Time Representation: A Taxonomy of Interval Relations *

Peter Ladkin
Kestrel Institute
1801 Page Mill Road
Palo Alto, Ca 94304-1216.

Abstract

James Allen in [*All2*] formulated a calculus of convex time intervals, which is being applied to commonsense reasoning by Allen, Pat Hayes, Henry Kautz and others [*AllKau, AllHay*]. For many purposes in AI, we need more general time intervals. We present a taxonomy of important binary relations between intervals which are unions of convex intervals, and we provide examples of these relations applied to the description of tasks and events. These relations appear to be necessary for such description. Finally, we provide logical definitions of a taxonomy of general binary relations between non-convex intervals.

Introduction

James Allen in [*All2*] formulated a calculus of convex time intervals, which is being applied to commonsense reasoning by Allen, Pat Hayes, Henry Kautz and others [*AllKau, AllHay*]. *Convex* intervals are intuitively those which have no gaps. The term *convex* comes from topology. Allen's calculus is a finite relation algebra in the sense of Tarski [*JoTa1, JoTa2, Mad1*]. It has 13 atoms, which Allen enumerates, and hence the algebra has 2^{13} elements. We refer to the elements of this algebra as *convex relations.* There are close relations between algorithms used by Allen [*Freu*], and work in representations of relation algebras [*Mad1, Com1*] We present some mathematical results on Allen's algebra in [*LadMad*]. Other ways of representing time in AI have been argued for in [*McDer1, McDer2*].

Here, we investigate the binary relations that can hold between intervals which are unions of convex intervals. We call such relations *non-convex relations.* These intervals consist intuitively of some (maximal) convex subintervals with convex gaps in between them. We start by discussing points-based and intervals-based representations of time. We then present a taxonomy of important binary relations between intervals which are unions of convex intervals, and we provide examples of these relations applied to the description of tasks and events. These relations appear to be necessary for such description. Finally, we provide logical definitions of a taxonomy of general binary relations between non-convex intervals.

The combinatorial explosion of possible binary relations between unions-of-convex intervals is dampened by considering only a subset of all possible relations. However, results of the author and Roger Maddux show that there are infinitely many relations definable in the algebra generated by these intervals [*LadMad*]. The notion of *convex interval* is definable in the algebra, as are the notions of *having exactly (greater than, less than) n maximal convex subintervals*, for each *n* [*LadMad*].

*This work was partially supported by RADC contract F30602-84-C-0109 and DARPA contract N00014-81-C-0582

Instants, Intervals and the Representation of Periods

In [Lad1], we discussed points-based and intervals-based ways of representing time.

Project management systems, amongst others, need a way of representing *periods* of time over which tasks happen, are scheduled, etc. There is a choice to be made between instant-based and interval-based representations of periods:

Instants are atomic, indivisible entities which do not overlap, and are usually partially or linearly ordered. The order is usually called *later than.* Instants have no duration.

This notion is used in the semantics of serial or concurrent programming languages with atomic instructions. Instants of time are identified with states of the system, and attached to these instants are propositions which describe the internal, non-temporal structure of the states.

Use of instants in this way can be referred to as *taking snapshots*, and this approach is often taken when the system to be modelled is clocked. All snapshots are then synchronised with the clock, and the problem of determining durations of periods is reduced to a count of clock interrupts.

To build periods from instants, we have to specify a range of instants, e.g.
$period(t_1, t_2) \equiv \{t : t_1 \leq t \leq t_2\}$. There is a question about whether to include endpoints, which we shall refer to again. A more complex, but useful, kind of period can be specified by taking (finite or arbitrary) *unions* of these basic *convex* periods.

We then have periods which can represent, say, the time during which a given process has control of the processor in a time-shared environment.

Intervals represent time periods directly. Intervals have *duration*, and are not necessarily indivisible. They are thus an abstraction from the properties of sets of time instants with measure. Thus, there are 12 ways that intervals may be related, excluding equality, e.g. *precedes, overlaps, contained in* [*All2*]. By contrast, instants can be related by only two, *earlier than, later than.*

To determine what structure we *need* in this context, we believe it is best to work with the abstraction directly. This position is argued in [*Lad1, All2, AllKau, AllHay*]. We consider the sets-of-points notion as one possible *interpretation* of intervals.

The use of intervals is not restricted to AI. [*Lam1*] defines an ordering on intervals (there referred to as *sets of events*), in order to prove the correctness of certain concurrency algorithms. Interval representations are also considered in [*vBen, Hum, Dow*]. There are mathematical constructions that convert point structures to interval structures and vice versa, e.g. the ***sets-of-points with measure*** construction [*vBen*]. See also [*LadMad*].

In particular, [*Hum*] resolves some supposed difficulties in the definition of truth of propositions over intervals.

Additional reasons we prefer to work with intervals are

- intervals provide a natural way of talking about duration, the length of time over which something happens, because they are an abstraction of the properties of periods of time
- interval ***notation*** is extensible: complexifying time structure doesn't lead to changes in syntax; whereas point notation isn't extensible, in that the number of points needed to specify a time structure varies with the complexity of the structure, e.g. we need $2 \times n$ points to specify the union of n convex periods
- there are unresolved difficulties with the endpoints of time periods, whether specified in point structures or interval structures. These need to be resolved before ***any*** implementation of time is attempted, but the difficulties are treated in a more ***ad hoc*** manner by points-based models of time. Notions such as ***temporal conjunction*** of propositions may be expressed in interval formulations [*Hum*], but not easily (so far) in standard points-based temporal logic. The interval approach allows all possible relations between endpoints, whereas such points-based approaches as [*McDer1, McDer2*] have to chose a convention which is then hard-wired into the semantics. In terms of [*All2*], the points-based approach has to choose between ***precedence*** (don't include the endpoint), or ***overlapping*** (include the endpoint), and usually rules out ***meeting*** for any intervals. See [*Hum*] for another example involving temporal conjunction.

Relation Primitives for Unions of Convex Intervals

Intervals which are unions of convex intervals occur naturally. For example, any recurring time period can be represented: we can regard the period *MONDAYS* as being composed of each individual Monday, *LABOR-DAYS* is likewise the union of convex intervals, consisting of each individual Labor day, the period of the regular weekly meeting with the boss is also a union of convex intervals, each of them the period of single meeting. These kinds of intervals seem to be among the most useful of the non-convex cases, and since we have reason to hope that knowledge gleaned from considering the convex case will transfer in part, we consider unions of convex intervals in detail. We develop further the definition of time units, which include examples such as the above, in [*Lad3*].

An interval which is a union of convex intervals looks like this

i

This interval ***i*** has three "parts", i.e. maximal convex subintervals, which we call
maxconsubints.

The Choice of Relation Primitives

There are too many discrete ways that unions of convex intervals may be related to each other. An exhaustive enumeration is infeasible, because

Theorem 1 ***The number of relations between unions of convex intervals is at least exponential in the number of maxconsubints.***

In fact, a much sharper result is true (see [*LadMad*]), but we intend only to establish infeasibility here.

Proof:

We prove the theorem by enumerating the relations between two intervals with n maxconsubints and using an inductive argument.

Consider two intervals which are the unions of 2 convex subintervals each. Suppose the first subinterval of each is entirely disjoint from the second subinterval of the other. Then each first maxconsubint ***precedes*** the second maxconsubint of the other. The intervals are related in 13^2 ways, including equality, since the first maxconsubints can be related in 13 ways, including equality, and so can the second maxconsubints. When we consider that the first maxconsubint of each may be related by other than ***precedence*** to the second maxconsubint of the other, e.g. they may ***overlap***, or ***meet***, we see there are more than 13^2 relations overall.

Now consider two intervals with $n+1$ maxconsubints, such that the first n maxconsubints of each interval all ***precede*** the final maxconsubint of the other. By the inductive hypothesis, there are more than 13^n ways the subintervals consisting of the first n maxconsubints may be related. The final two maxconsubints may be related in 13 ways, and therefore the total number of possible relations is more than 13^{n+1}. Again, when we consider that the final two maxconsubints may be related by ***overlaps*** or ***meets*** to the penultimate maxconsubint of the other, we notice there are many more relations than just those we enumerated in the proof.

Hence we have established the base and the inductive steps, and we draw the conclusion of the induction.

End of Proof.

To avoid the combinatorial explosion implied by the theorem, our basic relations don't depend on the number of maxconsubints. It is intuitively plausible that we don't need relations that depend on the number of maxconsubints for expressing properties of time periods associated with actions, tasks, events or propositions.

However, the relation algebra generated by the relations we consider is still infinite, and still enables us to define the class of intervals with exactly n subintervals, for each n [*LadMad*].

The approach we take will generalise the convex relations, by introducing ***functors*** that generate non-convex relations from convex relations, by enumerating new subclassifications of relations that weren't there in the convex case, and by enumerating the various relations that arise from considering just the first and last maxconsubints.

Additionally there is one relation, ***bars***, which is not obtained by generalising the convex case in some way.

We obtain the following relations between unions of convex intervals:

- those generated by the functors *mostly, always, partially, sometimes*, and *disjunction* from convex relations (*always* may be defined in terms of *mostly*)
- ***contains***
- ***disjoint from***, which splits into:
  - ***precedes*** and ***follows***
  - ***meets*** and ***is met by***
  - ***intermingles with***, which splits further into:
    - ***disjointly-contains*** and ***disjointly-contained by***
    - ***disjointly-overlaps*** and ***disjointly-overlapped by***
    - ***begins preceding*** and ***begins following***
    - ***ends preceding*** and ***ends following***
    - ***surrounds*** and ***is surrounded by***
- ***strictly intersects***, splitting into
  - ***begins after*** which splits into
    - ***begins in***
    - ***begins with***
    - ***begins following***
  - ***ends before*** which splits into
    - ***ends in***
    - ***ends with***
    - ***ends preceding***
  - ***begins before*** and ***ends after***, with the corresponding case splits
  - ***begins at*** and ***ends at***
- ***bars***, which is a new relation not generalised from convex relations

We give below the definitions of the relations, and follow with examples to show they are naturally occurring. We conclude that these relations are necessary for expressiveness in a calculus of intervals which are unions of convex intervals.

These relations are not atomic (not disjoint as sets of interval pairs), and some of them are definable from others. For example, ***surrounds*** is the conjunction of the relations ***begins before*** and ***ends after***, i.e. ***begins before*** $\wedge$ ***ends after***, where $(i\ (R \wedge S)\ j) \equiv ((i\ R\ j) \wedge (i\ S\ j))$

The Definition of the Relation Primitives

The Intended Calculus

We intend that these relations will be manipulated algebraically, that is, by considering only derived relations defined in the relation algebra generated by these relations [*JoTa2, Mad1, LadMad*]. We are not concerned with first-order definability, since we don't intend to use a first-order theorem-proving approach, and thus we have many more relations than we would need if we were using a first-order language approach. We believe the payoff is in the simpler structure of an algebraic theory.

Informal Definitions of the Primitives

Define a *component* of a non-convex interval to be a maxconsubint.

We shall speak of the ***n'th component of*** i and the ***n'th component of*** j, where i and j have finitely many components, as a ***matched pair*** of subintervals. In case of i and j having infinitely many components, we assume without defining it a one-to-one function that matches the "closest" components. This function may, in fact, be rigorously defined [*LadMad*].

$R^{\cup}$ is the converse relation to R; i.e. $(i\ R^{\cup}\ j)$ iff $(j\ R\ i)$.

We draw example intervals i and j for each relation. We represent them on different lines, but they are intended to be intervals on the same time line. It would be better to use two colors.

The *relation functors* are:

- **mostly**: i ***mostly*** $R\ j$, where R is a convex relation, if, for every component of j, there is a component of i that is R to it. This allows the possibility that there are other components of i, but not of j.

  E.g. i ***mostly meets*** j

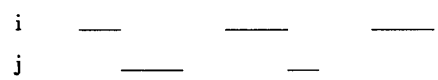

- **always**: i ***always*** $R\ j$, where R is a convex relation, iff matched pairs of components of i and j are R to each other. Alternatively, a form of the definition that will work for both finite and infinite unions of convex intervals:- every component of i R some component of j, and every component of j $R^{\cup}$ some component of i. ***Always*** is definable from ***mostly***: i ***always*** $R\ j$ iff i ***mostly*** $R\ j$ and j ***mostly*** $R^{\cup}\ i$, where $R^{\cup}$ is the converse relation to R.

  E.g. i ***always meets*** j

- **partially**: i ***partially*** $R\ j$, where R is a convex relation, iff some pairs of components of i and j are R to each other, and all others are disjoint. This allows the possibility that the disjoint intervals may *meet*

  E.g. i ***partially meets*** j

- **sometimes**: i ***sometimes*** $R\ j$, where R is a convex relation, iff some pairs (at least one pair) of components of i and j are R to each other

  E.g. i ***sometimes meets*** j

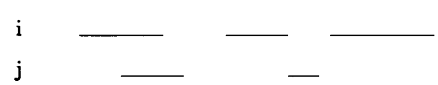

- **disjunction**: $i\ R \vee .... \vee Q\ j$ iff every pair of components is related by R or .... or Q or ***precedes*** or ***follows***

  E.g. i ***meets*** $\vee$ ***contains*** j

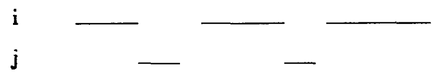

Many of the convex classifications generalise directly to the non-convex case. However, some of the convex classifications get new subclassifications in the non-convex case, notably ***disjoint from***, which obtains a new category of ***intermingles***, which itself has subclassifications, and ***strictly intersects***, which obtains many new subcategories. Some of the new subcategories are valid for both ***intermingles***, which is a category of ***disjoint***, and ***strictly intersects***, which is a different category.

- **contains**; unchanged from the convex case; *i contains j* iff every component of *j* is contained by some component of *i*

  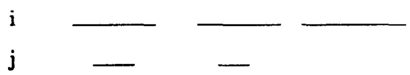

- **disjoint from**, which is a symmetric relation, and can be classified into:
  - **precedes,** as in the convex case; precedes is antisymmetric, and *i precedes j* iff all subintervals of *i* precede all subintervals of *j*

    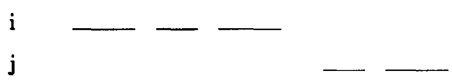

  - **follows**, the inverse of *precedes*
  - **meets**, antisymmetric; *i meets j* iff the final component of *i* meets the first component of *j*

    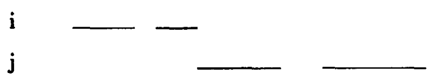

  - **is met by**, the inverse of *meets*
  - **intermingles with**; new in the non-convex case, symmetric, and itself has subclassifications enumerated below
- **strictly intersects**, which has new subclassifications generated by the ***relation functors***, as well as the new subcategories enumerated below

We now enumerate the subclassifications of *strictly intersects*.

- *i* **begins after** *j*; which is split into the mutually exclusive cases:
  - *i* **begins in** *j*; the leftmost component of *i* is overlapped by a component of *j*.

    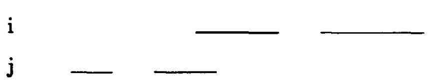

  - *i* **begins with** *j*; the leftmost component of *i* is overlapped by the leftmost component of *j*

    

  - *i* **begins following** *j*; some component of *j precedes* all components of *i*

    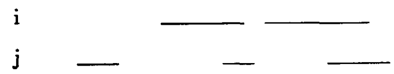

- *i* **ends before** *j*; split into the mutually exclusive cases:
  - *i* **ends in** *j*; the rightmost component of *i* overlaps a component of *j*

    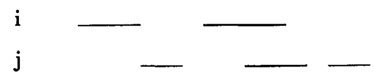

  - *i* **ends with** *j*; the rightmost component of *i* overlaps the rightmost component of *j*

    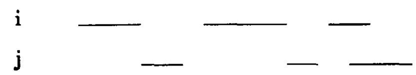

  - *i* **ends preceding** *j*; some component of *j succeeds* all components of *i*.

    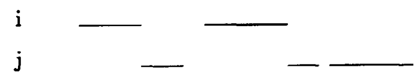

- the converses **begins before** and **ends after**, which are also split into cases, giving **begins preceding** and **ends following,** and the converses of the other four relations. We can't find any useful names for these other four at present
- *i* **begins at** *j*; the leftmost component of *i* starts the leftmost component of *j*

  

- *i* **ends at** *j*; the rightmost component of *i* finishes the rightmost component of *j*

  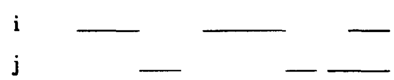

We enumerate the subclassifications of *intermingles with*:

- *i* **disjointly-contains** *j*; equivalent to *i is disjoint from and surrounds j*

  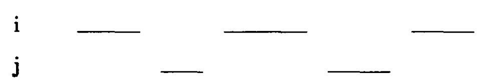

- *i is* **disjointly-contained** *in j*; the converse of *disjointly contains*

- *i* **disjointly-overlaps** *j*; equivalent to *disjoint and begins preceding and ends preceding*

- *i is* **disjointly-overlapped** *by j*; the converse of *disjointly-overlaps*

Finally, we note there are certain classifications that are valid in both the *intermingling* and the *intersecting* cases:

- **begins preceding, begins following, ends preceding** and **ends following**, with the definitions given in the *intersecting* case valid also for the *intermingling* case
- *i* **surrounds** *j*; equivalent to *i begins before and ends after j* in the *intersecting* case, and *disjointly-contains* in the disjoint case. We illustrate the *intersecting* case:

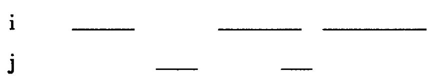

- *i is* **surrounded by** *j*; the converse of *surrounds*

Additionally, there is a polyadic relation that is of some importance. We illustrate it for two intervals, and it should be clear how to generalise it to many. In our calculus, we only consider the binary case [*LadMad*].

- *i* **bars** *j*; the union of *i* and *j* is convex. *Bars* is a symmetric relation, and commutative in the general case

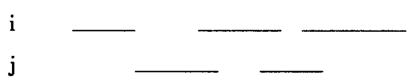

Examples of Relations Between Non-Convex Intervals

We illustrate the relations by examples drawn from general processes, procedures, tasks and occurrences. We conjecture that the most useful applications of the calculus will be in the areas of task description and management, action theory and process theory.

The reader may observe that many of the relations above are converses of relations already included in the enumeration. We include enough examples to show that the relations are useful and natural for descriptive purposes. Each example has the form of expressing a relation between two tasks, events or actions. Suppose P is such a creature. Then we attach to P the interval *int(P)* [*Lad3*]. Then, for two tasks, P and Q, we can consider the interval relation *R(int(P), int(Q))* between the associated intervals. This relation R appears above each example.

- *mostly meets*:
  - after the committee has reviewed the market on Monday, the brokers may act to buy the desired stock.
  - when designs are finalised, the programmers assigned may be immediately available for the implementation task
- *always meets*:
  - after the committee has reviewed the market on Monday, the brokers always act to buy the desired stock.
  - when designs are finalised, the programmers assigned should be immediately available for the implementation task
- *always overlaps*:
  - investigation of the system crash starts before system service is restored, and continues afterwards.
  - preparations for performing the task should be initiated while the design team is finishing the detailed description
- *(overlaps* $\vee$ *contains)*:
  - investigation of the system crash starts before system service is restored, and sometimes continues afterwards.
  - these tasks are always concurrent:
work on the processor configuration and the distributed system design task
- *partially contains*:
  - if you need to cross the road, you do so while there's no traffic.
  - task 34 should only be worked on when Fred and Mary are available for it
- *partially meets*
  - when system service degrades, sometimes we need to reboot
  - the distributed system design task may need to be followed by a feasibility study
- *begins in*:
  - the emergency procedures were introduced and used during the company reorganisation
  - the implementation should be commenced while the system is being configured
- *ends in*:
  - the final reorganisation before dissolution occurred during last year's first financial crisis
  - the implementation should be completed while the system is in the test stage
- *begins with*:
regular system backups were started during the first time the system lost a drive because of a head crash
- *ends with*:
communication with other machines ceased with the last failure of the main processor of the gateway

- ***begins at***:
  tasks A and B are independent starting tasks for the project
- ***ends at***:
  finishing a project at the deadline date

Relation Primitives for General Non-Convex Intervals

General non-convex intervals correspond to arbitrary sets of time-points, in a points-based model of time.

To define the relations logically, we can either assume that certain relations are primitive, that a notion of ***subobject*** is primitive, or that ***part of*** is primitive (subobjects are ***parts of*** their superobjects, and ***contains*** and subobjects are interdefinable also, as indicated below).

The notion of ***subobject*** may be introduced by definition also from interval ***operators***. For example, if we have a notion of ***intersection*** of intervals, which is natural in certain representations such as that of a set of points, we may stipulate that the result of any ***intersection*** of intervals is a subobject of all those intervals. For another example, the argument intervals of a ***union*** operation on two intervals are subobjects of the result of the union. We do not consider operators on intervals in this paper, and prefer to avoid them where possible, since the addition of operators vastly complicates the algebra, which would be no longer simply a relation algebra in the sense of Tarski.

We have the following classification of general non-convex intervals:

- ***contains***
- *disjoint from*, which splits into:
  - *precedes* and *follows*
  - ***meets*** and ***is met by***
  - ***intermingles with***, which splits further into:
    - ***disjointly-contains*** and ***disjointly-contained by***
    - ***disjointly-overlaps*** and ***disjointly-overlapped by***
    - ***begins preceding*** and ***begins following***
    - ***ends preceding*** and ***ends following***
- ***strictly intersects***, splitting into
  - ***begins after***
  - ***ends before***
  - ***begins before*** and ***ends after***, with the corresponding case splits
  - ***begins at*** and ***ends at***
- ***bars***

We shall refer to the basic object of time (point, interval or whatever) as a ***time object***. For example, if you prefer points-based time notions, then you will represent intervals as sets of points, and your basic time object will be a set of points. We assume at present that there is no ***null*** object.

Allen and Hayes [*personal communication, AllHay*] are able to define the convex relations from one primitive, in a first-order manner. Van Benthem [*vBen*] uses two. We have more, to enable us to keep the logical form of the definitions to a statement with at most two bounded quantifiers. We give English descriptions of the definitions, but it should be obvious how to translate them into a first-order logical language with the declared primitives.

Given a basic time object, we define its ***subobjects*** as those objects which ***are contained in*** it. So, for example, for the set-of-points notion, the subobjects of i are precisely the subsets of i. For the convex interval notion, subobjects are convex subintervals, and for the unions-of-convex-intervals notion, subobjects are unions of convex subintervals. We assume that in a given ontology of intervals, either the subobjects are precisely defined, or the notion of *containment* is primitive. They are interdefinable, as indicated below.

If ***subobject*** is a primitive notion, or, alternatively, *containment* is a primitive relation, and ***precedes*** and ***meets*** are also primitive relations, we can give the following definitions of the relations. We assume that the conditions in the subclassifications are conjoined with the conditions in the apropriate superclassification, so that we may avoid repeating parts of definitions; e.g. ***i disjointly contains j*** is to be read as ***(i is disjoint from j)*** $\wedge$ ***some subobject of i is*** ....

i contains j: j is a subobject of i

i precedes j: a primitive relation

i meets j: also primitive

i is disjoint from j: i and j have no common subobjects.

> We note that this definition is adequate only because there is no ***null*** object, which would have to be a subobject of every object.
>
> **i disjointly contains j:** some subobject of i *precedes* all subobjects of j and some subobject of i *follows* all subobjects of j
>
> **i disjointly overlaps j:** some subobject of i *precedes* all subobjects of j and some subobject of j *follows* all subobjects of i
>
> **i begins preceding j:** some subobject of i *precedes* all subobjects of j
>
> **i begins following j:** some subobject of i *follows* all subobjects of j
>
> **i ends preceding j:** some subobject of j *follows* all subobjects of i
>
> **i ends following j:** some subobject of i *follows* all subobjects of j

i strictly intersects j: (also known as *i overlaps j*) i and j have a common subobject, and neither *i contains j* nor conversely.

> Notice this relation is symmetric. We can turn this relation into an antisymmetric relation, as *overlaps* is for Allen, with our primitives, by asserting that the subclassifications below are mutually exclusive and exhaustive of the ***strictly intersects*** relation.
>
> However, we would consider this move to be a claim about the structure of a particular interval model. For example, if one were to consider axiomatising the structure of closed sets of real numbers, one might want to assert that *meets* is the empty relation (since two closed sets can only meet

by intersecting at a point, which is a closed set); that *precedes* is a dense (partial) order (any two non-intersecting closed sets may be separated by a closed set); etc.

We prefer to leave the assertion of exhaustiveness as a structure axiom if it is needed. Thus we must allow *strictly intersects* the luxury of symmetry.

i begins after j: some subobject of j *precedes* all subobjects of i

i ends before j: some subobject of j *follows* all subobjects of i

i begins before j: some subobject of i *precedes* all subobjects of j

i ends after j: some subobject of i *follows* all subobjects of j

i begins at j: there is no subobject of i that precedes all subobjects of j, and symmetrically for j and i

i ends at j: there is no subobject of i that follows all subobjects of j, and symmetrically for j and i

For the case of *bars*, we did not include a monadic predicate in our language for selecting *convex* intervals. It is obvious that we would need such a predicate, whether primitive or defined, in order to define *bars*, however this doesn't solve all the problems.

For example, if we were to provide a predicate for *convexity*, we might try to define:

i bars j: there is a convex object k such that the subobjects of k are exactly the subobjects of i and j

Such an attempt doesn't work: consider arbitrary sets of real numbers, with subobjects being subsets, and convex objects being convex sets.

Let $A = [a, b)$, and $B = [b, c]$, where $a < b < c$. Then intuitively A *bars* B since $A \cup B$ is $[a, c]$ which is convex. However, let x, y be such that $a < x < b < y < c$. $[x, y]$ is a subobject of $A \cup B$, but isn't a subobject either of A or of B.

It's probable that *bars* has to be a primitive relation.

We note that primitives for interval-based notions of time are discussed extensively in [vBen] and [Hum]. We note further that our proposed classification is much finer-grained than the primitives in these works, but that all are first-order definable from the primitives used therein. We have mentioned before that our motives for such profusion are algebraic. We note that [vBen] also provides an extensive discussion of the types of interval structure that may be obtained in different domains of application.

Acknowledgements

We thank Tom Brown, Allen Goldberg, Pat Hayes, Richard Jullig, Wolf Polak, Bob Riemenschneider, Richard Waldinger, and the referees for discussion and comments.

Bibliography

All1 : Allen, J.F., *Towards a General Theory of Action and Time*, Artificial Intelligence 23 (2), July 1984, 123-154.

All2 : Allen, J.F., *Maintaining Knowledge about Temporal Intervals*, Comm. A.C.M. 26 (11), November 1983, 832-843.

AllKau : Allen, J.F. and Kautz, H., *A Model of Naive Temporal Reasoning*, in Hobbs, J.R. and Moore, R.C., editors, *Formal Theories of the Commonsense World*, Ablex 1985.

AllHay : Allen J.F. and Hayes, P. J., *A Commonsense Theory of Time*, in Proceedings IJCAI 1985, 528-531.

Com1 : Comer, S.D. *Combinatorial Aspects of Relations*, Algebra Universalis 18, 1984, 77-94.

Dow : Dowty, D.R. *Word Meaning and Montague Grammar*, Reidel, 1979.

Freu : Freuder, E.C., *Synthesizing Constraint Expressions*, Comm. A.C.M. 21 (11), November 1978, 958-965.

Hum : Humberstone, I.L., *Interval Semantics for Tense Logic: Some Remarks*, J. Philosophical Logic 8, 1979, 171-196.

JoTa1 : Jonsson, B. and Tarski, A., *Boolean Algebras with Operators I*, American J. Mathematics (73), 1951.

JoTa2 : Jonsson, B. and Tarski, A., *Boolean Algebras with Operators II*, American J. Mathematics (74), 1952, 127-162.

Lad1 : Ladkin, P.B., *Comments on the Representation of Time*, Proceedings of the 1985 Distributed Artificial Intelligence Workshop, Sea Ranch, California.

Lad2 : a self-reference.

Lad3 : Ladkin, P.B., *Primitives and Units for Time Specification*, Proceedings of AAAI-86 (this volume).

LadMad1 : Ladkin, P.B. and Maddux, R.D., *The Algebra of Time Intervals*, in preparation.

Lam : Lamport, L., *On Interprocess Communication Part I: Basic Formalism*, Distributed Computing, to appear.

Mad1 : Maddux, R.D., *Topics in Relation Algebras*, Ph. D. Thesis, University of California at Berkeley, 1978.

McDer : McDermott, D., *A Temporal Logic for Reasoning about Actions and Plans*, Cognitive Science 6 (2), April-June 1982, 71-109.

McDer2 : McDermott, D., *Reasoning about Plans*, in Hobbs, J.R. and Moore, R.C., editors, *Formal Theories of the Commonsense World*, Ablex 1985.

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