Path Consistency in a Network of Non-convex Intervals

Robert A. Morris and William D. ShoafF and Lina Khatib
Florida Institute of Technology
Computer Science Department
150 W. University Blvd.
Melbourne, FL 32901
U. S. A.

Abstract

Reasoning about time often involves incomplete information about periods and their relationships. Varieties of incompleteness include uncertainty about the *number* of objects involved, the *distribution* of a set of temporal relations among these objects, and what can be called the *participation* of a set of objects in a temporal relation. A solution to the problem of representing and reasoning about incomplete temporal information of these kinds is forthcoming if a restricted class of non-convex intervals (called *n-tntervals)* is added to the temporal domain of discourse. An n-interval corresponds to the common sense notion of a recurring period of time with a (possibly) unspecified number of occurrences. In this paper, we formalize a representation for temporal reasoning problems using n-intervals. The language of the framework is restricted in such a way that tractable techniques from constraint satisfaction can be applied. Specifically, it is demonstrated how the problem of determining path-consistency in a network of binary n-interval relations can be solved.

1 Introduction

Common sense knowledge about the occurrence of events is often incomplete. The kinds of incompleteness of interest here is expressed in the following sentences:

1. John sometimes drives to Orlando.
2. John's ten trips to Orlando this month twice preceded his phoning Phyllis.
3. Some of John's ten trips to Orlando preceded her phoning Phyllis.

The first example exhibits *number* incompleteness: it is not clear how many driving events there are, hence how many periods of time are involved, but that at least one occurs. The second example illustrates *distribution* uncertainty; we don't know which phoning intervals overlapped with which meetings. The third example exhibits, in addition to distribution uncertainty, what can be called *participation* uncertainty; it is not known how often the relation between the trips and the phonings occurred.

Uncertainty of this kind is removed by adding quantitative information. Compare (l)-(3) with "John's *ten* trips to Orlando this month *twice* preceded his phoning Phyllis: the *first* time and the *third* time". Adding the first quantity eliminates number uncertainty; the second quantity ("twice") eliminates participation uncertainty; the last two quantities together eliminate distribution uncertainty.

Temporal knowledge incompleteness makes it difficult to build a reasoner that adequately deals with time. In this paper, the concern is to formalize a representation which addresses the problem of reasoning with number, distribution, and participation uncertainty. To accomplish this aim the remaining sections

1. Define a class of binary temporal relations between pairs of n-intervals;
2. Construct an algebra for formalizing reasoning problems involving n-interval relations; and
3. Demonstrate how path consistency can be applied to a knowledge base of assertions about n-interval relations, recast as a relation network, in the sense of [Van Beek, 1990].

2 Underlying Model of Time

Time is considered here to be a linear order on a domain consisting of points identified with the real numbers under the ordering $<$. The collection of intervals over this domain is the set of finite sequences of the form $I = \langle I_1^-, I_1^+, I_2^-, I_2^+, \ldots, I_n^-, I_n^+ \rangle$, representing a gapped interval with n convex components, or *subintervals*, starting at I_1^- and ending at I_n^+. We call such a sequence an *n-interval* [1], with the ordering $I_1^- < I_1^+ < I_2^- < \ldots < I_n^+$. $I_i = \langle I_i^-, I_i^+ \rangle$ represents the ith convex component of I.

In addition to the ordering constraints on n-intervals, there is structure imposed by the fact that n-intervals are the temporal component of the representation of knowledge about *recurring events.* This complexity is represented by introducing binary relations between convex parts of n-intervals.

[1] This term is used differently by Ligozat [Ligozat, 1990], viz., to denote an interval consisting of n points.

First, time is viewed as being partitioned into segments, or blocks, each representing a set of occurrences of recurring intervals. Blocks come in various sizes; common examples include days, weeks, months, years, etc. Formally, a temporal partition can be viewed as a function from a convex interval R, which we will call a *reference frame,* to a set of convex intervals. Being a partition, no two items in this set overlap, and each segment of R is in one of the members of the set. For example, let *Weeks(R)* = the set of convex intervals in R with a duration of (at most) one week, starting at R_1[2]. From this notion, where r is a partition, we define an equivalence relation $ASSOC_\tau(I_i, J_j)$, which is true of pairs of intervals /, and Jj if and only if /, and J, are in the same member of $\tau(R)$. For example, if $ASSOC_{Weeks}(I_i, J_j)$, then I_i and J_j occur in the same week.

In addition to associated intervals within a block, the model views periods of time to be *correlated.* Correlation is an equivalence relation between pairs of convex parts of recurring events that holds as a result of binary relations holding between them. For example, consider the sentence "Meetings precede lunches twice a week". There are two distinct groupings here: the first is indicated by the term "weeks", which partitions the reference frame into weekly blocks of associated intervals. The second grouping pairs convex parts of the recurring events "meetings" and "lunches". This pairing phenomenon is represented in the model by the relation of correlation. Further examples of correlation are provided below.

James Allen [Allen, 1983] defined a set of primitive binary relations on convex intervals by enumerating all ways in which two ordered pairs of real numbers can be related. There are 13 such primitives, which we collect together in the set ACR: before (6), meets (m), during *(d),* overlaps (o), starts (s), finishes (/), their inverses (after (bi), met by (mi), etc), and equals (=). From these, he defined a class of 2^{13} interval relations by considering all possible disjunctions of primitives; we call this set CR. Members of CR can be depicted as finite disjunctions of the form "$R_1 \vee R_2 \vee \ldots R_n$", where $R_i \in \mathbf{ACR}$. As constraints in a relation network (defined below) the same relation is depicted as the set $\{R_1, R_2, \ldots, R_n\}$.[3]

Since temporal relations in such a knowledge base are binary, it is possible to represent the knowledge as a network, where each arc represents one of the temporal relations. More precisely (from [Van Beek, 1990]):

Definition 1. A *Network of Binary Relations* is a set X of m variables $\{X_1, \ldots X_m\}$, a domain D_i of possible values for each variable, and binary relations $R_{i,j}$ between variables. An *instantiation* of the variables in X is an m-tuple representing an assignment of elements of D_i to each X_i. A *consistent instantiation* of a network is an instantiation that satisfies all the relations between the variables. A network is *consistent* if such an instantiation exists; otherwise, it is *inconsistent*. A relation $R_{i,j}$ is *feasible* with respect to a network if there exists a consistent instantiation of the network where $R_{i,j}$ is satisfied. For each edge in a network, the set of feasible relations is the set consisting of all and only the relations that are feasible.

Definition 2 A *(Convex) Interval Binary Relation Network* (IRN) is a binary relation network whose variables represent intervals, and whose binary relations are members of **CR**. An atomic IRN is an IRN each of whose edges are labeled with a single member of **ACR**. A *solution network* (scenario) is a consistent atomic IRN.

A cheap, useful inference technique used in examining systems of interval relations is that of *path consistency* [Van Beek, 1990][4]:

Definition 3 A network is *path consistent* if and only if for every triple of variables (i, j, k), $\forall x, \forall z, xR_{i,k}z \Rightarrow \exists y \in D_j\ xR_{i,j}y \wedge yR_{j,k}z$.

A path consistency algorithm such as the one found in [Allen, 1983], is an approximate technique for determining the set of feasible relations in an IRN. It performs intersection and composition of relations, and assumes that a inverse operation on relations is defined. Intersection of relations is just intersection of sets. In the temporal domain, a *Composition Table* ([Allen, 1983], not reprinted here) defines composition between elements of **ACR**. In the general case, the results of composing two elements of **CR** is just the union of the result of the pairwise composition of its atomic elements. Finally, for each relation $R_{i,j}$, its *inverse* ($R^{-1}_{i,j}$), is defined as the relation $R_{j,i}$ that satisfies $IR_{i,j}J \iff JR_{j,i}I$. The path consistency algorithm repeatly performs composition over triangles of edges in a relation network, until no more changes to any edges are performed. Changes occur when a composition of relations on a pair of edges results in a change (a "refinement") of the relation in the third edge.

To summarize, the underlying model for reasoning with recurring events has the following components:

1. A set of n-intervals $I, J, K \ldots$, each consisting of an ordered set of convex parts called subintervals;
2. A function for retrieving the xth convex part of any n-interval I, the result of which is denoted by I_x.
3. A binary equivalence relation *correlation*, (COR), defined between convex intervals; in addition to being an equivalence relation, if $COR(I_x, J_y)$ and $COR(I_x, J_z)$ then $y = z$.
4. A binary equivalence relation *association*, abbrevi-

[2] If it is not possible to divide R into equal segments of one week duration, there will be a segment at the end of R of duration less than one week.

[3] A word on notation used throughout: following convention, *R, S* etc. stand for members of CR. Lower case letters from i,,;,... are used to specify arbitrary convex intervals; as noted, upper case letters from /, J,... specify n-intervals. Single subscripts (**e.g.** R_k) distinguish members of ACR. Double subscripts are used (**e.g.** $R_{i,j}$) to emphasize that R is a relation between i and *j*\ these are often omitted.

[4] Freuder [Freuder, 1978] generalizes path consistency to k-consistency: a network is k-consistent if, given any instantiation of any $k-1$ variables satisfying all the relations among the variables, there exists an instantiation of the k-th variable such that the k values together satisfy all the relations. Path consistency corresponds to 3-consistency.

ated as ASSOC; [5].

5. The set CR, based on the set of Allen primitives ACR

In what follows, we extend the IRN framework for the representation of n-intervals and their relations. It is demonstrated how the method of path consistency can be applied to the problem of determining the set of feasible relations in a n-interval relation network.

3 A Class of TV-Interval Relations

In this section, a language is constructed for specifying a collection of binary temporal relations between pairs of n-intervals, where a specification is a set of assertions describing n-interval relations.

This language regiments fragments of natural language discourse (in this case, English) involving the application of temporal adverbs "sometimes" "only", "always", "always and only", and their negations, "never", "not always", etc. These adverbs can be viewed as operators on elements of CR to make assertions about recurring or repeating events. For example, applying "always" to the convex interval relation "before or meets" results in the n-interval relation "always before or meets".

An adequate semantics for temporal specifications requires sensitivity to what we will call co-designation. For illustration, contrast the following three fragments:

1. John sometimes goes to work before calling his dad. Then, he misses (i.e. the call overlaps) the meeting.
2. Joan sometimes goes to work before calling her dad. Otherwise, she calls her dad first.
3. Faculty meetings sometimes precede seminars. Those meetings overlap with lunch.

In the first passage, the adverb "then" serves to establish a connection with the previously introduced occurrence of going to work. In the second, "otherwise", serves to introduce a temporal relationship between a different occurrence of going to work and that of calling. Informally, co-designation can be viewed as a relation between two assertions, in which, informally, each succeeds in "picking out" a common object (in this case, the same period of time). Thus, there is co-designation between the sentences in fragments 1 and 3, but not in 2.

The case of "sometimes" thus illustrates the need for a mechanism for representing co-designation. This is also the case with contexts involving "always", "only" and "always and only", but these cases are more complicated. To say that cocktails always follow faculty meetings implies a cocktail event after every faculty meeting event. Dually, to say that cocktails only follow faculty meetings is to say that for every cocktail event, there is a faculty meeting preceding it. Finally, to assert that cocktails always and only follow faculty meetings is to express a one-to-one correlation between the two events.

[&]The association relation is not discussed extensively in this paper, and does not figure explicitly in the interpretation of the formal language defined in this paper. It is intended to serve within more general systems for reasoning about recurring events, and is introduced here for the reader to become acquainted with this broader framework.

Assertions involving any of these operators will also be sensitive to co-designation relationships. However, in certain of these contexts, co-designation is being implicitly established with any subinterval of an n-interval. An example will motivate this idea. If I say "Faculty meetings only meet cocktails", and also "Sometimes, voodoo chanting begins faculty meetings", one can infer a temporal relation between voodoo chanting and cocktails (viz., the former are sometimes before the latter), independently of the distribution of the chanting among the faculty meetings. To draw this inference, there needs to be a mechanism for expressing the added degree of uncertainty expressed by the italicized phrase.

With this complexity in mind, a language for expressing n-interval relations is now defined.

Definition 4. The language **NRL** consists of the following: A set of terms for designating n-intervals $I, J \ldots$; a set of *basic operators*: "Π" ("always"), "Ω" ("only"), "Θ" ("always and only"), and "Σ" ("sometimes"); a symbol for *negation* of basic operators (e.g., $\bar{\Sigma}$), the result is a set of operators for "not always" ($\bar{\Pi}$), "not only" ($\bar{\Theta}$), "not always and only" ($\bar{\Omega}$), and "not ever" ($\bar{\Sigma}$); finally, a set of *i-terms* (*index terms*), consisting of either: a variable (taken from $x, y, \ldots$), a constant (taken from $a, b, \ldots$) or a functional expression with one argument of the form $c_i(t)$, ("the ith correlate of t") where i is an integer and t is an index term.

This language can be used to specify temporal knowledge in the form of a conjunction of assertions about n-interval relations expressed in the form $IOP_{t,t'}R_{i,j}J$, where I and J designate n-intervals, t and t' are i-terms and $OP \in \{\Sigma, \bar{\Sigma}, \Pi, \bar{\Pi}, \Omega, \bar{\Omega}, \Theta, \bar{\Theta}\}$. This set will be called **OPS**. The set of all n-interval relations which are the extensions of relational predicates of the form $OP_{t,t'}R_{i,j}$ will be included in a set called **NCR**. Finally, in the sentence of the form "I $OP_{t,t'}$ R J", we call t the "I i-term", and t' the "J i-term".

The interpretation of single sentences in **NRL** containing the basic operators are provided by the following translations into first order logic. It is assumed that I is an X-interval and J is an Y-interval and $1 \leq x \leq X$ and $1 \leq y \leq Y$, and that $R_i \in$ **ACR**. μ and ν stand for variables, $c(\mu)$ is a functional expression whose meaning can be expressed as "The correlate of μ", and μ' is a constant index term. "$COR(\mu, \nu)$" is true of correlated intervals.

- $I\,\Pi_{c(\nu),\nu}\{R_1, \ldots, R_k\}J$ iff $\forall\nu\exists\mu, COR(\mu,\nu) \wedge I_\mu\{R_1, \ldots, R_k\}J_\nu.$ (always)
- $I\,\Omega_{\mu,c(\mu)}\{R_1, \ldots, R_k\}J$ iff $\forall\mu\exists\nu, COR(\mu,\nu) \wedge I_\mu\{R_1, \ldots, R_k\}J_\nu.$ (only)
- $I\,\Theta_{\nu,\mu}\{R_1, \ldots, R_k\}J$ iff $I\,\Pi_{c(\mu),\mu}\{R_1, \ldots, R_k\}J \wedge I\,\Omega_{\nu,c(\nu)}\{R_1, \ldots, R_k\}J.$ (always and only)
- $I\,\Sigma_{\mu',\nu'}\{R_1, \ldots, R_k\}J$ iff $\exists\mu\exists\nu, COR(\mu,\nu) \wedge I_\mu\{R_1, \ldots, R_k\}J_\nu.$ (sometimes)

Restricting extensions of relational predicates in **NRL** to correlated intervals represents the tendency in discourse to restrict attention to correlated events. For example, if one says "faculty meetings sometimes precede cocktails", one usually intends to express a relation between a faculty meeting and the *stngle* happy hour in closest proximity following the meeting. The close proximity between the two occurrences is typically necessary and sufficient to establish their correlation.

I-terms are used to express both correlation and co-designation. The latter is depicted by co-indexing, i.e., subscripting with the same term. For example, let GoingToWork, Calling, Meetings, Seminars, and Lunches each denote n-intervals associated with the recurring events going to work, calling one's dad, meetings, seminars, and lunches, respectively. Then 1-3 above can be regimented as follows:

1. **GoingToWork** sometimes$_{a,b}$ before **Calling**. **Calling** sometimes$_{b,c}$ overlaps **Meeting**.
2. **GoingToWork** sometimes$_{a,b}$ before **Calling**. **Calling** sometimes$_{c,d}$ before **GoingToWork**.
3. **Meetings** sometimes$_{a,b}$ precede **Seminars**. **Meetings** sometimes$_{a,c}$ overlap **Lunches**.

In specifying relations involving "always", as in "I always before J", the I i-term needs to be indexed, since it may be necessary to refer later to either "the same I" or "some other I" (i.e., other than the ones related to J). On the other hand, the J i-term can be a variable, since every J is after I, and therefore there won't be a need to speak of "other Js". To make this distinction, we introduced the functional expression "$c_i(t)$", indicating a functional relationship to the variable i-term; e.g. $\Pi_{c_1(x),x}$. The case of "only" is the dual of "always"; here, the J variable needs to be indexed. Finally, the case of "always and only" can be handled with only variables. [6]

Finally, the meaning of assertions involving the negation of a basic operator is captured using the notion of the *complement* of a relation. The complement of an element of **CR**, R (R'), is **ACR** $-R$. Thus, "never-R" is "always-R', and "not-always-R" is "sometimes-R'", etc. For example, to say that "Dinners never precede, meet, or overlap cocktails" is to say "Dinners always do anything but precede meet, or overlap cocktails".

4 An TV-interval Algebra

In this section, a set of operations for manipulating a specification written in NRL is defined. First, a set of operations for an *n-interval relation algebra* is defined, specifically, the operations of *inverse, intersection* and *composition.* Secondly, rules governing substitution of i-terms are briefly discussed.

An algebra is.a set together with one or more operations on that set, where the set is closed under those operations. The operations to be defined on NCR are inverse, composition, and intersection. These operations allow for the formalization of the reasoning tasks involved in evaluating specifications.

First, the proper definition of intersection on n-interval relations requires the introduction of a form of conjunction of relations. To see this, consider contexts involving "always-R" and "sometimes-S". For example, compare "Dinners sometimes end or overlap with speeches", with "Dinners always end with speeches; the latter is a refinement of the former. On the other hand, the case involving the pair "Dinners always end or overlap with speeches" and "Dinners sometimes end with or are met by speeches" is more complicated: there is added information in the latter assertion that can't be completely refined. The $\oplus$ operator is introduced to express what in English is expressed by "and furthermore"; the result of intersection in this example is "Dinners always end or overlap with speeches and furthermore they sometimes end with speeches".

Definition 5. (N-Interval Algebra) **NCR** is the underlying set for an N-interval algebra. It consists, first, of 4×2^{13} basic relations that result from applying one of four operators $\Sigma, \Pi, \Omega, \Theta$ to one element in **CR**. Secondly, if $OP \in \{\Pi, \Omega, \Theta\}$, $R, S_1, \ldots S_k \in$ **CR** and $S_i \subset R$, and $S_i \cap S_j = \emptyset$ for all $i, j = 1 \ldots k$, then **NCR** contains the relation $OP(R) \oplus \Sigma(S_1) \oplus \ldots \oplus \Sigma(S_k)$.

By generalizing the $\oplus$ operator in this way, n-interval relations can be defined which are sucessive refinements of other n-interval relations. For example, "Always before or meets or after and (furthermore) sometimes meets and (furthermore) sometimes before" is more specific than "Always before or meets or after and (furthermore) sometimes meets".

The operators defined on this set are inverse, intersection and composition. Due to space limitations, attention is restricted to defining these operations on basic relations. Extending the definitions to fit the case of relations defined using the $\oplus$ operator is straight forward, but somewhat tedious.

First, generalizing the notion of inverse to n-interval relations is straight forward. When R is in **CR** we have:

1. $(\Pi R)^{-1} = \Omega R^{-1}$
2. $(\Theta R)^{-1} = \Theta R^{-1}$
3. $(\Sigma R)^{-1} = \Sigma R^{-1}$
4. $(\Omega R)^{-1} = \Pi R^{-1}$

For example, if cocktails always follow or meet faculty meetings, then faculty meetings only precede or are met by cocktails.

Intersection is defined among members of **NCR** of the form $OP_{t,t'}R$ and $OP'_{t,t'}S$; Table 1 contains rules governing intersection. We use $\sqcap$ to distinguish n-interval relation intersection from convex relation intersection ($\cap$).

$$
\begin{array}{lcll}
1.\ \Pi(R) \sqcap \Theta(S) & = & \Theta(R \cap S); & \\
2.\ \Omega(R) \sqcap \Theta(S) & = & \Theta(R \cap S); & \\
3.\ \Omega(R) \sqcap \Pi(S) & = & \Theta(R \cap S); & \\
4.\ \Pi(R) \sqcap \Pi(S) & = & \Pi(R \cap S); & \\
5.\ \Omega(R) \sqcap \Omega(S) & = & \Omega(R \cap S); & \\
6.\ \Sigma(R) \sqcap \Sigma(S) & = & \Sigma(R \cap S); & \\
7.\ \Theta(R) \sqcap \Theta(S) & = & \Theta(R \cap S); & \\
8.\ \Theta(R) \sqcap \Sigma(S) & = & \left\{ \begin{array}{ll} \Theta(R) & \text{if } R \subseteq S; \\ \emptyset & \text{if } R \cap S = \emptyset; \\ \Theta(R) \oplus \Sigma(R \cap S) & \text{otherwise} \end{array} \right. &
\end{array}
$$

[6]The reader will have noticed the similarity between the functional index and conventions for eliminating existential quantification in the process of producing clausal forms for first-order sentences, when the existential quantifier is in the scope of a universal quantifier.

$$9.\ \Pi(R) \sqcap \Sigma(S) = \begin{cases} \Pi(R) & \text{if } R \subseteq S; \\ \emptyset & \text{if } R \cap S = \emptyset; \\ \Pi(R) \oplus \Sigma(R \cap S) & \text{otherwise} \end{cases}$$

$$10.\ \Omega(R) \sqcap \Sigma(S) = \begin{cases} \Omega(R) & \text{if } R \subseteq S; \\ \emptyset & \text{if } R \cap S = \emptyset; \\ \Omega(R) \oplus \Sigma(R \cap S) & \text{otherwise} \end{cases}$$

Table 1. Intersection of N-interval relations

The rules collectively expresses the fact that "always and only" is as specific or more specific than "only", "sometimes" or "always". For example, to say that dinners always follow cocktails is more specific than to say that they sometimes do. Similarly, to say that they always and only follow cocktails is more specific than to say that they always do.

Table 2 defines composition between non-convex interval relations of the form $OP_{t,t'}R$ and $OP'_{t'',t'''}S$, where R and S are in **CR**, and $OP \in$ **OPS**. For example, if I is always and only (Θ) before J and J always (Π) starts K, then using composition one infers that I is always and only before (before-convex-composed-with-starts) K.

$\odot$	$\Pi(S)$	$\Omega(S)$	$\Sigma(S)$	$\Theta(S)$
$\Pi(R)$	$\Pi(R \circ S)$	$\Sigma(R \circ S)$	$\Sigma(R \circ S)$	$\Pi(R \circ S)$
$\Omega(R)$	$\Sigma(R \circ S)$	$\Omega(R \circ S)$	$\Sigma(R \circ S)$	$\Omega(R \circ S)$
$\Sigma(R)$	$\Sigma(R \circ S)$	$\Sigma(R \circ S)$	$\Sigma(R \circ S)$	$\Sigma(R \circ S)$
$\Theta(R)$	$\Pi(R \circ S)$	$\Omega(R \circ S)$	$\Sigma(R \circ S)$	$\Theta(R \circ S)$

Table 2. Composition of N-interval Relations

As noted above, negation operates on each of the four basic non-convex relation operators to form new operators. Because negation can be defined in terms of relation complement, contexts involving them can be reduced to contexts without them. Hence, the operations defined in this section can be applied to all members of **OPS**.

Finally, in order for a reasoner to effectively apply these operations on relations in **NRL**, it is necessary to define rules for i-term substitution. We do this informally. In general, i-term substitution is applied to two pairs of i-terms T and T', and occurs prior to performing either intersection or composition. First, any i-term can be substituted (uniformly within a pair) for a variable i-term. In the case of intersection, if one i-term in T is the result of a substitution with another i-term in T', then restrictions on correlation relationships constrain the value of the other i-term in T to be identical to value of the other i-term in T'. For example, intersection can be performed on relations $\Pi_{c_1(x),x}b$ and $\Sigma_{a,c}b, m$, by making the appropriate substitutions; the result of the intersection is the refinement $\Pi_{c_1(x),x}b$. In the case of composition between $OP_{t,t'}R$ and $OP'_{t'',t'''}$, for the result to be defined it is necessary either that t' and t'' be the same term, or that one of them is a variable. The examples in the next section provide other illustrations of the substitution process.

5 Consistency In N-interval Networks

A knowledge base of assertions about n-intervals can be represented as a network of binary relations defined as follows:

Definition 6. A TV-interval relation network (NRN) is a network of binary relations whose variables (nodes) represent n-intervals and whose arcs link pairs of nodes. The arcs are labeled with a single element of NCR.

Example 1. I always Overlaps J. J sometimes occurs before K. K always meets L. L is sometimes before I.

The NRN representation is found in Figure 1 [7]. The figure shows the edges added as the result of composition involving nodes I, J, and K and K, L and J. The knowledge base is consistent.

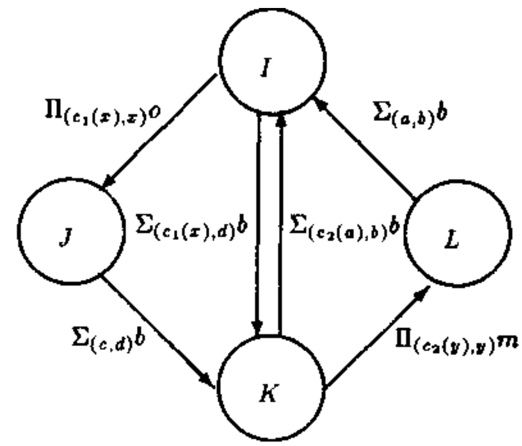

Figure 1. N-Interval Network For the Specification of Example 1.

Example 2. I is sometimes during J. Some (other) times, I meets J. The same J sometimes occurs before K. K is never after I. The network representation is found in Figure 2. The knowledge base is inconsistent. Composition will establish that I is sometimes before K; but taking the inverse of the arc between K and I will establish that I should *never* be before K.

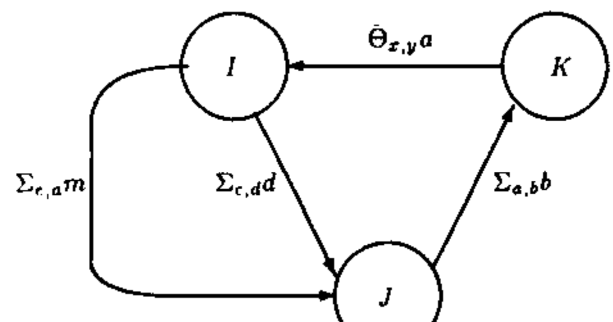

Figure 2. N-Interval Network for the Specification of Example 2.

Determining path consistency for NRN networks involves:

[7] Not all the information has been represented; for example, the inverse relation between J and I (only overlapped by) has not been represented. In addition, pairs of n-intervals for which no initial relation has been specified (e.g. between I and K) can be joined by an arc indicating this fact. Let ACR be the set of all atomic Allen relations; then complete lack of information in an n-interval network can be depicted by the relation $I_i \Sigma_{x,y} ACR_{i,j} J_j$. (In English: some part of I_i has some temporal relation to some part of J_j).

1. Possible additions to the number of apparent edges in the network as the result of composition; hence
2. Dealing with possibly multiple arcs between the same two vertices; (i.e., an NRN is a multigraph)\ and
3. Two kinds of relation composition (0 and o), intersection, inverse, complement, and i-term substitution.

For an n-interval relation network, the worst-case behavior of a path consistency algorithm is determined by both the number of nodes and the number of edges, both the edges initially specified, and those added as the result of composition. The latter can be estimated by considering the number of distinct values of x and y among the ^relations, as the following proof sketch demonstrates.

Theorem. Path consistency for NRN networks can be determined in $0(k^3m^3)$ time, where k is the number of distinct values of initial xy-indices, and m is the number of nodes in the network.

Proof (sketch): In an m-node NRN there are m(m—1)/2 pairs of nodes. For each pair (J, J) there are up to $0(k^2)$ distinct xy-indexed temporal relations (arcs) from / to J. Given a node pair (I, J) and an arc A from I to J there are m—2 intermediate nodes K where composition can be used to form an alternate path of length 2 from I to J. For this node K there are up to a maximum of k edges from / to A' that could qualify for composition. Since the composition operation itself has constant time complexity, the time complexity of the path consistency algorithm on an NRN is $0\{k^3m^3)$.

Contrasting this result with the complexity of path consistency for convex interval networks ($0(rn^3)$ [Van Beek, 1990]), it should be noted that because clustering of convex intervals is involved, there will in general be fewer variables in an NRN than in a convex interval network. In addition, k will tend to be small in practice.

6 Related Work

The research presented here extends current efforts in developing interval-based systems for reasoning about time by allowing for an explicit representation of intervals with gaps. The purpose of this extension is to expand the set of properties and relations that can be attributed to periods of time for applications such as static scheduling, natural language processing, and temporal database querying.

The impetus for this work was provided by Peter Ladkin's initial discussion [Ladkin, 1986], and by advances made by Ligozat [Ligozat, 1990]. Our work, as noted, has been also inspired by the success of constraint-based approaches to time [Mehri, 1991]. Finally, we were motivated in part by issues raised and discussed in [Koomen, 1989] which led to similar ideas of clustering intervals and their relations as a heuristic to aid in the reasoning process. The results here should also be compared to earlier work by two of the authors [Morris and Al-Khatib, 1991].

7 Summary

The problem of reasoning about certain kinds of incomplete knowledge about time was discussed. The mechanism chosen for the solution uses a non-convex interval representation, and a class of non-convex relations constructed by considering only groups of correlated convex intervals within a single temporal block. It was demonstrated that path consistency techniques can be applied to problem solving involving networks of n-interval relations.

References

[Allen, 1983] James Allen. Maintaining Knowledge About Temporal Intervals. In Ronald J. Brachman and Hector J. Levesque, editors, Readings in Knowledge Representation, pages 510-521, Morgan Kaufman Publishers, San Mateo California, 1985.

[Kautz and Ladkin, 1991] Henry A. Kautz and Peter B. Ladkin. Integrating Metric and Qualitative Temporal Reasoning. Proceedings of the Ninth National Conference on Artificial Intelligence, pages 241-246, Boston, Massachusetts, August 1990. American Association for Artificial Intelligence.

[Freuder, 1978] Eugene C. Freuder. Synthesizing Constraint Expressions. Communications of the ACM, 21(1978), pages 24-32.

[Koomen, 1989] Johannes A. G. M. Koomen. Localizing Temporal Constraint Propagation. Ronald J. Brachman, et. al., editors. Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning, KR'89, pages 198-202, Toronto, Ontario, Canada, May, 1989. Morgan Kaufman Publishers, San Mateo, California.

[Ladkin, 1986] Peter B. Ladkin. Time Representation: A Taxonomy of Interval Relations. Proceedings of the Third National Conference on Artificial Intelligence, pages 360-366, Philadelphia, Pennsylvania, 1985. American Association for Artificial Intelligence.

[Ligozat, 1990] Gerard Ligozat. Weak Representation of Interval Algebras. Proceedings of the Eighth National Conference on Artificial Intelligence pages 715-720, Boston, Massachusetts, August 1990. American Association for Artificial Intelligence.

[Mehri, 1991] Combining Qualitative and Quantitative Constraints in Temporal Reasoning. In Proceedings of the Ninth National Conference on Artificial Intelligence, Anaheim, California, July, 1991, pages 260-265. American Association for Artificial Intelligence.

[Morris and Al-Khatib, 1991] Robert A. Morris and Lina Al-Khatib. An Interval-Based Temporal Relational Calculus For Events With Gaps. The Journal of Experimental and Theoretical Artificial Intelligence, 3(1991), pages 87-107.

[Van Beek, 1990] . Peter Van Beek. Exact and Approximate Reasoning about Qualitative Temporal Relations. Technical Report TR 90-29, University of Alberta, Edmonton, Alberta, Canada.