Binary String Relations: A Foundation for Spatiotemporal Knowledge Representation

Delis Vasilis[1,2] and Hadzilacos Thanasis [2]

[1]Computer Engineering and Informatics Dept.
University of Patras, Greece

[2]Computer Technology Institute
Patras, Greece

tel: +3061 273496, 221834 fax: +3061 222086

{delis, thh}@cti.gr

ABSTRACT

The paper is concerned with the qualitative representation of spatiotemporal relations. We initially propose a multiresolution framework for the representation of relations among 1D intervals, based on a binary string encoding. We subsequently extend this framework to multiple dimensions, thus allowing the description of spatiotemporal relations at various contexts. The feasible relations at a particular resolution level are inherently permeated by a poset structure, called *conceptual neighbourhood*, upon which we propose efficient relation inferencing mechanisms. Finally, we discuss the application of our model to spatiotemporal reasoning, which refers to the classic problems of satisfiability and deductive closure of a set of spatiotemporal assertions.

Keywords

Spatiotemporal relations, spatiotemporal reasoning, conceptual neighbourhoods.

1 INTRODUCTION

Spatiotemporal information is a first class citizen in modern DBMS. For the effective manipulation of such entities and relationships we need appropriate representation and reasoning frameworks, which both have traditionally been two important research areas in the DB and AI communities. A major focus of AI research has been the design of sound and complete temporal and spatial reasoning systems, while DB research has been concerned with the incorporation of such systems as vital components in query languages. This fruitful interaction between the two fields has yielded an abundance of models and frameworks.

Qualitative representation of spatial knowledge, has dominated the spatial DB literature (see [29], [8] for thorough treatments), mainly due to the following factors:

- the input/output of spatial processes is often qualitative in nature, while the precision of quantitative representations is not always desirable.
- the processing of qualitative knowledge is cheaper.
- such representations facilitate spatial query processing, as human spatial reasoning is qualitative in nature.

All the available models, like *2D strings* for symbolic image encoding [7], *symbolic arrays* [28] for hierarchical symbolic image representation, etc. in essence concern the representation of *spatial relations* among spatial entities, which are often categorised into *topological*, *directional*, and *qualitative distance* relations. In particular, the most widely adopted among such models concern the encoding of *binary* spatial relations, like the *four-* and *nine-intersection models* for topological relations [9], the logic-based *RCC-calculus* [30] for the formal definition of topological relations between regions, *cardinal direction* models [10], to name just a few common ones. Potential applications include computer vision [4], Geographic Information Systems and Image Databases [29, 31], motion planning [22], image and multimedia similarity [25, SK97], etc. (see [29] and [20] for extensive comparative reviews on different representation and reasoning frameworks). The main drawback of the above models is their relatively limited flexibility. They all allow for the representation of fixed sets of relations, thus restricting the potential range of applications. To our knowledge, no existing model provides varying granularity levels in the description of relations. Moreover, querying and reasoning mechanisms that exploit such models usually focus on a subset of spatial relations (like topological and directional, directional and distance or hybrid combinations thereof) [5, 25, 31], limiting again their applicability. No existing model offers a unified treatment of all three types of spatial relations.

On the other hand, temporal representation and reasoning is an important recurring research topic mainly in AI community (see for example [32], [3], [16]) with several applications like natural language processing [2], planning [14], plan recognition [15], diagnosis [26], knowledge-based systems [17], etc. Undoubtedly, Allen's temporal interval algebra [1] was the most influential work. A few researchers [23, 12] have suggested that this algebra could be extended to represent spatial relationships. Such attempts again have limited applicability as they allow the definition of fixed relation sets. It seems that a uniform relation model satisfying the following requirements is elusive:

- generic enough to capture both spatial and temporal aspects in the same framework.
- formal, yet practical (easy to understand and use, providing efficient inferencing and query processing).
- allowing the description of relations at varying levels of detail, thus flexible enough to meet the needs of many applications.

In this paper we propose a unified relation model that satisfies the above. We adopted two assumptions: i) there is no a priori complete set of spatiotemporal relations (there is *north-east* between *north* and *east*, while *overlap* could be further refined to *strong* or *weak*, etc.), and ii) the primitive element underlying any qualitative representation of space(time) should be the ability to discern between empty or non-empty spatial(temporal) occupancy. So we defined *resolution schemes* as models for the representation of 1D relations at tuneable detail levels. Such a relation is determined by the emptiness or not of a set of intersections between an interval and some predefined regions of interest. The distinguishable relations at a particular resolution scheme form complete relation sets which have an inherent poset structure, whose lattice is called *conceptual neighbourhood* and has the property that similar relations are closer to each other than non-similar ones. This framework for representing 1D relations can be easily extended to ND, allowing us to model and reason on spatial and spatiotemporal relations (e.g. in 2D one can define a complete relation set comprising the topological, directional and distance relations allowed at the particular resolution scheme) between N-dimensional rectangular approximations of the actual database objects. The above are presented in Section 2, which draws from material from [PMD98]. In that paper the relation framework was informally introduced and used in combination with spatial indexing methods for the efficient processing of spatial queries. However, in this paper the relation model is rigorously formulated for the first time (resolution schemes and conceptual neighbourhoods are formally defined). A good understanding of the model is essential for the reader in order to address the second contribution of this work, which is the application of our model to spatiotemporal reasoning, which is treated in Section 3.

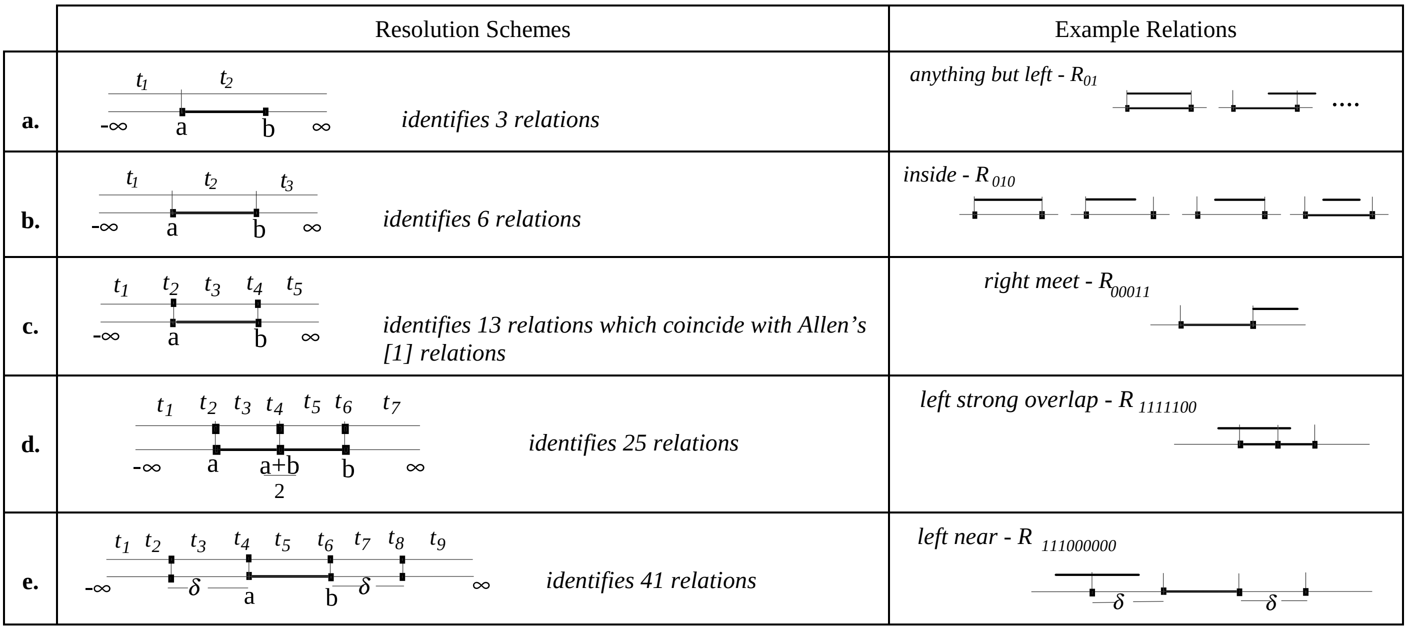

Figure 1 Example Resolution Schemes

The distinguishable relations at a particular resolution scheme allow for the definition of relation algebras which are Boolean algebras augmented with the operators of converse and composition, for the computation of which algorithmic procedures are described. On the negative side is that the deductive closure of a set of assertions in such algebras (and thus their satisfiability) is NP-complete, which is already known for Allen's algebra. On the positive side is that these problems are polynomial for practically useful restrictions of the algebras, which can be elegantly characterised in our framework.

Section 4 concludes by summarising the results of this work.

2 BINARY STRING RELATIONS

We shall initially focus on 1D and consider probably the most primitive way for representing relations, by defining *reference points*. The simplest model with a single reference point, already allows for the definition of only a few 1D relations (*left*, *right*, etc.). At the other extreme we can have an infinite set of reference points possibly regularly distributed along the 1D axis (e.g. a coordinate grid); this would allow for an infinite number of distinct relations, which is too much for qualitative considerations. In between the two extremes, a good starting point for discussing 1D relations is to define an ordered pair of two reference points, which corresponds to a convex interval (a point-set is convex if the points of a line segment that connects any two points belong to the set). We shall only consider relations between convex intervals with non-zero finite duration (zero-length intervals obscure the model's uniformity and complicate some of its properties and proofs, so they are not treated in this work).

Given such a *reference* interval $[a,b]$, we can identify five potential regions of interest: 1.$(-\infty,a)$ 2.$[a,a]$ 3.(a,b) 4.$[b,b]$ 5.$(b,+\infty)$. Given an arbitrary interval $[c,d]$ (which we call *primary*), its relation to $[a,b]$ can be determined by the five intersections of $[c,d]$ with each of the above five regions, the emptiness of which are modelled by five binary variables t_1, t_2, t_3, t_4, t_5, ("0" corresponds to an empty intersection while "1" corresponds to a non-empty one). For example, 11000 corresponds to the *left-meet* relation (—__). The different relations that can be defined in this way coincide with Allen's [1] interval relations (*before, after, overlap, during, starts, finishes,* their six converse relations and *equals*). The choice of regions of interest can be arbitrary.

Let $Y=<X_1, X_2, \ldots, X_n>$ be a partition of $(-\infty,\infty)$. Given an interval X, let $X.l$ and $X.r$ represent its left and right endpoint (regardless of closure at endpoints), respectively (in case X has zero length, $X.l=X.r$). Also, let $seq(Y)=<x_1, x_2, \ldots, x_m>$ be the ordered sequence of the elements of $\{X_1.l, X_1.r, X_2.l, X_2.r, \ldots, X_n.l, X_n.r\}-\{-\infty,\infty\}$ (e.g. $seq((-\infty,a), [a,a], (a,b), [b,b], (b,+\infty)) = <a, b>$).

Definition 1. Assume a reference interval $[a,b]$. A *resolution scheme* (or simply *scheme*) $Y(a,b) = <X_1, X_2, \ldots, X_n>$ is a partition of $(-\infty,\infty)$, where $seq(Y(a,b)) = <x_1, x_2, \ldots, x_m>$. For any reasonably practical scheme a and b must be included in $seq(Y(a,b))$ i.e. $\exists k\ \exists p$, $k<p\leq m$: $x_k=a$, $x_p=b$. The following also hold:

1. $\forall i\ i<k$, $x_i=a-c_i$ or $x_i=a-\lambda_i(b-a)$, where c_i, $\lambda_i > 0$ and $c_1>c_2>\ldots>c_{k-1}$ and $\lambda_1>\lambda_2>\ldots>\lambda_{k-1}$
2. $\forall i$, $k<i<p$, $x_i=a+c_i$ or $x_i=b-d_i$ or $x_i=a+\lambda_i(b-a)$ or $x_i=b-\mu_i(b-a)$, where c_i, λ_i, d_i, $\mu_i > 0$, λ_i, $\mu_i < 1$ and $d_{k+1}>d_{k+2}>\ldots>d_{p-1}$ and $\mu_{k+1}>\mu_{k+2}>\ldots>\mu_{p-1}$ and $c_{k+1}<c_{k+2}<\ldots<c_{p-1}$ and $\lambda_{k+1}<\lambda_{k+2}<\ldots<\lambda_{p-1}$
3. $\forall i$, $i>p$, $x_i=a+c_i$ or $x_i=a+\lambda_i(b-a)$, where c_i, $\lambda_i > 0$ and $c_{p+1}<c_{p+2}<\ldots<c_m$ and $\lambda_{p+1}<\lambda_{p+2}<\ldots<\lambda_m$

A relation R over Y is a binary string $t_1t_2\ldots t_n$, $t_i\in\{0,1\}$, complying with the following constraints: i) it contains exactly one substring of consecutive "1"s and ii) there is at least one t_i="1", in a position i such that X_i is a non-zero length interval. Given a primary interval $[x, y]$, $R([x, y], [a, b]) = t_1t_2\ldots t_n$, where t_i="1" iff $X_i\cap[x,y]\neq\emptyset$ and t_i="0" otherwise, i=1..n. We shall equivalently use the notation $R_{01100}([x, y], [a, b])$ to indicate that $R([x, y],[a, b])$=01100.

The intuition behind the above is the following: In order to distinguish several *left* or *right* relations with respect to a reference interval, one can establish appropriate regions of interest by defining points at either absolute distances from the interval ($a-c$, $b+c$) or at distances relative to the interval's length ($a-\lambda(b-\alpha)$, $b+\lambda(b-\alpha)$). The same holds for the interval's interior points, although in this case the choice of relative distances is likely to be more useful in practice as opposed to absolute values (e.g. "*strong overlap* is at least 50% coverage" vs. "*strong overlap* is coverage of at least 5 seconds").

For example, the resolution scheme $\{(-\infty, a-100), [a-100, a-100], (a-100, a-50), [a-50, a-50], (a-50, a), [a, a], (a, a+1/3(b-a)], (a+1/3(b-a), a+2/3(b-a)), [a+2/3(b-a), b), [b, b], (b, b+50), [b+50, b+50], (b+50, b+100), [b+100, b+100], (b+100, +\infty)\}$ uses symmetrical distances and can among others identify three distance relation levels, *very far* (>100), *far* (<100, >50), *close* (<50), and three overlap levels, *weak overlap* (up to 33.33% coverage), *medium overlap* (greater than 33.33% and less than 66.66% coverage) and *strong overlap* (at least 66.66% coverage).

According to the definition above it can be trivially proved that any relationship between two convex intervals can be uniquely expressed by a binary string relation at a particular resolution scheme and also that any valid relation string corresponds to a feasible interval relationship. The constraints imposed on the strings guarantee that no relation between a non-convex primary interval and the reference one can be expressed (the "1"s are consecutive) and also that zero length primary intervals are not allowed (by the obligatory non-empty intersection with at least one non-zero length region of interest).

The distinguishable relations at a particular scheme are called *primitive* relations. Figure 1 shows examples of a few more schemes which are discussed below.

– $(-\infty, a)$ $[a, +\infty)$: a very coarse scheme which distinguishes left disjoint relations from anything else (Figure 1.a - the figure portrays different interval relationships that correspond to the same example relation string).
– $(-\infty, a)$ $[a, b]$ $(b, +\infty)$: a coarse scheme where meet (at endpoints) relations can not be distinguished (Figure 1.b - the figure portrays different interval relationships that correspond to the same example relation string). This example shows that the endpoints of regions of interest need not necessarily constitute zero-length regions by themselves, unless meeting or not at the endpoints should be treated as two different relations. This is not the case in this scheme, as Figure 1.b shows.
– $(-\infty, a)$ $[a, a]$ $(a, (a+b)/2)$ $[(a+b)/2, (a+b)/2]$ $((a+b)/2, b)$ $[b, b]$ $(b, +\infty)$: a scheme which refines overlap relations (Figure 1.d).

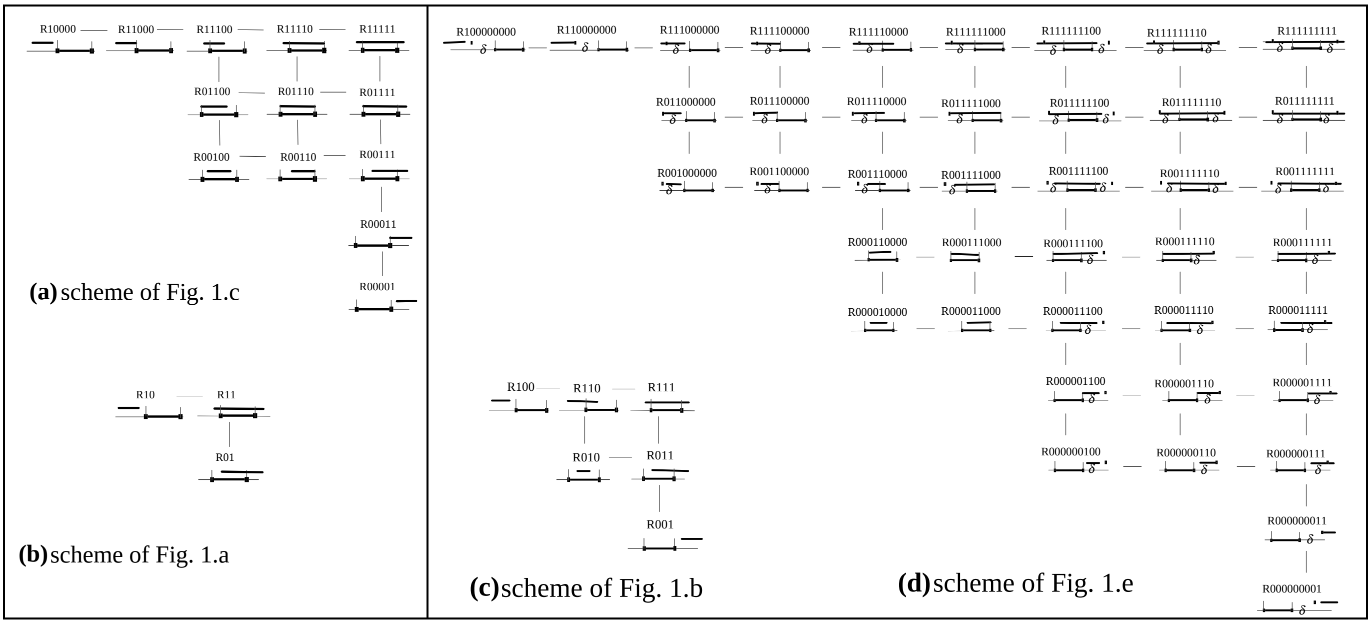

Figure 2 Example Neighbourhood Graphs

– $(-\infty, a-\delta)$ $[a-\delta, a-\delta]$ $(a-\delta, a)$ $[a, a]$ (a, b) $[b, b]$ $(b, b+\delta)$ $[b+\delta, b+\delta]$ $(b+\delta, +\infty)$: a distance enhanced scheme, where *near* is defined as being in a distance of up to δ and *far* otherwise (Figure 1.e). Left and right distances need not be symmetric (e.g. in applications to non-isotropic spaces).

A resolution scheme determines the detail level in the description of relations, which is affected by the number of bits associated with the scheme: the more detail we need to capture in a relation's representation, the greater the number of bits (regions of interest) we have to use. This refinement need not be uniform. For example one can focus on a particular relation class (e.g. overlap relations) and identify subclasses (like weak overlap, strong overlap, etc. - Figure 1.d) or subdivide the regions $(-\infty, a)$, $(b, +\infty)$, in order to capture qualitative distances (Figure 1.e).

In general, if n is the number of bits used by the resolution scheme, the number of distinguishable relations is $n(n+1)/2 - k$, where k is the number of bits assigned to zero-length regions of interest. In order to count the valid binary relation strings, we can fix the starting point at some bit then we can put the ending point at the same or some subsequent bit. There are n choices if we fix the first point to the leftmost bit, $n-1$ if we fix it to the second from the left, and so on. The total number is $n(n+1)/2$ from which we subtract the k invalid strings that contain exactly one "1" that corresponds to a zero-length region of interest. For Allen's scheme in Fig. 1.c ($n=5, k=2$) we get 13 relations, while for the distance enhanced scheme (n=9, k=4, see Fig. 1.e) we get 41 relations.

The advantages of the model so far are i) the simple representation of relations at varying resolutions levels and ii) its expressiveness, in the sense that given the new notation, the corresponding 1D relationship can be easily inferred, and vice versa.

2.1 Conceptual Neighbourhoods

Encoding the set of 1D relations at a particular scheme as binary strings gives it a poset structure. For two relations $X=x_1x_2\ldots x_n, Y=y_1y_2\ldots y_n$, $x_i, y_i \in \{0,1\}$,

$$X \le Y \text{ iff}$$

$$rm(X) \le rm(Y) \wedge lm(X) \le lm(Y)$$

where $rm(X)$, $lm(X)$ return the position of the rightmost and leftmost "1", respectively, in the binary string of a relation X. The resulting ordered structure is a distributive lattice. Every pair of primitive relations has a greatest lower bound and a least upper bound which are computed as follows:

$LUB(X,Y) = r_1r_2\ldots r_n$, where

r_i=1, for $max(lm(X), lm(Y)) \le i \le max(rm(X), rm(Y))$

0 otherwise, and

$GLB(X,Y) = r_1r_2\ldots r_n$, where

r_i=1, for $min(lm(X), lm(Y)) \le i \le min(rm(X), rm(Y))$

0 otherwise.

Intuitively, this partial order reflects the inherently assumed west-east direction of the 1D line. Thus, $X \le Y$ means that Y represents a configuration where the primary interval is somehow shifted to the right with respect to X's configuration. Several lattices for various schemes are depicted on Fig. 2, with a specific horizontal and vertical arrangement of edges (the greatest element is right-down while the least element is up-left).

A property of such lattice graphs is that similar relations are closer to each other than dissimilar ones (two relations are linked through an edge if they can be transformed to each other by a continuous minimal deformation of the primary interval). Freksa [13] defined the graph of Fig. 2a to capture the concept of similarity among Allen's relations, using the term *conceptual neighbourhood*, which we adopt in this paper for the lattices defined at any resolution scheme. Therefore, an obvious similarity measure for two relations is their distance in the corresponding neighbourhood graph, the calculation of which is facilitated by the binary encoding. We just have to count the minimal number of "0"s that have to be replaced with "1"s in order to make the two relation strings identical (essentially the number of minimal deformations that separate these relations). For example, $d(R_{11\underline{00}0}, R_{\underline{00}110})=4$ (the underlined "0"s are the ones to be replaced). The corresponding pseudo-code is given below.

```
INT distance(relation R1, relation R2)
R = R1 OR R2;  /*bitwise OR */
d = 0;
FOR i:= leftmost_1(R) to rightmost_1(R) DO {
        IF R1[i]=0 THEN d++;
        IF R2[i]=0 THEN d++; }
RETURN (d);
```

The neighbours of a relation *R* (termed according to their spatial arrangement in the neighbourhood graph) can be also efficiently computed: *right*(*R*) can be derived from *R* by inserting a "1" from the right, i.e., finding the first "0" after the rightmost "1" and replacing it by "1". *left*(*R*) can be therefore derived from *R*, by pruning a '1" from the right, i.e. replacing the rightmost "1" by a "0". Similarly, *up*(*R*) can be derived from *R* by inserting a "1" from the left while *down*(*R*) can be derived by pruning the leftmost "1". Of course a derived neighbour is accepted if it confronts to the string validity constraints. Since there are only four potential neighbours of a relation (just four possible minimal interval deformations: a left or right shift of either interval's endpoints), the above rules automate the construction of neighbourhood graphs. Given a resolution scheme one can start by any valid relation string and recursively construct the whole graph. This computation, as well as distance calculation are significant advantages against other relation models where neighbourhood graphs have to be defined manually and distances have to be pre-computed and stored in tables.

2.2 Extension to Higher Dimensions

The relation between two objects in *N* dimensions corresponds to the combination of the *N* 1D relations. Thus an *N*-dimensional relation can be naturally defined as a *N*-tuple of 1D relations between the projections of the ND objects on each of the dimensions, e.g. $R_{11000\text{-}11100} = (R_{11000}, R_{11100})$. Therefore, relations at each of the dimensions need not necessarily be defined at the same resolution scheme. Assuming the distance-enhanced scheme for the spatial dimensions, and Allen's scheme for the temporal one, Figure 3 illustrates a 3D composition, by giving the $R_{x\text{-}y\text{-}time}$ representations of some indicative relations between the five rectangles and the central gray one. Given an ND real world relationship, the corresponding ND relation can be derived by establishing the N 1D relations between respective 1D projections for each of the dimensions, and vice versa.

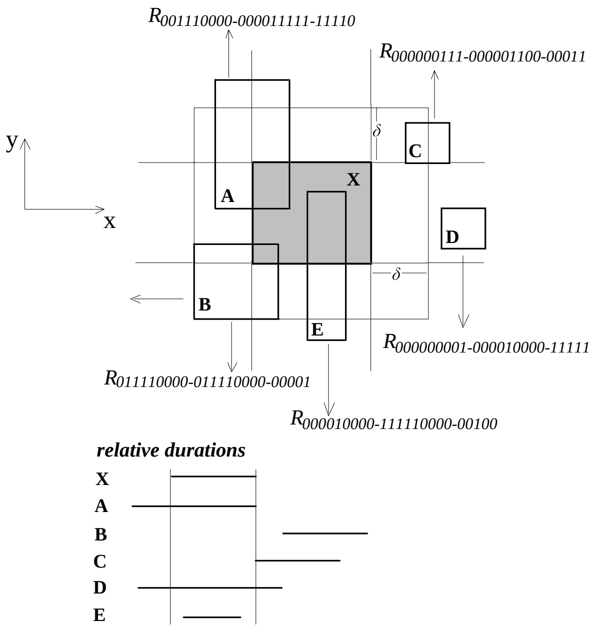

Figure 3 A Spatiotemporal Composition

There are $S_1 S_2 \ldots S_n$ distinct ND relations, where S_i is the number of distinguishable relations at the *i*-th dimension. The inference mechanisms described for 1D (distance and neighbour computation) can be analogously extended to ND. Since in 1D every relation has 4 potential neighbours, in ND every relation has maximum 4**N* neighbouring relations, derived by replacing each of its constituent 1D relations R_{ik} with its 1D neighbours, say $R_{ik}{}^{j}$, j=1..4:

$$neighbour(R_{i1\text{-}i2\text{-}\ldots ik\ldots\text{-}iN}) \in \{R_{i1\text{-}i2\text{-}\ldots ik^j\ldots\text{-}iN} \mid j=1..4,\ k=1..N\}$$

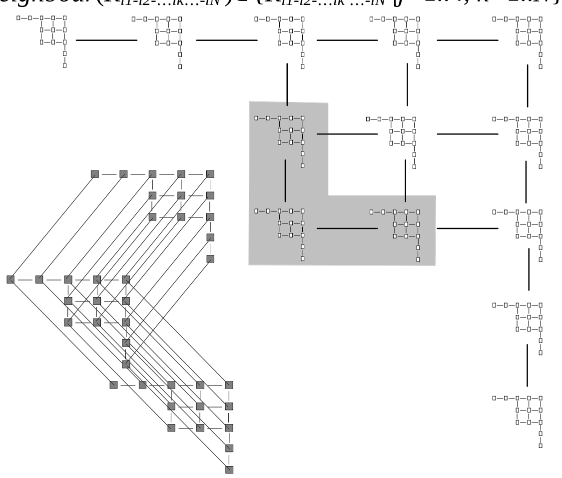

Figure 4 Example 2D Neighbourhood Graph

ND conceptual neighbourhoods are "fractal" graphs (graphs whose nodes are graphs, etc.). Figure 4 shows the 2D neighbourhood, where Allen's scheme is utilised for both dimensions. In this graph, 13 conceptual neighbourhoods corresponding to one dimension are linked, forming a higher level conceptual neighbourhood for the other dimension (each node in the big neighbourhood graph is a small neighbourhood graph). As the magnified section of the graph shows, in addition to the conceptual neighbours with the same *x* value, each relation R_{xj} has four more potential neighbours; i.e., the relations with the same *j* value and the *x* values related as shown in the higher level graph.

According to an ND neighbourhood graph, the distance between two ND relations is the sum of the pair-wise distances between the corresponding constituent 1D relations, i.e.
$d(R_{i1\text{-}i2\text{-}\dots\text{-}iN}, R_{j1\text{-}j2\text{-}\dots\text{-}jN})=d(R_{i1}, R_{j1})+d(R_{i2}, R_{j2})+..+d(R_{iN}, R_{jN})$.

3 SPATIOTEMPORAL REASONING

As with Allen's 13 relations, the *n* distinct relations at a particular resolution scheme allow for the definition of an interval relation algebra. The 2^n elements of the algebra are relations that may exist between convex 1D intervals, which are represented as sets of primitive relations, semantically corresponding to their disjunction. The *n* singleton elements represent *exact* relations, while any other set denote some degree of *uncertainty* in the definition of relations - e.g. the relation {10000, 11000} at the scheme of Fig. 1c corresponds to the temporal relationship *before or meets.* At the two extremes, the set consisting of all primitive relations (universal relation) denotes the "unknown" relation while ∅ corresponds to the empty relation.

Given a set of relations between interval variables (in other words, a set of assertions in the algebra), a fundamental reasoning problem is to determine their satisfiability. This problem is equivalent to determining the deductive closure of these assertions (in fact there is a polynomial mapping between them). For Allen's algebra these problems are shown to be NP-complete [33]. Thus, satisfiability for any interval algebra containing Allen's relation (algebras for every resolution scheme finer than the one in Fig. 1c) is NP-complete.

The above problems constitute the central theme of the *temporal reasoning* research topic, extensively studied in the AI community as a special case of *constraint satisfaction problems* (CSPs) [24], whose main concepts are described in the sequel.

3.1 Basic Constraint Satisfaction Concepts

A *network of binary relations* (or *binary constraint network* - BCN) is defined as a set of variables $X=\{X_1, X_2, \dots, X_n\}$, a domain D_i for each variable and a set of binary relations (constraints) between variables that indicate the allowed combinations of instantiations of the variables. An *instantiation* of the variables in *X* is an *n*-tuple $<x_1, x_2, \dots, x_n>$, representing an assignment of $x_i \in D_i$ to X_i. A *consistent instantiation* of a network is an instantiation of the variables such that the relations between variables are satisfied. A network is *inconsistent* if no consistent instantiation exists.

In our case the variables X_i are convex intervals (ordered pairs of real numbers) and the binary relations R_{ij} between the variables X_i and X_j are elements of an interval relation algebra at a particular resolution scheme. Such algebras provide the theoretical framework for the study of BCNs [18]. They are Boolean algebras with two additional operators:

- *relation converse* ($^{-1}$): $R(a,b)\equiv R^{-1}(b,a)$, where *R* a relation and *a, b* intervals.
- *relation composition* (*): $R_1(a,b)*R_2(b,c)\equiv R(a,c)$, where *R* is the least restrictive relation between *a* and *c* permitted by relations R_1 and R_2 (e.g., in Allen's scheme $R_{10000}(a,b)*R_{11111}(b,c)=R_{10000}(a,c)$).

This framework is easily extended to ND: a network of ND relations among ND objects is equivalent to *N* networks of 1D relations among intervals, each corresponding to the relations between the 1D projections of objects for each of the dimensions. If an instantiation of X_i and X_j satisfies R_{ij} then only one of the disjuncts R_k in R_{ij} is satisfied. Such networks can be obviously represented as labelled directed graphs whose nodes correspond to interval variables. The label C_{ij} of an edge (*i, j*) is the set of all the primitive relation disjuncts of R_{ij}. We trivially assume that C_{ii} is the equality relation.

An element $R_k \in C_{ij}$ is *feasible* if there exists a consistent instantiation of the network where R_k is satisfied. The set of feasible relations of a graph edge is called the *minimal labels.* The correspondence between constraint networks and temporal reasoning is clear. Computing the deductive closure of a set of assertions in an interval relation algebra is equivalent to finding all the minimal labels in a constraint network. The satisfiability of a set of assertions can be also reduced to the minimal labels problem (the set is satisfiable unless a minimal label is empty).

Mackworth [21] defined three consistency properties of constraint networks. A network is *node consistent* iff for every node *i*,

$$\forall\, x\, (x\in D_i) \rightarrow x\, R_{ii}\, x$$

We have already assumed that relation networks are node consistent. A network is *arc consistent* iff for every edge (*i, j*),

$$\forall\, x\, (x\in D_i) \rightarrow \exists y\, (y\in D_j) \wedge\, x\, R_{ij}\, y$$

which means that for every instantiation of X_i there exists an instantiation of X_j such that R_{ij} is satisfied. In our case, relation networks are always arc consistent. Finally, a network is *path consistent* iff for every triple (*i, k, j*) of nodes,

$$\forall x \forall z \;\; x\, R_{ij}\, z \rightarrow \exists y\, (y \in D_k) \wedge x\, R_{ik}\, y \wedge y\, R_{kj}\, z$$

which means that for every instantiation of X_i and X_j that satisfies the direct relation R_{ij}, there exists an instantiation of X_k such that R_{ik} and R_{kj} are also satisfied.

Freuder [11] generalised the above, defining *k-consistency*. A network is *k*-consistent if given any instantiation of any *k*-1 variables satisfying all the direct relations among them, there exists an instantiation of any *k*-th variable such as the *k* values satisfy all the direct relations among the *k* variables. According to this definition, node, arc, and path consistency correspond respectively to one-, two- and three-consistency. *Strong k-consistency* is *i*-consistency for all $i \leq k$. Notice that an ND relation network is consistent (under any consistency type) iff the *N* corresponding 1D networks are consistent.

Path consistency does not guarantee global consistency (strong *n*-consistency for a network of *n* intervals) for the full Allen's algebra, and thus for any relation algebra at a finer scheme, although it does so for useful subsets thereof, as we will see in the sequel. Path consistency is achieved by *constraint propagation*, a well known $O(n^3)$ in the number of intervals algorithm (see [1], [32]). The main idea of the algorithm is that a constraint C_{ij} on an edge (*i, j*) is propagated to any path of length two that the edge participates. Thus, the label on an edge (*k, j*) may be restricted by $C_{ki}*C_{ij}$, and so does the label on an edge (*i, k*) by $C_{ij}*C_{jk}$. This is applied recursively until no updates are possible, guaranteeing that for any $i, j, k \leq n$, $C_{ik} \subseteq C_{ij}*C_{jk}$.

3.2 Convex Relations

Although the full algebra at any practically useful resolution scheme is shown to be almost assuredly intractable (unless P=NP), we can restrict it to a still useful subclass and gain in complexity.

Definition 2: A *convex relation* (the term was coined by Ligozat [19] in a different framework though) is an interval in the relation lattice at a particular resolution scheme, i.e. a relation of the form [*R*, *S*] where *R*, *S* are primitive relations. Therefore, an ND relation is convex if its 1D constituent relations are convex.

Intuitively, convex relations capture a continuous uncertainty in our knowledge of a particular relationship. They are more reasonable approximations than arbitrary disjunctive elements of the full algebras, since in practical cases they are likely to arise due to inexact/ambiguous observations (e.g. when asking for directions, it is more likely to get an answer "the post office is two or three blocks away", than "the post office is two blocks or ten blocks away").

Viewing each interval as a point in the $x<y$ half-plane, a 1D relation in represented by a 0D, 1D or 2D region in the half-plane (see Fig. 5). Convex relations have an elegant geometrical characterisation, which is a convex point-set, as we prove in the following lemma.

Lemma 1: *Assume a reference interval [a,b], a primary interval [x,y] and a resolution scheme* $Y(a,b)=<Y_1,Y_2,\ldots,Y_m>$, *where* $seq(Y(a,b))= x_1x_2\ldots x_n$. *A relation R between [x,y] and [a,b] is convex iff its geometric representation can be defined by the conjunctive expression*

$$(a_1\, \sigma_1\, x\, \sigma_2\, a_2) \wedge (a_3\, \sigma_3\, y\, \sigma_4\, a_4) \wedge (x<y),$$

where σ_i *is either* $<$ *or* $\leq$, $a_i \in \mathbf{R}$ *and* $a_1<a_3$, $a_2<a_4$, $a_1 \leq a_2$, $a_3 \leq a_4$.

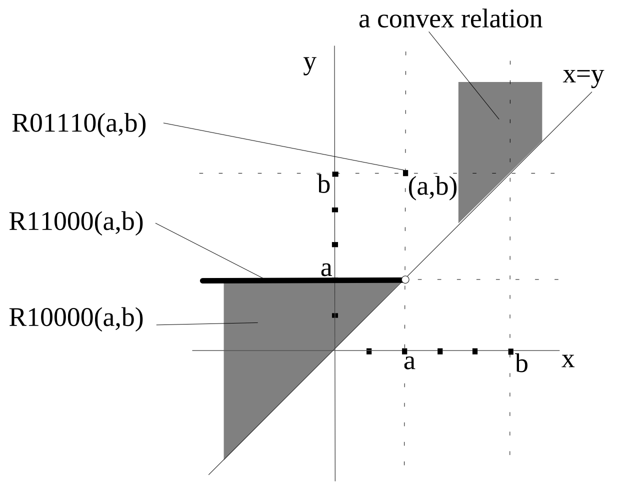

Figure 5 Geometric Representation of Interval Relations

Proof. The key idea is that, according to the above expression, *x* and *y* are constrained by convex interval domains, say *dom*(*x*) and *dom*(*y*). If *R* is convex then it is a relation interval $[R_1, R_2]$ at the corresponding lattice. Let Y_i,Y_j, be the leftmost and rightmost regions of interest whose corresponding bit is set to "1" in R_1 and Y_k,Y_l, be the leftmost and rightmost regions of interest whose corresponding bit is set to "1" in R_2. Since *R* can be any relation between R_1 and R_2, it can be expressed by the conjunction:

$$(Y_i.l\; \sigma_1\; x\; \sigma_2\; Y_k.r) \wedge (Y_j.l\; \sigma_3\; y\; \sigma_4\; Y_l.r) \wedge (x < y),$$

where σ_1, σ_2, σ_3, σ_4 are either $\leq$ or $<$, depending on whether Y_i, Y_k, Y_j, Y_l is a zero-length interval or not, respectively. By definition, $Y_i.l<Y_j.l$, $Y_i.l\leq Y_k.r$ (equality holds when $R_1=R_2$ and $Y_i=Y_k$ is a zero-length interval), $Y_j.l\leq Y_l.r$, $Y_k.r<Y_l.r$.

For the inverse, assume that *x* and *y* are constrained by *dom*(*x*) and *dom*(*y*), given respectively by the expressions $(a_1\; \sigma_1\; x\; \sigma_2\; a_2)$, $(a_3\; \sigma_3\; y\; \sigma_4\; a_4)$. Let again Y_i,Y_j, $i,j\leq m$, be the leftmost and rightmost regions of interest respectively whose intersection with *dom*(*x*) is non-empty, and Y_k,Y_l, $k,l\leq m$, be the leftmost and rightmost regions of interest respectively whose intersection with *dom*(*y*) is non-empty. Now consider the relation strings

$R_1=b_1b_2\ldots b_m$ where b_q="1" for $i\leq q\leq k$ and "0" otherwise,

$R_2=b_1b_2\ldots b_m$ where b_q="1" for $j\leq q\leq l$ and "0" otherwise. Since $a_1<a_3$, $a_2<a_4$, $a_1\leq a_2$, $a_3\leq a_4$, it follows that $i\leq j\leq k\leq l$.

Therefore, according to the definitions in Section 2, $R_1 \leq R_2$, while there are no other relations R_3, R_4 whose corresponding points could possibly belong to the region $(a_1\ \sigma_1\ x\ \sigma_2\ a_2) \wedge (a_3\ \sigma_3\ y\ \sigma_4\ a_4) \wedge (x<y)$, such that $R_3 \leq R_1$, $R_2 \leq R_4$. Also, the convexity of $dom(x)$, $dom(y)$ ensures that all the relations R, $R_1 \leq R \leq R_2$, are valid representatives of the conjunction. Thus, according to Definition 2, $R=[R_1, R_2]$ is convex. □

In order to be used for temporal reasoning problems, the class of convex relations has to qualify as a relation algebra, thus be closed under the operations of set theoretic union and intersection (which trivially holds as every convex relation is a set of primitive relations), relation converse and relation composition [18].

Lemma 2: *Convex relations are closed under converse.*
Proof. Assume a reference interval $[a,b]$, a primary interval $[x,y]$, a resolution scheme $Y(a,b)=\langle Y_1, Y_2, \ldots, Y_m \rangle$, where $seq(Y(a,b))=x_1x_2\ldots x_n$ and a primitive relation $R([x,y],[a,b])$. According to Lemma 1, R being convex implies that x, y belong to convex domains. The key idea is that while x, y belong to continuous ranges with respect to a and b, a and b belong to continuous ranges as well relatively to x and y, thus a and b have convex domains with respect to $Y(x,y)$. The following elaboration on this idea in fact outlines an algorithm for computing the converse of a convex relation. Let $R=[R_1, R_2]$, Y_i and Y_j be the regions of interest that correspond to the leftmost and rightmost "1" of R_1, respectively, and Y_k and Y_l be the regions of interest that correspond to the leftmost and rightmost "1" of R_2, respectively. Then, according to Lemma 1 R can be expressed by the conjunction

$$(Y_i.l\ \sigma_1\ x\ \sigma_2\ Y_k.r) \wedge (Y_j.l\ \sigma_3\ y\ \sigma_4\ Y_l.r) \wedge (x < y)$$

We also adopt the following notation: $\uparrow a$ indicates a value asymptotically converging to a from the left ($x<a$, $\lim x=a$), while $\downarrow a$ indicates a value asymptotically converging to a from the right ($x>a$, $\lim x=a$). We must also provide it with some common sense operations ($\uparrow a + \uparrow b = \uparrow(a+b)$, $\uparrow a + \downarrow b = (a+b)$, etc.)

We shall use this notation in order to define $Y(x,y)$ by establishing points as functions of a and b. According to whether σ_i represents a strict inequality or not, the minimum and maximum values for x and y are $Y_i.l$ or $\downarrow Y_i.l$, $Y_k.r$ or $\uparrow Y_k.r$, $Y_j.l$ or $\downarrow Y_j.l$, $Y_l.r$ or $\uparrow Y_l.r$, respectively. Let R_3 be the relation between $[a, b]$ and $[max(x), max(y)]$ and R_4 be the relation between $[a, b]$ and $[min(x), min(y)]$. Using arguments similar with the ones in the inverse part of the proof of Lemma 1, it follows that $R^{-1}=[R_3, R_4]$. □

In the example in Figure 6 we assume the resolution scheme of Figure 1.d and compute the converse of $R_{1110000}([x, y], [a, b])$, which is a primitive relation (a special case of convex relations).

Figures 6(i) and 6(ii) depict $Y(x,y)$ for the maximum and minimum values of x and y respectively. In the second case, where $min(x)$ tends to $-\infty$, we replaced $-\infty$ (since it can not participate in algebraic expressions) with $-M$ where M is a sufficiently large positive integer. This trick together with the asymptotic notation introduced above are gadgets whose purpose is to make the positions of a and b relatively to $seq(Y(x,y))$ (thus R^{-1}) clearly identified (e.g. if $x = \uparrow a$ then we know that $a>x$). For this example we see that $(R_{1110000})^{-1} = [R_{0011111}, R_{0000111}]$. An intuitive explanation is that although $[x, y]$ weakly overlaps $[a, b]$ from the left, $[a, b]$ can weakly overlap, extend exactly until $(x+y)/2$ or strongly overlap $[x, y]$ from the right, depending on the length and relative position of $[x, y]$.

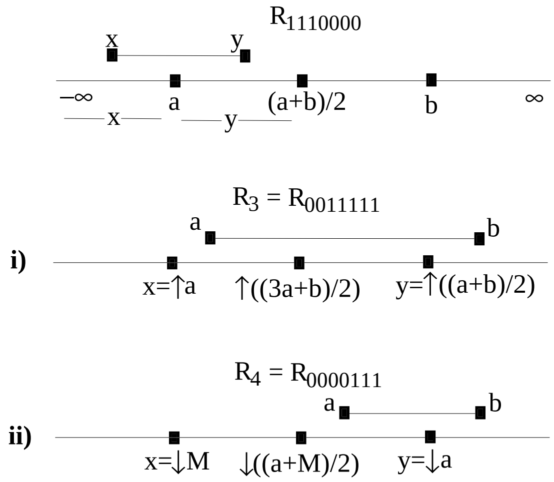

Figure 6 Example of Converse Computation

Lemma 3: *Convex relations are closed under composition.*
Proof. Assume the relations $R_1([a_1, b_1], [a_2, b_2])$ and $R_2([a_2, b_2], [a_3, b_3])$ at a resolution scheme Y. Convexity, according to Lemma 2, implies that $R_2^{-1}([a_3, b_3], [a_2, b_2])$ is convex as well. Therefore, according to Lemma 1, a_1, b_1, a_3 and b_3 are constrained by convex interval domains whose endpoints belong to $seq(Y(a_2, b_2))$. By assessing the relative positions of a_1 and b_1 with respect to a_3 and b_3 we can show, reasoning in a way completely analogous to the proof of the previous lemma, that $R_1 * R_2$ is the convex relation $[R_3, R_4]$ (see Fig. 7). □

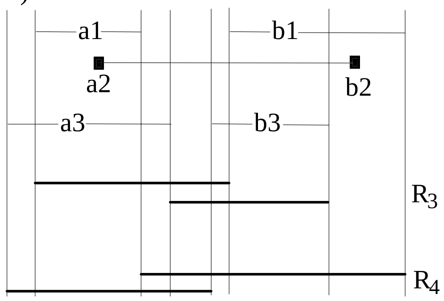

Figure 7 Example of Composition Computation

The proofs of the two previous lemmas indicate the algorithms for the computation of the converse and composition of convex relations, thus for primitive relations as well. Therefore these algorithms can compute converse and composition of any relations (sets of primitive relations): i) the converse of a set of primitive relations is the union of the individual converses, and ii) the composition of two sets of primitive relations is the union of the results of the pair-wise compositions of primitive relations, as composition distributes over union. The latter is a significant advantage against previous approaches, in which composition of relations has to be computed in advance by hand and stored in tables (see [13] for the 13×13 composition table for Allen's relations). In our case, this obviously would be extremely difficult if not impossible, even for moderately fine resolution schemes (the distance-enhanced resolution scheme would require 41^2 such computations). Also, converse and composition in ND is the result of the combination of the *N* 1D converses and compositions, respectively.

The following result is a generalisation of a result presented in [32] which proves tractability for a subset of Allen's algebra (all relations which can be expressed as conjunctive expressions where each conjunct relates any two of the four interval endpoints with one of the operators <, ≤, >, ≥, <=>, =), which coincides with the convex relations allowed in Allen's resolution scheme.

Theorem 1: *The satisfiability of the algebra of convex relations at any resolution scheme is polynomial.*

Proof. We will prove by induction on the number of nodes of a convex 1D relation network with *n* intervals that path consistency (thus the polynomial constraint propagation algorithm) guarantees strong *n*-consistency.

Basis. The claim is true for *k*=1, 2, 3, since convex relation networks are always arc consistent, trivially assumed to be node consistent, and path consistent according to the hypothesis, thus strong 3-consistent.

Inductive step. We assume strong (*k*-1)-consistency. This implies that any set of variables $X_1, X_2, \ldots, X_{k-1}$ can be consistently instantiated to values $[s_i, e_i]$ such that $[s_i, e_i]\ R_{ij}\ [s_j, e_j]$, $i, j=1, \ldots, k-1$. To prove *k*-consistency we must show that there exists at least one instantiation of any variable $X_k=[x_k, y_k]$ such that

$$[x_k, y_k]\ R_{ki}\ [s_i, e_i] \quad i=1..k-1 \quad (1)$$

Let *Y* be the resolution scheme. Since R_{ki}, $i=1, \ldots, k-1$, are all convex relations, according to Lemma 1 x_k and y_k are constrained by convex intervals whose endpoints are functions of s_i ane e_i. Thus, R_{ki} can be expressed as

$$(f_1(s_i, e_i)\ \sigma_1\ x_k\ \sigma_2\ f_2(s_i, e_i)) \wedge (f_3(s_i, e_i)\ \sigma_3\ y_k\ \sigma_4\ f_4(s_i, e_i)) \wedge (x_k<y_k),$$

where $f_j(s_i, e_i) \in seq(Y(s_i, e_i))$ and $\sigma_j \in \{<, \leq\}$, $j=1..4$, $i=1..k-1$.

Proving that there is at least one instantiation for $[x_k, y_k]$ that satisfies (1) is equivalent to proving that the intersection of the above 2(*k*-1)+1 ($x_k<y_k$ is repeated in every expression so it counts once) convex regions -which we call atomic- is non-empty. We make use of Helly's theorem [6] which states that given a family of at least *n*+1 convex sets in $\mathbf{R}^n$, all the sets have a point in common iff every *n*+1 sets have a point in common. In our case *n*=2, so it suffices to prove that any three atomic regions have a common point. There are several cases, depending on the corresponding three expressions.

<u>Case 1</u>

$f_1(s_i, e_i)\ \sigma_1\ x_k\ \sigma_2\ f_2(s_i, e_i)$

$f_3(s_j, e_j)\ \sigma_3\ x_k\ \sigma_4\ f_4(s_j, e_j)$

$f_5(s_h, e_h)\ \sigma_5\ x_k\ \sigma_6\ f_6(s_h, e_h)$, for some $i, j, h \leq k-1$, $\sigma_g \in \{\leq, <\}$, $g=1..6$. In this case only one variable is restricted by convex sets, so by Helly's theorem is suffices to show that any two such sets have a point in common. However, any two of the above three expressions have at least a point in common because they correspond to a triangle in the network (e.g. the first two ones correspond to the triangle (X_k, X_i, X_j)) for which there is a consistent instantiation due to path consistency, assumed in the hypothesis. This case is equivalent to three expressions restricting y_k.

<u>Case 2</u>

$f_1(s_i, e_i)\ \sigma_1\ x_k\ \sigma_2\ f_2(s_i, e_i)$

$f_3(s_j, e_j)\ \sigma_3\ x_k\ \sigma_4\ f_4(s_j, e_j)$

$f_5(s_h, e_h)\ \sigma_5\ y_k\ \sigma_6\ f_6(s_h, e_h)$, for some $i, j, h \leq k-1$. In this case, the atomic regions defined by the first two expressions have a common point, using arguments similar with the ones in the first case. In fact their intersection is a zone parallel to *y* axis which obviously intersects with the region defined by the third expression, which is a zone parallel to the *x* axis. This case is equivalent with two expressions restricting y_k and one expression restricting x_k.

<u>Case 3</u>

$f_1(s_i, e_i)\ \sigma_1\ x_k\ \sigma_2\ f_2(s_i, e_i)$

$f_3(s_j, e_j)\ \sigma_3\ x_k\ \sigma_4\ f_4(s_j, e_j)$

$x_k < y_k$, for some $i, j \leq k-1$. Similarly, the first two expressions define a zone which intersects with the $x_k<y_k$ half-plane. This case is equivalent with two expressions restricting y_k and $x_k<y_k$.

<u>Case 4</u>

$f_1(s_i, e_i)\ \sigma_1\ x_k\ \sigma_2\ f_2(s_i, e_i)$

$f_3(s_j, e_j)\ \sigma_3\ y_k\ \sigma_4\ f_4(s_j, e_j)$

$x_k < y_k$, for some $i, j \leq k-1$. Again, the first two expressions define regions which have at least a common point, as they correspond to the (X_k, X_i, X_j) triangle which is consistent due to path consistency. Obviously, the common point(s) belong to the $x_k<y_k$ half-plane. □

3.3 Reasoning in Multiple Dimensions

As with every property of ND relations, the satisfiability of an ND relation network is equivalent to *N* satisfiability problems, one for every 1D relation network that corresponds to the relations for each of the *N* dimensions. The ND relation network is satisfied if all the *N* corresponding 1D relation networks are satisfied. The significance of the previous theorem is obvious for *N*=2 (application to Geographic Information Systems), *N*=3 (application to reasoning with volumes or multimedia -viewed as 2D snapshots over time- entities), *N*=4 (application to reasoning with moving volumes).

One may notice that in several ND applications, *N*>1, overlapping relations may not be allowed for some of the dimensions. This is the case with cadastral applications, reasoning with solid volumes, etc. In such cases the only feasible 1D primitive relations at any resolution scheme are refinements of equal, before, after, meet and meet^{-1} relations. Even so, the curse of intractability is still with us. In [32] reasoning with the full algebra of Allen is proven to be NP-complete with the use of a reduction from graph colouring that utilises only equal, before, after, meet and meet^{-1} primitive relations. This fact intensifies the significance of convex relations which combine adequate expressiveness to be used in practice with tractability.

4 CONCLUSIONS

Spatiotemporal information is increasingly becoming an integral part of databases and knowledge bases. The need for its effective and efficient manipulation has motivated a lot of research efforts in the DB and AI areas. These have shown that formal frameworks for the qualitative representation of spatiotemporal relations can often be effective means to deal with certain spatiotemporal retrieval and reasoning problems.

Our work follows this path and proposes a unified framework for the definition of spatiotemporal relations, mainly inspired by two well established models, Egenhofer's intersection model [9] and Allen's interval algebra [1]. The former provided the assumption that relationships in space and time can be qualitatively described by the emptiness or not of intersections (coincidence of spatial and temporal occupancy) while the latter indirectly suggested that different amounts of detail can be captured in the description of a relation between intervals, depending on the definition of appropriate 1D regions of interest.

We formalised these ideas and developed a model for the representation of 1D relations at varying resolution levels. This is by itself an important advantage compared to previous approaches which have limited or no capabilities of tuning the relations' detail level. This 1D model, along with its efficient neighbourhood inference capabilities can be uniformly extended to ND, which renders it naturally applicable to spatiotemporal similarity problems (like image and multimedia similarity), where the queried content is the spatiotemporal arrangement among the objects in a spatiotemporal scene. This has been addressed elsewhere [PMD98] where it has been shown that our model can be seamlessly combined with available indexing techniques (R-trees in the particular study) to handle such problems efficiently.

Additionally, the framework is enriched with powerful reasoning capabilities. Two algorithmic procedures for the computation of the converse and composition of relations are described. The automatic computation of these operators is again an advantage against previous models in which relation inferencing through composition had to be manually defined. At last, a useful tractable class of relation algebras is identified, namely convex relation algebras, for which deductive closure (and equivalently, satisfiability) can be computed by the polynomial constraint propagation algorithm. Providing the above features, the relation reasoning framework could be successfully used as the core of a spatiotemporal inference engine.

Aknowledgements

Part of this work was done during a first author's visit to the Hong Kong University of Science and Technology, supported by DAG96/97.EG36 Hong Kong RGC, and also by a PhD studentship from Bodosakis Foundation. The cooperation with Dr Papadias during the visit contributed significantly to the shaping of these ideas.

REFERENCES

[1] Allen, J., "Maintaining Knowledge About Temporal Intervals", *CACM*, 26(11), 1983.

[2] Allen, J., "Towards a General Theory of Action and Time", *Artificial Intelligence*, 23, 1984.

[3] Allen, J., “Time and time again: The many ways to represent time,” *International Journal of Intelligent Systems*, 6(4), 1991.

[4] Ballard, D., Brown, C., "*Computer Vision*", Prentice Hall, 1984.

[5] Bruns, T.H., Egenhofer, M.J., "Similarity of Spatial Scenes", *7th International Symposium on Spatial Data Handling*, 1996.

[6] Chvatal, V., *"Linear Programming"*, W. H. Freeman.

[7] Chang, S.K., Shi, Q.Y., Yan, C.W., "Iconic Indexing by 2D Strings", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(3), 1987.

[8] Egenhofer, M.J., “Spatial Relations: Models and Inferences”, *Tutorial 2, International Symposium on Large Spatial databases (SSD'97)*, 1997.

[9] Egenhofer, M.J., Herring, J., “Categorizing Binary Topological Relationships Between Regions, Lines and Points in Geographic Databases”, *Technical Report*, Department of Surveying Engineering, University of Maine, Orono, ME, 1991.

[10] Frank, A., "Qualitative Spatial Reasoning about Distances and Directions in Geographic Space", *Journal of Visual Languages and Computing*, 3(4), 1992.

[11] Freuder, E.C., "Synthesizing Constraint Expressions", *CACM*, 21, 1978.

[12] Freksa, C., "Conceptual Neighbourhood and its Role in Temporal and Spatial Reasoning", *IMACS Workshop on Decision Support Systems and Qualitative Reasoning*, 1991.

[13] Freksa, C., "Temporal Reasoning based on Semi Intervals", *Artificial Intelligence*, 54, 1992.

[14] Hogge, J., "TPLAN: A Temporal Interval-Based Planner with Novel Extensions", *Technical Report UIUCDCS-R-87*, University of Illinois, 1987.

[15] Kautz, H. A., "A Formal Theory of Plan Recognition", *PhD Thesis*, *Technical Report 215*, University of Rochester, 1987.

[16] Kautz, H. A., "Temporal Reasoning", *MIT Encyclopedia of Cognitive Science (forthcoming)*, http://www.research.att.com/ ~Kautz/papers-ftp/index.html.

[17] Koubarakis, M., Mylopoulos, J., Stanley, M., Borgida, A., "Telos: Features and Formalisation", *Technical report KRR-TR-89-4*, Department of Computer Science, University of Toronto, 1989.

[18] Ladkin, P., Maddux, R., "The Algebra of Binary Constraint Networks", *Kestrel Institute Technical Report KES.U.88.9*, 1988.

[19] Ligozat, G., "On Generalised Interval Calculi", *International AAAI Conference*, 1991.

[20] Li, J., Ozsu, T., Szafron, D., "Spatial Reasoning Rules in Multimedia Management Systems", *International Conference on Multimedia Modelling (MMM)*, 1996.

[21] Mackworth, A.K., "Consistency in Networks of Relations", *Artificial Intelligence*, 8, 1977.

[22] McDermott, D., Davis, E., “Planning Routes Through Uncertain Territory”, *Artificial Intelligence*, 22, 1984.

[23] Mukerjee, A., Joe, G., “A qualitative Model For Space”, *International AAAI Conference*, 1990.

[24] Montanari, U., “Networks of Constraints: Fundamental Properties and Applications to Picture Processing”, *Information Science*, 7, 1974.

[25] Nabil, M., Ngu, A., Shepherd, J., "Picture Similarity Retrieval using 2D Projection Interval Representation", *IEEE TKDE*, 8(4), 1996.

[26] Nokel, K., "Temporal Matchning: Recognizing Dynamic Situations from Discrete Measurements", *11th IJCAI*, 1989.

[27] Papadias, D., Mamoulis, N., Delis, B. "Algorithms for Querying by Spatial Structure", *VLDB Conference*, 1998.

[28] Papadias, D., Sellis, T., "Spatial Reasoning Using Symbolic Arrays", *International GIS Conference - From Space to Territory: Theories and Methods of Spatio-temporal Reasoning in Geographic Space*", LNCS 639, 1992.

[29] Papadias, D., Sellis, T., "Qualitative Representation of Spatial Knowledge in Two-Dimensional Space", *VLDB Journal*, 3(4), 1994.

[30] Randell, D.A., Cui, Z., Cohn, A., "A Spatial Logic Based on Regions and Connection", *Third International Conference on Principles of Knowledge Representation and Reasoning*, 1992.

[31] Sistla P., Yu, C., Haddad, R., "Reasoning About Spatial Relationships in Picture Retrieval Systems", *VLDB Conference*, 1994.

[32] van Beek, P., Cohen, R., "Exact and Approximate Reasoning about Temporal Relations", *Computational Intelligence*, 6, 1990.

[33] Vilain, M., Kautz, H., van Beek, P., "Constraint Propagation Algorithms for Temporal Reasoning: A Revised Report", appears in *Readings in Qualitative Reasoning about Physical Systems*, (eds) de Kleer and Weld, Morgan Kaufmann, 1989.