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.

<u>Models</u>

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

<a,b> before <c,d> 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.

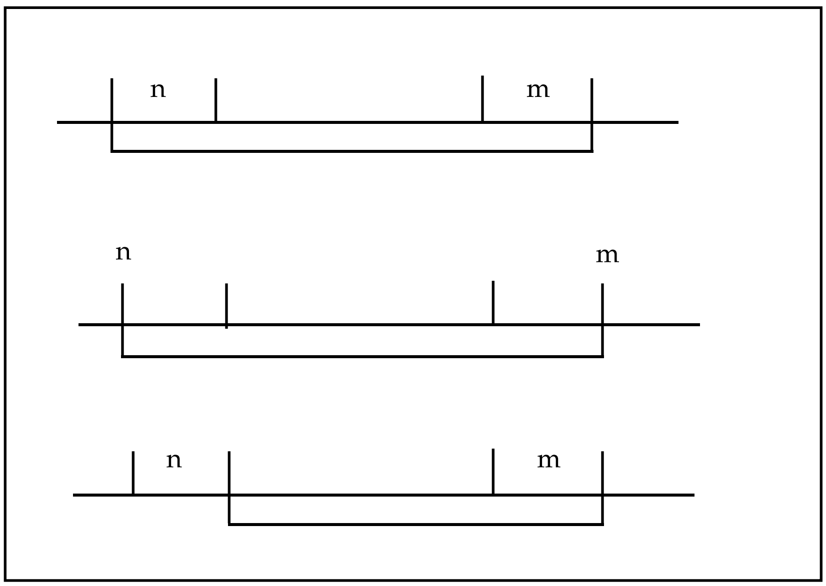

Figure 1. Three ways to label end-ticks