Hierarchical Reasoning about Direction Relations

Dimitris Papadias
Artificial Intelligence Research Division, Institute for Applied Information Technology, GMD, D-53754 Sankt Augustin, Germany.
dimitris@nathan.gmd.de

Max J. Egenhofer
National Center for Geographic Information and Analysis, and Department of Spatial Information Science and Engineering, and Department of Computer Science, University of Maine, Orono, ME 04469-5711, USA.
max@spatial.maine.edu

Jayant Sharma
Oracle Corporation, 110 Spit Brook Rd, ZK02-1/019, Nashua, NH 03062, USA.
jsharma@us.oracle.com

ABSTRACT

Spatial reasoning is an important area of geographic information systems (GIS) and spatial databases research. This paper deals with reasoning about direction relations (east, northeast) in spatial hierarchies. We assume a database that stores the direction relations between objects in the same geographic region and propose algorithms for the inference of relations between objects in different regions. We present two types of inference: the first one uses the relations of ancestor regions in the hierarchy, while the second one is based on compositions of spatial relations and path consistency. For both types we provide inference rules, illustrate examples, and study the computational complexity. Although we use a specific set of relations for demonstration purposes, the algorithms are applicable to any set of direction relations provided with appropriate inference rules.

1. INTRODUCTION

The hierarchical representation of space has a strong psychological motivation (Hirtle and Jonides, 1985) and numerous computational advantages that have been exploited in a number of areas including data structures (Guttman, 1984; Samet, 1989) and wayfinding (Car and Frank, 1994). This paper studies hierarchical reasoning about direction relations in spatial databases that cannot resort to coordinate-based representations. Such situations are typically found in narratives, trip reports, and scientists' field notes. We assume that there only exists relative information about the objects within a region and inclusion relations (i.e., the hierarchical structure). The goal is to infer the direction relations between objects located in different regions and to detect potential inconsistencies.

As an example consider that you have the information that location X_1 is east of X_2 in the map of region A_1, and that X_2 is east of X_3 in the map of (neighbouring) region A_2. In addition you learn that X_3 belongs also to region A_3, which is northeast of A_1. The above data contain an inconsistency about the relation between X_1 and X_3: from their relation with X_2, we can infer through transitivity that east(X_1,X_3); using the relation of their ancestor regions (northeast(A_3,A_1)) we can also infer that northeast(X_3,X_1) which contradicts the previous inference. Such inconsistencies may occur by combining spatial knowledge from different sources and alternative representations like images, topographic surveys or verbal descriptions (for an extended discussion see Frank, 1992). Spatial inference mechanisms are essential for explicating relations and enforcing consistency.

The rest of the paper is organized as follows: Section 2 defines direction relations between points and regions, and describes spatial databases preserving directions. Section 3 discusses the retrieval of explicit relations and outlines a framework for the computation of the cost. Section 4 presents an algorithm for the inference of the relation between points using the relations of their ancestor regions. Section 5 describes a complementary form of inference that uses chains of common points and achieves path consistency for the whole database. Section 6 concludes with comments.

2. DIRECTION RELATIONS AND SPATIAL DATABASES

Various types of direction relations have been used to match different needs that range from cognitive modelling (Herskovits, 1986) to image similarity retrieval (Lee et al., 1992) and from navigation (Holmes and Jungert, 1992) to user interfaces (Roussopoulos et al., 1988; Papadias and Sellis, 1995). Although in this paper we use a set of *projection-based* definitions, the proposed methods are applicable to other types of directions (for a discussion on alternative types see Frank, 1992; Hernandez, 1994).

According to the projection-based model, the relation between two points is determined by the position of the primary object with respect to the projection lines from the reference object to the coordinate axes (Freksa, 1992; Papadias and Sellis, 1994; Nabil et al. 1995). In this way, nine mutually exclusive relations can be defined between points:

NorthWest: $NW(P_1,P_2) \equiv (X(P_1)<X(P_2)) \wedge (Y(P_1)>Y(P_1))$,

RestrictedNorth: $RN(P_1,P_2) \equiv (X(P_1)=X(P_2)) \wedge (Y(P_1)>Y(P_1))$,

NorthEast: $NE(P_1,P_2) \equiv (X(P_1) > X(P_2)) \wedge (Y(P_1)>Y(P_1))$,

RestrictedWest: $RW(P_1,P_2) \equiv (X(P_1)<X(P_2)) \wedge (Y(P_1)=Y(P_1))$,

SamePosition: $SP(P_1,P_2) \equiv (X(P_1)=X(P_2)) \wedge (Y(P_1)=Y(P_1))$,

and the converse relations: $RE(P_1,P_2) \equiv RW(P_2,P_1)$, $SW(P_1,P_2) \equiv NE(P_2,P_1)$, $RS(P_1,P_2) \equiv RN(P_2,P_1)$, $SE(P_1,P_2) \equiv NW(P_2,P_1)$.

U denotes the *universal* relation (the disjunction of all primitive relations) and $\varnothing$ the *empty* relation (the relation that arises during inconsistencies). The above relations form a relation algebra and can be used for relation-based reasoning. They constitute the set of *high resolution* relations; we also define a set of *low resolution* relations using disjunctions: $N = NW \vee RN \vee NE$, $E = NE \vee RE \vee SE$, $S = SW \vee RS \vee SE$, $W = NW \vee RW \vee SW$, $SL = RW \vee SP \vee RE$ (SameLevel), and $SH = RN \vee SP \vee RS$ (SamewidtH). The projection-based definitions are applied for regions in an analogous way. There are 13 mutually exclusive relations between intervals in 1D space (Allen, 1983). If we extend to 2D space we get the 169 primitive relations between regions of Figure 1.

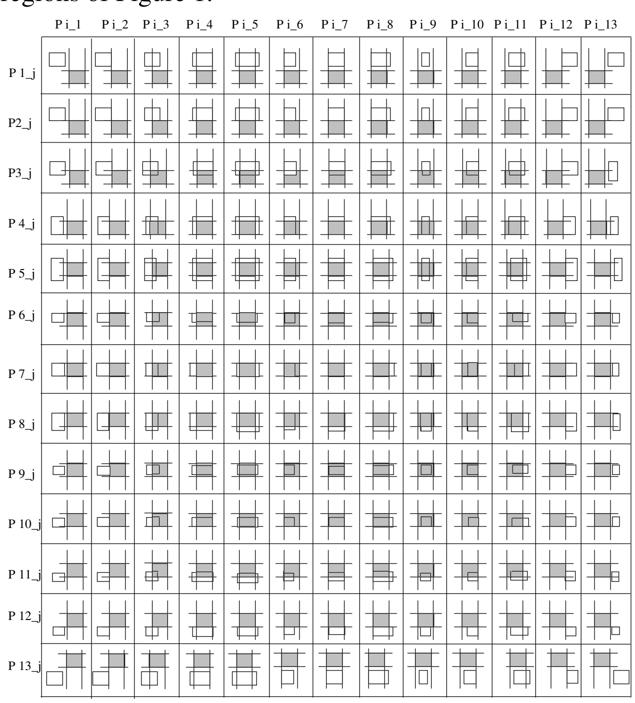

Figure 1 Projection relations in 2D space

A number of *relation-based* systems that store only the above type of relations and discard other forms of spatial information (such as shape, distance and topological relations) have been proposed. Chang et al. (1987) designed *2D strings* for iconic indexing in image databases. A 2D string is a pair of one-dimensional strings that represent the symbolic projections of the objects on the x and y axis. Glasgow and Papadias (1992) developed *symbolic arrays*, which are nested array structures that preserve direction relations among the distinct parts of complex spatial entities at different levels. Most previous work, however, has focused on the representation and processing of explicit relations and the proposed systems do not include mechanisms for inference and inconsistency checking.

Let *DB* be a spatial database of maps (or more generally, spatial representations) each corresponding to a distinct region. For every map there is a relation-based representation (2D string, symbolic array, a relational table, or a set of binary predicates) that stores the relations between all pairs of objects in the region. The objects in the map can be either points or regions, but not both (the regions that contain points are called *leaf regions*). Each pair of objects in a map is related by a primitive direction relation explicitly represented.

We use the notation $A \vdash R(X_1,X_2)$ to express that objects X_1 and X_2 are related by relation R in the map of region A. The hierarchy is represented by pointers to next-level areas (IN relation). $DB \vdash IN(X_i,A_j)$ denotes that object (point or region) X_i is a part of (therefore, totally contained in) the next level region A_j. IN* is the transitive closure of IN: $DB \vdash IN^*(X_i,A_j) \equiv DB \vdash IN(X_i,A_j) \vee \exists A_k [DB \vdash IN^*(X_i,A_k) \wedge DB \vdash IN^*(A_k,A_j)]$. IN* needs not be explicitly represented, but can be computed by a recursive function that traverses the hierarchy bottom-up and marks all the ancestors of an object in the hierarchy. For demonstration, we use the example of Figure 2a; Figure 2b illustrates the hierarchy and the relations explicitly represented.

R, R_1, R_2 ... denote relation variables between points, and r, r_1, r_2 ... between regions. The general problem is to retrieve the (explicit or implicit) relation between any pair of points P_i and P_j in the database: $DB \vdash R(P_i,P_j)$. There are three cases regarding the direction relations between points: *explicit retrieval*, *inference through regions,* and *inference through points*. In the next sections we discuss algorithms that extract the relation between all pairs of points and detect inconsistencies. For each case we provide rules of inference, describe examples, and obtain formulas for the cost.

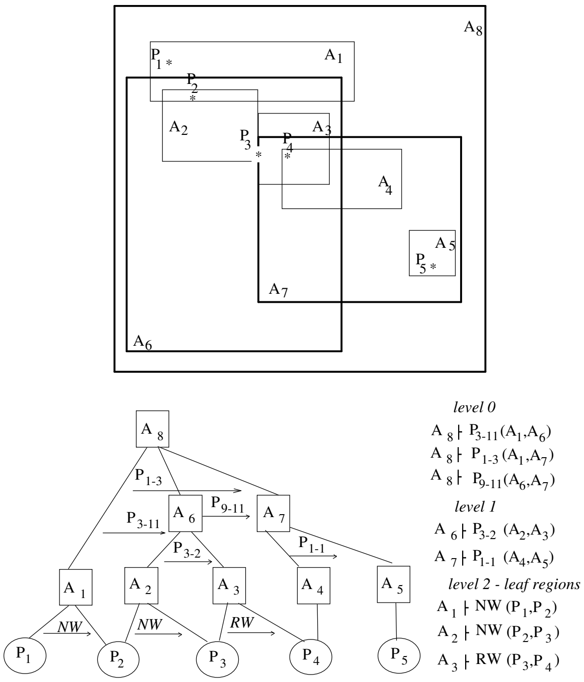

Figure 2 Example of hierarchical relation-based structure

3. EXPLICIT RETRIEVAL

According to explicit retrieval, the relation between points P_i and P_j is R if there is a leaf region A in which the two points are related by R: $\exists A(A \vdash R(P_i,P_j)) \Rightarrow DB \vdash R(P_i,P_j)$. Inconsistencies in this case arise when P_i and P_j exist together in multiple maps and their relations in these maps are different (an inconsistency of this form would be: $A_1 \vdash NW(P_1,P_2)$ and $A_2 \vdash NE(P_1,P_2)$). The following algorithm performs explicit retrieval by retrieving all leaf regions and examining the relations between all pairs of points in them. An *initialization* process assigns *U* to the relation between each pair of distinct points (*SP* for identical points). All the algorithms assume that information in each region is *node* ($A \vdash SP(P_i,P_i)$) and *arc consistent* ($A \vdash R(P_i,P_j) \Leftrightarrow A \vdash converse(R(P_j,P_i))$) and work only on the pairs (P_i,P_j) for which i<j.

```
Explicit_retrieval
for each (leaf) region A_k
 retrieve A_k;
 for each P_i such that DB⊢ IN(P_i,A_k)
   for each P_j such that DB⊢ IN(P_j,A_k) and i<j
       get the relation R' : A_k ⊢ R'(P_i,P_j);
       R(P_i,P_j)= R(P_i,P_j)∩R';
       if R=∅ then return INCONSISTENCY;
       else R(P_j,P_i)=converse(R(P_i,P_j));
  end-for
 end-for
end-for
```

In order to obtain formulas for the cost of the algorithms we make the following simplifications (although such simplifications may not apply for real applications, they provide a good measure for the expected cost in most cases). Each region contains *k* objects (points or other regions). Each object belongs to *m* regions in the upper hierarchy level, except for the region at the top (0 level) that does not belong to any region, and the objects at level 1 that belong only to the top-level region. It is always the case that $k/m>1$ and in regular applications $k/m>>1$. *N* is the total number of points in the database and *h* is the height of the hierarchical structure. We assume that there is a buffer that stores the $N(N-1)/2$ relations between all pairs of points.

The cost is a function of the *number of map retrievals* because such operations require access to secondary storage (i.e., retrieval of the disk pages that contain the map). This is common practice in database literature where indexing methods are compared on the number of accessed pages from the disk (Guttman 1984; Papadias et al., in press). In the case of explicit retrieval we have to retrieve all leaf regions. Due to the fact that leaf regions store all points and their copies[1], their number is mN/k, resulting in the same number of map retrievals. Unlike explicit retrieval, which is straightforward, the other two cases require inference mechanisms that potentially search large parts of the database.

4. SPATIAL INFERENCE THROUGH REGIONS

In inference through regions, the relation between P_i and P_j is inferred from the relations between their ancestor regions. The notation r→R means that when the relation r holds between two regions, then the relation R holds between all pairs of points in the regions. For instance, if two regions are related by projection relation $P_{1\text{-}1}$, the relation between any two points, each belonging to one region, is NW ($P_{1\text{-}1} \rightarrow NW$). Inference through regions can be described as: $[\exists A_k\ \exists A_l\ (DB \vdash IN^*(P_i,A_k) \wedge DB \vdash IN^*(P_j,A_l) \wedge DB \vdash r(A_k,A_l)) \wedge (r \rightarrow R)] \Rightarrow DB \vdash R(P_i,P_j)$. Inconsistencies in this case arise when the relation between P_i and P_j in some map is inconsistent with the relation between some of their ancestors regions. As an example consider: $A \vdash RN(P_1,P_2)$, $DB \vdash IN^*(P_1,A_1)$, $DB \vdash IN^*(P_2,A_2)$, and $DB \vdash P_{1\text{-}1}(A_1,A_2)$.

When the projections of two regions are disjoint on both axes (projections $P_{1\text{-}1}$, $P_{1\text{-}13}$, $P_{13\text{-}13}$, and $P_{13\text{-}1}$ in Figure 1), then high resolution information can be inferred for both south-north and west-east directions. However, not all projections allow such inferences about the

---

[1]Since an object appears *m* times in the next hierarchical level, we say that it has one original representation and m-1 copies.

P	1	2	3	4	5	6	7	8	9	10	11	12	13
1	NW	NW∨RN	N	N	N	N	N	N	N	N	N	NE∨RN	NE
2	NW∨RW	NW∨RN∨RW∨SP	N∨SL	N∨SL	N∨SL	N∨SL	N∨SL	N∨SL	N∨SL	N∨SL	N∨SL	NE∨RN∨RE∨SP	NE∨RE
3	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
4	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
5	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
6	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
7	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
8	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
9	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
10	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
11	W	W∨SH	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	E∨SH	E
12	SW∨RW	SW∨RS∨RW∨SP	S∨SL	S∨SL	S∨SL	S∨SL	S∨SL	S∨SL	S∨SL	S∨SL	S∨SL	SE∨RS∨RE∨SP	SE∨RE
13	SW	SW∨RS	S	S	S	S	S	S	S	S	S	SE∨RS	SE

Figure 3 Direction relation between points implied by the relation of ancestor regions (r→R)

relations between points. When the projections are disjoint on only one axis, low resolution relations about this axis can be derived, but information on the other axis is lost. Figure 3 summarizes the relations that can be derived about points given the projection relation between regions. The entries with *U* correspond to overlapping projections on both axes (in this case no conclusion can be drawn about the relations between points).

Initially the relation between any pair of points is given by explicit retrieval. Inference through regions retrieves one by one all non-leaf regions A and gets the relation r: A├ r(A_k,A_l) for all pairs of regions IN A. Let R' be the relation implied by r: r→R' (according to the rules of Figure 3). If R' ≠ *U*, the relation between all pairs of points P_i such that DB├IN*(P_i,A_k), and P_j such that DB├ IN*(P_j,A_l) is updated to R(P_i,P_j)=R(P_i,P_j)∩R'.

```
Inference_through_regions
for each non-leaf region A
 retrieve A;
 for each region A_k such that DB├IN(A_k,A)
  for each A_l such that DB├IN(A_l,A) and k<l
    get the relation r : A ├ r(A_k,A_l);
    lookup R': r → R';
    if R'≠ U then
     for each point P_i such that DB├IN*(P_i,A_k)
      for each point P_j such that DB├IN*(P_j,A_l)
        R(P_i,P_j)=R(P_i,P_j)∩R' ;
        if R(P_i,P_j)=∅ then return INCONSISTENCY
             else R(P_j,P_i)=converse(R(P_i,P_j));
      end-for
     end-for
    end-for
  end-for
end-for
```

For demonstration of the algorithm we use the example of Figure 2. Figure 4a illustrates the explicit relations between all pairs of points in the form of a constraint network. First A_6 is retrieved and the relation between A_2 and A_3 is found to be $P_{3\text{-}2}$. Since $P_{3\text{-}2}$→W∨SH, the relation between P_2 (which belongs to A_2) and P_4 (which belongs to A_3) is refined to *U*∩(W∨SH) = W∨SH (Figure 4b). The relation between P_3 and the other points of A_2 and A_3 remains unchanged, because NW ∩ (W∨SH) = NW and RW ∩ (W∨SH) = RW (for (P_2,P_3) and (P_3,P_4) respectively). Then A_7 is retrieved and the relation NW between P_4 and P_5 is inferred, because the ancestor regions of the two points (A_4 and A_5) are related by $P_{1\text{-}1}$ and, $P_{1\text{-}1}$ → NW (Figure 4c). After the retrieval of A_8 (the last non-leaf region) the network takes its final form of Figure 4d. From A_8├ $P_{1\text{-}3}$(A_1,A_7) and $P_{1\text{-}3}$→N, the relation North is inferred between all points of A_1 and the ones *IN** A_7, resulting in N(P_1,P_4), N(P_1,P_5), N(P_2,P_5), and NW(P_2,P_4)∨ RN(P_2,P_4) (the last relation is obtained by N∩(W∨SH)). The relations $P_{3\text{-}11}$(A_1,A_6) and $P_{9\text{-}11}$(A_6,A_7) do not allow any inferences, because $P_{3\text{-}11}$ → *U* and $P_{9\text{-}11}$ → *U*.

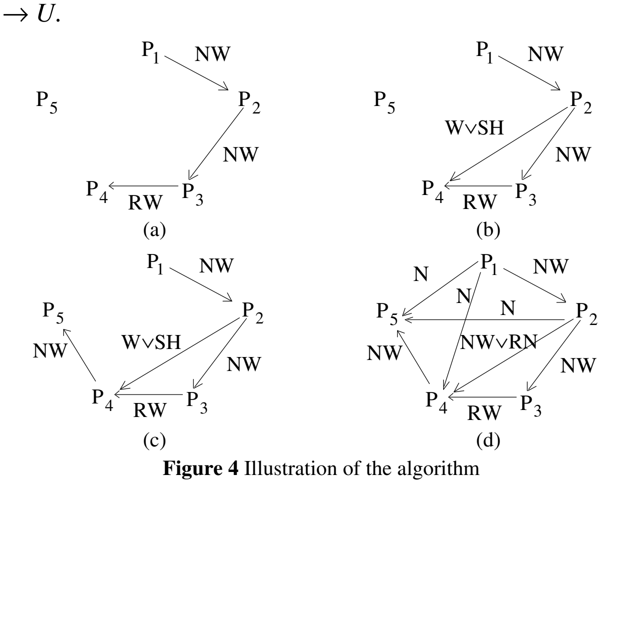

Figure 4 Illustration of the algorithm

In order to measure the cost of inference through regions we need to calculate the number of non-leaf regions, because all these regions are retrieved. There is only one node at level 0, k nodes at level 1, and k^2 at level 2. Out of these k^2 nodes, k^2/m correspond to objects and the rest to copies. Level 3 contains k nodes for each original node of the previous level resulting in a total of k^3/m nodes, out of which only k^3/m^2 are original and represented at level 4. Similarly, at level h-1 there are k^{h-1}/m^{h-2} nodes that correspond to actual leaf regions. Since the number of leaf regions is mN/k, we obtain a formula for h (Equation 1):

$$\frac{k^{h-1}}{m^{h-2}} = \frac{m \cdot N}{k} \Rightarrow h = \left\lceil \log_{(k/m)}\left(\frac{N}{m}\right) \right\rceil \qquad (1)$$

The number of non-leaf regions (and therefore the number of map retrievals) is the sum of original regions from level 0 to level h-2. Substituting the height of equation 1 we get an approximation for the cost of inference through regions:

$$1 + \sum_{i=1}^{h-2} \frac{k^i}{m^{i-1}} \cong m^2 \, \frac{N-k}{k(k-m)} \qquad (2)$$

Because the algorithm generates the permitted relations for all pairs of points, it needs to be performed only once and its results can be stored for future use. The above algorithm produces fast and high resolution relations in many situations; however, for overlapping projections with multiple common points (as in Figure 2) further refinements are possible by using the common points.

5. Spatial Inference Through Points

Inference using common points can be formulated as a path consistency problem in a network of binary direction constraints. Each constraint in the network is a disjunction of primitive relations and represents the permitted relations between a pair of points after explicit retrieval and inference through regions have taken place (e.g., Figure 4d). According to this form of inference, the relation between P_i and P_j, which belong to different maps, is derived through a chain of common points by composition of spatial relations: $[\exists P\ (DB \vdash R_k(P_i,P) \wedge DB \vdash R_l(P, P_j)) \wedge (R_k{*}R_l = R)] \Rightarrow DB \vdash R(P_i,P_j)$. The composition constraint $R_k{*}R_l$ is computed by forming the cross products of the primitive relations that comprise R_k and R_l, composing each resulting ordered pair by looking up the results in the composition table and taking the union of the resulting sets. Inconsistencies arise when different relations are inferred by different chains of points, or when the inferred relation contradicts the results of explicit retrieval or inference through regions (e.g., $A_1 \vdash NW(P_1,P)$, $A_2 \vdash NW(P,P_2)$ and $A_3 \vdash RS(P_1,P_2)$).

Figure 5 describes the rules that are applied in order to produce the possible direction relations between P_i and P_j when their relation with respect to a third point P is known. A very important point has to do with the type of constraints that appear in the network. If any disjunction of primitive relations is allowable, then the detection of all inconsistencies is NP-Complete even for point networks (Van Beek and Cohen, 1990) and path consistency (which is polynomial) does not suffice. However, in the problem that we study here, we start with a set of *33 relations* (*U*, ∅, 9 primitive, 6 low-resolution, and 16 relations of the form NW∨RN that may appear during inference through regions - Figure 3), which is closed under composition and intersection (see Sharma 1996). Path consistency suffices for this case and exponential algorithms are not needed to enforce satisfiability (for more details see Papadias and Egenhofer, 1996).

A number of path consistency algorithms have been proposed (Allen, 1983; Mackworth and Freuder, 1985). The following one is a variation modified for the current

	NW(P,P$_j$)	**RN(P,P$_j$)**	**NE(P,P$_j$)**	**RW(P,P$_j$)**	**SP(P,P$_j$)**	**RE(P,P$_j$)**	**SW(P,P$_j$)**	**RS(P,P$_j$)**	**SE(P,P$_j$)**	***U*(P,P$_j$)**
NW(P$_i$,P)	NW	NW	N	NW	NW	N	W	W	*U*	*U*
RN(P$_i$,P)	NW	RN	NE	NW	RN	NE	W	SH	E	*U*
NE(P$_i$,P)	N	NE	NE	N	NE	NE	*U*	E	E	*U*
RW(P$_i$,P)	NW	NW	N	RW	RW	SL	SW	SW	S	*U*
SP(P$_i$,P)	NW	RN	NE	RW	SP	RE	SW	RS	SE	*U*
RE(P$_i$,P)	N	NE	NE	SL	RE	RE	S	SE	SE	*U*
SW(P$_i$,P)	W	W	*U*	SW	SW	S	SW	SW	S	*U*
RS(P$_i$,P)	W	SH	E	SW	RS	SE	SW	RS	SE	*U*
SE(P$_i$,P)	*U*	E	E	S	SE	SE	S	SE	SE	*U*
***U*(P$_i$,P)**	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*	*U*

Figure 5 Composition table for high-resolution relations

problem. All pairs of points whose relation is not U are inserted into a queue. Then every pair is popped from the queue and the corresponding relation is used to refine the relation between the popped points and all the other points that co-exist with them in some region. The pairs of points whose relation is refined are pushed in the queue for propagation of the update through the network.

```
Inference_through_points

for each point P_i
 for each point P_j such that i<j
  if R(P_i,P_j)≠U then push-queue(P_i,P_j);
 end-for
end-for

while not-empty-queue
 pop(P_i,P_j);
 for each region A_l such that DB├ IN(P_i,A_l)
  retrieve A_l;
  for each P_k such that DB├IN(P_k,A_l) and k≠i≠j
   R_t(P_k,P_j)=R(P_k,P_j) ∩ (R(P_k,P_i)*R(P_i,P_j));
   if R_t=∅ then return INCONSISTENCY;
    else if R_t(P_k,P_j)⊂ R(P_k,P_j) then
      R(P_k,P_j)= R_t(P_k,P_j);
      R(P_j,P_k)=converse(R(P_k,P_j));
      if not in-queue(P_k,P_j) then push(P_k,P_j);
  end-for
 end-for
for each region A_m such that DB├ IN(P_j,A_m)
 retrieve A_m;
 for each P_k such that DB├ IN(P_k,A_m) and k≠i≠j
  R_t(P_i,P_k)=R(P_i,P_k) ∩ (R(P_i,P_j)* R(P_j,P_k));
  if R_t=∅ then return INCONSISTENCY;
    else if R_t(P_i,P_k)⊂ R(P_i,P_k) then
       R(P_i,P_k)= R_t(P_i,P_k);
       R(P_k,P_i)=converse(R(P_i,P_k));
       if not in-queue(P_i,P_k) then push(P_i,P_k);
  end-for
 end-for
end-while
```

For demonstration of the algorithm, we use the configuration of Figure 2. Figure 6a illustrates the network and the queue after explicit retrieval and inference through regions have been applied. First the pair (P_1,P_2) is popped and all the regions that contain these points are retrieved. P_3 co-exists with P_2 in region A_2 and its relation with P_1 is updated according to: $R(P_1,P_3) = R(P_1,P_3) \cap (R(P_1,P_2) * R(P_2,P_3)) = U \cap (NW*NW) = NW$. Because the new relation is a refinement of the previous one ($NW \subset U$), the pair (P_1,P_3) is pushed into the queue for propagation. The new network and the state of the queue at this phase are shown in Figure 6b. Then the pair (P_1,P_4) is popped from the queue, the regions A_1, A_3, and A_4 are retrieved, and the relations between the points (P_2,P_4), and (P_1,P_3) are updated. However, the network does not change at this stage because, $R(P_2,P_4) = R(P_2,P_4) \cap (R(P_2,P_1)*R(P_1,P_4)) = (NW \vee RN) \cap (SE*N) = NW \vee RN$, and $R(P_1,P_3) = R(P_1,P_3) \cap (R(P_1,P_4)*R(P_4,P_3)) = NW \cap (N*RE) = NW$. Similarly the pair (P_1,P_5) does not alter the network, while the pair (P_2,P_3) produces: $R(P_2,P_4) = NW$. The remaining pairs update the network in the same fashion; the final state after the termination of the algorithm is illustrated in Figure 6c.

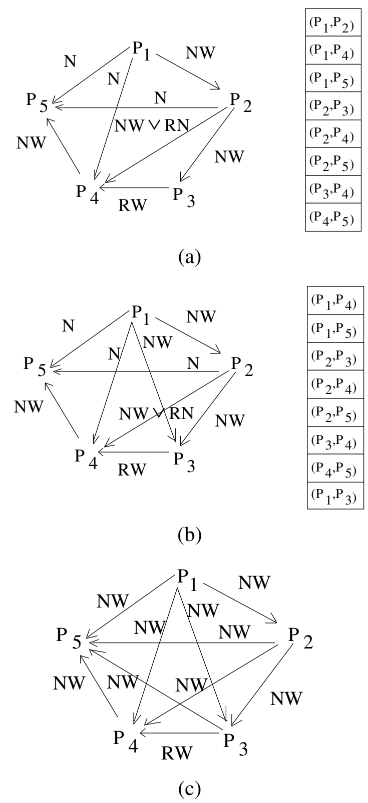

Figure 6 Illustration of the algorithm

In order to find the cost of inference through points we start with the observation that only 32 different constraints may appear in the network (and $\varnothing$ in which case the algorithm terminates with an inconsistency). A constraint imposed by inference through regions or explicit retrieval may be refined several times until it reaches its final state at the end of path consistency. The maximum number of refinements for any constraint is *four* (for details see Papadias and Egenhofer, 1996). For example, a constraint between two points may initially be U and become $N \vee SL$, then $NW \vee RN \vee RW \vee SP$, then $NW \vee RN$, and finally NW.

There exist $N(N-1)/2$ distinct pairs of points in the database and each may be pushed into the queue up to four times. Each time a pair is popped from the queue, $2m$ map retrievals are performed to retrieve the points that are related with the popped points in some region. Therefore, inference through points requires *4mN(N-1)* map retrievals in the worst case, which makes it significantly more expensive than inference through regions.

6. Discussion

In the previous sections we argued that first explicit retrieval obtains the relations between pairs of points that exist in the same region, then inference through regions generates additional constraints imposed by the relations between the ancestor regions, and finally inference through points takes advantage of common points to produce further refinements. The order in which explicit retrieval and inference through regions are performed is not important. As long as the content of the database remains unaltered they generate the same result independently on which operation is performed first. On the other hand, inference through points has always to be performed at the end, otherwise it may not produce all relations.

Unlike topological relations where there is a set of widely used relations in both research literature and commercial products (Egenhofer and Franzosa, 1991; Medeiros and Cilia, 1995), universally accepted definitions do not exist for direction relations. Although we have dealt with a set of projection-based direction relations often found in the literature, the methods of the paper are not relation-specific but can be applied to higher dimensions and other types of directions (for an example see Papadias and Egenhofer, 1996). In general, what is needed for the application of the algorithms is (1) a set of direction relations for points and one for regions, (2) rules for the inference of the relation between points given the relation between ancestor regions and, (3) composition rules. Notice, however, that depending on the choice of the relation set, inference through points may become exponential.

Hierarchical inference mechanisms are necessary to complement other qualitative spatial reasoning methods (Mukerjee and Joe, 1991; Sharma et al., 1994; Grigni et al., 1995; Scarponcini et al., 1995). Even in a single system, data about the same or overlapping areas but from different sources are stored separately. This information may be incomplete or inconsistent, and inference mechanisms are required to explicate relations and remove inconsistencies. As interoperability issues are solved, heterogeneous spatial databases and open GIS will soon become a reality. Such systems will store huge amounts of spatial data in various formats and of variable quality. Users will query the systems requiring fast and accurate results (and not contradictory answers of the form "A is north and south of B"). Spatial inference mechanisms will play an important role for the detection of inconsistencies in the data and the integration of the different systems.

Acknowledgements

This work was partially supported by grants from the National Science Foundation under grant number No. SES 88-10917 for the NCGIA. Dimitris Papadias is currently supported by an ERCIM fellowship. Max Egenhofer's work is further supported by the National Science Foundation under grant number IRI-9309230 and IRI-9613646, Rome Laboratory under grant number F30602-95-1-0042, the Scientific and Environmental Division of the North Atlantic Treaty Organization, and by a Massive Digital Data Systems contract sponsored by the Advanced Research and Development Committee of the Community Management Staff and administered by the Office of Research and Development.

References

Allen, J.F. “Maintaining Knowledge about Temporal Intervals”. Communications of the ACM 26 (11), 832-843, 1983.

Car, A., Frank, A. “General Principles of Hierarchical Spatial Reasoning: The Case of Wayfinding”. In the Proceedings of the 6th International Symposium on Spatial Data Handling, 1994.

Chang, S.K., Shi, Q.Y. Yan C.W. “Iconic Indexing by 2-D String”. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-9, no 3, 413-428, 1987.

Egenhofer, M., Franzosa R. (1991) “Point Set Topological Relations”. International Journal of Geographic Information Systems, 5, 161-174, 1991.

Frank, A.U., “Qualitative Spatial Reasoning about Distances and Directions in Geographic Space”, Journal of Visual Languages and Computing, Vol. 3, 343-371, 1992.

Frank, A. U. “Qualitative Spatial Reasoning: Cardinal Directions as an Example”. International Journal of Geographic Information Systems Vol. 10 (3), 269-290, 1996.

Freksa, C. “Using Orientation Information for Qualitative Spatial Reasoning”. Proceedings of the International Conference GIS - From Space to Territory: Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Pisa, Italy. Springer Verlag LNCS.

Glasgow, J.I., Papadias, D. “Computational Imagery”. Cognitive Science, 16, 355-394, 1992.

Grigni, M., Papadias, D., Papadimitriou, C., "Topological Inference", Proceedings of the International Joint Conference of Artificial Intelligence, 1995.

Guttman, A., "R-trees: a Dynamic Index Structure for Spatial Searching", Proceedings of ACM SIGMOD International Conference on Management of Data, 1984.

Hernandez, D., "Qualitative Representation of Spatial Knowledge", Springer Verlag LNAI, 1994.

Herskovits, A. "Language and Spatial Cognition". Cambridge University Press, 1986.

Hirtle, S, Jonides, J. "Evidence of Hierarchies in Cognitive Maps". Memory and Cognition 13, 208-217, 1985.

Holmes, P.D., Jungert, E. "Symbolic and Geometric 2Connectivity Graph Methods for Route Planning In Digitized Maps". IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-14, no 5, 549-565, 1993

Lee, S-Y., Yang, M-C., Chen, J-W. "Signature File as a Spatial Filter for Iconic Image Database". Journal of Visual Languages and Computing, 3, 373-397, 1992.

Mackworth, A, Freuder, E. "The Complexity of some Polynomial Network Consistency Algorithms for Constraint Satisfaction Problems". Artificial Intelligence, 25, 65-74, 1985.

Medeiros, C.B., Cilia, M. "Maintenance of Binary Topological Constraints Through Active Databases". Proceedings of 3rd ACM Workshop on GIS, 1995.

Mukerjee, A., Joe, G., "A Qualitative Model for Space". Proceedings of the National Conference on Artificial Intelligence (AAAI), 1991.

Nabil, M., Shepherd, J., Ngu, A. "2D Projection Interval Relationships: A Symbolic Representation of Spatial Relationships", Proceedings of 4th Symposium on Large Spatial Databases (SSD), Springer Verlag LNCS, 1995.

Papadias, D., Sellis, T., "Qualitative Representation of Spatial Knowledge in Two-Dimensional Space", Very Large Data Bases Journal, Special Issue on Spatial Databases, Vol 3(4), 479-516, 1994.

Papadias, D., Sellis, T. "A Pictorial Query-By-Example Language". Journal of Visual Languages and Computing, Special Issue on Visual Query Systems, Vol. 6, No. 1, 53 -72, 1995.

Papadias, D., Theodoridis, Y., Stefanakis, E., "Multi-dimensional Range Query Processing with Spatial Relations", To appear in Geographic Systems. Available via ftp from: http://nathan.gmd.de/persons/dimitris.papadias//publications.htm

Papadias, D., Egenhofer, M., "Algorithms for Hierarchical Spatial Reasoning". NCGIA TR-96-2 University of California, Santa Barbara, 1996. Available via ftp from: ftp://ncgia.ucsb.edu/pub/Publications/tech_reports/96/96-2/96-2.ps

Roussopoulos, N., Faloutsos,C., Sellis, T. "An Efficient Pictorial Database System for PSQL". IEEE Transactions on Software Engineering, Special issue on Image Databases, 639-650, 1988.

Samet,H. "The Design and Analysis of Spatial Data Structures". Addison Wesley, 1989.

Scarponcini, P., St. Clair, D.C., Zobrist, G.W "An Inferencing Language for Automated Spatial Reasoning about Graphic Entities", Proceedings of 4th Symposium on Large Spatial Databases (SSD), Springer Verlag LNCS, 1995.

Sharma, J., Flewelling, D., Egenhofer, M., "A Qualitative Spatial Reasoner". In the Proceedings of the 6th International Symposium on Spatial Data Handling, 1994.

Sharma, J., "Integrated Spatial Reasoning in Geographic Information Systems". Ph.D. Thesis, Department of Spatial Information Science and Engineering, University of Maine, 1996.

van Beek, P., Cohen, R. "Exact and Approximate Reasoning about Temporal Relations". Computational Intelligence 6, 132-144, 1990.