r ?y @l)))))

;;BASIC-syntax
(predicative predicative chained)
```

3.1. Simple Point Axioms

These axioms simply describe an ordering of timepoints. Intervals, etc. are not mentioned. The chief interest of this is to act as a definitional support for later, more complex, theories.

All the quantifiers in the axioms in this section should be understood to be restricted to `timepoint` when the axioms are used in a broader context. This restriction is omitted here for clarity.

```
;;LP-syntax
(predicative timepoint)
```

There is one relation, `before`, between timepoints. It is a chained relation:

```
;;before-LP-syntax
(chained before)
```

It should follow from `before-LP-syntax` and BASIC-SYNTACTIC-TOOLS that

```
(chained before)
```
and hence follow from
```
(before a b c)
```

that `(before a b)` and `(before b c)`.

Right now (10/10/94) I am not certain that this does follow in KIF , in fact.

Ordering axioms: `before` is transitive, irreflexive and linear:

```
;;before-LP-trans
(forall (?x ?y ?z)(=> (and (timepoint ?x ?y ?z)(before ?x ?y ?z))
                      (before ?x ?z)))

;;before-LP-irreflex
(forall (?x)(=> (timepoint ?x)(not (before ?x ?x))))

;;before-LP-order
(forall (?x ?y)(=> (timepoint ?x ?y)
                   (or (= ?x ?y) (before ?x ?y) (before ?y ?x))))
```

These axioms are sufficient to guarantee that all timepoints lie on a single line. However, it allows that line to be either finite or extremely infinite, as in the nonstandard models described below.

Time is infinite in both directions:

```
;;before-LP-infinite-past
(forall (?x)(=> (timepoint ?x)
                (exists (?y)(and (timepoint ?y)(before ?y ?x)))))

;;before-LP-infinite-future
(forall (?x) (=> (timepoint ?x)
                 (exists (?y) (and (timepoint ?y)(before ?x ?y)))))
```

This pretty much accounts for the overall shape of the temporal universe of points, although it will be useful later to define the 'before-or-equal' relation:

```
;;bbefore-LP
(defrelation bbefore (?x ?y) :=
(and
(chained bbefore)
(<=> (bbefore ?x ?y)(or (before ?x ?y)(= ?x ?y)))
```

These axioms constitute the basic theory of point-ordering, called LINEAR-POINT, or LP.

The fine structure can be described in two different ways, as dense or discrete.

Fine-structure axioms:

```
;;LP-dense
(forall (?x ?y) (=> (timepoint ?x ?y)
                    (=> (before ?x ?y)
                        (exists (?z)(and (timepoint ?z)
```

```
                                    (before ?x ?z ?y))))))
```

This asserts that timepoints are dense. Adding this to the definition of 'before' produces the theory DENSE-LINEAR-POINT, or DENSE-LP

An alternative (and inconsistent) assumption is that there is an atomic spacing of timepoints which allow no closer divisions; ie, that there simply are no points between two adjacent atomic 'ticks'. This is a bit more complicated to express:

```
;;DLP-discrete

(forall (?x)(=>(timepoint ?x)
        (and
         (exists (?y)
                 (and (timepoint ?y)
                      (before ?x ?y)
                      (not (exists (?z)(and (timepoint ?z)
                                            (before ?x ?z ?y)))))
          (exists (?y)
                  (and (timepoint ?y)
                       (before ?y ?x)
                       (not (exists (?z)(and (timepoint ?z)
                                             (before ?y ?z ?x))))))))))
```

Adding this to LP produces DISCRETE-LINEAR-POINT, or DLP.
Notice that DLP-discrete already has the infinity assumptions built into it, so the infinity axioms are redundant in DLP

Van Benthem (1983) shows that DENSE-LP and DLP are *syntactically complete*. That is, any assertion using this (admittedly rather restricted) vocabulary which is true in all first-order models of these axioms can be deduced from these axioms. They entail all there is to be entailed; all sentences (written in this vocabulary) which are consistent with them, are already provable from them.

So there's nothing more to be said about linearly ordered times without extending the vocabulary in some way.

Models

The obvious models are the rationals **Q** of DENSE-LP and the integers **N** of DLP. The real line **R** is also a model of DENSE-LP, but it is not easy to distinguish **Q** from **R** within a first-order theory.

However, there are also many nonstandard models. For example, consider the rational plane with

$\langle a,b\rangle$ before $\langle c,d\rangle$ iff $a<c$ or $a=c$ and $b<d$.

This amounts to **Q** copies of **Q** ordered in sequence, and it is a model of DENSE-LP. Similarly, **N** copies of **N** ordered in sequence is a model of DLP. These first-order theories are not sufficiently powerful to state that all pairs of points are only a finite distance apart; they can describe the 'local' structure of the line, but all they can say about the 'global' structure is that it is a total order, and to fully capture the structure of the line one needs to say more than this, as the existence of these models shows.

The key problem is how to express ‘finite’. In the standard models two points can only be a finite distance apart, and this cannot be expressed in this (or any other) first-order vocabulary. For example, an axiom which says that the distance between two points is always an integer is still true in the nonstandard models.

Exactly what counts as a 'point' is also not clearly specified by these simple theories. Consider for example the model of DLP created from the standard model by simply removing the closed unit interval around zero, so that

a before b iff $(a < b < -1)$ or $(a < -1$ and $b > 1)$ or $(1 < a < b)$

This line with a hole in it (or, with the zero point swollen into an interval) is indistinguishable from the standard model from within this theory. It is order-isomorphic to the rational line.

Theories which more exactly specify the standard models can be got by adding a little second-order expressiveness, provided of course that these second-order quantifiers are understood to have the standard interpretation, ie varying over *all* properties. Adding the following axiom to DENSE-LP specifies the integers up to isomorphism. (This is Dedekind's principle of continuity, asserting that whenever a property changes with time, there must be a 'dividing point' which separates the times when it is true from those when it is not (taken in this form from Van Benthem 1983).) When added to the DLP theory this still does not quite nail down the standard interpretation, since **R** copies of **R** is still a nonstandard model, but it forces the models to be suitably dense in timepoints, ruling out the mere rationals and guaranteeing that there are no gaps.

Dedekind:
```
(forall (?p)(=>
   (and
       (exists (?x)(?p ?x))
       (exists (?x)(not (?p ?x)))
       (forall (?x ?y)(=>(and(?p ?x)
                            (not(?p ?y)))
                       (before ?x ?y)))
    )
    (exists (?z)(forall (?u)
                  (and  (=>(before ?u ?z)
                            (?p ?u))
                        (=>(before ?z ?u)
                            (not (?p ?u))
                   )
    )))
```

When interpreted in the integers, the ?z asserted to exist is the last integer when ?p is true (or the first where it is false); when interpreted in the reals, it is the 'dividing' point where the intervals of ?p's truth and falsity meet one another.

To see how this rules out the nonstandard models, consider a model of DLP consisting of two copies of **N**, and let **p** be the property of being in the first half of this double-line model. Then there is no single point which ‘divides’ the points satisfying **p** from those that do not: but the Dedekind axiom requires such a point to exist.

This property **p** is not expressible in the language, however, so if we interpret the second-order quantifier in the Henkin sense, ie as varying over all relations which can be named using the vocabulary of the theory, then the Dedekind axiom no longer has the semantic force necessary to eliminate nonstandard models. The importance of this lies in the fact that there is no computational way to distinguish a ‘classical’ semantics from the Henkin semantics, since second-order logic is complete relative to the Henkin interpretation. Dedekind is essentially second-order: a first-order transcription of it would not carry the semantic force needed to rule out the nonstandard models of out axioms.

We could take bbefore as the basic relation and define `(before ?x ?y)` as `(and (bbefore ?x ?y)(not(bbefore ?y ?x)))`. This has the advantage of not needing to use equality. However, there are now some rather unintuitive models, eg one in which the `bbefore` relation is circular and the `before` relation is therefore everywhere false. These correspond to nonstandard models of equality which allow equivalence classes of indistinguishable individuals. If the logic has equality, then asserting the connection is sufficient to explicitly eliminate these peculiar interpretations:

```
(forall (?x ?y)
    (=> (bbefore ?x ?y ?x)  (= ?x ?y)))
```

A useful extension to the vocabulary is provided by skolemising `DLP-discrete`, which gives the functions “next” and “previous”:

```
;;next-DLP
(deffunction before
(forall (?x)(=>(timepoint ?x)
        (and
         (before ?x (next ?x))
         (not (exists (?z)(and (timepoint ?z)
                               (before ?x ?z (next ?x)))))
))))

;;previous-DLP
(deffunction previous
(forall (?x)(=>(timepoint ?x)
        (and
         (before (previous ?x) ?x)
         (not (exists (?z)(and (timepoint ?z)
                               (before (previous ?x) ?z ?x)))))
))))
```

3.2 Alternative: nonlinear time

If we simply omit `before-LP-linear` then timepoints can be only partially ordered. Then for example one model is the real plane with <x,y> before <u,v> just when x<u, which allows every point to have infinitely many immediate successors and predecessors. Often we want to insist that time only branches in the future direction. To do this, replace the axiom with:

```
;;before-BP-order
(forall (?x ?y ?z)
        (=> (and (timepoint ?x ?y ?z)(before ?y ?x) (before ?z ?x))
            (or (= ?z ?y) (before ?z ?y) (before ?y ?z))
        ))
```

This allows a 'forest' of branching structures. To restrict to a single tree add:

```
;;before-BP-tree
(forall (?x ?y)
```

```
        (or (= ?z ?y) (before ?z ?y) (before ?y ?z)
            (exists (?z) (and (before ?z ?x) (before ?z ?y)))))
```

which forces all timepoints to fit into a single tree-structure.

If we make these changes and omit the backward-infinity axiom (to allow the common interpretation in planning systems in which a 'start' time is considered for the planning process), the resulting theory is christened BRANCHING-POINT. Just as in the linear case, this can be extended by assuming either density or discretness. The density axiom LP-dense works here as before, but the discreteness axioms needs to be stated slightly differently.

```
;;BP-discrete
(forall (?x)(=>(timepoint ?x)
   (and (exists (?y)
           (and (timepoint ?y)
                (before ?x ?y)
                (not (exists (?z)
                           (and (timepoint ?z)(before ?x ?z ?y))))))
        (=> (exists (?y)(and (timepoint ?y) (before ?y ?x)))
            (exists (?y) (and (timepoint ?y)
                              (before ?y ?x)
                              (not (exists (?z)(and (timepoint ?z)
                                              (before ?y ?z ?x)))))))
)))
```

Adding BP-discrete to BRANCHING-POINT gives the theory DISCRETE-BRANCHING-POINT or DBP.

Models

Possible models include any suitably large tree-structured graph. In particular, a Herbrand universe of terms provides a model, in which `before` denotes the subterm relation. (Notice this would no longer be true if we add the backwards infinity axiom, since there are always terms with no subterms.)

Any model of the linear theory is also a model of the corresponding branching one, of course, so the nonstandard “stretched” orderings are still possible here. However, there are also other more exotic possibilities. For example, consider the set of finite sequences of integers, and interpret `before` as the initial-subsequence relation. Then the timeordering is like the integers, satisfying BDP, but every time has infinitely many immediate successors, so the branching rate is infinite. It is not clear whether this should be regarded as “nonstandard”, however. I do not know how to specify simply that the branching rate is finite without describing some definite branching pattern.

If we add the backward infinity axiom then infinite-branching models are still possible. For example, consider infinite sequences of integers which are not isomorphic to a subsequence of themselves. Think of these as directed backwards in time, so each represents an infinite history reaching to the present. Now let `before` denote the tail-subsequence relation. Again there is infinite branching, but now each time also has an infinite past.

3.3 Alternative: Situation calculus time

The interesting part of the situation calculus notation from our perspective is how it describes timeordering. Times –situations – are partially ordered by the structure of the '`do`' terms which can be generated by instantiation from those which occur in the action axioms. The set of possible such terms amounts to a subset of the Herbrand universe of the language.

Situation calculus axiomatisations succeed as a basis for planning only if their ways of referring to situations are restricted to this subset. For example, suppose that we wish to assert an optimistic claim that somewhere in the future of any situation, it will rain pennies from heaven. It would be natural to write something of the form

```
****   (forall (?s) (exists (?t) (and (before ?s ?t)
                                   (rains-pennies-from-heaven ?t)) ))
```

But when skolemised, this would completely trivialise the task of budgetary planning: just `do` the skolem function. The situation theory implicitly assumes that functions from situations to situations represent actions that can be performed, not just assertions about times that may exist. Because of this restriction, we could connect the usual situation-calculus notation to the present theories by asserting that one time is before another just when there is a sequence of actions which take the first to the second. This can be stated in KIF using sequence variables, as follows.

```
;;done-BDP
(deffunction done
(forall (?s ?a @l)(and
                    (= (done ?s ?a @l) (do ?a (done ?s @l)))
                    (= (done ?s ?a) (do ?a ?s)) ))
)
```

That is, `done` takes a situation and a sequence of actions and returns the situation resulting from do-ing those actions in that sequential order. The following axiom can then be regarded as an alternative definition of the point-ordering relation <u>before</u> in terms of the <u>do</u> relation.

```
;;before-do-BP
(forall (?x ?y) (<=>(exists (@l)(= ?y (done ?x @l)) )
                   (before ?x ?y)))
```

Notice the biconditional, which makes this rather a strong claim;in particular, makes it inconsistent with the density assumption. (If the biconditional is replaced with a simple forward conditional then the theory is consistent with density, but then this axiom cannot be regarded as a definition of `before`.) Transitivity now follows from properties of sequences. Irreflexivity however is not trivial: it essentially asserts that there are no 'null' actions. Backward-infinity must be rejected. Forward-infinity can be added, but it is consistent only when the preconditions are such that it is always possible to apply an action. (This would be satisfied, for example, by having two actions each of which undoes the effect of the other.)

<u>3.4 Variations</u>

Essentially the same ideas can be expressed in different ways.

<u>a. Cohistorical (replaces 3a/b)</u>

`Cohistorical` is a chained relation meaning 'on the same time-line':

```
(defrelation cohistorical (?x ?y) :=
            (and
            (<=> (cohistorical ?x ?y)
                  (or(= ?x ?y)(before ?x ?y)(before ?y ?x)) )
            (<=> (cohistorical ?x ?y @l)
                   (and (cohistorical ?x ?y) (cohistorical ?y @l)))
```

The linearity axiom can then be expressed more neatly by :

```
(forall(?x ?y)(cohistorical ?x ?y)))
```

and the tree axiom for branching time by

```
(forall (?x ?y ?z)(=> (and (before ?y ?x) (before ?z ?x))
                        (cohistorical ?y ?z)))
```

b. Timelines

Still another variation introduces timelines explicitly.

We use a relation on between a timepoint and a timeline.

```
;;on-Timeline-1
(forall (?x ?y)(=> (timepoint ?x ?y)
        (<=>
         (exists (?h) (and (timeline ?h)(on ?x ?h) (on ?y ?h)))
         (or (= ?x ?y) (before ?x ?y) (before ?y ?x)) )))
```

Now, to assert linearity we can just claim that timelines do not overlap:

```
;;on-Timeline-2
(forall (?h ?k)
        (=> (and (timeline ?h ?k)(exists (?x)
                   (and (timepoint ?x)(on ?x ?h) (on ?x ?k)))
            (= ?h ?k) ))
```

This allows several timelines to exist, but they can have no connection to one another. We can capture the branching structures by insisting that overlapping is only possible in the past:

```
;;on-Timeline-3
(forall (?h ?k ?x)
        (=>  (and (timeline ?h ?k)(timepoint ?x)(on ?x ?h) (on ?x ?k))
             (forall (?y )(=> (timepoint ?y)
                     (=>(before ?y ?x)
                        (<=>(on ?x ?h) (on ?y ?k))
                        ))
                     )
             ))
```

Several authors have argued that since time itself is linear, the apparent branching of alternative futures should be thought of as just one kind of hypothetical reasoning about alternatives, having no particular connection with time. The right way to think of the branching futures axioms, on this view, would be as alternative world-lines in which only the future is allowed to vary, presumably on the grounds that less is known about it than about the past. This intuition can be described reasonably well here by distinguishing two kinds of existential claim. A quantification over a timeline represents a claim of possibility, but a quantification over a timepoint within a timeline is a temporal statement.

c. tick-tock

Discrete time can be explicitly identified with the integers by assuming a function `tick` and its inverse `tickof`:

```
;;TT-syntax
(and (forall (?n)
             (<=> (integer ?n)
                  (timepoint (tick ?n))
                  (forall (?x) (<=> (timepoint ?x)
                                    (integer (tickof ?x)) )))))

;;tickof-TT
(and (forall (?x) (= ?x (tick (tickof ?x)))
             (forall (?n) (= ?n (tickof (tick ?n)))  )))
```

The discreteness axiom follows from this (by identifying the ?y variables with (tick n-1) and (tick n+1)), and the 'infinity' axioms can obviously be omitted. The useful functions `next` and `before` can be defined:

```
;;next-TT
(defunction next (?x) := (tick (+ 1
                                     (tickof ?x))))

;;before-TT
(forall (?x ?y) (<=> (before ?x ?y)
                     (< (tickof ?x) (tickof ?y))))
```

and the earlier ordering can be defined using the conventional arithmetic vocabulary in the obvious way:

```
  (forall (?x ?y)(<=> (before ?x ?y)
                      (lessthan (tickof ?x)(tickof ?y))))
```

Models

This style of axiomatisation depends crucially upon the arithmetical terms 'integer' , 'plus' and 'lessthan'. If we take the meaning of these to be independently established, then the only models of the tick-tock language are isomorphic to the integers, ie standard models. While there is no first-order theory which can make such a guarantee, writers of axioms often simply assume that such language is available.

If we do not make this assumption, then of course many nonstandard models exist.

Some apparently unimportant, but potentially confusing, variations are possible between different conventions for how intervals are described in terms of 'ticks'. The pair <tick(n),tick(m)> can be thought of as identifying an interval in at least three different ways, as shown in figure 1. The first thinks of the interval as containing these as its end-ticks, so that the intervals <...,tick(m)> and <tick(m+1),...> meet one another with no space between. The second thinks of tick as identifying the spaces between atomic moments of time, so that the intervals <...,tick(m)> and <tick(m),...> meet one another. We will use this second convention.

However, these axioms are also consistent with the idea that a timepoint is an 'atomic' interval, by using the third convention illustrated, where an interval is identified by its last time-tick, so that the interval would be described as <tick(n-1),tick(m)>. In the first case, the interval <tick(m),tick(m)> contains a single clocktick; in the second and third cases such a construction is meaningless, and the single-tick moment would be described as <tick(m-1),tick(m)>. The difference between the second and third cases corresponds to the difference between thinking of timepoints as lying between moments, or as being atomic moments themselves. Much of the literature on temporal databases fails to distinguish between timepoints and moments, using the term “instant”ambiguously for both. As we show later, there is a coherent theory which supports this ambiguity, but it is not consistent with conventional real analysis.

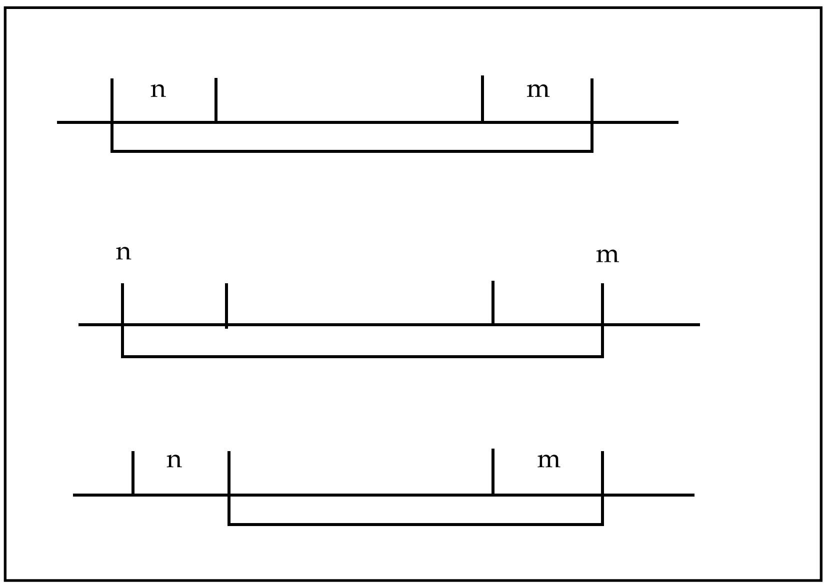

Figure 1. Three ways to label end-ticks

4. Interval theories

The axioms in this chapter all quantify over intervals of time. Intervals are not totally ordered, and a wide variety of relationships might hold between two intervals even if we assume, as I usually will in this section, that time is linearly ordered. Allen (1984) lists thirteen relations; equality and the following six plus their inverses (ie with arguments reversed) precedes, meets, overlaps, starts, finishes and during. This is often taken to be a standard set, but we will also consider theories which use subsets of these and which use different sets of relations. Any relation between intervals (on the line) can be defined in terms of these six, but smaller sets also suffice to define them all. We will give a theory expressed entirely in terms of meets within which all the relations can be defined, for example. These thirteen relations are all the relations which completely specify the relative orderings of the endpoints of the intervals. This suggests a way to define them within the timepoint theories, and we also present a theory based on this intuition.

The distinction between dense and discrete time also runs through these theories. Most of them can be extended by axioms which restrict to one or the other case, but they are also often consistent with the idea that time is dense in some places and discrete in others.

4.1 Intervals as information

An interval might be thought of as giving partial information about the location of a point: as an approximation to a point, in a sense. Two intervals might be sufficiently distinct that they establish the relative ordering of the points they contain: following Allen, we will call this relation `precedes` (although we think of it as more like the disjunction of his *precedes* and *meets* relations.) On the other hand, one interval might surround the same point as another but more precisely; let us call this `finer`. Both are orderings, but `finer` is partial and `precedes` is total.

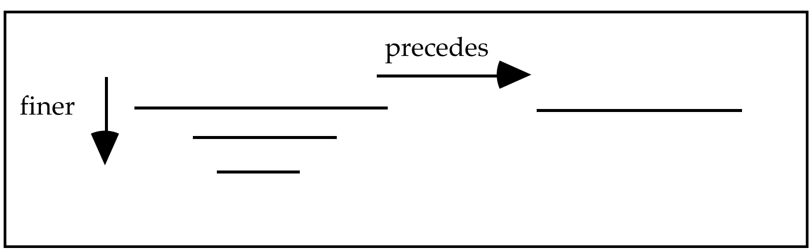

The following axioms constitute a theory called APPROXIMATE-POINT, since it treats intervals as approximations to points.

```
;;AP-Syntax
(chained precedes finer)

;;AP-Trans-1.
(forall (?x ?y ?z) (=> (precedes ?x ?y ?z)
                       (precedes ?x ?z)))

;;AP-Trans-2
(forall (?x ?y ?z) (=> (finer ?x ?y ?z) (finer ?x ?z)))

;;AP-reflex
```

```
(forall (?x) (and (not (precedes ?x ?x)) (finer ?x ?x)))

;;AP-finer-asym
(forall (?x ?y)(=> (finer ?x ?y ?x)
                   (= ?x ?y)))
```

If two intervals are not clearly separate, then they must somehow intersect. That is, there must be a shorter interval contained in them both. This leads to the axiom corresponding to the total-order assumption for points:

```
(forall(?x ?y)(or (precedes  ?x ?y)
                  (precedes  ?y ?x)
                  (exists (?z)(and  (finer  ?z ?x)
                                    (finer  ?z ?y)))
      ))
```

It is easier to state this using an intermediate concept of ‘not clearly distinguishable from’, or `ncdf`:

```
;;AP-ncdf
(defrelation ncdf (?x ?y) :=
  (exists (?z) (and (finer ?z ?x) (finer ?z ?y))))
```

`ncdf` plays the role here that equality does in the linear-point theory (which might be called an exact-point theory):

```
;;AP-orthogonal-1
(forall (?x ?y)
        (or (ncdf ?x ?y) (precedes ?x ?y) (precedes ?y ?x)))
```

Comparing this with `before-LP-linear` shows how intervals can be interpreted naturally as approximations to points.

The relation `ncdf` is not transitive, since `(ncdf a b)` and `(ncdf b c)` are consistent with `(precedes a c)`:

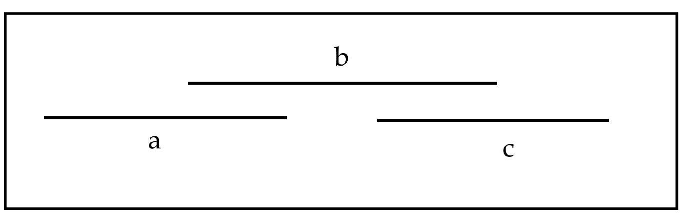

However, under these circumstances there must be two subintervals `e` and `f` of `b` such that `(precedes e f)`, ie the overlapping interval can be separated into two subintervals which are order-distinguishable.

We need a few more axioms which establish the connection between `precedes` and `finer` :

```
;;AP-orthogonal-2
(forall (?x ?y) (not (and (finer ?x ?y) (precedes ?x ?y))))

;;AP-separation
(forall (?x ?y ?z) (=>
                    (and  (finer ?x ?y)
```

```
                              (precedes ?y ?z))
                       (precedes ?x ?z) )
          )
```

It now follows that the disjunction in AP-orthogonal is exclusive: for if `(finer ?z ?x)` and `(precedes ?x ?y)`, then `(precedes ?z ?y)` by separation; but `(finer ?z ?y)`, contradicting the second orthogonality condition.

Notice that of Allen's relations, precedes, overlaps , and during are the only ones which do not specify that two endpoints are identical, so we might expect them to correspond to these relations. Following the claim that we can only have approximate information, let us say that meets, start, and finish are always false. Then we could read `ncdf` as meaning the disjunction of equality and the relations {`o,d,di,oi`}.

The infinite extension of the timeline can be asserted in the obvious way:

```
;;AP-infinite-past
(forall (?x) (exists (?y) (precedes ?y ?x)))

;;AP-infinite-future
(forall (?x) (exists (?y) (precedes ?x ?y)))
```

However, this still does not guarantee that any large intervals exist. For example, the set of all rational intervals less than unit length would satisfy the axioms so far. One way to assert this is to claim explicitly that there is an interval which is less precise than any two intervals:

```
;;AP-large
(forall (?x ?y) (exists (?z) (and (finer ?x ?z) (finer ?y ?z))))
```

Models

> This theory, in spite of its claims to infinite extension, can be straightforwardly interpreted as being entirely about the open subintervals of the unit real interval. The axioms guarantee that there there will always be a future, but they say nothing about how long that future may last. This can only be achieved by talking about the durations of intervals, as in section 6.

Density, which in a sense is the opposite of AP-large, is also straightforward:

```
;;AP-dense
(forall (?x) (exists (?y) (and (finer ?y ?x) (not (finer ?x ?y)))))
```

Although this way to describe density is fairly standard, its practical utility is not so obvious. The axiom rules out non-dense models by insisting that finer and finer distinctions must be possible. If we skolemise this and think of it computationally, it can be interpreted as saying that finer *measurements* are always possible. In practice this is not usually the case, and there is a definite limit to the precision with which measurements can be made, a 'grain size' past which it is not possible to discriminate separate points; or, equivalently, a size below which intervals seem like points. We could just say that the world is therefore 'really' discrete, as is usually done in temporal databases, where times are usually reckoned as integer counts of a clock-tick (Snodgrass et al 1994). But this is unsatisfactory for several reasons. We would like to be able to work with the assumption of density up to a point, as it were; and second, discreteness, as it is usually described, actually is too rigid. For example, it is quite possible for two clocks, both at our limits of discernment, to be beating out of phase with each other; but this is impossible if *time itself* is reckoned to be discrete.)

Anyway, its about the only apple in the shop, so we will adopt it as the density axiom. Given this, the Allen relation of meets can be defined using the idea that the meeting place can be contained in finer and finer intervals. If two intervals meet, and a third overlaps the meetingpoint, then if the line is dense it is always possible to find a subinterval which also overlaps that point and hence is *notclearlydistinguishablefrom* both the touching intervals:

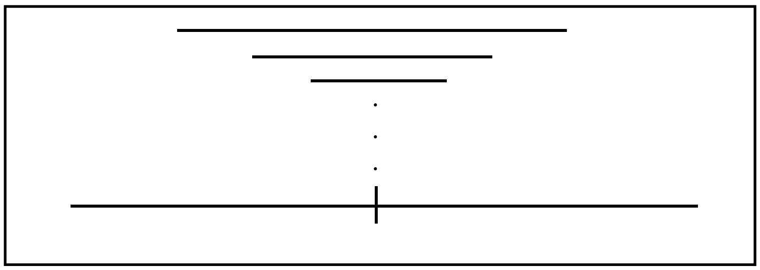

```
;;meets-AP-dense
(defrelation meets (?x ?y) :=
             (and (precedes ?x ?y)
                  (forall (?z)
                          (=> (ncdf ?x ?z ?y)
                              (exists (?u)
                                      (and (finer ?u ?z)
                                           (ncdf ?x ?u ?y)))))))
```

Discreteness can be defined by following the intuitions behind those in the point axioms. First define a notion of an atomic clock-tick:

```
;;moment-ADP
(defrelation moment (?x) :=
             (not (exists (?y) (and (finer ?y ?x)
                                    (not (= ?x ?y))) )))
```

or equivalently

```
(defrelation moment (?x) := (forall (?y)(=>(finer  ?y ?x)
                                           (= ?x ?y)) ))
```

Discreteness could now be stated by the claim that all intervals have a next moment and a previous moment, following the definitions in the linear-point case; if, that is, we could express 'meets'. To simply assert that there is a next *interval*, which works for point orderings, is not enough here, as this is true even on the rational line. The definition given for the dense case does not work here, but a simpler alternative is available:

```
;;meets-ADP-discrete
(defrelation meets (?x ?y) :=
             (and (precedes ?x ?y)
                  (not (exists (?z) (precedes ?x ?z ?y))) ))
```

And now we can state discreteness by

```
;;ADP-discrete
(forall (?x)
        (and (exists (?y) (and (meets ?x ?y) (moment ?y)))
             (exists (?y) (and (meets ?y ?x) (moment ?y)))))
```

Models

Open intervals on the rational line **Q** or the real line **R** are probably the most intuitive models. The axioms do not distinguish between open and closed intervals, and can be interpreted in either way. For example, consider the universe of all open rational intervals and define `precedes` in the obvious way by (a,b) `precedes` (c,d) just when b is less than or equal to c, and `finer` as the weak subinterval relation. The same kind of nonstandard orderings ($\mathbf{Q} \uparrow \mathbf{Q}$, ie **Q** copies of **Q**, etc.) are also models here, and for the same reason. They now take on some extra style, since there are now intervals which span entire copies of the rational line. In the nonstandard models there are also special nonstandard intervals which do not have endpoints, such as the interval in $\mathbf{Q} \uparrow \mathbf{Q}$ consisting of all rationals <A,b> for some fixed A.

AP does not requre *all* set-theoretic intervals to be in the universe, however. For example, it is satisfied on the integers. Indeed, a model of AP can be made from any model of the point theories by selecting <u>any</u> infinite subset E of points, and letting the universe of intervals be all pairs of points <a,b> with $a,b \in E$ and a before b. Then two intervals meet just when they share an endpoint, and moments exist wherever E is not dense.

4.2 Intervals meeting in the glass continuum

Returning now to the idea of the continuum as consisting of intervals which meet *at* points, we assume only a single relation, `meets`, which is chained but not transitive. This theories can be regarded as axiomatisations of the intuition of the glass continuum.

First we specify that when two timeintervals meet defines a definite temporal location. This axiom 'ties together' the various ways a meetingplace of several timeintervals could be specified.

```
;;IM-syntax
(predicative timeinterval)

;;meets-IM-place
(forall (?i ?j ?k ?m)
        (=> (and (meets ?i ?k) (meets ?j ?k) (meets ?i ?m))
            (meets ?j ?m)  ))
```

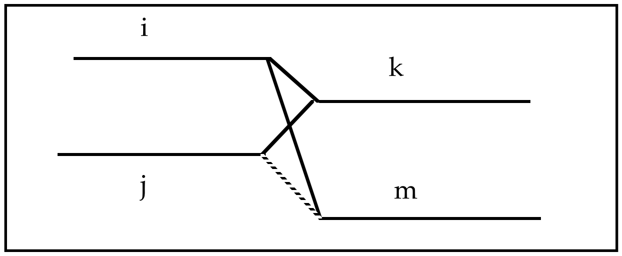

Time is infinite:

```
;;meets-IM-infinity
(forall (?i) (exists (?j ?k) (meets ?j ?i ?k)))
```

Until now this has been an essentially arbitrary assumption, but here it takes on extra importance since proofs often need those extra intervals to nail down the position of a meeting-point. Notice however that being infinite, as noted earlier, does not guarantee infinite duration unless further assumptions are made.

That time is ordered can be stated by following the intuition that a meetingplace is an endpoint, and insisting that these points be ordered:

```
;;meets-IM-total-order
(forall (?i ?j ?k ?m)
        (=> (and (meets ?i ?j) (meets ?k ?l))
            (or (meets ?i ?l)
                (exists (?n) (or (meets ?i ?n ?l)
                                 (meets ?k ?n ?j) ))
                )))
```

If the 'or' here were exclusive disjunction this would suffice, but we need to add the following to prevent time becoming circular:

```
;;meets-IM-line
(forall (?i ?j) (not (meets ?i ?j ?i)))
```

The transitivity of the timeorder corresponds to the assumption that two meeting intervals form a single larger interval:

```
(forall (?i ?j ?k ?m)(=> (meets ?i ?j ?k ?m)
                             (exists (?n)(meets ?i ?n ?m)) ))
```

It is often convenient to replace this by its skolem form, using an explicit addition function on intervals:

```
(deffunction plus
(forall (?i ?j ?k ?m) (<=> (meets ?i ?j ?k ?m)
                           (meets ?i (plus ?j ?k) ?m)
                           ))
)
```

(This is not a logical definition as it does not specify a value for `(plus a b)` when `a` and `b` do not meet.)

This provides a neat way to define the notion of a `moment`:

```
;;moment-IM
(defrelation moment (?i) :=
  (forall (?j ?k) (not (= ?i (plus ?j ?k)))
   ))
```

These axioms constitute a small but surprisingly powerful theory within which all other binary temporal relations on intervals can be defined.

Density and discreteness can be defined in obvious ways:

```
;;IM-dense
(forall (?i) (not (moment ?i)))

;;IM-discrete
(forall (?i) (and
              (exists (?j)
                      (and (meets ?i ?j) (moment ?j)))
              (exists (?j)
                      (and (meets ?j ?i) (moment ?j)))
              ))
```

Models

These axioms have models which reflect several of the intuitions which motivate the glass continuum, all of which can be regarded as in some sense “standard” models, although they are not mutually consistent. First consider the dense case.

One model interprets intervals as open connected subsets of the rational line **Q** (ie open intervals in the conventional mathematical sense of 'interval'), asserts that (a,b) `meets` (c,d) just when b=c, and defines `plus` as the interior of the set-theoretic union of the closures of the intervals being added together. Notice then that the sum of two intervals contains the point where they meet. In general, if more complex operations are defined, they are modelled by always taking the interior of the corresponding set-theoretic operation applied to the closures of the arguments. The intersection of two meeting intervals is always the interior of a single point, ie empty; so two meeting intervals are disjoint.

This model amounts to a response to the intuitive puzzle discussed earlier which says that the point of division is in *neither* half of the split interval, but just somehow vanishes when they are considered seperately.

Alternatively, a dual model is provided by closed connected subsets of **Q**, with subsets meeting just when they share an endpoint, and the sum of two intervals being defined to be the closure of the set-theoretic union of their interiors. This means that two intervals which meet share the meeting point, but their interval-overlap is the closure of the interior of a singleton, ie the empty set. Thus there is no shared interval to play the role of `?n` in the definition of `overlap`.

This answers the puzzle by saying that the point of division is in *both* halves, but still insists that the two halves have an empty intersection. (The point might be thought of as identifying the surfaces of the two pieces of glass filament, which were in the same place before the break.)

A third way to model this theory thinks of every point on the line as having two halves, one facing to the past and one to the future. Suppose we have two distinguished copies of **Q**, and write <a| for a member of one copy and |a> for the other: call these the past and future "halves" of the timepoint <a|a>. An interval denotes a pair <a|b> where a is less than b; <a|b> `meets` <b|c>. The meeting point <b|b> is not an interval and is therefore "invisible"to the theory. While this may seem unnecessarily baroque, it supports the common naive intuition, when faced with the midpoint puzzle, that the midpoint itself must somehow be divided into two equal parts.

Note that these are all different from the standard model for the point continuum, in which a closed interval can only `meet` an open interval. They all refuse to distinguish between two kinds of interval, for example.

These theories can also be interpreted as referring to the point continuum, however. For example, one model is provided by the set of semiclosed rational intervals (a,b] where (a,b] `meets` (b,c]. Here the meeting-point clearly belongs to the first interval, and a timeinterval is interpreted directly as a mathematical interval.

Nonstandard models analogous to the nonstandard models of the point ordering theories also exist, of course. In fact, any model of a point-ordering theory can be extended to a model of the corresponding (discrete or dense) interval-meeting theory by defining intervals, and the `meets` relation, in any one of the three ways just described.

Analogous models of the discrete theory can be constructed on the integers, but here simpler models are possible. One defines `meets` as the relation between [n,m] and [m+1,k]. Here integers label moments [n,n], and meeting-points have to be thought of as <u>between</u> integers. Another model consists of the integer intervals [n,m] with n < m and defines `meets` as the relation between [n,m] and [m,k]; here integers label points and a moment is an interval [n, n+1]. A nonstandard model here might be provided by **Q** copies of **N**, which, as Van Benthem points out, is "locally" discrete even though having a dense global structure.

Models of the basic theory can consist of moments inserted into other regions of density. Almost any combination is possible: in particular, nonstandard orderings such as **Q** followed by **N**, or **Q** with several copies of **N** inserted into it.

One class is of special interest, which I will call *point-moment* models. Consider a continuous interpretation such as **Q** or **R** with some selected subset M of isolated points (that is, for any p,q∈ M, there must be a point r with p < r < q), and define an interval as a pair <a,a> where a ∈ M, or <a,b> when a<b ; and `meets` be true of <a,b> and <b,b> (and of <b,b> and <b,c>) when b ∈ M, but of <a,b> and <b,c> otherwise. M can be any subset of isolated points ; it can be finite or infinite, and its members can be arbitrarily close to one another. Moments are identified by the members of M.

In such a model, a moment is being interpreted as a point, which is possible precisely because a moment has no internal structure. Point-moment models mix together some aspects of density and some of discreteness, so not surprisingly they are not models of either of the extended theories. (Not of density because moments exist, and not of discreteness because intervals with endpoints not in M – which must exist since M is isolated – do not satisfy the discreteness axiom.) These models are discussed further in section 5.3.1.

Constructing points in models of the intervals.

Since this theory is so firmly based on the concept of an interval, it may be surprising that any model of it contains things that can be regarded as timepoints. This is less trivial than it sounds, since in many models mathematical points (eg rational numbers) are not suitable interpretations of timepoints, especially in models based on the glass continuum. The interest of this result is that these theories are sometimes claimed to be 'point-free', in the sense that they are described purely in terms of intervals *sui generis*, rather than considering intervals as being defined by endpoints or consisting of sets of points. But such a claim has to taken with caution, since it is always *possible* to interpret this theory as talking about intervals constructed from points.

Suppose M is any model of these axioms with universe (of intervals) U, and consider pairs <B,A> (for 'Before' and 'After') of subsets of U with the property that every member of B `meets` every member of A; call such a pair a focus. The idea here is that a focus locates the meeting-point. Suppose a focus is *maximal* when the sets B, A are as large as possible, so that B contains every interval in U that meets a member of A, and vice versa. Then maximal foci (called 'nests' in Hayes & Allen 1987) can serve as points. The ordering relation `before` of section 3.1 can be defined thus: <B1,A1> is before <B2,A2> just when A1 ∩B2 is nonempty. The interval in the intersection is the interval between the points:

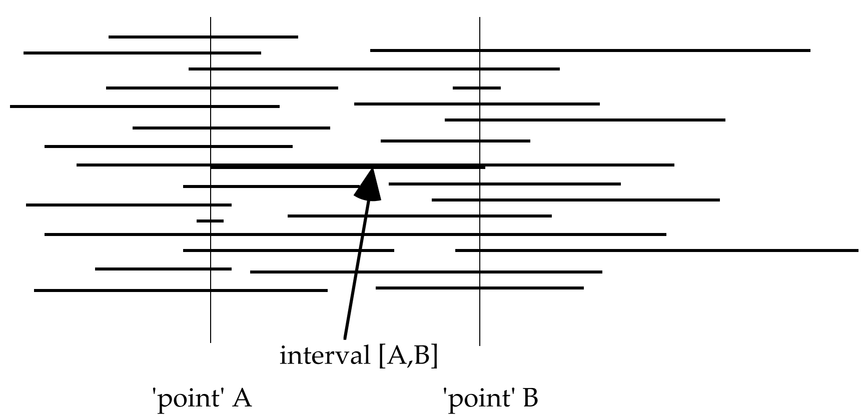

The function `beginof` from intervals to points is defined by: (beginof i) is the point <B,A> with i ∈A; similarly for endof but i ∈ B. That these are unique follows from the axioms and the assumption of maximality. It is now straightforward to show that with these interpretations, all the axioms of the linear point theory LP and all the interval-endpoint definitions given below are satisfied. Thus, any model of the interval-meeting axioms can be interpreted also a model of LP . This works equally well for "nonstandard" models.

This filter construction, originally due to A. N. Whitehead, may seem unintuitive: a point is very small, but these maximal foci are very large. But when one considers that the role of the focus is to isolate a point as precisely as possible, then the idea that it might take an infinite amount of information to isolate something infinitely small makes the filter seem somewhat more compelling.

4.3 Thirteen Relations

Probably the most familiar temporal theory is often described as being based on the thirteen relations described earlier. In this section I will use the compact notation used by Allen (1984) and extended by Freska(1992) where each relation is denoted by its intial letter , followed by '`i`' to indicate the inverse.

The thirteen relations are an exhaustive and mutually exclusive set:

```
;;TT-list
(forall (?x ?y)
    (=> (timeinterval ?x ?y)
        (xor (p ?x ?y)
             (o ?x ?y)
             (m ?x ?y)
             (s ?x ?y)
             (d ?x ?y)
             (f ?x ?y)
             (= ?x ?y)
             (fi ?x ?y)
             (di ?x ?y)
             (si ?x ?y)
             (mi ?x ?y)
             (oi ?x ?y)
             (pi ?x ?y) )
     ))
```

Notice the use of exclusive-or here. This axiom can be written rather more tediously without it.

```
;;TT-inverses
(forall (?x ?y)
        (and (<=> (p ?x ?y) (pi ?y ?x))
             (<=> (m ?x ?y) (mi ?y ?x))
             (<=> (o ?x ?y) (oi ?y ?x))
             (<=> (s ?x ?y) (si ?y ?x))
             (<=> (d ?x ?y) (di ?y ?x))
             (<=> (f ?x ?y) (fi ?y ?x)) )
        )
```

The various relations between these are often, following Allen, summed up in a transitivity table. This 13 x 13 array gives the set of possible relations that can hold between x and z when one relation holds between x and y and the second between y and z. We summarise the table here as a variadic relation `ttable` on relations. The first two are the table coordinates, the rest are the possible entries at that point. This is expressed in the following second-order assertions, written here in pseudo-KIF notation:

```
4.3.3   (forall (?r ?s ?e @l)(<=>
              (ttable ?r ?s ?e @l)
              (forall (?x ?y ?z)(=>(and (?r ?x ?y)(?s ?y ?z))
                                (disjoin ?x ?z ?e @l) )) ))

4.3.4   (defrelation disjoin (?x ?z ?e @l) :=
                    (or   (?e ?x ?y)
                          (disjoin ?x ?y @l))
        )
```

Unlike the Dedekind axiom discussed earlier, however, these are not essentially second-order quantifiers. These 'second-order' variables are really only schematic placeholders for first-order relations: indeed, for our purposes, they need only refer to the twenty-six realtions in our current list. These axioms can therefore be quite reasonably transcribed into KIF by using the 'holds' relation between a relation and its arguments:

```
;;table-TT
(defrelation ttable (?r ?s ?e @l) :=
         (forall (?x ?y ?z) (=> (and (holds ?r ?x ?y)
                                     (holds ?s ?y ?z))
                                (disjoin ?x ?z ?e @l) )) ))

;;disjoin-TT
(defrelation disjoin (?x ?z ?e @l) :=
             (or (holds ?e ?x ?y)
                 (disjoin ?x ?y @l))))
```

The table is then summarised in the following rather long conjunction:

```
(and  (ttable p p p)
      (ttable p m p)
      (ttable p o p)
      (ttable p fi p)
      (ttable p di p)
      (ttable p si p)
      (ttable p s p)
      (ttable m p p)
      (ttable m m p)
      (ttable m o p)
      (ttable m fi p)
      (ttable m di p)
      (ttable o p p)
      (ttable o m p)
      (ttable fi p p)
      (ttable p d p m o s d)
      (ttable p d p m o s d)
      (ttable p f p m o s d)
      (ttable p oi p m o s d)
      (ttable p mi p m o s d)
      (ttable m si m)
      (ttable m s m)
      (ttable m d o s d)
      (ttable m f o s d)
```

```
(ttable m oi o s d)
(ttable o d o s d)
(ttable o d o s d)
(ttable fi d o s d)
(ttable m mi fi = f)
(ttable o p p)
(ttable o m p)
(ttable o o p m o)
(ttable o fi p m o)
(ttable o di p o m fi di)
(ttable o si m fi di)
(ttable o s o)
(ttable fi s o)
(ttable o oi o fi di si = s d f oi)
(ttable fi m m)
(ttable fi o o)
(ttable fi di di)
(ttable fi si di)
(ttable di fi di)
(ttable di di di)
(ttable di si di)
(ttable si fi di)
(ttable si di di)
(ttable fi f f = fi)
(ttable di p p m o fi di)
(ttable si p p m o fi di)
(ttable si s si = s)
(ttable s p p)
(ttable s m p)
(ttable d p p)
(ttable d m p)
(ttable f p p)
(ttable s o p m o)
(ttable s fi p m o)
(ttable s di p m o fi di)
(ttable s si si = s)
(ttable d o p m o s d)
(ttable d fi p m o s d)
(ttable f m m)
(ttable f o o s d)
(ttable f fi f = fi)
(ttable oi p p m o fi di)
(ttable mi p p m o fi di)
(ttable oi m o fi di)
(ttable oi o o fi di si = s d f oi)
(ttable mi m si = s)
)
```

(This ordering tries to illustrate the natural grouping of entries in the table. Some entries are omitted: in particular, those involving equality follow by ordinary logical principles; some entries provide no information; and the entire table is symmetric so only half of it is necessary.) More discussion of the topological structure of the table can be found in (Freska 1992), who develops a very intuitive iconic notation.

The table can be thought of as defining an algebra on subsets of the set AR = {`p,m,o,s,d,f,fi,di,si,oi,mi,pi,=`} of thirteen relation *names*,. Temporal reasoning can then be performed by computing multiplication in this algebra, where an empty string indicates a contradiction. This has been a very influential technique for interval reasoning. Vilain and Kautz (1986) have shown that computing closure in this algebra is an NP-complete problem in general, although useful polynomial time algorithms have been developed for special cases (Vilain, Kautz & Van Beek 1989)

This algebraic approach to interval reasoning is rather different in spirit than the simple "reduction to endpoint" definitions given below. For example, consider the relationship between two intervals `?i` and `?j` which is that the first begins before the second ends, ie `(before (beginof ?i)(endof ?j))` . To express this in the algebraic style one must consider which of the thirteen relationships is consistent with this arrangement and list them all explicitly. The answer is `{p, m, o, s, d, =, fi, si, fi, di, oi}`, ie all eleven of these are consistent with that constraint on the endpoints. This awkwardness of the disjunctive-list notation is due to the fact that each particular interval relation commits itself to the relative positions of both endpoints, so a lack of commitment can only be expressed by a disjunction.

Several entries recur in the table several times: the sequences `(p m o fi di)` and `(o s d f = fi di si oi)` for example. Freska notes that several of these have natural interpretations as weaker constraints on interval endpoints. For example, `(p m o fi di)` is the relation of beginning earlier than, with no constraint on the relations between the ends of the intervals; while the second disjunction is the relation `ncdf` in section 4.1. The thirteen relations can be defined from these by conjunction rather than disjunction. (For example, if we also define *exact* to mean that the two intervals' endpoints are somewhere exactly aligned (this is `(m s f = si fi mi)`), then meets is *beginsearlierthan* and *not ncdf* and *exact*.) This suggests the possibility of finding such a set of weaker relations which is also closed under transitivity, which would greatly simplify the task of computing transitive closure.

One might have hoped for a simpler summary of the relations between such an intuitively acceptable set of temporal relationships. They can be appropriately defined in simpler theories. A theory needs to provide enough axioms to conclude that these relations are an exhaustive and exclusive set of interval relations, and to derive the transitivity table. We give two ways of doing this.

4.4 Thirteen relations in terms of meets

These definitions use one basic formal trick, which is to force two intervals to have a common endpoint by hypothesising that a third interval exists which meets them both, or which they both meet. This is rather like the familiar juggling feat of holding boxes in the air by clamping them between two boxes held in the hands.

The intuitive content of these definitions is probably best shown by drawing the linear patterns of intervals named in the quantifiers. `Overlaps` is the most complicated case.

```
;;precedes-IM
(defrelation precedes (?i ?j) :=
              (exists (?k) (meets ?i ?k ?j)))

;;overlaps-IM
(defrelation overlaps (?i ?j) :=
                   (exists (?k ?m ?n ?o ?p)
                      (and (meets ?k ?m ?n ?o ?p)
                           (meets ?m ?j ?p)
                           (meets ?k ?i ?o) )))
```

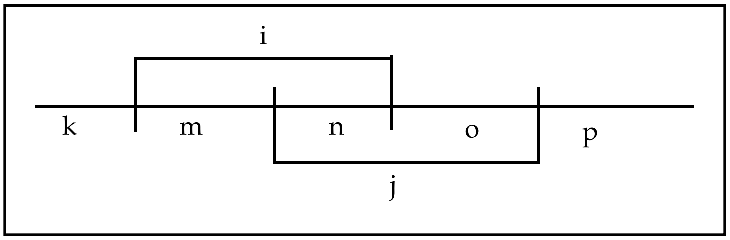


```
;;starts-IM
(defrelation starts (?i ?j) :=
            (exists (?k ?m ?n)
                    (and (meets ?k ?i ?m ?n)
                         (meets ?k ?j ?n) )))

;;during-IM
(defrelation during (?i ?j) :=
            (exists (?k ?m ?n ?o)
                    (and (meets ?k ?m ?i ?n ?o)
                         (meets ?k ?j ?o) )))



;;finishes-IM
(defrelation finishes ( ?i ?j) :=
            (exists (?k ?m ?n)
                    (and (meets ?k ?m ?i ?n)
                         (meets ?k ?j ?n) )))
```

Allen & Hayes (1987) show how the entire 13 x 13 transitivity table of these relations and their inverses can be derived within the theory of ‘meeting’ described earlier plus these definitions.

5. Points and Intervals

In this section we consider putting together the ideas of point and interval into combined theories.

5.1 Thirteen relations in terms of endpoints

The thirteen relations can be defined directly, following their intuitive meanings, by using the point ordering relation `before` applied to the endpoints of the intervals. To do this we introduce two functions `beginof` and `endof` from intervals to points. The inverse function `between` is also useful. These axioms and definitions, added to the theory LINEAR-POINT, constitute a theory we will call ENDPOINTS:

```
;;beginof-endof-EP
(forall (?i)(=> (timeinterval ?i)
                       (and (timepoint (beginof ?i) (endof ?i))
                            (before (beginof ?i) (endof ?i))
                       )
        ))

;;between-EP
(forall (?p ?q)(<=> (before ?p ?q)
                   (and (= ?p (beginof (between ?p ?q)))
                        (= ?q (endof (between ?p ?q))))))

;;precedes-EP
(defrelation precedes (?i ?j):=
             (and (timeinterval ?i ?j)
                  (before (endof ?i) (beginof ?j))))

;;overlaps-EP
(defrelation overlaps (?i ?j):=
             (and (timeinterval ?i ?j)
                  (before (beginof ?i) (beginof ?j) (endof i))))

;;starts-EP
(defrelation starts (?i ?j):=
         (and (timeinterval ?i ?j)
              (= (beginof ?i) (beginof ?j))
              (before (endof ?i) (endof ?j)) ))
;;
;;during-EP
(defrelation during (?i ?j):=
          (and (timeinterval ?i ?j)
               (before (beginof ?j)
                       (beginof ?i)
                       (endof ?i)
                       (endof ?j)) ))

;;finishes-EP
(defrelation finishes (?i ?j):=
         (and (timeinterval ?i ?j)
              (before (beginof ?j) (beginof ?i))
              (= (endof ?i) (endof ?j)) ))
```

The entries in the transitivity table can be straightforwardly derived within ENDPOINTS. However, we can also take these axioms, add them to the theory INTERVAL-MEETING, and regard `between-EP` as a definition of `before`; and then all the axioms of LINEAR-POINT are derivable. So this theory provides a kind of bridge between the point-ordering view of time and the interval-meeting view, from the glass continuum perspective.

Notice that `beginof-endof-EP` means that there cannot be a single-point interval in this theory. If we allowed single-point intervals then the transitivity table would no longer be derivable. For example, it would then be possible to have three intervals such that I meets J meets K and I meets K. Moreover, one interval could both meet and start another. Later we will consider a variation in which this restriction is removed, and which reconciles these apparently unintuitive consequences.

<u>Models</u>

> All models of LP , including the nonstandard ones, extend immediately to models of EP by interpreting intervals in any of the three ways mentioned in section 4.2 above. These axioms fit most naturally into the glass continuum, since they make no distinction between open and closed timeintervals; but as noted there, it is possible to interpret timeintervals as half-closed intervals in the point continuum.
>
> Some care is necessary in considering models over the rationals. We cannot take the universe to be *all* rational intervals, since there are rational intervals with no rational endpoints. The universe of models of EP must be made up from intervals with rational endpoints. Perhaps not surprisingly, therefore, when we insist that intervals are attached to their ends, we can only have countably many of them.

<u>5.2 Open and closed intervals: the point continuum</u>

Although the thirteen relations can be interpreted in the point continuum, difficulties arise if we combine the mathematical intuitions from real analysis with these axioms. For example, Galton (1990) argues that Allen's axiom for 'property negation', when applied to a smoothly moving body, seem to imply that it is never at <u>any</u> position (since it is never at any position for a whole subinterval, `(pnot (at ?x))` is always true throughout any interval). This argument however assumes that "interval"means mathematical interval, and indeed Galton goes on to develop a variation on Allen's theory which is as directly interpretable in the point continuum as Allen's is in the glass continuum. We give a variation here which is equally expressive but somewhat less complex.

Galton distinguishes two kinds of proposition: a *state of position* and a *state of motion*. Paradigmatic examples are respectively, a ball being at rest and a ball being in motion. Galton however is careful to distinguish being at rest from having zero velocity. For example, a ball tossed in the air has zero velocity at the single timepoint when it reaches the top of its parabola, but it is not *at rest* there; this is a state of motion but not of position.

Galton gives many axioms relating states of position and motion, using a complicated variation of the 'holds' notation discussed in section 2 earlier. However, since states of position are those things that can hold true during a closed interval, while states of motion must hold during open intervals, the two classes of proposition must be exclusive, and to distinguish them it is sufficient to distinguish the kinds of interval during which they hold. If we further think of an interval as consisting of the points it contains, we can easily define an equivalent language which has as primitive just the notion of `holds at` a point. As described earlier, with this simple interpretation 'holds' can be transparently removed from the language, and such axioms as

```
**     (<=> (holds-at (not ?p) ?i) (not (holds-at ?p ?i)) )
```

become tautologous.

The distinction between states of motion and position can then be reinterpreted as differences between different kinds of relation. Some can only be relativised to open intervals, others only to closed. For example, the property of *moving* is taken to be a relation between an object and a closed interval, while *resting* becomes a relation to an open interval. When something starts to move, the two intervals of its resting and moving `meet`, but the latter contains the meeting point, where the motion has begun but the velocity is zero. For the old example of the light, either one or the other of *being lit* and *not being lit* has to be a state of motion; or else, perhaps more plausibly, they can both be states of rest, but then the single meetingpoint is where a special state of motion predicate – <u>goingout</u> or <u>comingon</u> , or perhaps simply <u>changingstate</u> – is true.

There are several ways to obtain a theory for open and closed intervals. One method is to define intervals as sets of points, and use set-theoretic comprehension principles to establish that appropriate intervals exist. We do not develop this standard mathematical approach here, but give selfcontained axiom systems which can be connected to set theory later, if required. There are in any case some complications in thinking of intervals literally as sets of points.

The most straightforward way is probably to talk of endpoints, as in the theory EP. I will use the relation `in` between a point and a containing interval (this is not the same as `during`) and `open` and `closed` as predicative relations on intervals.

```
;;PC-in-syntax
(forall (?x ?y)(=> (in ?x ?y)
                   (and (timepoint ?x)
                        (timeinterval ?y)) ))

;;open-close-in-PC
(forall (?i)(=> (timeinterval ?i)
        (or (and (open ?i)
                 (not (closed ?i))
                 (forall (?p)(<=> (in ?p ?i)
                                  (before (beginof ?i) ?p (endof ?i))
                 ))  )
            (and (closed ?i)
                 (not (open ?i))
                 (forall (?p)(<=> (in ?p ?i)
                                  (bbefore (beginof ?i) ?p (endof ?i))
                 ))  ))))
```

So far nothing establishes that intervals actually exist, but we can assert this directly.

```
;;begin-end-PC-1
```

```
(forall (?p ?q)(=> (timepoint ?p ?q)
        (or  (before ?q ?p)
             (exists (?i)
                     (and (timeinterval ?i)
                          (closed ?i)
                          (= (beginof ?i) ?p)
                          (= (endof ?i) ?q)
                     )
              ))
        )
)
```

The rather awkward disjunction here is to allow the case where ?p and ?q are the same point. The similar axiom for open intervals reads more naturally:

```
;;begin-end-PC-2
(forall (?p ?q)(=> (timepoint ?p ?q)
        (<=>  (before ?p ?q)
              (exists (?i)
                      (and (timeinterval ?i)
                           (open ?i)
                           (= (beginof ?i) ?p)
                           (= (endof ?i) ?q)
                       )
              ))
        )
)
```

The skolem form of this axiom provides the useful function `between` from a pair of points to the open interval between them.

```
;;between-PC
(defunction between
(forall (?p ?q)(=> (timepoint ?p ?q)
        (<=> (bbefore ?p ?q)
             (and (timeinterval (between ?p ?q))
                  (= ?p (beginof (between ?p ?q))
                  (= ?q (endof (between ?p ?q))  ))))))
)
```

The biconditional here makes the very strong assumption that every pair of points defines an interval and vice versa. Weaker “bindings” between interval and points could also be considered, for example by replacing this by a simple implication, or restricting the quantifier to a special subset of ‘endpoints’. For example, one might want to consider a dense theory of points but only allow a discrete universe of intervals, so that some points had no interval between them. I will not explore such ideas in detail, however.

Every open interval has a closed interval with the same endpoints. The skolem form of this statement provides the function `closure`. This axiom would be trivial in the glass continuum, but needs to be explicitly stated here:

```
;;closure-PC
(deffunction closure
(forall (?i)(=> (timeinterval ?i)
        (and  (timeinterval (closure ?i))
```

```
              (closed (closure ?i))
              (= (beginof ?i) (beginof (closure ?i)))
              (= (endof ?i) (endof (closure ?i)))
        )))

)
```

It follows that an interval is closed just when it is equal to its own closure, as expected.

A moment is an interval which is as short as possible, ie has no points inside it :

```
;;moment-PC
(defrelation moment (?i) :=
       (forall (?p)(<=> (in ?p ?i)(or (= ?p (beginof ?i))
                                      (= ?p (endof ?i)) ))))
```

The definitions of the thirteen relations used in 4.5 above now need to be reconsidered. (As they were inspired by the glass continuum, this is perhaps not surprising.) Only a closed interval can meet an open one, and vice versa; and a subinterval can start or finish only an interval of the same kind. The concept of being the same kind as, ie <u>also closed or also open</u>, or `acoao`, will be useful here:

```
;;acoao-PC
(defrelation acoao (?i ?j)(=> (timeinterval ?i ?j)
             (or  (and (open ?i) (open ?j))
                  (and (closed ?i) (closed ?j))
                  )))
```

Then the three cases where endpoints are supposed to be precisely aligned can be rewritten thus:

```
;;meets-PC
(defrelation meets (?i ?j) :=
          (=> (timeinterval ?i ?j)
              (and (not (acoao ?i ?j))
                  (= (endof ?i) (beginof ?j)) )))

;;starts-PC
(defrelation starts (?i ?j) :=
         (=> (timeinterval ?i ?j)
             (and  (acoao ?i ?j)
                   (= (beginof ?i) (beginof ?j))
                   (before (endof ?i) (endof ?j))
                   )))

;;finishes-PC
(defrelation finishes (?i ?j) :=
        (=> (timeinterval ?i ?j)
            (and  (acoao ?i ?j)
                  (before (beginof ?j) (beginof ?i))
                  (= (endof ?i) (endof ?j)) )))
```

The other Allen relations are defined just as in ENDPOINT.

If these axioms are added to the basic point-order theory LINEAR-POINT, the resulting theory is a sketch of the usual mathematical account of the rational line or the real continuum.

It does not have continuity assumptions, however, since these are essentially not first-order. Whether they are needed is an interesting, and ultimately empirircal question. If this theory is being used in a scientific or engineering domain then certain *consequences* of continuity are certainly useful, and may be essential: notably, for example, the ijntermediate-value theorem and the existence of solutions to certain classes of equations. However, these can often be stated directly. In any case, I do not here consider this issue (which goes beyond the narrow question of temporal effectiveness) in detail.

Two open intervals cannot meet, so if two open intervals `(between a b)`and `(between b c)` share an endpoint, then there is a single-point closed interval just separating them, so that `(between a b) meets [b,b] meets (between b c).`These single-point intervals are the moments.

Moments are closed, and every closed single-point interval is a moment, as we would expect. In stark contrast to the axioms in 5.1, these single-point moments are everywhere: by `between-end-PC-1`, <u>every</u> point occupies one. The biconditional in `between-end-PC-2` ensures that a single-point open interval is impossible.

The chief utility of moments until now has been to state a discreteness condition. But because a moment now can consist of nothing but a point, axioms which assert that intervals have adjacent moments will be true even in the real line, thus failing completely to guarantee an underlying discrete structure of seperate time-ticks. We cannot insist that a moment `meet` another moment since this is provably impossible; being both closed, their meeting is excluded by the acoao condition. The correct way to state density or discreteness is in terms of points rather than intervals, since the intervals in this theory are defined by the locations of their endpoints. We can simply add the appropriate definitions in the extended theories DENSE- or DISCRETE- LINEAR-POINT.

<u>Models</u>

> These axioms have a curious consequence in a discrete model of time, where there are points `p,q` ordered by `before` but with *no* points between them; therefore there are open timeintervals `(between p q)` which contain no timepoints at all. The discrete timeline consists of these empty open intervals interleaved with single-point moments containing the timepoints. This rather strange (although consistent) picture arises from combining two rather different intuitions: the universal quantifier in 5.2.6 is suggested by a vision of smooth continuity, which the discreteness axiom explicitly denies.
>
> This would be a contradiction in a set-theoretic model, where these would all be the empty set and hence be identical, forcing time to be circular in a particularly pathological fashion. However, a set-theoretic development would not need these axioms; and we do not need to identify these empty moments, since they are distinguishable by their endpoints (which, being open, they do not contain). For example, one discrete model of these axioms is provided by the integers where `timepoint` is true of the even integers, a `timeinterval` is a connected sequence of integers which is `open` when its ends are odd and `closed` when they are even, and `meets` is simply adjacency. The 'empty' open intervals are then the odd integers.
>
> The density axiom produces more intuitive results, and this extension of the theory has the usual models. With the density assumption, a moment must contain only a single point which is both its end and its beginning.

Galton's distinction between states of motion and rest is now simply the distinction between the kind of interval they can be said to hold during. They both `hold-in` an interval just when they `hold at` all the points `in` the interval; but if the interval is closed, then they are a state of motion, and if open, a state of rest. Things that `hold-in` single-point moments (as opposed to holding at a point) must therefore be states of motion.

5.3 The Vector Continuum: placing points in glass

The different intuitions about the continuum give rise to very different axioms. However, it is possible to combine some of the most useful features of both of them into a single consistent framework.

First, we simply add the notion of a point to the description of the glass continuum given in 4.2, ie the theory INTERVAL-MEETING. To preserve the intuition, however, the basic relation between points and intervals is not that of containment – the `in` relation of 5.1 – but that of a *meeting-point*: points are the places where intervals meet. We will express this by a three-way relation `meets-at` between two meeting intervals and a point.

The most natural relation between interval and point is expressed here by the relation meets-at which is true when a point is the meeting-place of two intervals:

```
;;VC-syntax-meets-at
(forall (?x ?y ?z)(=>(meets-at ?x ?y ?z)
                     (and (timeinterval ?x ?z) (timepoint ?y))))

;;meets-at-VC
(defrelation meets-at (?i ?p ?j) :=
               (= ?p (endof ?i) (beginof ?j)))
```

It follows therefore that `(meets ?i ?j)` if and only if `(exists (?p)(meets ?i ?p ?j))`.

If we now define the timepoint ordering :

```
 (defrelation before (?p ?q) :=
          (exists (?i ?j ?k)
               (and (meets-at ?i ?p ?j)(meets-at ?j ?q ?k)))  )
```

then its properties follow from the axioms of INTERVAL-MEETING: that is, LP can be derived within IM+these three axioms. (Notice again the importance of the infinity assumption , which is now necessary to establish that every interval has endpoints.) This theory, like ENDPOINTS, therefore provides a bridge between the point and interval ways of conceptualising the glass continuum. To emphasise the point; if we define `beginof` and `endof` correctly, the definitions of the thirteen relations in the style of the interval theory IM, or that of the endpoint theory EP become provably equivalent:

```
(defrelation beginof (?i) := (exists (?j)
                   (meets-at ?j (beginof ?i) ?i)))

(defrelation endof (?i) := (exists (?j)
                   (meets-at ?i (endof ?i) ?j)))
```

The resulting theory is still openminded about density or discreteness, and can be pushed in either direction by adding suitable extension axioms, eg by extending LP or IM.

A more interesting alternative, however, allows the idea of intervals with a direction. These will be useful when talking of durations, since it is natural there to think of a negative duration as a 'debt' of time; and it results in a very elegant simplification of many of the axioms, overcoming some artificial restrictions which have been necessary in order to make the axioms have the needed conclusions.

Consider the axiom `beginof-endof-EP`:

```
(forall (?i)(=> (timeinterval ?i)
                       (and (timepoint (beginof ?i) (endof ?i))
                            (before (beginof ?i) (endof ?i))
                       )
        ))
```

and consider the effect of weakening it to remove the restriction that the beginning of an interval is `before` its end.:

```
(forall (?i)(=> (timeinterval ?i)
                (timepoint (beginof ?i) (endof ?i))
        ))
```

This allows intervals with are 'pointed backwards' as well as interval with idenical beginning and endings, ie single-point intervals. The lack of single-point intervals is a distinct weakness of the 'glass continuum' theories such as ENDPOINTS, compared to the POINT-CONTINUUM, but this suggests an alternative way to incorporate them.

```
;;moment-VC
(forall (?x) (<=> (moment ?x)
                  (= (beginof ?x)(endof ?x))))
```

Here then, an interval is defined by *any* two points. If the beginning is `before` the end, then we will say that the interval is `forward`, and given any interval we will have a function `back` which 'reflects' it.

```
;;forwards-VC
(defrelation forwards (?i) := (before (beginof ?i)(endof ?i)))
```

```
;;back-VC
(deffunction back (?i) = (between (endof ?i)(beginof ?i)))
```

That `(back (back ?x)) = ?x` then follows from the definition of `between` in `ENDPOINTS.`

Actually this 'mirror' analogy is somewhat misleading, since a backward interval should not be thought of as *going* backwards in time - time itself does not move, of course - but rather as a debt of time or an amount owing. For example, a sequence of meeting intervals represents a longer interval which is their sum; to include a backward interval in such a sequence would simply be to *diminish* the total time it represented:

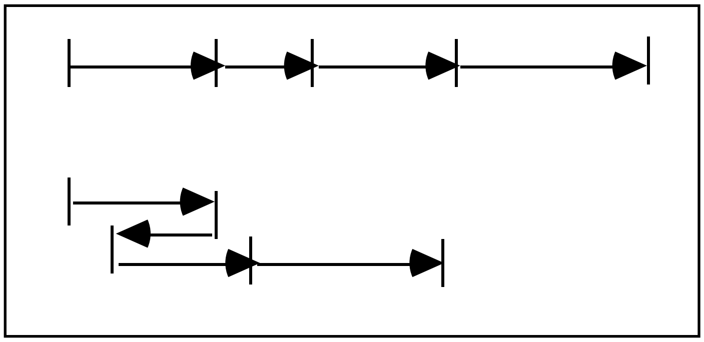

This is rather like observing that some event is earlier than expected.

Similarly, a single-point interval - a moment - added to another interval should not affect its duration at all. There seems now, in fact, to be little distinguish a moment from a timepoint, other than our insistence that the two categories are disjoint. The motivation for such an insistence comes from the idea that intervals, but not points, are the things during which propositions are true, or when events happen. This intuition was fundamental to Allen's development and has motivated others; and it was precisely this (considered more as a bug than a feature) which led Galton to explicitly deny the glass continuum intuition and return to the point continuum. There seems to be little more than apoint when the ball is motionless at the top of its parabola, for example.

This all suggests that we take the unusual step of allowing a moment to be both an interval and a point. The categories are disjoint everwhere else, but here they overlap:

```
;;VC-moment
(forall (?x) (<=> (moment ?x)
                  (and (timepoint ?x)
                     (timeinterval ?x))))
```

Moments now have some interesting properties. Following the axioms in `INTERVAL-MEETING` and the definitions in `LINEAR-POINT`, it is easy to show that a moment meets itself: indeed, that it meets itself at itself. This seems extremely unintuitive until we observe that any interval meets its own back-reflection. We can think of a moment not as the limiting case of two copies of itself placed one before the other - which is clearly impossible - but rather as the limiting case of an interval meeting its own reflection:

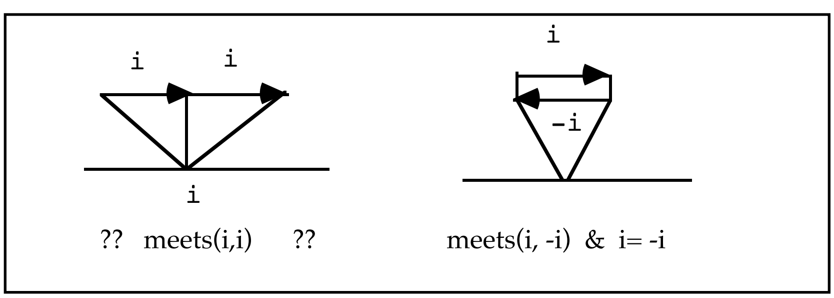

It now also follows that if two (nonmomentary) intervals meet at a moment, then they also meet. This means that a moment can be placed 'between' two meeting intervals without blocking their meeting, in marked contrast to the situation in the point continuum.

In many ways this seems more natural. It allows us to describe a pattern of meeting intervals without being overly concerned about whether or not the meeting-places are worthy of mention. Deciding that a meetingplace is a moment (and hence that something can be true there but nowhere else, as in the tossed ball) would not require massive updating of a set of assertions, rewriting all the 'meets' to be 'precedes', re-evaluationg the open/closed status of intervals, and so forth, but can be transparently added, and is inconsistent only with an explicit denial of its status as momentary..

This provides the most satisfying framework for describing the glass continuum. A moment now is a both a timepoint and a timeinterval. Since it meets other intervals, it can also take part in the other interval relationships defined in INTERVAL-MEETING. It is (when considered as a point) both the `beginof` and `endof` itself (when considered as an interval). In some ways therefore it acts similarly to the one-point intervals in the point continuum. For example, we could distinguish the interval `between` two such point-moments from the result of adding them to its ends, and regard this addition as something like the operation of closure in the point continuum. But many intervals need not end in such moments, so 'closure' might be rarely possible: and in any case, these 'closed' intervals do not behave any differently than their 'interiors' in how they relate to their neighbors. Points here are still not as <u>substantial</u> as they are in the point continuum, even when given the status of being intervals.

This also fits quite well with the intuition of an interval as giving information about the exact time of a timepoint. Since a moment is a point, it has no other points within it, hence it is an interval which identifies a point as precisely as possible. It represents the limits of our abilities to measure exactly when something happens. In the dense case, this would seem naturally to be the point itself.

5.3.1 Approximate-meeting

Another way to motivate this way of describing time comes from the following essentially model-theoretic argument. Return for the present to the older idea of moment as a shortest possible interval, but *not* a point, as described in the theory INTERVAL-MEETING. Let us say that two intervals *approximately meet* if they are separated only by a moment:

```
;;A
(defrelation ameets (?i ?j):=
                    (exists (?k)
                          (and (moment ?k)(meets ?i ?k ?j)) ) )
```

and then allow this as one way of meeting:

```
;;B
(defrelation mmeets (?i ?j):=
                    (or (ameets ?i ?j)(meets ?i ?j)) )
```

(The name should suggest a momentary hesitation) Now, if we assume that two moments cannot meet:

```
;;C
(forall (?i ?j) (=>
            (and (moment ?i)(moment ?j))
                  (not (meets ?i ?j)) ))
)
```

then `mmeets` satisfies the axioms of INTERVAL-MEETING (IM) . That is, if we rewrite the IM axioms with `mmeets` in place of `meets`, the resulting theory – call it AIM – is deducible within IM + A,B,C. (Hayes & Allen 1991)

(The reason for the no-meeting requirement can be seen by considering the 'place' axiom with `ameets` instead of `meets`. The assumptions allow this to occur:

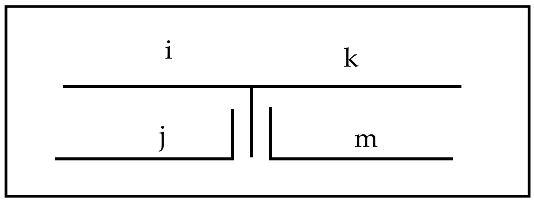

But now the intervals j and m do not even approximately meet, since there are two moments between them. If tiny errors are allowed to add up, they become large enough to notice.)

This means that `meets` could consistently have been interpreted this way all along (assuming still that moments cannot meet). Any model of IM + A,B,C is automatically a model for AIM; and any model of AIM is also one of IM where `meets` is sometimes interpreted as `ameets`. Therefore we can merge the relations `meets` and `ameets` into a single relation without confusion, retaining all truths. If we call this combined relation `meets`, then this amounts to saying that some meetings can be at moments without violating the theory. The point-moment models described in section 4.2, in which a subset of points in the model were singled out as standing for moments, were indeed such interpretations, with the 'isolation' of the moment-points corresponding exactly to the no-meeting axiom C. The effect of this merging can be described as identifying moments with points.

It would be nice if we could simply conjoin all the axioms in INTERVAL-MEETING and ENDPOINTS, but it is now necessary to adapt some of the earlier axioms slightly, since point-moments will no longer serve to establish a space between meetingpoints:

```
;;beginof-endof-VC (generalises EP)
(forall (?i)(=> (timeinterval ?i)
                (and (timepoint (beginof ?i) (endof ?i)))))

;;between-VC (generalises EP)
(forall (?p ?q) (and (= ?p (beginof (between ?p ?q)))
                     (= ?q (endof (between ?p ?q))))))

;;plus-VC
(deffunction plus
(forall (?i ?j)(=>(meets ?i ?j)
                  (= (plus ?i ?j)
                     (between (beginof ?i)(endof ?j)))))
)
```

The axioms of INTERVAL-MEETING adapt quite well to the generalisation which allows backwards and pointlike intervals. For example, the ‘place’ axiom is still true even when some of the intervals mentioned are backwards:

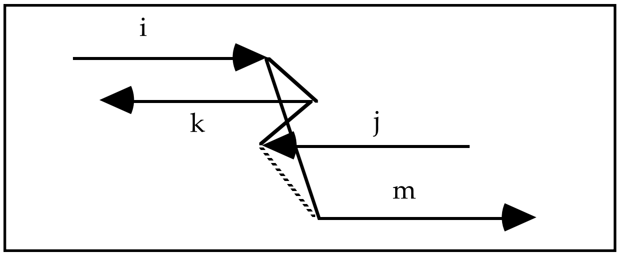

The definitions of the Allen relations given in ENDPOINTS work here perfectly well. The definitions in INTERVAL-MEETING, however, need to be restricted to forward intervals, or they fail to make the relevant distinctions. Here, the mere existence of an interval is not sufficient to ensure the past-to-future ordering of its endpoints.

6. Durations

A duration is a property of an interval, or equivalently of a pair of points defining the ends of the interval. Since the duration of a point (including a point-moment) is zero, we can here be somewhat more careless about the distinctions between different views of the continuum, and between open and closed intervals. Any theory of duration ought to apply to both views equally well; and, except where noted, these axioms can be added to any of the earlier theories. The conclusions can often only be reached in one of the more comprehensive theories, however. The distinction between dense and discrete time is often important since in dense time, a moment has no duration, but in discrete time it must have some.

6.1 Basic properties of durations

It would be acceptable to assume immediately that durations were, say, real numbers; but in the spirit of the earlier sections I will develop the theory with the minimal assumptions necessary. Whatever durations are, some things seem clear. Durations can be compared; durations can be added together (since the duration of two meeting intervals is the sum of their durations); there is a zero duration (which is the duration of a point); and finally, clocks measure duration by counting , so it must be possible to multiply durations by integers. These three basic assumptions are embodied in the constant `zero`, the chained transitive relation `less`, and the functions `add` and `times` used in these axioms. The type predicate is `timeduration`, and `duration` is a function from intervals to timedurations.

```
;;DU-syntax
(and (predicative timeduration)
     (forall (?x ?y)(and
                (=> (exist (?z) (= ?z (mult ?x ?y)))
                    (and (integer ?x) (timeduration ?y)))
                (=> (exist (?z) (= ?z (add ?x ?y)))
                    (timeduration ?x ?y))
                (timeduration zero)  ))


;;mult-DU
(deffunction mult
(forall (?d ?n ?m)
        (and
         (= (mult 0 ?d) zero)
         (= (mult 1 ?d) ?d)
         (= (mult (+ ?n ?m) ?d))
            (add (mult ?n ?d) (mult ?m ?d)))
        ))
)

;;add-DU
(deffunction add
(forall (?d ?e @f)
        (and (= (add zero ?d) ?d)
             (= (add ?d ?e) (add ?e ?d))
             (= (add ?d (add ?e @f)) (add (add ?d ?e) @f))))
)
```

Several different functions satisfy this axiom. Since in general there is no notion of a unit duration, addition cannot be defined recursively in terms of a successor function.

```
;;less-DU
(forall (?d ?e)(<=> (positive ?d)(less ?e (add ?d ?e))))
```

The relation positive cannot be defined purely in the duration theory, but requires the notion of a forward interval, which in turn is defined in terms of the basic before relation on timepoints. Without reference to before, times,and hence timedurations, are completely symmetric with regard to the direction of time.

```
;;numberin-DU
(deffunction numberin
(forall  (?d ?e)
         (=>
          (exist (?n)(and (integer ?n) (= (mult ?n ?d) ?e)))
          (= (mult (numberin ?d ?e) ?d) ?e)))
)
```

`Numberin` is a useful function when we can be sure that copies of one duration fit exactly the another.

6.2 Intervals, points and durations

The most basic facts about the duration of intervals is that they add up properly:

```
;;duration-DU-plusadd
(forall (?i ?j)
        (=> (meets ?i ?j)
            (= (duration (plus ?i ?j))
               (add (duration ?i) (duration ?j)) ))))
```

and that points have no duration:

```
;;duration-DU-zeroduration
(forall (?x)(<=> (timepoint ?x) (= (duration ?x) zero)))
```

Examples of durations include

```
;;DU-examples
(and  (timeduration year)
      (timeduration week)
      (timeduration day)
      (timeduration hour)
      (timeduration minute)
      (timeduration second) )
```

We might also want to insist that nontrivial moments cannot have zero duration. In VC, since moments are timepoints, this covers them. If not, however, the following might seem a reasonable way to do it:

```
???    (forall (?i:timeinterval)(=> (= (duration ?x) zero)
                                     (moment ?i) ))
```

but in simple discrete theories it is false, while in simple dense theories it says nothing , since moments do not exist. A suitable axiom for general use talks about points:

```
(forall (?p ?q)(<=>
            (before ?p ?q)
            (not (= (duration (between ?p ?q)) zero)) ))
```

The biconditional guarantees that moments in ENDPOINTS have zero duration, for example, and that one-point closed intervals in the point-continuum theory also do.

(One might wish to allow non-pointlike intervals with zero duration, by weakening the biconditional in `duration-DU-zeroduration` to a simple conditional. These things would complicate the description of clocks, however, so here we assume that they are impossible.)

If time is totally ordered and we select a certain fixed starting timepoint, then any other point is uniquely defined by the duration of the interval between it and the start time. So timepoints can be <u>identified</u> by durations:

```
;;duration-DU-rigid
(forall (?i ?j)
        (=> (and (timeinterval ?i ?j)
                 (= (beginof ?i) (beginof ?j))
                 (= (duration ?i) (duration ?j)) )
            (= (endof ?i) (endof ?j)) ))
```

Notice that this does not claim that the interval itself is *uniquely* identified, since open and closed intervals have the same duration in the point continuum. However, in theories where endpoints identify intervals this will of course suffice. Also, rigidity does not hold in branching time models, where the clock or calendar time of a timepoint fails to uniquely specify it since there are many alternative timelines all going at the same rate, as it were. A convenient variation on the rigidity axiom uses the function `from`, from a point and a duration to the point that much later:

```
;;from-DU
(deffunction from
(forall (?p ?d) (= ?d (duration (between ?p (from ?d ?p)))))
(forall (?p ?d)(=> (and (timepoint ?p)(timeduration ?d))
                   (timepoint (from ?p ?d))))
)
```

Backwards and reflected intervals have negative and negated durations:

```
;;positive-DU
(defrelation positive
(forall (?p ?q)
        (<=> (before ?p ?q)
             (positive (duration (between ?p ?q))) ))

;;DU-back
(forall (?i)(= zero (add (duration ?i)(duration (back ?i)))))
```

Some other useful functions include a function which totals a sequence of durations

```
;;total-DU
(deffunction total
(forall (?d @s)
        (and  (= (total ?d) ?d)
              (= (total ?d @s)
                 (add ?d (total @s)) )))
```

```
)
```

and one that forms a sequence of the durations of a sequence of intervals, a kind of duration-on-sequences function:

```
;;dduration-DU
(deffunction  dduration
(forall (?x @l)(=> (timeinterval ?x)
                   (and (= duration ?x)(dduration ?x))
                        (= (dduration ?x @l)
                        (listof (duration ?x) (dduration @l)))))))
```

6.3 Simple Clocks

A simple clock is characterised by a `start`, which is a timepoint, and a `beat`, which is a duration. It works by counting the number of beats between its start time and the time being measured. A clock makes a continuous timeline seem discrete by mapping every time to a `clocktick`.

```
;;simpleclock-CL
(forall (?c) (=>(simpleclock ?c) (and (duration (beat ?c))
                                 (timepoint (starttime ?c)))))
```

```
;;clocktick-CL
(defrelation clocktick (?p ?c) :=
            (exists (?n)(and (integer ?n)
                             (= (duration (between (starttime ?c) ?p))
                                (mult ?n (beat ?c)) ))))
```

hence

```
 (forall (?n ?c)
      (clocktick (from (times ?n (beat ?c)) (starttime ?c)) ?c))
```

Since intervals can be backwards, a clock 'tells' the time even before its starttime, but such clocktimes are negative.

A (simple) clock time is the time of a point as measured by the clock, which is equal to the time of the immediately preceding clocktick:

```
;;simpleclocktime-CL
(defunction simpleclocktime (?p ?c)
  (and  (integer (simpleclocktime ?p ?c))
        (less (mult (simpleclocktime ?p ?c) (beat ?c))
              (duration (between (starttime ?c) ?p)))
        (not (less (mult (+ 1 (simpleclocktime ?p ?c)) (beat ?c))
                   (duration (between (starttime ?c) ?p))  ))
        )
  )
```

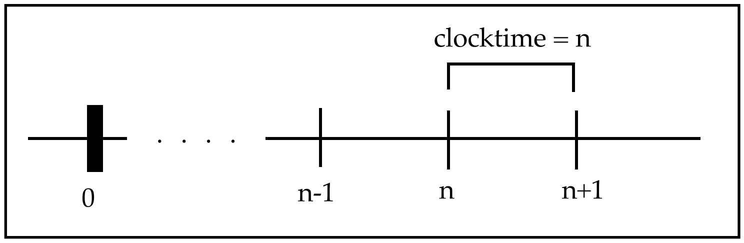

This is the best we can do in general, since no clock can precisely locate every point in a dense timeline. In the discrete timelines, however, we might hope for more accuracy.

Until now nothing has insisted that every moment of a discrete time must have the same duration, and indeed this would not be an appropriate claim to make in all theories (Situation-Calculus Time for example). However, now seems the right time to make this insistence. If we call this universal atomic amount of time `quantum`, then we have simply:

```
(forall (?i) (<=> (moment ?i)
             (= (duration ?i) quantum) ))
```

(In some continuous time theories, this means that `quantum` is `zero`; in others, it says nothing since there are no moments.)

We can reasonably expect that a clock-beat is some definite number of quanta, since when time is discrete there are no other times available for the clock to tick at. Indeed, if time is discrete then every duration is made of quanta:

```
(forall (?d)(<=> (timeduration ?d)
                (exists (?n)
                        (= ?d (mult ?n quantum)) ) ))
)
```

and so it follows that a clock which could beat at the pulse-rate of the universe could indeed serve as a universal clock:

```
;;quantumclock-Q
(defrelation quantumclock (?c) :=
             (and (simpleclock ?c)(= (beat ?c) quantum))))

;;simpleclocktime-Q
(forall (?c ?p)
        (=>  (and (quantumclock ?c)(timepoint ?p))
             (= (mult (simpleclocktime ?p ?c) (beat ?c))
                (duration (between (starttime ?c) ?p))))))
```

Clocks are shift-invariant. This is easier to say in KIF than in English:

```
;;CL-shift-invariance
(forall (?c ?d)
       (=> (and (= (beat ?c) (beat?d))
                (clocktick (starttime ?d) ?c))
```

```
            (forall (?p)(=>(timepoint ?p)
                           (= (simpleclocktime ?p ?c)
                              (+ (simpleclocktime ?p ?d)
                                 (simpleclocktime (starttime ?d) ?c)))))
        ))
```

Models

Again, the semantic weight of these axioms is closely tied to the possible interpretations of the arithmetic terms that occur in them. If `integer` really refers only to integers, then all the nonstandard models of the timeline are ruled out and these axioms, when added to any of the earlier discrete-point theories, have only standard interpretations. However, nonstandard models of arithmetic adapt perfectly well to give nonstandard clocks, which might be called pink-rabbit clocks, since they beat forever and then keep on going.

6.4 Calendars

A calendar is a fixed system of timeintervals and subintervals which divide the timeplenum into separate, identifiable pieces. A clock defines a calendar, but not all calendars can be defined that way, since a calendar need not be shift-invariant.

One way to think of a calendar is that it provides a way to make continuous time feel like discrete time at a certain scale. Thus we can talk of the next year, next month, next minute etc., and combine these together to refer to the third hour of the second day of the ninth week in 1995.

A *simple calendar* is a clock with a finite sequence of durations which add up to its beat. Years divided into months and days divided into hours are examples. The beat of a simple calendar is called its *scale*., and the sequence of durations is its *rhythm* For example, our conventional year-scale calendar's rhythm is the sequence <31 days, 28days, 31days,30days,31 days,30days,31 days,31 days,30 days,31 days,30 days,31 days>. (This description insists that this is true even in a leap year, by the way: the 29th of February is always a shift-interval rather than a part of the calendar pattern.)

```
;;simplecalendar-CL
(defrelation simplecalendar
(forall (?c)
        (=> (simplecalendar ?c)
            (and (simpleclock ?c)
                 (= (beat ?c) (total (rhythm ?c))))))
)
```

Since KIF sequences are finite, such pathological examples as an oscillator approaching infinite frequency are ruled out. A rhythm defines a sequence of intervals between successive clockticks.

```
;;dates-CL
(deffunction dates
(forall (?c ?p) (=> (and (simplecalendar ?c) (clocktick ?p ?c))
                    (and (= (beginof (first (dates ?c ?p)) ?p)
```

```
                (meets (dates ?c ?p))
                (= (rhythm ?c) (dduration (dates ?c ?p))  )
                )
             )))
```

So the dates of a year are the months, and the dates of a day are the hours. (This axiom should be rewritten for the point continuum, since the simple requirement of meeting may be too simplistic. It would lead for example to some hours being open and some closed. The proper thing to say there is that there is a moment between them.)

6.5 Correcting and adjusting Clocks

Clocks are prone to corrections of various kinds. A clock may be adjusted to be faster or slower, and clocks can be given deliberate hiccups as in leap years. Since leap years are regular, one could define a four-year clock to take them into account, but an alternative approach is to introduce the notion of a correction. Since our definition of clock insists that they never change, we have to describe making a correction as shifting to a different, but usually closely related, clock. For example, leap years and leap seconds re-set the clock's `starttime` slightly:

```
;;shift-CL
(defrelation shift (?c ?d ?e) :=
(and (= (beat ?c) (beat ?d))
     (= ?e (duration (between (starttime ?c) (starttime ?d))))))
```

The shifting interval need not be measurable by the clock. It can be much smaller than the beat of the clock, as in usually the case with such minor corrections.

Correcting the rate just amounts to re-setting the beat; but as this happens at a particular time, we usually implicitly consider the start-time also to have shifted to keep things straight.

```
;;adjust-CL
(defrelation adjust (?c ?d ?e ?p) :=
(and  (= (beat ?d) (add (beat ?c) ?e))
      (= (simpleclocktime ?p ?c)(simpleclocktime ?p ?d))
      )
)
```

It follows that when a `clock` is adjusted, its `starttime` moves to the timepoint which it would have to have been in order to have arrived at this time at the new rate!

Adjust makes a clock slower. The usefulness of allowing negative durations and backward intervals is shown by the fact that to make a clock faster is simply to make it slower by a negative amount.

We can now consider an intuitive 'clock' to be a series of simple clocks, each one differing from the previous one by a correction of some kind happening at a point. In the point continuum, the changes occurring during single-point closed intervals. In the glass continuum, the changes happen at the moments where the longer intervals meet each other. Either way, time would be ambiguous only for a moment.

Since this correcting might go on for ever, the sequence of clocks might be infinite, so cannot be described as a KIF sequence. It is simplest to define a function. Every clock has an associated set of points, called the `resettimes`, totally ordered by `before`. If this is empty then the clock is simple clock. We need the idea of the next reset time, provided by:

```
;;nextone-CL
(deffunction nextone
  (and  (member (nextone ?p ?s) ?s)
        (before ?p (nextone ?p ?s))
        (forall (?x) (=> (member ?x ?s)
                         (not (before ?p ?x (nextone ?p ?s)))))
  ))
```

I will now slip into second-order syntax for a moment, in order to define this predicate on functions:

```
;;clock-CL-1
(forall (?c)
        (<=> (clock ?c)
             (and (forall (?x) (=> (member ?x (resettimes ?c))
                                   (timepoint ?x) ))
                  (forall (?p)
                          (=> (member ?p (resettimes ?c))
                              (simpleclock (value ?c ?p))))
                   )))
```

and then of course the time as measured by a clock is the clocktime on the simple clock currently in use:

```
;;clock-CL-2
(forall (?p ?c)
        (or
         (and (empty (resettimes ?c))
              (simpleclock ?c)
              (= (clocktime ?p ?c) (simpleclocktime ?p ?c))  )
         (forall (?q )
                 (<=>
                  (and (member ?q (resettimes ?c))
                       (before ?q ?p (nextone ?q
                                              (resettimes ?c))))
                  (= (clocktime ?p ?c)
                     (simpleclocktime ?p (value ?c ?q)) )))))
```

6.6 Compound clocks and calendars

Days are a particularly simple form of calendar in which the inner beats are themselves regular clockticks. It is worth giving this class a special name, since everything becomes so much simpler in this case. Let us say that one simple clock `fits` another if all its clockticks are also clockticks of the second:

```
;;fits-cl
(deffunction fits)
(forall (?c ?d )
```

```
        (<=> (fits ?c ?d)
             (forall (?p) (=> (clocktick ?p ?c) (clocktick ?p ?d)))))
```

Then it follows that the first one's beat is an integral number of beats of the second. A sequence that long of identical beats of the second clock would serve as the rhythm of a calendar whose beat was the first clock. This is exactly the relationship between minutes and seconds, days and hours, and years and weeks. That these clocks all fit together so nicely is the reason why any time can be defined as a certain number of milliseconds after the birth of Christ. In discrete time, every clock fits onto the quantum clock, but in dense time arbitrarily fine misfits of rate can occur.

We can now define a compound clock to be a sequence of clocks, each of which fits the next , which restarts at the beginning of each beat of the first. The basic relation is between a simple clock and a (not simple) clock which gets reset at midnight:

```
;;dividesup-CL
(defrelation dividesup
(forall (?c ?d)
        (<=> (dividesup ?c ?d)
             (and (forall (?p)
                          (and (<=> (clocktick ?p ?c)
                                    (and (member ?p (resettimes ?d))
                                         (= ?p
                                            (starttime (value ?d ?p)))
                                         )
                                    )
                               )
                          ))))
)
```

and then:

6.6.3
```
;;compoundclock-CL
(defrelation compoundclock)
(forall (?c ?d @s)
        (<=> (compoundclock ?c ?d @s)
             (and  (dividesup (value ?c ?d))
                   (compoundclock ?d @s))))
```

So each of the slower clocks has to inherit the same corrections, if any, that are made to the faster ones. This means for example that adding a leap second at midnight on the 17th of July makes the enclosing week, year or century also shift forward by that second, if the whole nested structure is asserted to be a compound clock.

There is almost no end to the temporal structures that could be defined and that might be useful. For example, this document has not considered intermittent intervals, or developed the idea of branching-time clocks, or transitivity tables for relativistic time. However, these ideas are clearly capable of considerable expressive power, and I hope that the careful comparative development might be of some utility to future temporal formalizers.

7. Acknowledgements

I am grateful to Richard Fikes , Tom Gruber and Peche Turner for their support, to Tom for helping me debug the axioms, and to BobMorris for finding errors in earlier drafts.

Bibliography

Allen, J. (1984) Maintaining Knowledge about temporal intervals, *Commun. ACM* **26** (11) 832-843

Allen, J.F. and Hayes, P.J. Moments and points in an interval-based temporal logic, *Computational Intelligence* 5,4, 1989. pp 225-238

Freska, C. (1992) Temporal reasoning based on semi-intervals. *AI Journal* **54** (1-2)199-227

Galton, A. (1990) A Critical Examination of Allen's Theory of Action and Time. *AI Journal* **42** (2-3) 159-187

Gruber , T. R. & Olsen, G. R. (1994) An ontology for engineering mathematics. In Doyle, Torasso, & Sandewall, Eds., *Fourth International Conference on Principles of Knowledge Representation and Reasoning* , Morgan Kaufmann.

Hayes, P.J. & Allen, J. (1987) Short Time Periods. *Proceedings IJCAI-87,* Milan, 981-983

Snodgrass, R. T. *et al* (1994) *The TSQL2 Language Specification*. Technical Report, University of Arizona: FTP.cs.arizona.edu

Van Benthem, J. F. A. K. (1983) *The Logic Of Time*. Synthese Library vol. 156. D. Reidel Publishing Company.

Vilain, M. & Kautz, H. (1986) Constraint propgation algorithms for temporal reasoning. *Proceedings AAAI-86*, Philadelphia, 377-382

Vilain, M., Kautz, H. & van Beek, P.G. (1989) Constraint propgation algorithms for temporal reasoning: a revised report. in: Weld, D & de Kleer, J., eds, *Readings in Qualitative reasoning about Physical Systems*, Morgan Kaufman.