

Towards the global GIS

Christopher Gold ^{a,b,*} and Mir Abolfazl Mostafavi ^a

a: Geomatics Department, Laval University, Quebec City,
Quebec, Canada G1K 7P4.

b: Department of Land Surveying & Geo-Informatics
The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong

Abstract

This paper outlines the concept of a “Global GIS,” and defines various aspects of its development, as well as various options and decisions that must be made. The emphasis is on the advantages and disadvantages of maintaining a global topological structure, and whether topology should be generated “on the fly” in response to a specific query. We first define what we mean by “space” in this context, followed by a description of topological structures and how we may use them in the context of graph traversal problems. We then describe some appropriate data structures. After mentioning some of the real-world problems associated with polygon construction problems, we touch on how graphs may represent change over time.

A global topological structure is then proposed which resolves some of the problems mentioned. This is based on incremental graph updating methods. Some unresolved problems associated with using global topological structures are then discussed, especially those associated with queries, such as polygon overlay, that generate a new local topological structure. We conclude with some preliminary rules, and suggest further work.

Keywords: GIS, topology, data structures, globe, Voronoi diagrams.

1. Introduction: the problem

This paper outlines some of the issues to be resolved in the development of a truly global GIS. The term “GIS” is used loosely, because our concepts of GIS functionality are based largely on our experience with existing systems, and what we would really like is to be able to simulate as

many properties of “real-world” space as possible, within our computing systems. In addition, the term “global” here implies both that the data set is very large, and that we are working on the globe itself: the sphere or ellipsoid. Within that three-dimensional framework we have the gravitational force that gives us our concept of “vertical”, which means that for most terrestrial applications we may reasonably restrict ourselves to the two “horizontal” dimensions, in other words the surface of our globe.

2. What is space?

This space is, for most purposes, continuous, but with local variations in energy and matter. Where the matter is a solid, we think of objects embedded in that space; otherwise we think of fields, where some attribute of energy or matter varies more-or-less continuously over space. Where abrupt boundaries exist we interpret what we observe as due to the existence of solid objects, embedded in our space, and can reason about their form and spatial relationships with respect to other objects. In most other cases we think of “fields” - some attribute varying over our space, with occasional observations of these attributes. In this case we need to express their spatial relationships, just as if the observations were themselves objects.

Because digital computers use discrete representations, and because people find it easier to classify information than to visualize it as continuously varying, we use various techniques to “discretize” (classify) this information, either in the attribute domain, or in the spatial domain, or in the time domain. Contour maps, for example, discretize the attribute (elevation) information, whereas political maps discretize the spatial component into polygonal regions. Most maps discretize time, with only occasional map updates. For further discussion of these issues, see Gold (1996).

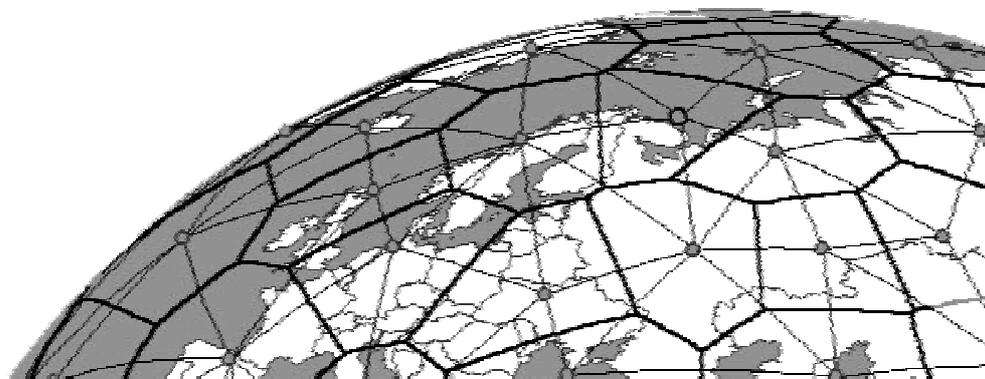


Figure 1. The earth and global topology.

Thus most of our information about space does not merely consist of defining objects, along with some arbitrary coordinate reference, but at a fundamental level requires some concept of spatial relationship (adjacency, etc.). The use of spatial relationships (a) avoids using arbitrary coordinate systems and (b) provides the tools for local navigation within the particular spatial framework of interest. Item (a) is important because local relationships and measurements are easier to detect than imaginary global coordinates (at least, this was true before the advent of GPS). Item (b) allows the navigation from neighbour to neighbour through a mesh giving the adjacency relationships between members of a selected class of objects. Fig. 1 illustrates the idea for a small point set.

3. What are topological structures?

In a global GIS one basic requirement is to have a set of spatial relationships linking our objects or observations in a “horizontal” sense - in other words, a topological structure. The key question is how to implement it, and at what level, in a “global” system. Moreover, both in theory and in practice, these topological links must be locally modifiable. On the theoretical level, we live in a world where spatial relationships change frequently, not least because we are navigating

in it. At the practical level, a global-sized data set may not have a topology that is built all at once in a "batch" fashion - it is too big. Indeed, at that scale it is clear that time comes to play a significant role - it will never be complete or up-to-date, and therefore local updating becomes an absolute necessity.

The idea of "topology" given above is that a set of spatial relationships exist between objects that are in some sense neighbours, and these spatial relationships, for the purpose of this paper, are expressed in terms of two local dimensions only. In mathematical terms, we can express this as a graph, with the objects being represented as nodes and the spatial relationships as edges connecting pairs of objects. (See Sedgewick, 1983.) We can go further: our nodes and edges are embedded in a two-dimensional surface. Usually this is adequate, but it does not allow for cases such as a road crossing a river, or overpasses and underpasses. As these are special cases, we will not include them in our current model. (Most GISs do likewise.) However, the Quad-edge data structure described below may be used to represent these cases if required.

If our graph is of this form, not permitting edge crossings, we refer to it as a planar graph, which has various additional properties. One property is that regions may be defined, as the portion of two-dimensional space enclosed by a closed cycle of edges. (We are usually concerned with a minimal region, where the cycle does not enclose any other nodes or edges - a single polygon, for example.) Another property is that all the edges connecting to a particular node can be given in some cyclic order - e.g. anticlockwise around the node. Finally, the number of regions, edges, or nodes may be determined from the other two entities (Euler's Formula: $\text{Regions} + \text{Nodes} - \text{Edges} = 2$, counting the exterior as a region). One example of a planar graph is a polygon map; another is a river network; yet another is a triangulation.

There are various ways of storing these various planar graphs, and much has been written about the various possible data structures used to store and manipulate them in the computer. They are all designed to allow the program to move from one entity (Region, Edge, Node) in the graph to

some adjacent entity. Examples of such structures are: Polygon/Arc/Node; Doubly-Connected-Edge-List (DCEL); the Winged-Edge structure; the Quad-Edge structure; and, for specific applications, the triangulation structure. Worboys (1995) gives a good summary. One way to look at the many possibilities is the PAN-graph (Gold, 1988). Here each entity in the graph (Polygons, Arcs, Nodes: equivalent to Regions, Edges, Nodes) may be connected together in various ways, indicating the pointer structure between adjacent graph entities. De Floriani et al. (1995) use the same method. While simple, this figure can readily indicate whether one can reach another entity from a given starting place, how quickly, and with how much storage. Given some knowledge of the kind of queries expected, an appropriate data structure may be defined.

One limitation of these approaches should be noted: they work only for connected graphs. If there are “islands” in a polygon map, for example, and the arcs between the nodes in the data structure graph correspond to the polygon edges in the map, then the presence of the island within some other polygon can not be detected without further processing. This shows up a weakness of this particular data structure. An alternative solution is to work with the “dual” of the original map. Here, each polygon may be represented as a point (the “capital city” perhaps). Then pairs of capital cities of polygons with common boundaries are connected (“roads” between capital cities of adjacent polygons). These form the redefined edges of the new graph. Finally, as before, regions are defined as closed cycles of edges: this gives redefined polygons around each original node. If the map is completely covered with polygons then the resulting graph is connected. In addition, the graph is now a triangulation, except for where four or more regions meet at a point (which can be fixed by adding zero length edges), or where there are islands. Thus, as in Gold (1988), the dual structure may be used as a more convenient representation of the adjacency relationships present in the polygonal map. For network maps the same approach may be used.

4. Traversing graphs

Within any of the data structures described above, and expressed by some form of PAN graph, we need to be able to navigate through the generated network in some way, in order to respond to various spatially-based queries. Examples of such queries include: finding an adjacent entity (Region, Edge, Node); walking along a particular path through the network to a specified destination (defined by location or object name); detecting all the objects within some zone of interest (defined by coordinates, or by some set of objects, e.g. a polygon boundary); or doing a complete traversal of the whole graph.

An example of the traversal of a whole graph can be seen in the process for the extraction of a set of polygons from a set of digitized arcs. In the usual approach (Burrough 1986) the digitized “spaghetti” is processed to find all the intersections among all the arcs. After clean up this gives a set of nodes, with arcs defined in (anticlockwise) order around each node. (In practice this process is extremely difficult to implement because of the usually messy results of manual digitizing.) Once this stage has been reached the resulting graph needs to have the region (polygon) information added. Starting at any arc, and examining the node at one end, the next arc clockwise is found. This process is repeated for the node at the other end of this arc, and this continues until we have traversed the polygon in anticlockwise order and have returned to the starting place. The additional relationships (Polygon/Arc, Polygon/Node, and their inverses) are added to the topological structure as required. This process has used Node/Arc and Arc/Node relationships, and the appropriate side of each arc is flagged as used. The opposite (unflagged) side of one of the arcs just processed is then used as a starting point, and the next polygon defined. This continues until the whole graph is processed.

Other examples of traversing the whole graph include the classic depth-first or breadth-first searches (Sedgewick 1983), which do not require any geometric information and can be guaranteed to visit every node of a connected graph, and the “Visibility Ordering” of a triangulation (Gold and Cormack, 1987, De Floriani et al. 1991). It is also possible to “walk”

through a triangulation, from some starting triangle to some destination location, by checking which triangle edges have the destination point on the “outside” edge, and then moving to that triangle and repeating the process, until the enclosing triangle is found (Gold et al., 1977).

The Quad-edge data structure (Guibas and Stolfi, 1985) allows similar navigation, even though all navigation is from edge to edge only. Its advantages are firstly that there is no distinction between the primal and the dual representation (it is symmetric with respect to edges and faces/polygons), and secondly that all operations are performed as pointer operations only, thus giving an algebraic representation to its operations. (By contrast, operations on triangles require that the three edges of the triangle be tested to establish where the previous triangle was connected. Similar situations arise with other polygon structures.) Its disadvantage is that storage costs are high by comparison with other methods. It is, however, very easy to implement, especially in an object-oriented environment. Fig. 2 shows the basic structure, with four branches for each edge of the graph being stored, one for each of the adjacent vertices or faces. Table 1 shows the object definition in Delphi (Object Pascal). Apart from a pointer to the node or face, each branch has a pointer “R” (Rot), connecting the four branches together in anticlockwise order, and a pointer “N” (Next), pointing to the next edge in anticlockwise order around the node or face. Also shown are some of the low-level operations using N and R: “Sym,” which moves from the current branch to the branch of the same edge that faces in the other direction; “Oprev” which moves to the next branch/edge clockwise around the vertex/face associated with the original branch; and “Lnext” (which equals Sym.Oprev) that finds the next branch/edge clockwise around the other end of the edge.

Figure 2. The Quad-Edge data structure.

Figure 3. The Splice Operation

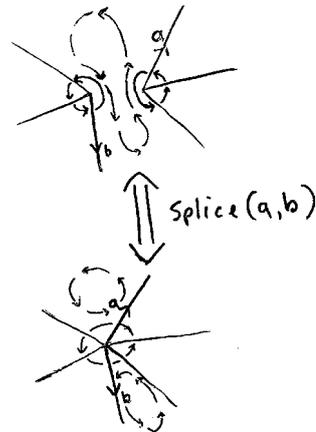
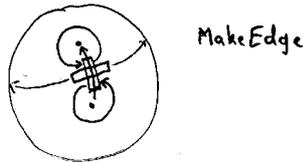
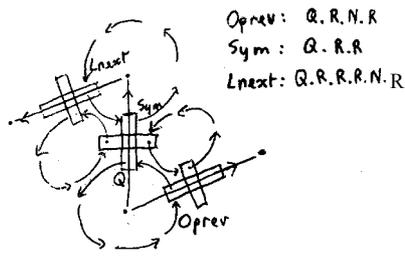


Table 1: Quad-Edge Code	
<pre> {This table gives an object-oriented version of Guibas and Stolfi's} {Quad-Edge data structure, its basic functions and example usage.} TQuad = class N : TQuad; {next edge anticlockwise} R : TQuad; {next 1/4 of edge} V : TPoint; {vertex} Index : Integer; {name* for debugging} end; class function TQuad.MakeEdge(Orig, Dest: TPoint) : TQuad; var Q0, Q1, Q2, Q3 : TQuad; begin {create four new 1/4 edges} Q0 := TQuad.Create; Q1 := TQuad.Create; Q2 := TQuad.Create; Q3 := TQuad.Create; {link the four parts} Q0.R := Q1; Q1.R := Q2; Q2.R := Q3; Q3.R := Q0; {link 0 & 2 to themselves, 1 & 3 to each other} Q0.N := Q0; Q1.N := Q3; Q2.N := Q2; Q3.N := Q1; {set pointers to vertices} Q0.SetVertex(Orig); Q2.SetVertex(Dest); Result := Q0; end; procedure TQuad.Splice(A,B : TQuad); {A, B: input Quad-Edges} var Alpha, Beta, An, Bn, Aln, Ben : TQuad; begin {get neighbouring edges: Alpha & Beta in Guibas & Stolfi} Alpha := A.N.R; Beta := B.N.R; An := A.N; Bn := B.N; Aln := Alpha.N; Ben := Beta.N; {reconnect the four pointers} A.N := Bn; B.N := An; Alpha.N := Ben; Beta.N := Aln; end; function TQuad.Sym : TQuad; {other end} begin Sym := Self.R.R; end; function TQuad.Oprev : TQuad; {next edge clockwise} begin Oprev := Self.R.N.R; end; </pre>	
<pre> function TQuad.Onext : TQuad; {next edge anticlockwise} begin Onext := Self.N; end; function TQuad.Lnext : TQuad; {next edge clockwise, other end} begin Lnext := Self.R.R.R.N.R; end; function TQuad.Rprev : TQuad; {next edge anticlockwise, other end} begin Rprev := Self.R.R.N; end; function TQuad.Vertex : TPoint; {read vertex} begin Result := Self.V; end; procedure TQuad.SetVertex(PtIn TPoint); {set vertex} begin V := PtIn; end; procedure TQuad.Delete; {disconnect and free an edge} begin Splice(Self, Self.Oprev); Splice(Self, Self.Lnext); Self.Free; end; function TQuad.Swap : Boolean; {swap a diagonal in a triangulation} var a, b : TQuad; begin Result := False; a := Self.Oprev; {get adjacent edges} b := Self.Lnext; if (a.Sym.Vertex <> b.Sym.Vertex) then begin Result := True; Splice(Self, a); {disconnect diagonal} Splice(Self, b); Splice(Self, a.Lnext); {re-connect diagonal} Splice(Self, b.Lnext); Self.SetVertex(a.Sym.Vertex); Self.Sym.SetVertex(b.Sym.Vertex); end; end; </pre>	

To show the simplicity of its use, Fig. 3 shows the two commands that are used to modify a graph: “Make-edge” to create a new edge on a manifold, and “Splice” to connect/disconnect Quad-edges together. In the simplest case, Splice connects two separate “Next” loops, joining the two nodes together, and at the same time splitting the “Next” loop around the common face. (The

process could equally well have split the “Next” loop around a node, and merged the loops around a face: Splice is its own inverse. See the original paper for more details.) Table 1 gives the code for these operations, as well as for three simple example operations using Make-edge and Splice: deleting an edge from a graph, inserting a point in a triangle, and switching the diagonal of a pair of triangles. These operations are sufficient to maintain a triangulation network for simple terrain modelling. Work is continuing on making these operations efficient in secondary (rather than just primary) storage.

5. Real-world problems with polygon construction

In reality, of course, topology construction in a traditional GIS is a complex operation. It normally consists of a global operation to find all line intersections between the input arcs, followed by various stages of cleanup: removing small dangling arc ends, merging nearby intersections into a single polygon node (see Pullar, 1994), etc. Only then may the resulting graph be traversed to create polygons. Additional problems are the detection of self-intersecting arcs, the generation of unwanted “sliver polygons” if arcs are nearly superimposed, and unconnected islands as described above. (In a traditional GIS, “topology” is largely thought of as a polygon construction problem.)

The nature of the intersection-detection operation indicates a global process rather than an incremental one, and this leads to the necessity of processing the whole “map” at once. Thus most systems are map-sheet based, allowing the processing of a manageable set of data at one time - even if only one arc has been changed from the previous time. Goodchild (1989) has discussed the problem of tiling large spatial data sets. This approach is clearly not appropriate for a global GIS, which would need an incremental algorithm, adding one arc at a time. (It should be noted that if the extent of an update can be well defined, then a “patch” may be created and sewn into the map sheet, thus speeding up the process. All the underlying problems remain, and it is not

problem here arises out of the complexities in the geometry operations, as the incremental update of a planar graph is not a problem.

Outside the traditional polygon map, there are many other applications whose topological structures deserve consideration. Networks (roads, rivers) possess a topological structure less complete than a polygon set; in addition, closed loops may not represent regions that need to be named, and unconnected nodes at the ends of arcs are permitted. Point data sets initially may be thought of as having no topological structure, but they are very often connected into a triangulation. These structures ought to be updateable on an incremental basis also.

Spatial operations that fall outside traditional GIS need to be considered also; in particular simulation of one form or another where the objects move in space and need dynamic update of their topology. This is a direct continuation of the update problem mentioned above, except that modifications are made as a result of the simulation of the process involved (the “future”), as opposed to the manual updating of the known “past.” A simple example would be atmospheric modelling, where particles representing packets of air would move around and interact with each other. This can be managed with the triangulation structure - in particular the Delaunay triangulation, which is the dual of the simple Voronoi diagram. This provides simple rules for determining the neighbouring points to any particular point, and how to update the topological structure as the points move (Roos, 1991). Another approach is that of reactive data structures (Van Oosterom, 1993). In the plane, there are many published examples of algorithms for creating the Voronoi/Delaunay structure – a good recent example is Devillers (1998). Okabe et al. (1992) have produced an excellent compendium of Voronoi algorithms, and they show that the concept of proximal regions has many applications in spatial analysis.

7. A Global Topological Structure

This example of atmospheric modelling naturally brings us to the question of moving our

topological structure onto the sphere (or ellipsoid). The transformation of a simple planar graph is very easy: take the “exterior” region, and consider this to be a finite region on the opposite side of the sphere to the viewpoint. For a triangulation, all exterior vertices may be connected to a point on the opposite side of the sphere. This simplifies management of the boundaries of the map, and provides, as Worboys (1995) calls it, the “perfect 2-space”.

Thus all displayed maps are considered to be projections of the globe - which is quite reasonable! Voronoi diagrams may be displayed (Figs. 5 and 6), as well as the dual Delaunay triangulation. Dynamic point movement may be managed as described above. The topological structure allows the navigation through the network to locate the enclosing triangle of, or the nearest object to, some specified location. Local updates may then be performed to insert or delete an object at that location. Watson (1988) and Lukatella (1989) have also developed methods for constructing the Voronoi diagram on the sphere, as have Augenbaum and Peskin (1985).

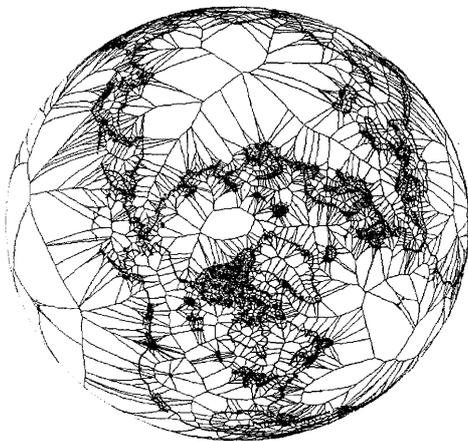


Figure 5. Orthographic projection, Voronoi cells of points on continent boundaries.

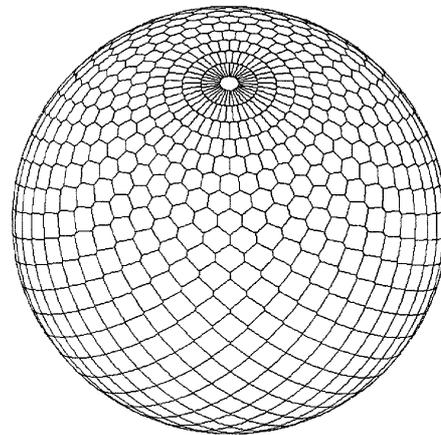


Figure 6. Regular Voronoi cells on the sphere

8. Construction of Voronoi diagram on the sphere

In this section we briefly describe the construction of the global topology based on the Voronoi diagram on the surface of sphere. As mentioned in previous sections, we used the direction

cosines in three dimensions as the basis of the coordinate system for representing the location of the objects on the surface of the sphere. Such a coordinate system is earth centred. This system is seamless and allows a continuous representation of the position of the objects on the surface of the sphere, as illustrated in Fig. 7.

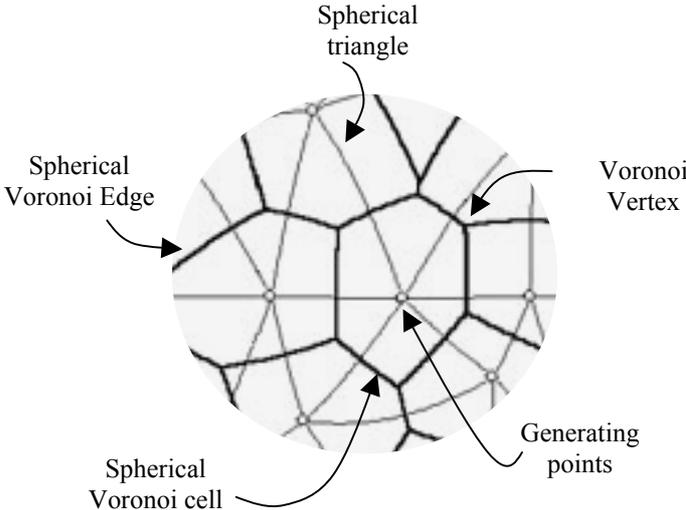


Figure 7. Portion of the Voronoi diagram and its components on the surface of unit sphere.

Consider a collection of N points $\{p_1, \dots, p_2\}$ on the surface of a unit sphere. A minimum of four points is needed to initialize the Delaunay triangulation and the Voronoi diagram on the sphere. For our purpose, we used six initial points that correspond to the intersection of the axes of the 3D Cartesian coordinate system with the surface of the sphere (Fig. 8). Using these initial

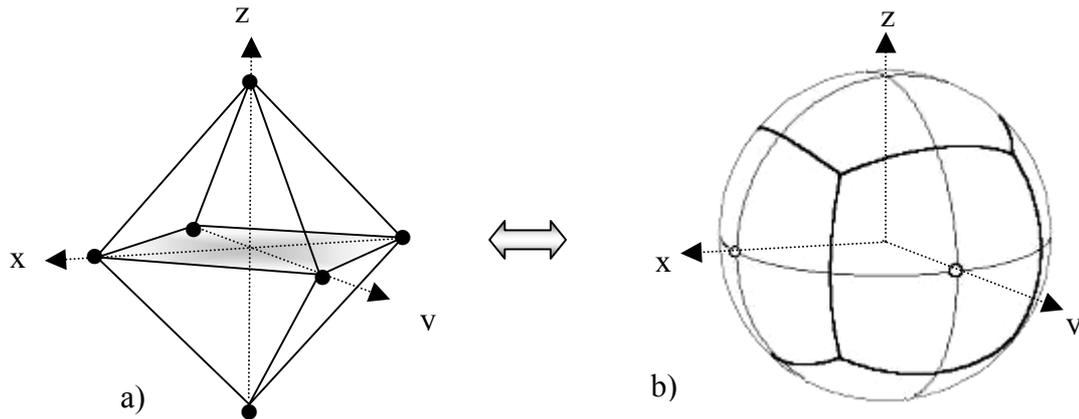


Figure 8. Initializing the Delaunay triangulation.

- a) Choice of six initial points at the intersection of the axes of the coordinate system with the surface of the unit sphere
- b) Projection of the eight triangles on the surface of the sphere and creation of the eight initial spherical triangles.

points we create eight triangles (four triangles on the northern hemisphere and four others on the southern hemisphere). The initially created triangles form a Delaunay triangulation on the sphere. This implies that all of the circumcircles of the triangles are empty. These triangles cover the surface of the globe and are used as the bases for the construction of the spherical Delaunay triangulation.

We use the incremental method to create the spherical Delaunay triangulation. This method allows local modification of the structure without reconstruction of the whole topology. The “Insert” and “Delete” operations for the Delaunay triangulation on the sphere are almost the same as in the planar case, but the optimization criterion is different. On the surface of sphere the shortest distance between two points is defined as the angular distance between these points

along the great circle formed by the origin of the sphere and the two points. With this definition, the circumradius of the triangle ABC is $r_{ABC} = \langle A, N \rangle$, where N is the circumcentre of the triangle (Fig. 9) and $\langle A, N \rangle$ is the distance between the points A and N on the surface of the sphere measured in radians.

The optimization procedure on the surface of the sphere consists of comparison of two angular distances. One is the circumradius of the triangle to be tested and the other is the angular distance between the circumcentre of the triangle and a fourth point on the surface of the sphere. If the latter is greater than the circumradius of the triangle, then the fourth point is outside the circumcircle - which implies that the triangle belongs to the Delaunay triangulation (Fig. 10). If the circumcircle of the spherical triangle ABC is not empty then it is not a valid Delaunay triangle. To maintain the Delaunay triangulation in this case we use the operation called “Swap” to change the diagonal of the spherical quadrilateral formed by the set of four points. The result of the “Swap” operation is two newly created triangles that should be tested again for the Delaunay criterion.

To compute the location of the circumcentre of a spherical triangle we use the following equation

$$\vec{N} = \frac{(\vec{B} - \vec{A}) \times (\vec{C} - \vec{A})}{|(\vec{B} - \vec{A}) \times (\vec{C} - \vec{A})|} \quad (1)$$

In this equation \times is the vector product and $|\cdot|$ denotes the norm of the vector. A , B , and C are three points on the surface of the unit sphere and \vec{N} is the position vector of the circumcentre of the spherical triangle ABC that passes through the origin and is reduced to the surface of the unit sphere. \vec{A} , \vec{B} and \vec{C} are the position vectors of the points A , B and C . We can see from the equation 1 that since the points A , B , and C can never be on the same straight line (because they are on the surface of the sphere) the equation is valid for all the triangles on the sphere.

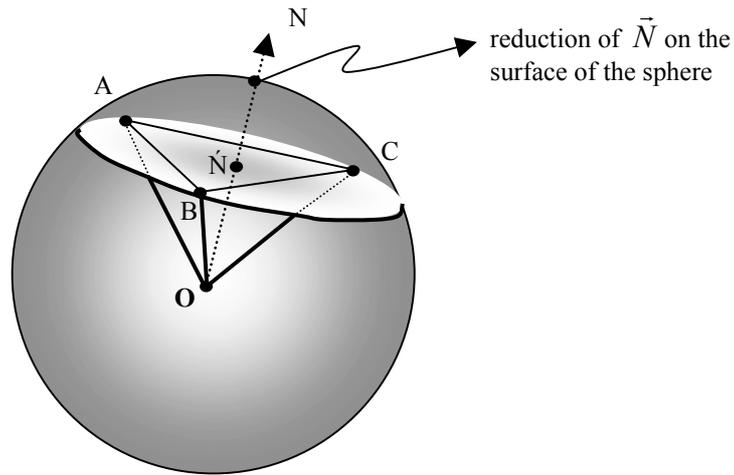


Figure 9. \dot{N} is the circumcentre of the triangle ABC and the point N is the centre of the spherical cap cut by the plan ABC .

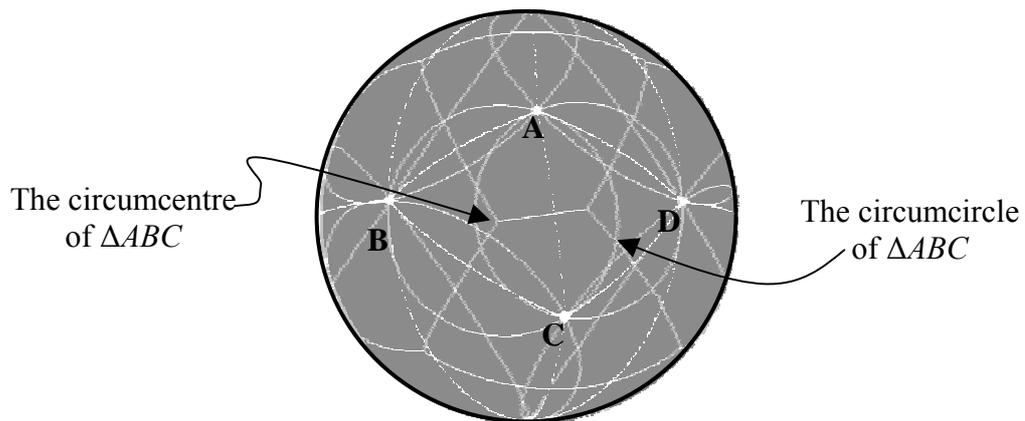


Figure 10. Circumcircle and Circumcentre of a spherical triangle.

The Voronoi diagram is the dual of the Delaunay triangulation. This implies that we can construct the Voronoi diagram from the Delaunay triangulation and vice versa. This is true for the Voronoi diagram on the plane and on the sphere. In the spherical case each Voronoi edge is an arc of a great circle between the circumcentres of any two adjacent spherical triangles (Fig. 7). For each triangle we connect its circumcentre to the circumcentres of its three adjacent triangles. The resulting diagram is the Voronoi diagram on the sphere (Figs. 11 and 12).

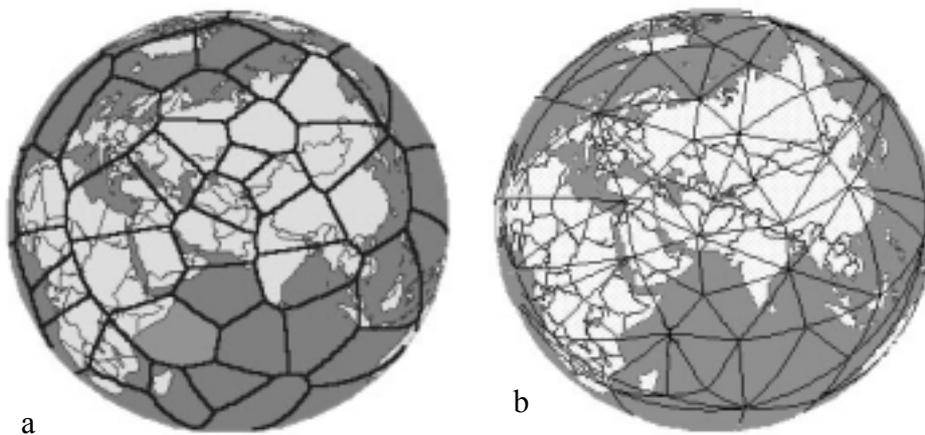


Figure 11. Sphere surface partitioning:
 a) Voronoi diagram
 b) Delaunay triangulation.

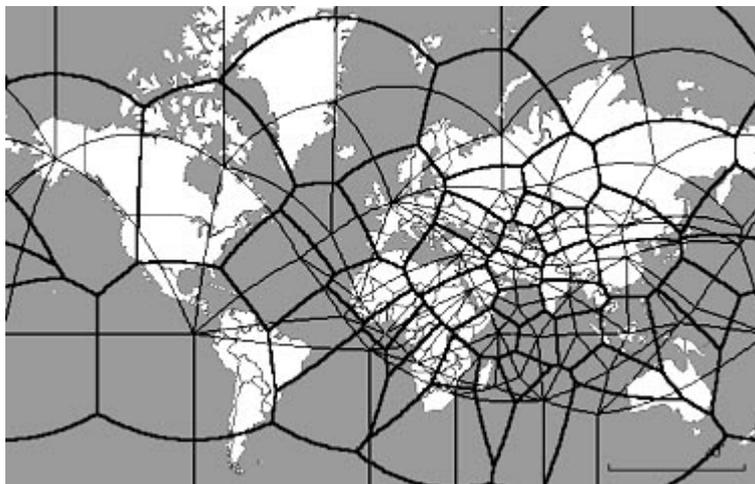


Figure 12. Stereographic projection of Voronoi diagram and Delaunay triangulation for a set of points.

9. Topology vs. attributes

Following topology construction issues, but related to them, are concerns about spatial queries.

Within a map-sheet-based system simple queries are global for the map sheet in most cases:

finding all objects (houses, polygons, etc.) with the desired attributes. This is implicit in a relational data base system: a particular table of attributes will be for the whole map. This is not feasible in a true global system: queries must be restricted first by some spatial component, and only then by attribute. The traditional approach would be by some map-sheet structure as described above, with all its complications.

Another alternative would be to use the global topological structure itself: first finding any object within the region of interest, and then expanding outwards to collect everything within that region. The first stage may be performed by walking through the network as described above. For successive queries this walk will usually be relatively short, as work continues on the region of interest. Nevertheless, this will not always be the case, and some pointer structure may be needed to avoid traversing half the globe to find some initial object within the new region of interest (see Gold, Remmele and Roos, 1995). The hierarchical structures of Dutton (1989, 1991, 1996), De Floriani and Puppo (1995) or Jones et al. (1992) could also be used at this stage. The second stage involves a simple graph traversal algorithm such as breadth-first-search, terminating at the region's boundaries, in order to examine all the objects of interest (see de Berg et al., 1997).

Given the premise that a global GIS contains many objects, and that the region of interest is small, this initial step reduces the number of objects examined to a reasonable set. Standard relational queries may then be performed on this reduced set. It should be remembered that typical relational queries are "global" in principle: the initial table contains the whole data set, with no spatial limits. Thus any subsequent query, whether as some form of coordinate window, or on some attribute, must start by examining the whole data set. Clearly, if the global data set is partitioned into map sheets as part of the fundamental data structure this problem is avoided, as queries are addressed to individual sheets. However, if a global data structure is used, with topological linkages, then the identification of some starting seed object followed by the regional search, as described above, avoids any requirement to examine the whole data set.

The introduction of time into the picture complicates things further. If it is a question of querying the database at some particular point in its updating process, rather than after the completion of some “snapshot” of the whole system, the approach just described remains unchanged. Queries that are constrained in space as well as time, and examine any of the past states of the region, remain a research topic (Mioc et al., 1999). However, if we consider that the number of time states remains small within a local region, then the above approach may be preserved and the lifespans of the objects may be treated as attributes. This, however, will not preserve the topological structuring at some past time: questions such as “When were these two objects adjacent?” can not be handled.

10. The Overlay Problem

Other questions arise when the traditional tools of polygon overlay and buffer zones are required. These are normally batch operations performed on the whole topological unit - usually the map sheet. In a global system the objects within the region of interest must be extracted by queries, as described above, and then the new topological structure created for that region. This last may be a batch operation, or based on local incremental topology modification. Nevertheless, the effect is to produce a new topological space (or structure) for the region of interest, for a particular query. Thus the utility of global topological spaces becomes less clear if analysis operations form new local structures that can not be integrated within the global parent structures.

The other old standby, the “dissolve” or “reclassify and merge” operation, is usually based on global reclassification, followed by local elimination of unwanted polygon boundaries. This does not necessarily require the generation of a new topological structure, except for archival purposes.

11. When is a Global Topological Space Useful?

This brings us to the key question of this paper: to what extent are global topological structures

appropriate in an ideal system? For our purposes, an ideal system would be organized on the globe (the sphere or ellipsoid), possess very large numbers of map objects, have some topological structure connecting these objects, and be capable of handling frequent small, local map updates. Due to the difficulties encountered in re-joining adjacent separate map sheets, it would be preferable to avoid this problem by having a seamless topology and coordinate system. (For the purpose of this exercise the problems of paging this data structure into secondary storage will be ignored, as being a separate problem.)

The primary question concerns the number of basic topological spaces that are required for the project under consideration. These are basically the same thing as “overlays”, as they can be thought of as planar (more precisely, plane) graphs: they may be connected (as in a polygon coverage, or a river network), or disconnected (as in point sets or buildings), but they may not conflict with one another by having intersections that are not represented by topological nodes.

Different overlays are used to separate different groups of objects. Each group must be topologically consistent, as just described, but may contain different classes of features as long as there is no superposition. For example, trees are not usually found on houses, so these may be grouped together in the same overlay, but cars are often found on roads, so these may not be grouped. See Molenaar (1998) for a more detailed treatment of spatial object modelling.

However, there is an important consideration: does the operation of classification of a set of objects in a single overlay require that the resulting classes be stored in multiple overlays? At one level they need not: they remain spatially non-overlapping after being reclassified. (Think of the reclassification of a choropleth map.) On another level they may: adjacency queries may be between buildings of the same type (such as police stations), rather than involving intervening buildings of another type. This issue may be resolved in practice by using more complex operations on the topological space than that of the basic adjacency properties of the simple graph. In our police station example the stations are not adjacent in terms of the buildings layer,

but they are accessible through that layer and the adjacency query may be resolved analytically.

Thus it appears that the answer to our question depends on how much local topology construction is required for queries (overlay, etc.). This in turn depends on the classification of object classes. If there are many types of objects, then the question becomes: do we have a separate global topology for each class? If yes, we can have adjacency type queries within each class, but local batch construction for inter-class queries. (Overlay is an inter-class query, buffer-zone is not.) If we need to classify our objects on the fly, (e.g. polygon reclassify and merge) then it is not obvious whether the results of this operation are good for local adjacency queries only, or need to be subdivided into their own topological spaces. At first sight it appears that, since the results can not be superimposed, they can be processed in a single topological space, as with our tentative approach to the police station problem above. (One basic rule for topological spaces is that objects in that space may not occupy the same locations: if that can possibly occur, then separate spaces are required.) Thus the primary question is: how many topological spaces are required? Can we make general rules, or do we need to perform semantic analysis for each specific application?

Conclusions

Within a single topological space a global system is appropriate. Where multiple topological spaces are required, then comparisons between them (often overlay operations) generate further spaces. These would be small-scale, one-off local operations, and would not need to be preserved beyond their use as a response to a query. This seems to be reasonably good for traditional polygonal maps, which are one representation of field-type data. Other fields would be continuous, with individual observations, and the single topological space would seem appropriate. For discrete objects a lot depends on the assumption of two-dimensionality, as with our examples of trees and houses, cars and roads. In both cases however there will be the need to

extract subsets of the operations in a single topological space in order to answer specific local queries: buffer zones around primary roads only, for example, or the distances between pubs or police stations. These are again one-off operations for specific queries.

The function of a global topological space therefore is to contain the objects of a particular (spatially non-overlapping) group. For simple cases, such as simulating atmospheric movement, this should be enough. In other cases several spaces should be maintained, and local topological structures generated for specific queries. The alternative is to subdivide the globe into map-sheets, creating many complications.

The starting point for any local operation on a topological space in traditional GIS is to find a map sheet containing some initial location. If a global structure is used, the first step is to find the closest object to the initial location, and then perform an outwards-search to cover the region of interest. This may be done by a walk through the structure (which could be long), or else by adding a simple hierarchical structure to give a close initial estimate. Having obtained a spatially restricted region, local topological spaces may be generated by overlay or other operations in order to respond to particular queries. A major technical question concerns the storage and retrieval of attribute information: where there are no permanent regions (map sheets) it is not obvious how to adapt to RDBMS-type structures.

The arguments given above would be appropriate for any appropriate topological structure that can be represented as a graph, possibly stored as Quad-edges, and can be updated using local operations. Nevertheless, the ideas were organized with the Voronoi structure in mind, (for example Gold and Condal, 1995), as it has been shown to be maintainable by local update, can handle time-varying information, and can readily be generated on the sphere. Thus the move towards a global GIS would require a locally-updateable topological structure, perhaps of that type. It would be necessary to maintain a topological space for each base class of objects/fields, defined on the basis that duplicate objects may not exist in the same space at the same location.

Topological queries between groups of classes, such as overlay, would require the generation of a temporary topological structure over the region of interest. For some single-space operations, such as flow simulation or interpolation, this would not be necessary.

Where temporary topological structures are required, it is less clear whether a global topological structure is necessary. On the positive side, it eliminates the worries over local coordinate systems and map edge-matching. On the negative side, maintenance might be excessive if most queries required local temporary topological structures. This would reduce the value of the global structure to data maintenance and extraction of the region of interest. Nevertheless, their benefits for the modelling of simple data types warrant experimentation with the approach, with extension to complex data types and queries if warranted at a later date.

Acknowledgments

This research was supported in part by an NSERC/AIFQ Industrial Research Chair in Geomatics at Laval University, Quebec.

References

- Augenbaum, J., and Peskin, C., 1985. On the construction of the Voronoi mesh on a sphere. *Journal of Computational Physics*, 59, 177-192.
- Burrough, P.A., 1986. *Principles of Geographical Information Systems for Land Resources Assessment*. Oxford University Press.
- de Berg, M., van Kreveld, M., van Oosterum, R. and Overmars, M., 1997. Simple traversal of a subdivision without extra storage. *International Journal of Geographical Information Science*, 11(4), 359-373.
- De Floriani, L., Falcidieno, B., Nagy, G. and Pienovi, C., 1991. On sorting triangles in a Delaunay tessellation. *Algorithmica*, 6, 522-532.
- De Floriani, L. and Puppo, E., 1995. Hierarchical triangulation for multiresolution surface description. *ACM Transactions on Graphics*, 14(4), 363-411.
- De Floriani, F., Puppo, E., Magillo, P., 1995. Technical aspects of spatial data. IN: Frank, A., (ed.), *Geographical Information Systems - Materials for a Post-Graduate Course - Vol. 2: GIS Technology*. Department of Geoinformation, Technical University of Vienna, 395-452.
- Devillers, O., 1998. Improved incremental randomized Delaunay triangulation. *Proceedings, 14th Annual Symposium on Computational Geometry*, 106-115.
- Dutton, G., 1989. Modelling locational uncertainty via hierarchical tessellation. In: M. Goodchild and S. Gopal (eds.), *Accuracy of Spatial Databases*, Taylor and Francis, 125-140.
- Dutton, G., 1991. Polyhedral hierarchical tessellations: The shape of GIS to come. *Geo Info Systems*, 1(3), 49-55.

- Dutton, G., 1996. Encoding and handling geospatial data with hierarchical triangular meshes. In Kraak, M.J. and Molenaar, M. (eds) *Advances in GIS Research II (Proceedings, SDH7, Delft, Holland)*. Taylor & Francis, London, 505-518.
- Fekete, G., 1990. Rendering and managing spherical data with sphere quadtrees. *Proceedings, IEEE Visualization '90*, IEEE Computer Society Press, 176-186.
- Gold, C.M., Charters, T.D. and Ramsden, J., 1977. Automated contour mapping using triangular element data structures and an interpolant over each triangular domain. *Computer Graphics*, 11, 170-175.
- Gold, C.M., 1988. PAN Graphs: an aid to GIS analysis. *International Journal of Geographical Information Systems*, 2, 29-42.
- Gold, C.M., 1996. An event-driven approach to spatio-temporal mapping. *Geomatica*, 50, 415-424.
- Gold, C.M. and Condal, A.R., 1995. A spatial data structure integrating GIS and simulation in a marine environment. *Marine Geodesy*, 18, 213-228.
- Gold, C.M. and Cormack, S., 1987. Spatially ordered networks and topographic reconstructions. *International Journal of Geographical Information Systems*, 1, 137-148.
- Gold, C.M., Remmele, P. and Roos, T., 1995. Voronoi diagrams of line segments made easy. In: Gold, C.M. and Robert, J.M. (eds.). *Proceedings, 7th. Canadian Conference on Computational Geometry, Quebec*, 223-228.
- Goodchild, M.F., 1989. Tiling large geographical databases. *Symposium on the Design and Implementation of Large Spatial Databases, Santa Barbara, California*, 137-146. Springer-Verlag, Berlin.
- Goodchild, M.F., Shiren, Y. and Dutton, G., 1991. Spatial Data Representation and Basic Operations for a Triangular Hierarchical Data Structure. Santa Barbara. NCGIA Tech. Paper 91-8.
- Guibas, L. and Stolfi, J., 1985. Primitives for the manipulation of general subdivisions, and the computation of Voronoi diagrams. *ACM Transactions on Graphics*, 4, 74-123.
- Jones, C.B., Ware, J.M. and Bundy, G.L., 1992. Multiscale spatial modelling with triangulated surfaces. *Proceedings, Spatial Data Handling 5, vol. 2*, 612-621.
- Lukatella, H., 1989. Hipparchus data structure: Points, lines and regions in a spherical Voronoi grid. *Proc. Ninth International Symposium on Computer-Assisted Cartography (Auto-Carto 9)*, Baltimore, 164-170.
- Mioc, D., Anton, F., Gold, C.M., and Moulin, B., 1999. "Time travel" visualization in a dynamic Voronoi data structure. *Cartography and Geographic Information Science*, 26(2), 99-108.
- Molenaar, M., 1998. *An Introduction to the Theory of Spatial Object Modelling for GIS*. Taylor and Francis, London.
- Okabe, A., Boots, B. and Sugihara, K., 1992. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley and Sons, Chichester.
- Otoo, E.J. and Zhu, H., 1993. Indexing on spherical surfaces using semi-quadcodes. *Advances in Spatial Databases. Proceedings, 3rd International Symposium, SSD'93, Singapore*, 510-529.
- Pullar, D., 1994. *A Tractable Approach to Map Overlay*. Ph.D. Dissertation, University of Maine at Orono, 134 p.
- Roos, T., 1991. *Dynamic Voronoi diagrams*. Ph.D. Dissertation, University of Würzburg, 213 p.
- Sedgewick, R., 1983. *Algorithms*. Addison-Wesley, Reading, Mass., 551 p.
- Van Oosterom, P., 1993. *Reactive Data Structures for Geographic Information Systems*. Oxford University Press.
- Watson, D.F., 1988. Natural neighbor sorting on the n-dimensional sphere. *Pattern Recognition*, 21(1), 63-67.

Worboys, M.F., 1995. GIS - A Computing Perspective. Taylor and Francis Ltd., London.

Figure 1. The earth and global topology.

Figure 2. The Quad-Edge data structure.

Figure 3. The Splice Operation

Figure 4. Direction Cosines

Figure 5. Orthographic projection, Voronoi cells of points on continent boundaries.

Figure 6. Regular Voronoi cells on the sphere

Figure 7. Portion of the Voronoi diagram and its components on the surface of unit sphere.

Figure 8. Initializing the Delaunay triangulation.

a) Choice of six initial points at the intersection of the axes of the coordinate system with the surface of the unit sphere

b) Projection of the eight triangles on the surface of the sphere and creation of the eight initial spherical triangles.

Figure 9. \hat{N} is the circumcentre of the triangle ABC and the point N is the centre of the spherical cap cut by the plan ABC .

Figure 10. Circumcircle and Circumcentre of a spherical triangle.

Figure 11. Sphere surface partitioning:

a) Voronoi diagram

b) Delaunay triangulation.

Figure 12. Stereographic projection of Voronoi diagram and Delaunay triangulation for a set of points.