CONSIDERATIONS FOR THE DESIGN OF A MULTIPLE REPRESENTATION GIS

David B. Kidner¹ and Christopher B. Jones²

¹ Department of Computer Studies University of Glamorgan Pontypridd Mid Glamorgan CF37 1DL, UK

> ² Department of Geography University of Cambridge Downing Place Cambridge CB2 3EN, UK

ABSTRACT

Maintenance of a multiple representation GIS using data from a variety of sources at differing scales requires update processing that can recognise equivalences and differences between new data and stored representations. Factors that may be taken into account when establishing equivalence and difference include a) measures of similarity of the location of new and stored geometry, having regard to positional error; b) comparison of classification and names; and c) comparison of geometric shape parameters. Decisions about update of representations of equivalent real-world phenomena may depend upon the capabilities of automatic generalisation procedures that could be used to derive one scale of representation from another. Access to data at different levels of detail, for answering queries and for processing updates, is facilitated by the use of multiresolution data structures.

INTRODUCTION

A multiple representation GIS allows the same real-world phenomena to be represented in different ways, at different scales and with different levels of accuracy. Many of the problems associated with their use refer to the changes in geometric and topological structure of digital objects which occur with the changing resolution at which those objects are encoded for computer storage, analysis and depiction (NCGIA, 1989, 1993).

Many organisations concerned with spatially-referenced information are faced with the need to integrate data from a variety of sources representing phenomena at a range of scales and in some cases at a range of points in time (NCGIA, 1989, 1993). Data received from a single source, such as a national topographic mapping agency may well be accompanied by indications of changes and updates from previous versions of the data. When working with multiple source data however, problems can arise in deciding which new data items should replace existing stored data and which should supplement existing data. When the same real world phenomena are represented at significantly different scales, decisions may need to be taken to determine whether multiple representations should be stored or whether a small scale representation could be automatically derived from a large scale representation using automatic generalisation procedures.

Ideally a GIS should include sufficient intelligence to recognise equivalences between spatial representations at different scales and different times and either automatically make decisions on appropriate update strategies, according to pre-specified rules, or else assist the user in making decisions by highlighting equivalences and differences between representations.

In a large database, efficiency in storage and access to multiscale and multiple representation data can be provided by multiresolution data structures providing progressive access to increasing levels of detail. If data for different themes and scales are to be combined in a flexible manner, it is also desirable to provide facilities for automated generalisation whereby appropriate levels of detail are selected and representations are modified to suit particular purposes and to avoid visual conflict between cartographic symbology.

An overview of a possible database architecture to support such a multiscale, multiple representation GIS was provided by Jones (1991), in which it was proposed that a deductive database architecture might be employed in combination with multiresolution data structures and specialised processors for performing update and scale-variable retrieval. In this paper we focus in more detail on some of the issues that arise in attempting automatic update of multiscale data that may be stored in multiple representations and with multiresolution data structures. Resulting strategies for update are being used to develop an experimental update processor for multiscale databases. This is being implemented using object-oriented programming with rule-processing.

STORAGE STRATEGIES FOR MULTIPLE REPRESENTATIONS

In a multiple representation database a many to many relationship may be established between real-world object descriptions, in the form of classifications and named and uniquely identified phenomena, and geometric-object descriptions consisting of sets of spatial primitives. Since real-world phenomena may be described at varying levels of generalisation, it is possible to envisage considerable complexity with multiple levels of overlapping, hierarchical, real-world object representations. Thus, for example, different types of overlapping administrative areas may refer to common lower level regions or topographic features. The different levels of each real world object hierarchy could refer to geometric objects that represented them spatially. Thus a high level real-world feature could have a relatively simple representation, such as a polygon, as well as the more detailed representations of its constituent parts. Individual items of geometry could be referred to from the various real-world objects that they were used to represent. Conversely, each geometric object can refer to the real-world objects that it is used to represent.

When extensive geometric objects are represented with relatively high accuracy, they may be subject to generalisation operators on retrieval. Efficiency in creating generalised representations can be gained by organising data by means of multiresolution data structures (e.g. van Oosterom, 1990, 1991; Becker et al, 1992; Jones and Abraham, 1986, 1987; Ware and Jones, 1992). Most of these data structures implemented to date make use of line or surface simplification procedures to categorise constituent vertices with different levels of scale significance.

In the Multi-Scale Line Tree (MSLT) of Jones and Abraham (1986, 1987), vertices of linear geometry are classified within error bands as defined by the Douglas algorithm (Douglas and Peucker, 1973). The resulting levels are spatially indexed using a quadtree scheme. Becker et al (1992) employ a related approach in which each level is spatially accessed by an R-tree type spatial indexing mechanism that accesses short sections of linear geometry. In the Reactive-tree, van Oosterom (1991) organises the vertices of entire line segments in a hierarchy based on classification of the vertices using the Douglas algorithm. Each complete line is given a priority and is spatially indexed with an R-tree type data structure. In the Multiscale Topographic Surface Database (MTSD), Ware and Jones (1992) integrate point, line and polygon objects within a hierarchically-structured terrain model based on constrained Delaunay triangulation. The vertices of objects embedded in the surface are classified by a combination of the Douglas algorithm, providing lateral error control, and the surface simplification algorithm used by De Floriani (1989) for vertical error control. In a recent development of the MTSD, flexibility in access is combined with considerable data storage savings by using the Implicit TIN approach to triangulated surface storage (Kidner and Jones, 1991), in which only vertex coordinates are stored explicitly, while the topology of the constrained triangulation is determined at the time of retrieval (Jones, Kidner and Ware). Using this approach, in which no explicit triangulation data are stored, the surface can be constrained with only those features that are relevant to the particular user query.

Multiresolution data structures of the sort referred to have been criticised for their dependence upon algorithms such as that of Douglas and Peucker (1973), which cannot be guaranteed to provide cartographically sound generalisations (Visvalingam, 1990). In practice, by building the data structure by means of algorithms that classify constituent vertices progressively according to their contribution to increasing locational accuracy, the data structure is able to provide efficient access to a representation that at least satisfies quantitative criteria relating to resolution. Generalisation may be regarded as a separate process that operates upon the retrieved geometry. Clearly the data structures enable as much of the stored geometry to be retrieved as is required by the generalising processor, which may then transform the retrieved data. The distinction is one that is made by Brassel and Weibel (1988) for example, who emphasise the importance of separating a Digital Land Model (DLM), based on survey data, from a Digital Cartographic Model, that represents a transformation of DLM data to a form suitable for effective visualisation. (This is not to argue however that there may not be algorithms that are preferable to that of Douglas and Peucker in the ranking order that is placed upon the vertices of a line.)

METADATA

Effective storage and maintenance of multiple representations will depend upon recording metadata relating both to real-world and geometric descriptions. If data originate from several sources, a record must be maintained of the source with regard to the types of classification and nomenclature schemes, along with dictionaries of codes used to describe the real-world attributes. 'Metadata relating to geometry include the positional accuracy of the survey method, the precision of recorded coordinates, the date of original survey, the coordinate system used, the geodetic datum, any transformations known to have been applied to the data and, specifically for data acquired by digitising existing maps, the map series, the scale of the map, the map projection, the positional error in digitising and the date of production of the map.

Appropriate interpretation of metadata relating to classification and nomenclature could depend upon the storage of rules that encode the hierarchical nature of individual classification schemes and the relationships between particular categories or types in different schemes. Rules of this sort are required to assist in determining the equivalences and differences between new data and stored data. Deductive databases can provide the facility to encode and execute such rules (Jones, 1993).

UPDATE SITUATIONS IN A MULTIPLE REPRESENTATION DATABASE

Maintenance of multiple versions of scale-variable data introduces considerable complexity to the data update procedures. Operations that may take place on update may be reduced in principle to those of adding a new representation and of deleting an existing representation. However, the word representation here must be interpreted as covering the range between detailed geometric coverages of a region of space and individual geometric primitives representing either localised phenomena or generalised descriptions of extensive phenomena. Furthermore, addition of data may involve attempts to unify the new data with existing data both geometrically, by edge matching, and semantically by recording situations in which the new geometry represents the same phenomena as does an existing representation.

Decisions about whether to add or delete scale-variable representations will depend on several factors. In a multiscale database, one of the most important of these factors is that of whether one representation can easily be derived from another. With the current status of generalisation software there is still very limited capacity to perform generalisation entirely automatically and thus there is often justification for storing versions of data at various levels of generalisation. It may be remarked though, that when integrating different datasets at retrieval time it would probably, never in practice be possible to have access to all possible pre-generalised representations, since the nature of generalisation may depend upon graphic interaction between specific items of data that may vary from one query to the next. At present there is certainly some limited capacity for entirely automatic generalisation, in particular for linear feature simplification and for scale-dependent selection of data, and this capacity could be exploited in a multiple representation GIS. To do so however could require that quantitative limits were placed on the ranges of scales over which generalisation could take place. This would allow update decisions to be taken to justify multiple representation.

Another major factor affecting the decision to store multiple representations is a knowledge of whether new data are equivalent to existing stored data. If data can be shown to be equivalent in that they represent the same real-world phenomena, then storage of multiple representations can be addressed from the previous point of view of whether or not existing data can be derived from other data automatically. If new data items can be proven to be separate in location from existing data then it may be assumed that they should be stored. As indicated earlier, the problem of establishing equivalence is most likely to arise when working with data from multiple sources, that may be at different levels of generalisation. Thus some datasets may record the presence of features that were omitted in another dataset, due to some process of feature selection. Equally, timevariant datasets may record differing features, such as administrative boundaries, roads and buildings, due to modifications or new developments.

The process of recognition of equivalence and difference is complicated by the fact that geometric location is always accompanied by uncertainty, the degree of which will vary between data recorded at different scales. Before considering methods of determining equivalence and difference we summarise below some of the factors that must be taken into account.

Differences in accuracy and generalisation between new data and stored data:

- new data may be less generalised than a stored representation
- new data may be equivalently generalised to a stored representation
- new data may be more generalised than a stored representation

The ability to derive one representation from another:

- new data may be derivable from a stored representation
- new data may be used to derive a stored representation
- new data may not be related to a stored representation

The locational relationships between new data and stored representations, all of which are established to some level of certainty:

- an entire new dataset may be separate from stored representations
- an entire new dataset may overlap a stored representation
- an entire new dataset may be edge-adjacent to a stored representation

- an entire new dataset may be partially edge-adjacent to a stored representation
- a new geometric primitive may be equivalent in location to an existing stored primitive
- a new geometric primitive may be separate from stored primitives

- a new geometric primitive may be edge-adjacent to stored primitives

The classification relationships between new and stored representations:

- the classification of new real-world objects may be the same as those of stored objects
- the classification of new real-world objects may be different from those of stored objects
- the classification of new real-world objects may be similar within an established level of certainty to stored objects.

TECHNIQUES FOR ESTABLISHING EQUIVALENCE AND DIFFERENCE

A key issue that emerges from consideration of the above factors is that of establishing equivalence of geometry that must be regarded as having inherent positional error. Since all stored geometry must be regarded as representing some real-world phenomena, methods for establishing equivalence can attempt to exploit both semantic and geometric data. Certain pieces of geometry may represent two or more real-world phenomena. Thus a single line segment could represent both the centre of a river and an administrative boundary. The same line would be less likely to represent a railway. Evidence for equivalences can be built up on the basis of multiple criteria that include the classification, the location and the shape characteristics.

Locational equivalence needs to encompass the situations of complete overlap, partial overlap (in which case the equivalent part must be identified) and adjacency, in which it must be determined whether two pieces of geometry are continuous, i.e. two parts of the same real-world phenomenon.

When comparing the location of geometric objects, it is essential that they are regarded as fuzzy in the sense that all coordinate data have associated errors. Error and uncertainty have always been a feature of cartographic information. However, the problem of handling error is compounded further within a multiple representation GIS. Source accuracy indices, metadata and functions are required to propagate these values to document the quality of operations (Lanter and Veregin, 1990). The uncertainty of digital map features may be characterised in terms of regions of locational probability around their edges, which are adaptable depending upon the level of certainty with which one wishes to access fuzzy objects (Dutton, 1992, Maffini et al, 1989). Brunsdon et al (1990) present a review of methods for

⁻ a new geometric primitive may be partially equivalent to stored primitives

handling error propagation, including the traditional 'epsilon' approach and a technique based on Monte Carlo simulation.

A variety of techniques are available to assist in determining whether two fuzzy spatial objects appear to be equivalent in location. The least computationally demanding method of determining separateness is that of comparison of the extents of the objects as indicated by minimum bounding rectangles. To give the method some reliability, the extents must be expanded to take account of the maximum error associated with the coordinates. If that test failed to establish separateness, in the case of a line, an extent oriented parallel to the 'anchor line' connecting start and end points could be used. If this resulted in overlap, then buffer zones could be created around each object and their intersections determined. The percentage overlaps relative to each object would constitute a measure of equivalence, though interpretation of this measure would depend upon several factors that include the relative dimension of the two objects and the magnitude of errors assumed for each object. Determination of the nature of partial equivalence of lines would require comparison of buffers associated with constituent vectors.

Analysis of levels of certainty of locational overlap using vector defined buffers could be expected to result in considerable computational overheads. These overheads could be reduced somewhat by working with simplified versions of the objects to be compared. Thus for example a line simplification algorithm might be used to reduce both objects to a similar level of detail. If stored data were represented by a multiresolution data structure, the simplified version of the line could be readily accessed, along with the corresponding error band. The method is thus comparable to that of operations upon strip trees (Ballard, 1981), though the assumption here is that buffer zones would take account of positional error, rather than just the error tolerance used in a simplification algorithm.

Alternative methods of location-based comparison could employ raster representations of the geometric objects. If rasterisation was performed to a relatively high resolution, then each pixel could be associated with a certainty measure, based on a probability distribution function centred on the stored geometry. When the raster representations were overlaid, these measures could then be combined to establish levels of certainty of overlap.

Locational comparison methods can provide valuable evidence of the equivalence of geometry. However, when the certainty level of comparisons is not very high, for example, due to a systematic locational shift, other methods, that characterise the geometry of objects, may also be considered to help accumulate evidence for similarity or difference. A major difference in geometric signature might be used to establish difference, without performing more detailed locational comparisons. Below are listed some examples of geometric parameters. Some of these are based on the line signature parameters proposed by Buttenfield (1991), and McMaster's (1986) measures for evaluating linear simplification.

- 1. Minimum bounding rectangle dimensions.
- 2. Line length.
- Anchor line length, measuring the Euclidian distance between the first and last coordinates of a line.
- 4. Bandwidth, based on the maximum perpendicular deviations of coordinates from the anchor line. Bandwidth may also be standardised by division by the anchor line length.
- Segmentation, which is the displacement along the anchor line of the point of maximum perpendicular distance.
- Error variance, which is the square root of the sum of the squared distances of line vertices from their corresponding anchor line, divided by the number points.
- Concurrence, which is related to the number of times that a line crosses its anchor line.
- Sinuosity ratio, which is the ratio between the line length and the anchor line length.
- Average angular change, which may be based on the angles between successive vectors defining a line.
- 10. Curvilinearity, which represents the tendency for successive line vectors to change direction positively or negatively.
- 11. Average separation of successive vertices in a line.

12. Number of vertices in a line.

- 13. Polygonal area.
- 14. Perimeter/Area ratio.

In attempting to use parameters of the sort listed above, it must be appreciated that some of them are scale-dependent. Thus when comparing lines from different scale representations, less, or no, emphasis should be placed on parameters that depended on the number of vertices, on measurements of line length or on frequency of directional change. It is however possible to compensate for scale differences. One approach is to apply factors that might be a function of line sinuosity or fractal dimension. Another is to simplify the more detailed representation in order to render the levels of generalisation comparable. In the case of line segments, this could be done with a line simplification algorithm or by retrieval from a multiresolution data structure, as indicated above in the context of locational comparison.

A RULE-BASED OBJECT-ORIENTED SYSTEM FOR MULTISCALE DATABASE UPDATE

Work is currently underway on the development of a rule-based system to control update in a multiscale, multiple representation database. It is implemented with an object-oriented programming system that includes rule processing. Programming objects that are employed include those of update dataset objects, containing both geometry and classification data; realworld objects equivalent to entries in a permanent database; and geometric objects also corresponding to permanently stored data. All of these object classes include slots to store relevant metadata derived from a permanent database. The constituents of dataset objects are matched against real world and geometric objects in an attempt to establish equivalences and differences between them and to create new real-world and geometric objects that may subsequently be loaded to the permanent database.

Methods are being implemented to calculate geometric shape parameters that will be stored, if calculated, within relevant slots of the geometric objects and the dataset objects. Locational comparison methods are also being implemented, based initially on vector-based techniques. Experiments are being carried out using rule processing to control execution of the location and shape descriptor methods and to compare classifications. The rules use results of the comparisons to accumulate measures of equivalence between components of new datasets and database representations and make decisions about database update operations of addition and deletion.

CONCLUSIONS

This research has been aimed at deriving new methods for handling multiple representations within a GIS. These aims are in accordance with the views of others tackling the same problem (NCGIA, 1989, 1993), namely, that two main areas in which research should be focused are a) database issues, such as the need to organise multiple topological and metrical versions for efficient access; and the implementation of linkages between multiple representations; and b) generalisation issues, such as formalising digital feature description models, flexible definitions of resolution for data sets, and rules for map generalisation. This paper addresses these issues with respect to maintaining a multiple representation GIS for update processing in particular.

The process of updating a GIS database using multisource data is complicated by problems of recognising the equivalence and difference between new data and that already stored. Determination of equivalence is subject to error due to locational inaccuracy at varying scales and to the possibility of differing classification schemes. A rule-based approach is being adopted in an experimental system to evaluate strategies for multiscale database update. The rules make use of methods that implement geometric procedures for comparing location and for comparing a variety of shape descriptors. The system is being developed in the context of a database capable of providing multiresolution access to stored geometry.

ACKNOWLEDGEMENTS

This research is supported by grant GR/F96288 from the Science and Engineering Research Council and by the Ordnance Survey, Southampton.

REFERENCES

- Ballard, D.H. (1981) "Strip Trees: A Hierarchical Representation For Curves", Comm. of the ACM, Vol. 24(5), May, pp. 310-321.
- Becker, B., H. Six & P. Widmayer (1991) "Spatial Priority Search: An Access Technique for Scaleless Maps", Proc. of the 1991 ACM SIGMOD, Denver, Colorado, May 29-31, (Eds: J. Clifford & R. King), ACM SIGMOD Record, Vol. 20(2), June, pp. 128-137.
- Brassel, K.E. & R. Weibel, (1988), "A Review and Conceptual Framework of Automated Map Generalization", Int. Journal of Geographical Information Systems, Vol. 2(3), pp. 229-244.
- Brunsdon, C., S. Carver, M. Charlton & S. Openshaw (1990) "A Review of Methods for Handling Error Propagation in GIS", Proc of the 1st European Conference on GIS (EGIS'90), Amsterdam, Netherlands, pp. 106-113.
- Buttenfield, B.P. (1991) "A Rule for Describing Line Feature Geometry", in Map Generalization: Making Rules for Knowledge Representation, (Eds: B.P. Buttenfield & R.B. McMaster), Longman, London, U.K., pp. 150-171.
- De Floriani, L. (1989) "A Pyramidal Data Structure for Triangle-Based Surface Description", IEEE Computer Graphics & Applications, March, pp. 67-78.
- Douglas, D.H. & T.K. Peucker (1973) "Algorithms for the Reduction of the Number of Points Required to Represent a Digitized Line or Its Caricature", The Canadian Cartographer, Vol. 10(2), December, pp. 112-122.
- Dutton, G. (1992) "Handling Positional Uncertainty in Spatial Databases", Proc. of the 5th Int. Symp. on Spatial Data Handling, Charleston, USA, pp. 460-469.
- Jones, C.B. & I.M. Abraham (1986) "Design Considerations for a Scale-Independent Cartographic Database", Proc. of the 2nd Int. Symp. on Spatial Data Handling, Seattle, Washington, July 5-10, pp.384-398.
- Jones, C.B. & I.M. Abraham (1987) "Line Generalisation in a Global Cartographic Database", Cartographica, Vol. 24(3), pp.32-45.
- Jones, C.B. (1991) "Database Architecture for Multi-Scale GIS", Proc. of the 10th Int. Symp. on Computer Assisted Cartography (Auto-Carto 10), ACSM/ASPRS, Baltimore, Maryland, March 25-28, pp.1-14.
- Jones, C.B., "Deductive Databases and the Representation of Hierarchies in Geographical Information", submitted for publication.
- Jones, C.B., D.B. Kidner & J.M. Ware, "The Implicit TIN and multiscale spatial databases", submitted for publication.

- Kidner, D.B. & C.B. Jones (1991) "Implicit Triangulations for Large Terrain Databases", Proc. of the 2nd European Conference on GIS (EGIS'91), Brussels, Belgium, pp. 537-546.
- Lanter, D.P. & H. Veregin (1990) "A Lineage Meta-Database Program for Propagating Error in Geographic Information Systems", Proc. of GIS /LIS'90, pp. 144-153.
- Maffini, G., M. Arno & W. Bitterlich (1989) "Observations and Comments on the Generation and Treatment of Error in Digital GIS Data", in Accuracy of Spatial Databases, (Eds: M. Goodchild & S. Gopal), Taylor & Francis, London, pp. 55-67.
- McMaster, R.B. (1986) "A Statistical Analysis of Mathematical Measures for Linear Simplification", The American Cartographer, Vol. 13(2), pp. 103-116.
- NCGIA (1989) "Multiple Representations", Scientific Report for the Specialist Meeting of the NCGIA Research Initiative 3, (Eds: B.P. Buttenfield & J.S. DeLotto), Dept. of Geography, SUNY at Buffalo, Feb. 18-21, Report 89-3, 87 pages.
- NCGIA (1993) "Multiple Representations", Closing Report for NCGIA Research Initiative 3, (Ed: B.P. Buttenfield), Dept. of Geography, SUNY at Buffalo, April, 27 pages.
- van Oosterom, Peter (1990) "Reactive Data Structures for Geographic Information Systems", Ph.D. thesis, Dept. of Computer Science, Leiden University, The Netherlands, 197 pages.
- van Oosterom, Peter (1991) "The Reactive-Tree: A Storage Structure for a Seamless, Scaleless Geographic Database", Proc. of the 10th Int. Symp. on Computer-Assisted Cartography (Auto-Carto 10), ACSM/ASPRS, Vol.6, Baltimore, Maryland, March 25-28, pp.393-407.
- Visvalingam, M. & J.D Whyatt (1990), "The Douglas-Peucker Algorithm for Line Simplification: Re-evaluation through Visualisation", Computer Graphics Forum, Vol. 10(3), pp. 225-235.
- Ware, J.M. & C.B. Jones (1992), "A Multiresolution Topographic Surface Database", Int. Journal of Geographical Information Systems, Vol. 6(6), pp. 479-496.