MULTI-SCALE SPATIAL DATABASE DESIGN FOR ONLINE GENERALISATION

Christopher B. Jones, Alia I. Abdelmoty, Michael E. Lonergan,
Peter van der Poorten and Sheng Zhou

School of Computing
University of Glamorgan
Tel +44 1443 482722
Fax +44 1443 482715
cbjones@glam.ac.uk

ABSTRACT

Internet-based access to a wide variety of spatially-referenced information services has introduced a strong motivation to develop interactive map interfaces that provide generalisation of data on demand. This paper presents a scheme for online generalisation that combines multi-scale data access with dynamic conflict resolution facilities that ensure that a displayed map is clearly legible. The approach focuses particularly upon the problem of modelling multi-scale objects and of maintaining topological consistency of linear and polygonal features across continuous scale ranges. Relevant techniques for topologically consistent generalisation of multiple features and for displacement-based proximal conflict resolution are introduced.

Keywords: map generalisation, multiresolution, topological consistency

1 Introduction

There are several situations in which it is desirable to generate a map very quickly at some arbitrary level of generalisation. Notable examples relate to the provision of information about tourist facilities, local government, retail outlets, the natural environment, cultural history and transport. The map should adapt its content and level of detail in direct response to user actions. At present, such requirements are typically met, if at all, by a compromise solution in which the map data are displayed at a small number of fixed scales (see for example the Microsoft Encarta Interactive World Atlas on CDROM, or the MapQuest internet site).

The alternative, of generalising maps on the fly, is a considerable technical challenge. In fact the whole field of automated generalisation is still in quite an immature state, irrespective of response times. Here we focus on two major problems that face the development of systems that are genuinely scale-variable, with a view to the provision of what may be termed, variously, *on-demand*, *online*, *on-the-fly* or *dynamic* map generalisation. The first problem is the design of multi-scale spatial data access schemes. The other is the design of online symbolisation and annotation facilities that can visually render the retrieved data and ensure that it is legible at the given display scale.

Multi-scale spatial data access schemes facilitate retrieval of geometric data across a range of resolution levels. Their design involves balancing access speed with storage requirements. This can be achieved with multiresolution data structures that organise the component geometry hierarchically in a non-duplicative manner. The approach enables increase in detail by the progressive addition of intermediate component geometry. In theory, an alternative approach to multi-scale data access is to apply generalisation procedures directly to a single large-scale representation. This however is not regarded as a practicable solution for on-demand generalisation in the near future, due to the potentially massive computational costs, for large datasets, and the lack of adequate generalisation procedures. There is a fairly long, though limited, history of non-duplicative multiresolution data access schemes (e.g. Jones and Abraham, 1996; Cromley, 1991; Becker et al, 1991; van Oosterom, 1994, 1995; Jones et al, 1994). Aspects of these schemes concerned with linear features are all based on the use of a sub-setting point selection procedure, notably that of Douglas and Peucker (1973), and may be regarded as inspired by hierarchical multiresolution data structures intended originally for image processing, exemplified by the strip tree (Ballard, 1981).

A significant limitation of these multiresolution access schemes is that they represent individual features that have been 'pre-generalised' independently of features with which they may be displayed. This can result in spatial conflicts between neighbouring features when displayed at smaller (derived) scales than their source representation. van Putten and van Oosterom (1998) recognise this problem and have addressed it in the context of the GAP-tree which provides multiple detail levels for areal partition. The focus there however is on the presence and absence of regional subdivisions rather than the representation of the component boundaries. The authors do refer to providing multiresolution access by means of the BLG-tree, but do not appear to address the problem of potential conflict between adjacent BLG-tree representations. Bertolotto and Egenhofer (1999) present a topologically consistent scheme for multiresolution access to multiple feature types, but do not deal with the problem of progressive multiresolution access to individual linear features. The progressive access scheme of Buttenfield (1999) proposes dealing with topological conflicts by incorporating Saalfield's (1999) line generalisation procedure. In Saalfield's approach topological conflicts are avoided by local re-instatement of points selected by the Douglas-Peucker algorithm. This is only a partial solution in that conflict is avoided by not generalising (or reducing the generalisation of) the line in the region of conflict.

The second problem, of online symbolisation and annotation, is concerned with the need for procedures that detect the presence of, and resolve, visual conflicts due to overlap or irresolvable closeness of map symbols and text. The assumption here is that retrieved geometry will be topologically consistent but that graphic conflicts may arise when the geometry is symbolised. Line symbols may be too close or may overlap, while area symbols may be too small. There is a need therefore for online conflict resolution procedures to ensure satisfactory levels of cartographic legibility. We will focus here on the problems induced by symbol conflicts (of overlap and proximity). Text legibility is achieved by a labelling postprocessor to be applied after symbol conflict resolution.

Resolution of graphic conflicts may be achieved by operations such as displacement, elimination and amalgamation. There are several examples of the implementation of these operators at a local level (e.g. Nickerson, 1988; Jones et al, 1995; Mackaness, 1994), but progress still remains limited in applying these operators to a complete map and handling the consequent propagation of conflict. Notable recent advances have however been made in global application of displacement by itself, using for example methods of least squares with constraints (Harrie, 1999), finite element analysis (Hojholt, 1998), simulated annealing (Ware and Jones ,1998) and multi-agent negotiation (Baeijs et al ,1996). This work may be regarded as in its infancy with very little progress on the combination of resolution operators. It may be noted that some commercial systems certainly enable multiple operators to be applied to a map in batch mode, but there is little facility for systematically handling the problems of conflict propagation that result.

In this paper we describe a scheme for multiple scale access and display. It integrates multiresolution storage techniques, based on prior application of generalisation algorithms, with procedures for online rendering and conflict resolution. The approach adopted provides a balance between pre-computation and online computation and addresses specifically the issue of maintaining topological consistency of generalised representations. We give an overview of the scheme in Section 2, before describing the maintenance of topological compatibility between map features in Section 3. Section 4 summarises an approach to pre-computation with topological consistency using constrained Delaunay triangulations. It is followed by a review of progress on procedures for online conflict resolution in Section 5 and some concluding remarks in Section 6.

2 Overview of multi-scale access scheme for online generalisation

In this section we present an overview of a design for a multi-scale spatial data access scheme for online map generalisation. Figure 1 illustrates the design, which falls into three major parts, the database in the upper enclosing box, the user interface module at the lower left and the online conflict resolution module in the lower right enclosing box. Database construction requires that topological relationships between component objects be correct at some specified maximum resolution. In practice such data would be obtained from a survey or data acquisition organisation that pre-processed the data to ensure validity. This dataset is then processed to 'priority-label' the component geometry.

Priority-labelling means that all spatial objects and their sub-objects have a designated range of resolution, that is used to prioritise retrieval with regard to the needs of generalisation. The term resolution is used here to refer to a minimum size, in map grid units, of discernable features in the dataset. This value is then scaled to determine whether the corresponding data should be retrieved for a particular map. At the finest level of detail all vertices are retrieved. In theory a single vertex could have multiple resolution ranges if the labelling procedure was not based on simple subsetting, and our design caters for that possibility. The process of priority labelling is performed using topology-maintaining procedures, which are the subject of subsequent sections. A consequence of this is that they may be applied using different sets of neighbouring feature types. This gives rise to the concept of *contexts*. Thus a particular priority labelling is associated with the set of features that were present when priority labelling was performed.

The multi-scale spatial database is composed of Multiresolution Spatial Objects (MSOs). An MSO corresponds to the representation of a spatial phenomenon at a single range of scales. Within this range the representation may undergo changes in geometric detail, corresponding typically to the prior application of (topology-maintaining) geometric detail reduction procedures that produce the context-specific priority labelling.

An MSO has a type that is application-specific. The multi-scale database supports representation of objects that change significantly in form, due to collapse and amalgamation operations, by including explicit links between the related MSOs. Thus an MSO resulting from amalgamation refers to the source MSOs that were amalgamated, while an MSO resulting from a collapse in dimensionality points to the higher dimension MSO from which it was derived. The linkages are two-way, denoted by source object links and derived object links respectively.

A simple MSO references a multiresolution geometry object, referred to as an mgeo-object, while a composite MSO references other MSOs. An mgeo-object may itself be composite, referring to other mgeo-objects, with primitive mgeo-objects consisting of an isolated point, a chain and a simple polygon. Higher level mgeo-objects refer to primitive objects in association with topological data recording connectivity and left/right adjacency. The vertices of primitive mgeo-objects are labelled individually with priority ranges, as indicated above. They are accompanied by ordering data, required to reassemble the objects at a specified level of detail. The ordering data may be explicit via sequence numbers or implicit via the pointers of a directed graph. These implementation details are not dealt with further in this paper (the references in Section 1 provide examples of existing techniques).

In the lower left containing box of figure 1 are the components of the user interface module. The user specification of a map is a set of parameter values defining the areal extent of the map, the set of feature classes (types) to be displayed, the required scale, and the symbolisation and associated display constraints of the required feature types. Such data may be a combination of pre-specified (default) values and transitory input values that might be modified repeatedly in a display session. Values of the user specification relating to areal extent, type and resolution are used to formulate a query on the multi-scale database. The enclosing box in the lower right of figure 1 contains the online conflict resolution functions. Evaluation of the user query results in the retrieval of the highest resolution subset of geometry from the relevant MSOs that most closely approximates the query scale value. The retrieved geometry is transferred to a temporary store of multiresolution map objects (MMOs). Each MMO corresponds to a single MSO and stores only that part of the geometry of the corresponding MSO that is needed to display the feature at the required scale.

The current user specification parameters then determine which MMOs are required. Note that because the presentation module supports changes in the user specification of the map, there cannot be assumed to be a one-to-one correlation between the objects in the MMO store and those features that appear on the map. Assemblage of the required MMOs leads to the creation of an internal map of MMOs. The internal map is an updatable spatial data structure. The contents of the internal map are analysed to identify any violation of cartographic constraints as recorded in the user specification. These constraints include minimum symbol sizes and minimum symbol separations. A process of conflict resolution is then applied to resolve any existing conflicts. The procedure operates upon and updates the internal map of MMOs. When a satisfactory level of conflict resolution has been achieved the internal map is labelled and its contents passed back to the interface module to be rendered on a graphics display using specified symbol characteristics.

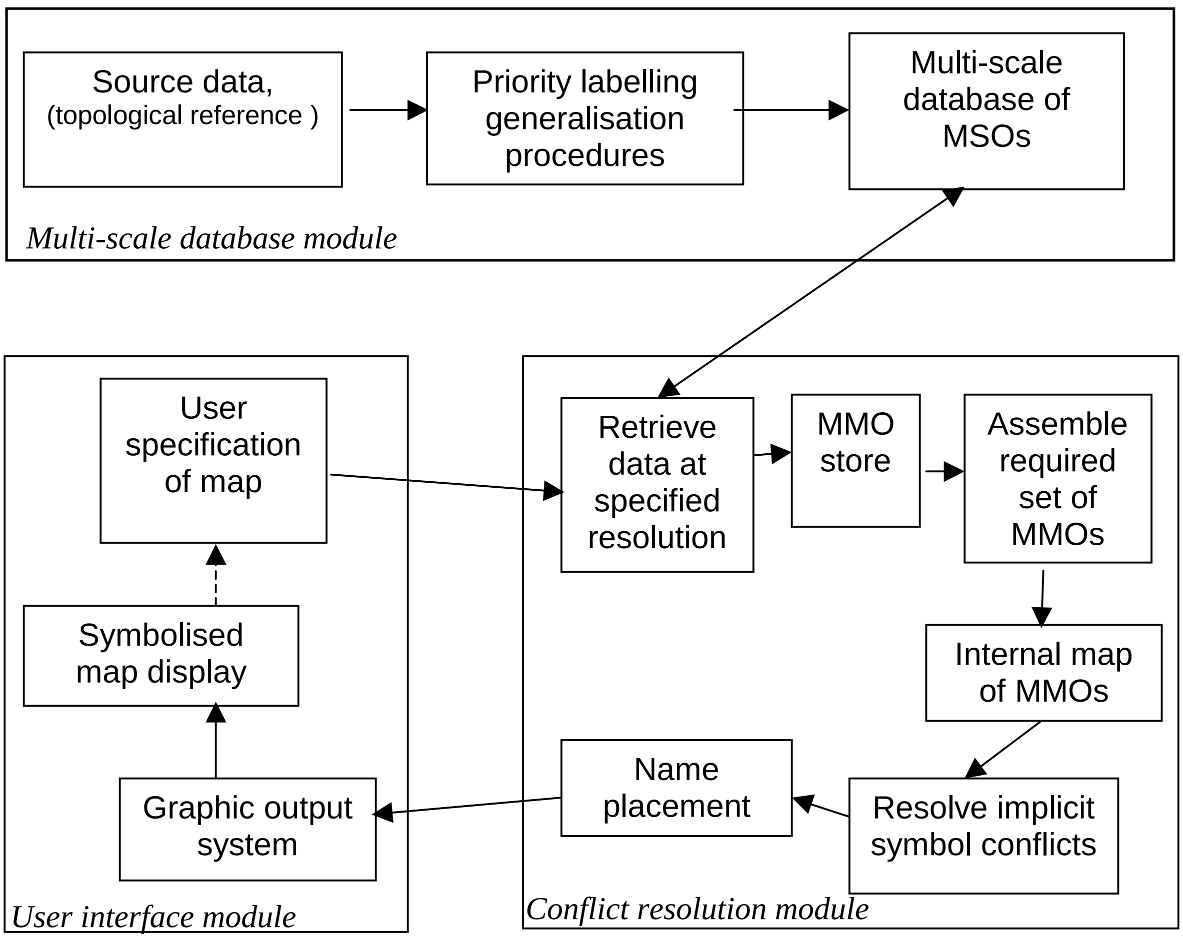

Figure 1. Overview of multi-scale access scheme for online generalisation

3 Neighbourhood Compatibility

The multi-scale database needs to support several levels of topological consistency. Data on topological consistency is supplemented with data that qualifies consistency with a minimum separation distance. The combination of topological consistency with the distance qualification is referred to as neighbourhood compatibility. We define three levels of compatibility. At the lowest level, referred to as *order-0* compatibility, an MSO has been pre-generalised to preserve internal topological consistency. For a point-referenced object this has no meaning, but for a linear feature or an areal feature with a linear boundary this level of compatibility guarantees that there are no cross-overs between originally disjoint lines at all levels of resolution at which the MSO may be represented (i.e. within its full resolution range). *Order-0* compatibility may be qualified with a distance value *d*, which records the minimum separation distance when different parts of the line approach each other. Lines displayed with a symbol width equal to this value will not overlap themselves. *Order-0* compatibility qualified by distance *d* is denoted *order-0:d* compatibility.

The next higher order of compatibility is between a feature and all other features of the same type. This is denoted by *order-1* compatibility. This means that for all available resolutions of representations of all features of the given type, no cross-overs are introduced. This degree of compatibility can again be qualified by the minimum separation distance between approaching parts of the same objects and any parts of disjoint objects. The distance for disjoint objects is simply the shortest distance between the two objects as measured for all parts of each of them. *Order-1* compatibility qualified by distance *d* is denoted *order-1:d*.

The next higher order of compatibility refers to topological consistency between a given feature and features of classes other than and including itself. This is referred to as *order-n* compatibility where $n>1$ indicates the number of types of feature (including itself) with which the given feature is compatible. Qualification with distance separations requires a distinction between distances for each pairing between the MSO and the respective compatible type. This is denoted *order-n:*$\{(t_i, d_i)\}$ where the set of tuples (t_i, d_i) indicates the compatible feature type and the minimum distance between the MSO and features of that type. In this case (t_1, d_1) refers to the relationship for the MSO's own type.

4 Compatibility Computation

The construction of a multi-scale database containing MSOs that conform to one or other of the above compatibility levels requires considerable overheads of pre-generalisation that maintain the respective topological consistency. Saalfield's (1999) procedure is a recent example of an approach to detecting and resolving conflict in the course of application of the Douglas Peucker algorithm. It uses convex hulls to assist in detection of conflicts, and resolves them, as indicated previously, by re-winding the Douglas-Peucker procedure to re-instate previously eliminated vertices until the conflict is resolved. Jones et al (1994) described an approach, also based on the Douglas-Peucker algorithm in which a more selective technique is adopted to identify only those previously eliminated vertices that are needed to resolve the conflict, rather than all those in reverse order until resolution is achieved. De Berg et al (1998) present a topologically consistent line generalisation algorithm based on that of Imai and Iri (1988) in which conflicts are avoided both with edges belonging to the feature being generalised and to isolated points in the neighbourhood.

One of the first papers to highlight the problem of topological inconsistency induced by line generalisation is that of Muller (1990). He presents a postprocessing approach whereby unacceptable closest approaches between parts of a line, resulting from application of the Douglas-Peucker algorithm, are detected and rectified using displacement vectors. He does not resolve topological conflicts, but the method is relevant here in that it is an example of a technique that might be of use in achieving specified levels of separation between topologically consistent features.

Those of the above methods that maintain topological consistency are intended for generalising an individual feature (in some cases in the presence of other fixed features), and hence could be used in principle to obtain a priority labelling of order-0. However, the generation of higher orders of compatibility requires a procedure for simultaneous generalisation of multiple features. In the current project this is achieved by applying triangulation-based techniques introduced in van der Poorten and Jones (1999). The method is based on creating a constrained Delaunay triangulation of the features to be generalised. Features are represented explicitly in an improved version of the SDS (simplicial data structure) scheme of Bundy et al (1995) (see also Ware et al, 1995 and Jones et al, 1995). The method is inherently topology maintaining. It is iterative and performs generalisation by successive elimination of line sub-features bounded by locally consecutive sequences of vertices. The sub-feature has an areal extent corresponding to a set of triangles that constitute a branch in the triangulation.

Figure 2 illustrates a constrained Delaunay Triangulation of a linear feature in which triangles have been shaded according to whether they a constitute leaf triangle at the tip of a branch (pale grey), a body triangle in the interior of a branch (not shaded) or a node triangle at the base of a branch (dark grey). A branch is composed of a leaf triangle and zero or more body triangles. Elimination of a feature is accompanied by its replacement by a single edge that forms the boundary between the feature and its associated node triangle. Since the sub-feature is bounded by edges belonging to a single linear object (not connected to any other object) and contains only free space, its elimination cannot induce topological inconsistency. Figure 3 illustrates the local effect of several iterations of the procedure for three linear features in close proximity. The result demonstrates that topological consistency has been maintained, but may be criticised for failure to maintain shape as well as might be expected.

Use of the approach is still the subject of experimentation and it is apparent that there are a number of options for enhancing the appearance of features undergoing generalisation by sub-feature elimination in this manner. For example, we have experimented with edge refinement that adds vertices for purposes of smoothing. It should be noted that the triangulation-based methods lend themselves to fast proximity measurement (Jones et al 1999). This assists in detecting proximal conflict (and in reporting the minimum separation distance of mutually generalised objects as required for distance qualified compatibility relations).

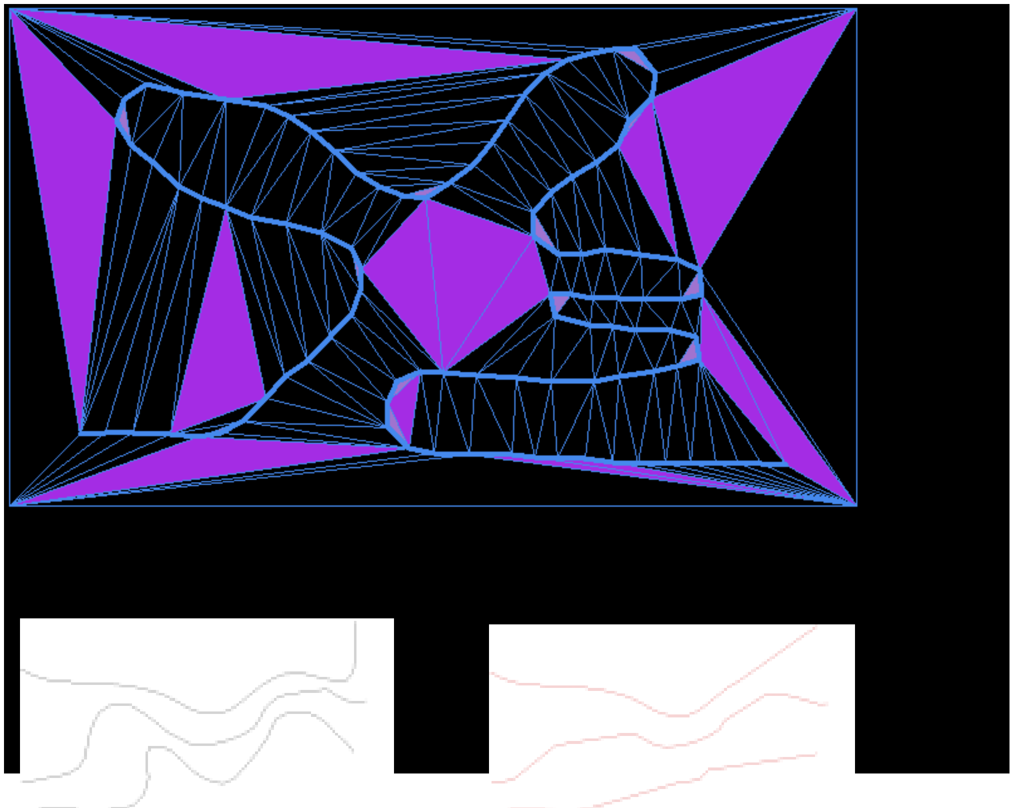

a b

Figure 3 Result of multi-feature topology preserving line procedure

5 Online Conflict Resolution

An integral component of the scheme is a procedure for resolving conflicts due to excessive symbol proximity. In the introduction we referred briefly to existing techniques based for example on least squares, finite element analysis, simulated annealing and multi-agent negotiation. In this project we are building upon and developing variations on the displacement-based dynamic conflict resolution procedures introduced in Ware and Jones (1998). The methods described there generated a set of candidate positions for each of the areal features under consideration. Gradient descent and simulated annealing search procedures were then applied to find individual candidate locations for each object that resulted in the least number of conflicts. A conflict was defined as the situation of two objects lying within the specified minimum tolerance distance of each other. Simulated annealing was found to be superior to the gradient descent method applied with regard to minimising the number of conflicts. A limitation of simulated annealing is that it is non-deterministic, since it depends upon random selection of candidate solutions, subject to probabilistic application of an evaluation function. It also involves examination of a potentially very large number of candidate solutions.

An alternative approach has been developed which is deterministic, greatly reduces the number of candidate solutions examined, but which incurs a higher level of computation to generate individual candidate solutions. It is an iterative method based on maximising the minimum distance between neighbouring objects (i.e. a *maxmin* approach). For each iteration one of the (two or more) objects that is closest to its neighbour is selected and a local displacement vector is calculated as a function of the 'forces' exerted by those neighbours that are closer than the minimum separation distance. The vector is derived from the maximum horizontal and vertical resolved components of the vectors representing the individual forces. The effect of the procedure is progressively to increase the size of the minimum separation distance on the map. In doing so it also reduces the number of conflicts as measured by the number object separations that are less than the threshold. The quality of the result is similar to that achieved by simulated annealing (Lonergan and Jones in press). The method is extensible in that it can be applied to displace 'plastic' linear and areal features. For all types of displacement however, when applied as described here, it may introduce misalignments of previously patterns of map features. Further work will experiment with various strategies for retaining pattern and structure in the course of displacement.

An extension of the existing procedures for conflict resolution is to introduce object elimination and object amalgamation. The problem is to decide which objects should be either eliminated or amalgamated, and which of the operations should be performed. A beneficial characteristic of the maxmin conflict resolution procedure is that it will generate clusters of objects, each of which is in proximal conflict with a neighbour. These clusters then become the natural focus for deletion and amalgamation operations. This is also the subject of ongoing research.

A major problem facing the implementation of online conflict resolution is the development of effective strategies for handling the multiresolution geometry that has been generated to the varying levels of neighbourhood compatibility. In the simplest case there would be a match between the user specification of feature types and a set of features types, i.e. a context, that had been pre-generalised as an ensemble (and the component geometry appropriately priority-labelled). If the separation distances corresponded to that of the symbol specifications then there would be little need for conflict resolution. When combining objects that have not been pre-generalised as an ensemble, topological consistency can only be assumed at the maximum resolution of objects with overlapping resolution ranges. In this case an online procedure could be applied to find the lowest acceptable resolution, pre-generalised versions that were topologically consistent, before applying topology-maintaining simplification procedures to derive the required level of detail.

6 Conclusion

This paper has presented a design for a cartographic system that supports applications requiring online map generalisation. The design has two major divisions, a multi-scale spatial database and an online presentation system with the latter subdivided into a user interface module and a conflict resolution module. The multi-scale database stores multiresolution spatial objects (MSOs) which correspond to map features that may be displayed over a specified resolution range. Each MSO references source

MSOs from which it may have been derived (in the case of object amalgamation) and derived MSOs of which it forms a part. The geometry of an MSO is priority labelled down to the vertex level to enable reconstruction of intermediate levels of generalisation of the object within its resolution range. Priority labelling of MSOs is performed to some specifed degree of compatibility with itself and its neighbours. Compatibility refers here to a combination of topological consistency and minimum separation distances between neighbouring objects. A technique for multiple feature pre-generalisation, using constrained Delaunay triangulations, has been introduced.

The online presentation system retrieves geometric representations from the multi-scale database that best approximate a user's requests. This geometry is then rendered using user-specified symbolisation. Legibility of the resulting map, with regard to minimum separation distances and minimum size objects, is ensured by means of an iterative conflict resolution procedure.

The design presented here is the subject of ongoing research. It is being implemented as a datablade in the Informix database and the initial prototype will serve as a workbench on which to experiment with several aspects of the design. Issues that are currently under consideration include strategies for hierarchical organisation of the priority-labelled geometry; determining an appropriate balance between pre-computation of compatible neighbours and online generalisation procedures; and extension of the existing online conflict resolution procedures to include elimination and amalgamation operators.

References

Baeijs, C., Y. Demazeau, et al. (1996). "SIGMA: Application of multi-agent systems to cartographic generalisation". *Agents Breaking Away, 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent Word, MAAMAW'96*, Einhoven, Springer, 163-176.

Ballard, D. (1981). “Strip trees: a hierarchical representation for curves.” *Communications of the ACM* 24: 310-321.

Becker, B., H.-W. Six, et al. (1991). "Spatial priority search: an access technique for scaleless maps". *ACM SIGMOD* 20(2): 128-137.

Bertolotto, M. and M. J. Egenhofer (1999). "Progressive vector transmission". *7th ACM Symposium on Advances in Geographic Information Systems*, Kansas City, MO, ACM Press.

Bundy, G. L., C. B. Jones, et al. (1995). "A Topological Structure for the Generalization of Large Scale Cartographic Data". *Innovations in GIS 2*. P. F. Fisher (ed), Taylor and Francis: 19-31.

Buttenfield, B. P. (1999). "Sharing Vector Geospatial Data on the Internet". *19th Conference of International Cartographic Association, Vancouver*, International Cartographic Association.

Cromley, R. G. (1991). “Hierarchical methods of line simplification.” *Cartography and Geographic Information Systems* 18(2): 125-131.

De Berg, M., M. van Kreveld, et al. (1998). “Topologically correct subdivision simplification using the bandwidth criterion.” *Cartography and Geographic Information Systems* 25(4): 243-257.

Douglas, D. H. and 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* 10(2): 112-122.
Harrie, L. E. (1999). “The constraint method for solving spatial conflicts in cartographic generalization.” *Cartography and Geographic Information Science* 26(1): 55-69.
Hojholt, P. (1998). "Solving local and global space conflicts in map generalisation using a finite element method adapted from structural mechanics". *8th International Symposium on Spatial Data Handling, Vancouver*, International Geographical Union.
Imai, H. and M. Iri (1988). "Polygonal approximations of a curve - formulations and algorithms". *Computational Morphology*. G. T. Toussaint (ed), Elsevier Science Publishers: 71-86.
Jaint Tcsstru". 1 Tf(l)-2.4(oI)1.1ci. (1 Tf(r)2.4(iAbrahiz)-7(a)]T(m(S)-486(P)-4.4(ol)-14(D)TJ/es)4.83aphic