MRDB APPLICATIONS FOR DATA REVISION AND REAL-TIME GENERALISATION

Mark HAMPE, Karl-Heinrich ANDERS, Monika SESTER

Institute of Cartography and Geoinformatics, University of Hannover, Germany
mark.hampe@ikg.uni-hannover.de
karl-heinrich.anders@ikg.uni-hannover.de
monika.sester@ikg.uni-hannover.de

ABSTRACT

This paper describes the ideas, advantages and methods maintaining a Multiple Representation Database (MRDB) for storing and serving spatial data. A Multi-resolution/representation-database (MRDB) can be described as a spatial database, which can be used to store the same real-world-phenomena at different levels of precision, accuracy and resolution. Furthermore these phenomena can be stored in different ways of presentation or symbolisation. There are several reasons for introducing an MRDB: On the one hand it allows for a multi-scale analysis of the data: Information in one resolution can be analysed with respect to information given on another resolution. A major reason for National Mapping Agencies to investigate and implement an MRDB is the possibility of propagating updates between the scales with the advantage that only the level with the highest resolution must be updated manually. Two projects using an MRDB will be presented. The EU-project GiMoDig develops methods for delivering geospatial data to a mobile user in real-time and for small displays. The other project, called WIPKA, aims at establishing a consistent representation of all topographic data sets in Germany (ATKIS) in an MRDB to enable an automatic incremental update process for the data sets.

1 INTRODUCTION

In a Multiple Representation Database (MRDB), different views on the same physical objects or phenomena can be stored and linked. This variety of sights can stem from different views of the world, different applications, as well as different resolutions. These lead to differences in the objects as such, i.e. in the semantics and in the geometry. In the paper we will concentrate on the design and maintenance of an MRDB and its use for two specific purposes. Real-time visualisation on small display devices is the main issue in the EU-project GiMoDig. The problem of automatic incremental update of topographic data sets of the National Mapping Agencies (NMA's) in Germany is one goal of the project WIPKA. First, we give an overview of principles and purpose of an MRDB, before presenting a proposal for the design. A main problem is the population of an MRDB – we will present approaches for introducing data into such a structure using generalisation operations and data matching. A concept for using the MRDB for update process is sketched in the end.

2 PRINCIPLES OF A MULTIPLE REPRESENTATION DATABASE (MRDB)

An MRDB can be described as a spatial database, which can be used to store the same real-world-phenomena at different levels of precision, accuracy and resolution [1, 2]. It can be understood both as a *multiple representation database* and as a *multiple resolution database*. In an MRDB, different views on the same physical objects or phenomena can be stored and linked. This variety can stem from different views of the world, different applications, as well as different resolutions. These lead to differences in the objects as such, i.e. in the semantics and in the geometry. Also the graphic representation can be taken into account, leading to geometric, semantic and graphic multiplicities [3].

There are two main features that characterise an MRDB (see fig. 1):

a) Different levels of detail (LoD's) are stored in one database and
b) The objects in the different levels are linked.

The first feature can be compared to the analogue map series of the NMA's: these maps of different scales exist separately, only implicitly linked by the common geometry. In the second case, however, individual objects are explicitly linked with each other and thus each object "knows" its corresponding objects in the other representations.

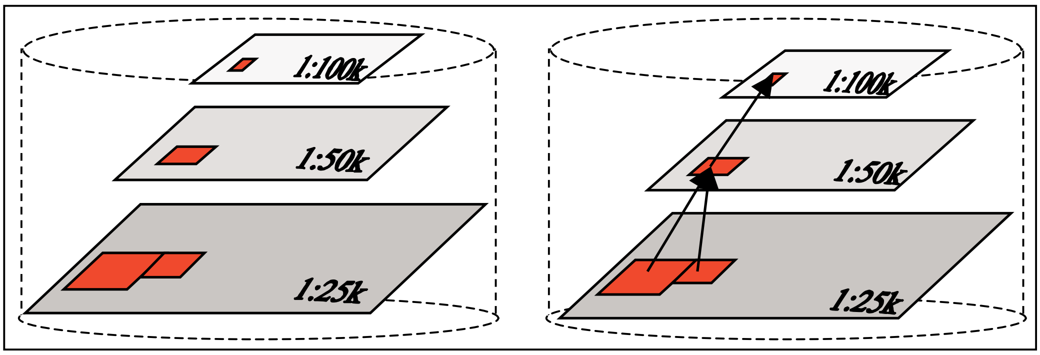

Fig. 1. a) MRDB with separate scales b) objects linked between the LoD's

There are several concepts of MRDB's which depend on the specific needs and requirements. Research projects aiming at combining real-time generalisation with multiple representation databases can be found in [4, 5, 6]. Vangenot et al. [7] describe modelling concepts which support not only the multi resolution view but also the different views on the object features like object types, attributes and their values. Kreiter [8, 9] describes the concept of an MRDB from the NMA's point of view. Cecconi [10] investigates the use of MRDB for the web mapping.

Important questions concerning the implementation of an MRDB are:

(a) The number of layers to be provided:
It depends on the application area, typically at factor two to four the appearance of spatial phenomena changes so dramatically, that intermediate layers have to be introduced.

(b) The necessity of links between layers:
The question whether explicit links should be established depends on the application; also the way these links are designed is of importance (uni-directional or bi-directional). If there is only the requirement to visualise the spatial data, separate scales may be sufficient (see Fig. 1a). In order to offer more functionality (i.e. GIS-analysis via different scales) it *is* necessary to provide the link (Fig. 1b).

There are several reasons for introducing an MRDB: On the one hand it allows for a multi-scale analysis of the data: Information in one resolution can be analysed with respect to information given in another resolution. Gabay and Sester [11] present an example where topographic data is linked with cadastral data. A topographic data set of lower resolution containing only settlement areas is queried concerning the buildings in that area – an information that can be derived from a more detailed cadastral data set, whose objects are directly linked. A major reason for National Mapping Agencies to investigate and implement MRDB is the possibility of propagating updates between the scales: the appealing idea is that the actual information only has to be updated in the most detailed data set, this new information can then be propagated through the links in MRDB to all the other scales [12, 13].

When designing an MRDB two important cases can be discerned:

a) Linking existing data sets of different scale or thematic contents by matching procedures.
b) Creating new data sets from existing ones, which then form new layers in the MRDB.

Concerning the first option, [1] classify three different stages for the design of the MRDB:

(a) *Correspondence between abstractions*: Database schemata translate phenomena of the real world into abstracted instances of databases by focusing only on relevant parts of these phenomena; integration of abstractions thus requires methods for schema integration on the semantic level.
(b) *Correspondence between individual objects* of different representations: Data models are required to describe the links between corresponding individual objects of the different representations.
(c) *Defining the matching process* between objects: in order to identify corresponding (homologous) objects and instantiate the corresponding links, two sets of geographical data must be searched for objects that represent the same real-world objects; methods for this purpose are subsumed under the term 'data matching' [5, 14].

Concerning the second option new data sets have to be created from existing ones. In order to do so, a function has to be known that allows this creation. In the case of deriving a lower scale data set from a higher resolved one, generalisation operations can be applied. The function immediately establishes also the links between corresponding objects. Consider for example the aggregation of two adjacent parcels of land to a new combined parcel in the lower resolution data set: links will be established between the high resolution parcels to the newly created one.

3 TWO PROJECTS DEALING WITH MRDB

3.1 The Project WIPKA

The project WIPKA is funded by the German Federal Agency of Cartography and Geodesy (BKG) in Frankfurt. It aims at establishing a consistent representation of all topographic data sets in Germany (ATKIS), as well as the option of also integrating additional data sets, e.g. from soil science or geology. An MRDB is used as a federated data model. The projects includes partners from Cartography and Geoinformatics (Munich, Hannover), Photogrammetry, Computer Science and Databases (Hannover). The tasks include the design of an MRDB [15], the homogenisation and consistent modelling of all topographic data sets, the population of the MRDB by real objects and their links, and finally a concept for the updating process. Concerning the creating of the links, both generalisation processes are needed (for the derivation of a new scale 1:50.000) and matching for seeking the correspondence between existing scales.

3.2 The EU-Project Gimodig

The EU-project GiMoDig, an acronym for „Geospatial Info-mobility Service by Real-time Data-Integration and Generalisation", aims at developing the spatial data delivery from national primary geo-databases for mobile use [4]. The project started in November 2001. The following partners are involved in this project:

- National Survey and Cadastre Denmark (KMS),
- National Land Survey of Sweden (LMV) and
- National Land Survey of Finland (NLS).
- Finnish Geodetic Institute (FGI) as coordinator,
- University of Hannover, institute for cartography and geoinformatics (ikg),
- Federal Agency for Cartography and Geodesy (BKG),

The main vision of GiMoDig is a mobile user, travelling within an European country and receiving on-line information of his or her environment on the mobile device. Even when crossing a border, the type of information presented does not change, thus having the situation, that seamless topographic information is available.

This vision requires on the one hand data from national databases, and the possibility of flexibly inspecting the data on the small displays, including zooming from overview to detail in real-time. Furthermore, the project deals with issues of user requirement analysis [16] and small-display cartography constraints. Also, the user must be able to state his requirements. Because of the limited display size and resolution it is important to transmit only the required information to the user. This will include the selection of the desired objects and features as well as the desired resolution of the presentation. Due to these requirements, one sub-objective of the project is the development of methods for generalising the graphic representation of geospatial data in real-time, to be suited for display of the data at varying scales on small, mobile devices. The presentation on the mobile display will be dependent on - and adaptive to - the special user requirements i.e. data resolution and content.

In automatic generalisation a considerable progress can be observed in recent years (see e.g. the results of Workshops of the ICA Commission on Generalisation (http://www.geo.unizh.ch/ICA/docs/mainlevel/home.html), resulting in efficient generalisation methods and algorithms that are applicable to perform scale transitions in given scale ranges. However, the processes involved going from a large scale to a small scale (say 1:10k to 1:Mio) are very complex. Thus, it is obvious, that (at least today) the generation and visualisation of ad hoc personalised products of spatial data in arbitrary scales on a mobile platform cannot be solved without pre-generalised datasets. Real-time generalisation can only be efficiently performed in small scale ranges and are restricted to operations of minor complexity, that can be solved completely automatically.

The second major issue in the project is the harmonisation of data sets of the partner National Mapping Agencies (NMA's) in order to allow for a seamless data provision. This presumes that (topographic) data sets of the different countries are analysed with respect to their contents and that a semantic harmonisation is performed in order to guarantee the same "world view" when traversing the borders. Also here, the necessity of providing different data sets with different views of the same physical entities, as well as means to link these data sets is given.

Thus, in order to solve the problems of data generalisation and harmonisation, the concept of an MRDB is used. The MRDB serves as a pre-generalised and pre-harmonised data structure with spatial data in given scales. To minimise the effort of computation work during the real-time generalisation process, the GiMoDig service selects a scale close to the desired scale requested by the mobile user. Based on this neighbouring scale, only small scale transitions are necessary, that can be handled in real-time. In this way the need for complex algorithms, for example displacement, can be minimised or even excluded.

4 THE DESIGN OF THE MRDB

In this section possible designs of a multiscale-database will be described; furthermore, the way it is implemented in the projects will be presented. In both projects the MRDB's are based on a *Federated Database System (FDBS)* [17, 18]. Starting point for a FDBS are several existing databases which should work together to provide a global service. A FDBS can be classified by three main characteristics: *Distribution*, *Heterogeneity,* and *Autonomy*. Distribution means that the data is stored on different database systems (DBS), which can run on a single computer system or in a distributed computing environment. Heterogeneity, because the DBS can use different database management systems (DBMS) and finally Autonomy, because the different DBS can be designed independently from each other and it is possible that each DBS runs independently from the *Federation Layer* (no changes of local application programs are necessary). Figure 2 shows the general architecture of a federated database system. The FDBS of both projects differ in the task which has to be provided by the Federation Layer. In the GiMoDig project the Federation Layer provides one global interface to different spatial distributed MRDBS's to unite them to one big MRDB. All these MRDB's have the same number of representation levels with the same scales. In contrast to GiMoDig the Federation Layer of the WIPKA FDBS is the MRDB. Every component DBS in WIPKA stores one specific representation layer (BaseDLM, DLM50, DLM250, or DLM1000). The Federation Layer integrates the component DBS's to an MRDB and provides the user interface to the MRDB. Therefore the Federation Layer contains a working database which stores all link information between the component DBS's. A database schema for this Federation Layer is described in [15].

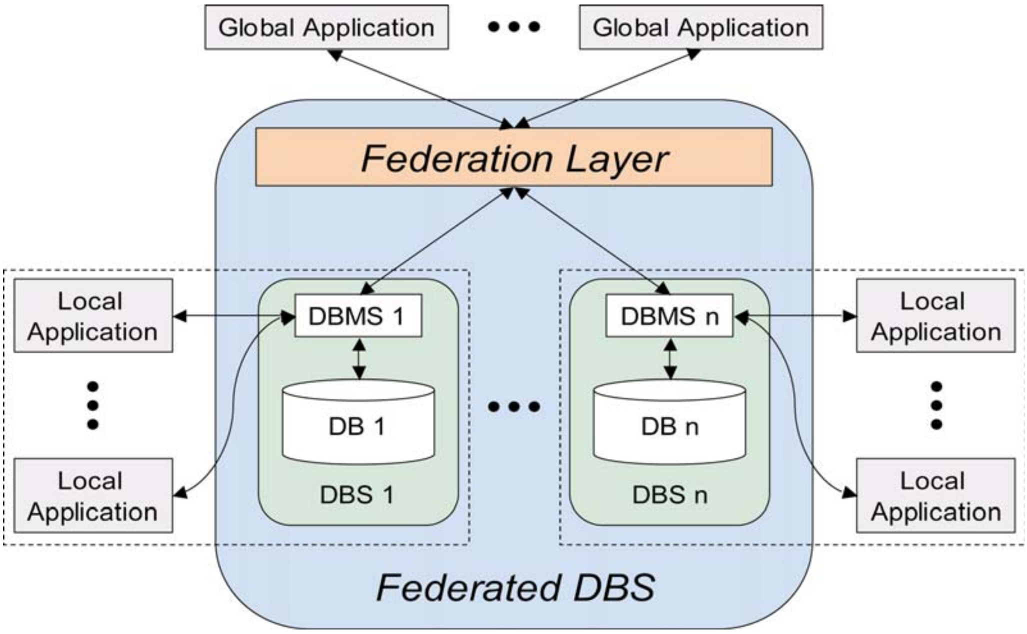

Fig.2: General architecture for federated database systems

Concerning the question of how to store the links between the corresponding objects, there are different approaches. Kreiter [8] analysed and identified different possibilities for linking the different levels of an MRDB. First he denotes the "attribute-variant": In this variant the whole MRDB will be stored in one dataset. This variant uses additional attributes to describe the different forms of appearance. The attributes denote the scale the object will appear and the geometric changes of the object while changing the LoD (see fig.3)

ID	dy25	dx25	dy50	dx50	dy100	dx100	scale25k	scale50k	scale100k	the_geom	attributes
10	0	0	-5	-3	-1	-10	1	1	1	LINESTRING(...)	...
20	0	0	10	5	3	5	1	1	1	LINESTRING(...)	...
30	0	0	-5	-6	-	-	1	1	0	LINESTRING(...)	...

Fig.3: "attribute variant" presenting an edge

The second variant is called "variant bottom-up". This possibility assumes two or more datasets of the same spatial phenomena. These datasets will be linked by using an additional attribute which refers to the corresponding objects in the following LoD (fig 4). This variant needs two additional columns, an ID for every object or part of an object and the ID of the linked object. A disadvantage of this model is the limitation to only one link per object. This link could also be stored in an extra table maintaining the ID's of the corresponding objects (see fig. 4b).

Another possibility to design an MRDB is to use the "variant top-down". This variant links the objects in the opposite direction. Using this alternative it would be best to use an extra table for storing the links. Because the top-down variant mainly has one-to-many-relations (like a built-up area contains many buildings) and the number of objects linked to the object in the lower level of detail highly differs, this would mean a high number of empty cells in the table.

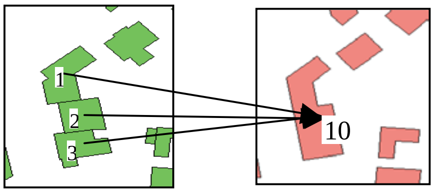

ID25k	link2level2	the_geom	attributes
1	10	MULTIPOLYGON(((3547844.797	...
2	10	MULTIPOLYGON(((3547944.667	...
3	10	MULTIPOLYGON(((3547444.785	...

ID25k	ID50k
1	10
2	10
3	10

Fig.4: "bootom-up variant", stored: a) as an attribute in the object table b) in an extra table

4.1 The MRDB-design in GiMoDig

The database-software used in GiMoDig is an open-source, unix-based and object-relational DBMS called *PostgreSQL* [19]. It supports almost all SQL constructs, including sub selects, transactions and user-defined types and functions. It will be used together with an extension called *PostGIS* [20], which adds support for geographic objects to the PostgreSQL object-relational database. *PostGIS* follows the OpenGIS "Simple Features Specification for SQL" [21]. The data are stored in a relational table. Every feature class will be stored in its own table. In our approach, the bottom-up variant is used. In fig. 5 we see an extract of a table, where three levels of detail are linked with each other.

ID	the_geom	ID50k	ID100k
0	SRID=-1;MULTIPOLYGON(((3547844.767 5805853.33	0	0
1	SRID=-1;MULTIPOLYGON(((3548593.272 5806036.09	0	0
2	SRID=-1;MULTIPOLYGON(((3548583.08 5806018.256	20	100
3	SRID=-1;MULTIPOLYGON(((3548582.631 5806069.61	30	100
4	SRID=-1;MULTIPOLYGON(((3548533.836 5806084.34	40	100
5	SRID=-1;MULTIPOLYGON(((3548553.841 5806053.16	50	100
6	SRID=-1;MULTIPOLYGON(((3548607.355 5806088.60	60	100
7	SRID=-1;MULTIPOLYGON(((3549097.043 5805889.74	70	220
8	SRID=-1;MULTIPOLYGON(((3549167.123 5805863.11	80	220
9	SRID=-1;MULTIPOLYGON(((3549174.006 5805846.30	90	220

Fig.5: Database schema storing spatial data and the link to the higher LoD's

In the case of the GiMoDig project the data will be accessed by WebFeatureService (WFS)-queries. Because of this the database queries should be limited, the less the better, the best solution would be to manage with one request. Therefore the best way will be to store the ID of the corresponding object as an additional attribute in an extra column. This solution has one disadvantage: if there is an n:1 relation we need more than one column for the link which means a lot of empty cells in the table. Although we are using the bottom-up-model, starting at the highest LoD, there appear some cases of an n:1 relation. For example when we are linking two datasets, stemming from different data sources, a street element can be composed of two or more elements in the lower LoD (see fig. 6). An alternative would be to store this

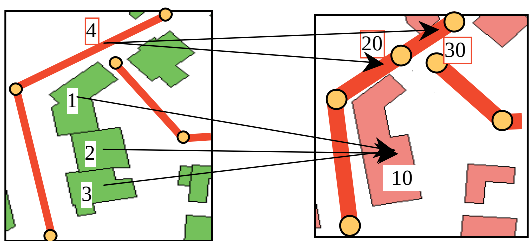

Fig. 6: possible relation between two LoD's in an MRDB

linking structure in an extra table. But in the case of GiMoDig this would mean to access two tables leading to two WFS-requests.

After clarifying the problem of how to store the data and the linking structure the second questions is how to acquire the corresponding objects. Two scenarios are possible:

a) There is one existing dataset (LoD), the following LoD has to be derived out of this dataset by generalisation.
b) There are two or more existing datasets which have to be combined to an MRDB by data matching.

In the next section the establishment of the links using these two principles will be described.

4.2 Establishing links between different levels in the MRDB

We will demonstrate the establishment of the links between different layers in the MRDB using two examples. In the first example, we show data sets of high to medium resolution, namely buildings, simplified buildings and built-up areas. In the second example we focus on linking medium to low resolution data sets, using built-up areas and higher aggregations thereof. The first example involves first a simplification of building outlines, followed by a matching of the generalised buildings with existing built-up areas, leading to three levels of detail. The second example mainly involves aggregation of area features.

4.2.1 High to medium resolution

We are concentrating on three LoD's including buildings and built-up-areas. The highest LoD consists of buildings, derived from an ALK-dataset of Hanover (fig. 7a). ALK (Automated cadastre) is a dataset created for the scales 1:500 to 1:5.000. The second level consists of generalised buildings (fig. 7b), created by simplifying the shape of the buildings. The third level contains built-up-areas (see fig. 7c). The transition from level 1 to level 2 is performed using a building simplification algorithms similar to [22, 23]. The main input parameter is the length of the smallest edge of the building to be retained. The links between the original and the generalised objects can be derived automatically from this generalisation process. In medium scales (1:25.000) the individual buildings are replaced by built-up areas. In this case, we use an existing layer of built-up areas given in the ATKIS-data set. The correspondence between individual buildings on level 2 and the build-up areas in level 3 are simply derived using a containment operation: a building is linked to the built-up area it is contained in.

Fig.7: Three stages of the MRDB: a) original buildings, b) generalised (simplified) buildings, c) built-up areas
Data from ALK® (courtesy of Mapping agency of city of Hannover) and ATKIS® Germany (courtesy of LGN, Hannover)

By generating these three levels of an MRDB we have pointed out two ways for receiving the links between the objects automatically:

a) The linking structure was derived from the generalisation process. The generalised objects (buildings) can be linked with the original objects buildings. The result is an 1:1 or 1:0 relation.
b) The linking information has to be derived by an data-matching process. To match the buildings with the built-up areas a point-in-polygon test has been implemented.

The result of these processes is an MRDB including three levels (fig. 7). All three levels have been linked with each other. Another possibility would have been to link one LoD only with the adjacent LoD. The decision if *all* layers should be linked or *only the adjacent ones* depends on the application. In case there are frequent requests for links between objects that are not adjacent, this information would have to be derived by several requests concerning the links in between.

4.2.2 Medium to low resolution

As an example for the medium to low resolution model generalisation we studied the automatic derivation of an ATKIS DLM50 (1:50.000) database from an ATKIS BaseDLM (1:10.000-1:25.000) database. The ATKIS DLM50 object catalogue describes the object types which should be captured in the DLM50 and which geometries types should be used to represent the real world objects. Therefore in general we have to deal with area, line, and point generalisation, but for our first studies we concentrated on the subject of area generalisation. These operations are implemented for the derivation of the DLM50, they are however generic enough to be transformed to representations of arbitrary scales.

The area generalisation process consist of three steps:

1. A reclassification of the object types.
2. An aggregation of adjacent areas with equal object type.
3. A shape generalisation.

The first step is needed because some object types of the ATKIS BaseDLM do not exist anymore in the ATKIS DLM50 landscape model. The second step is needed because in the ATKIS DLM50 the minimum size criteria for capturing of certain object types as area objects has increased. The third step handles all cases of area objects which are still not big enough after the second step to be captured as areas in the ATKIS DLM50 model. In such cases according to the ATKIS DLM50 object catalogue these areas has to be represented by a point or have to be eliminated completely. In both cases one has to establish a reclassification and an additional aggregation step to these areas. This additional reclassification and aggregation step can be done in different ways. Four possible solutions are shown in figure 8.

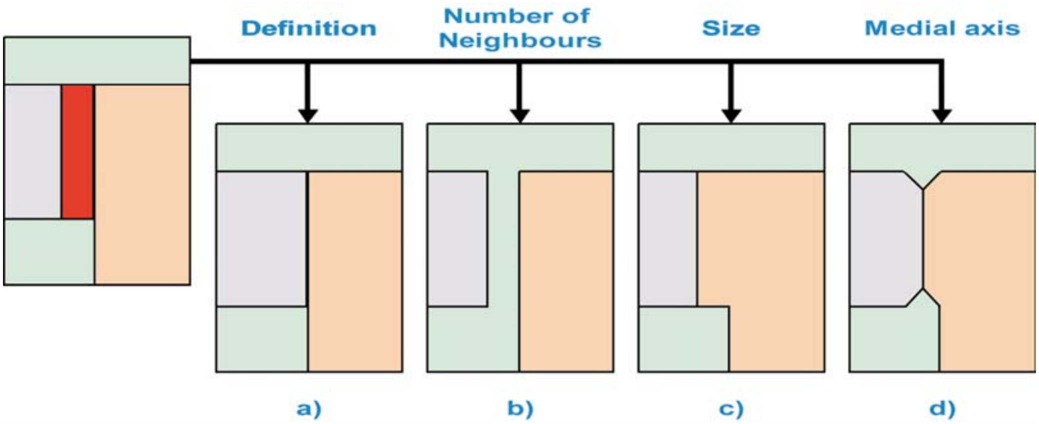

Fig.8: Four different ways to replace an area (red) by adjacent neighbour areas.

The replacement of an area can be done “by definition” which means that a priority list of new object types is given that describes which new object type has to be used to replace the old object type relative to the adjacent areas (fig.8a). E.g., if an area with object type *farmland* has to be replaced then a possible priority list can be: 1. grassland, 2. garden area, 3. area without vegetation, and so on. That means if an adjacent area has the type grassland then the area will be reclassified as grassland. If no adjacent grassland can be found then one has to look for an area without vegetation. If there is no such area then may be there is an area of type indefinable area, etc. The priority list has to make sure that always a new object type can be found. Another way is to choose the most frequent object type of the adjacent neighbour areas (fig.8b) or to choose the object type of the largest adjacent neighbour area (fig.8c). A more sophisticated approach is to compute the medial axis of the area which has to be replaced and to increase all adjacent areas according to the computed medial axis (fig.8d) [24]. The maximum number of equal neighbours and the maximum size approach have the same drawback, that this maximum number must not be unique. E.g., an area can have as many neighbours of object type X as of object type Y. It is also possible that an area has more than one adjacent neighbour areas with the same size but different object types. These ambiguities can be solved in the most cases by combining the criteria's (number and size of), but in general there can be still ambiguous cases. The approach with the medial axis takes all neighbour areas in account and increases all neighbour areas relative to the shape of the area which has to be replaced. The drawback of this approach is that it is more complicated to be implemented and it can produce artificial shapes, like the two small triangles in figure 8d). We decided to use the first approach because the result is clearly defined without ambiguity and can be implemented easily. The drawback is that one has to define for each object type a priority list. Figure 9 shows an example for this type of area aggregation. During this aggregation process the links between the original areas of medium resolution and the new areas of low resolution are stored in a link table.

Fig. 9: Left before and right after the area aggregation and reclassification process.

5 CONCEPT FOR UPDATING DATA IN AN MRDB

The goal of the WIPKA project is to implement an MRDB to enable an automatic incremental update process of all representation levels in the MRDB. The idea is to update only the lowest level with the highest geometric and semantic resolution (in our case is that the ATKIS BaseDLM). All other levels should then be updated automatically at most as possible to make the data revision process faster and more efficient and consistent. We are now developing a message passing system which will allow us to control the update process by propagating the changes through all levels according to the stored MRDB links and to trigger appropriated generalisation operators. We are in the early stage of these project, we can only describe the rough concept which we want to implement. An MRDB can be seen as an noncyclic directed Graph (Figure 10 shows this situation schematically) which provides information about the relationship between objects. These relationships can be used to propagate updates bottom up through the representation levels. In our concept we define three main types of update events which are propagated through the network:

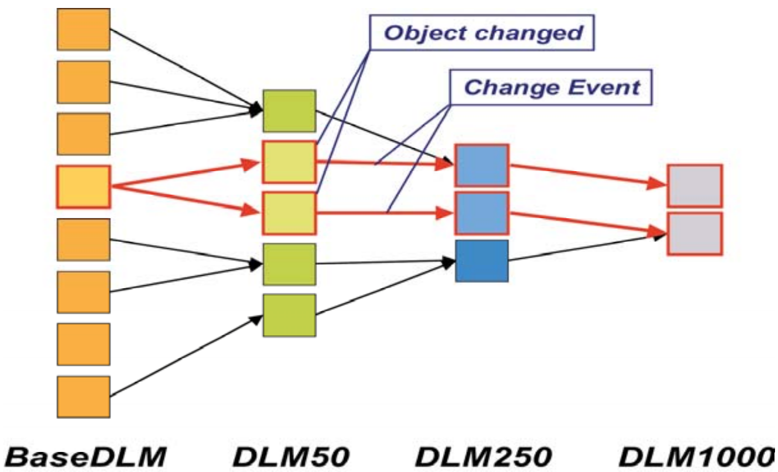

Fig.10: Update propagation in the MRDB.

1. Insert: A new object was created.
2. Remove: An object was deleted.
3. Change: An object has changed. These changes can be divided up in:
    - Change Attribute: Only attribute values are changed.
    - Change Geometry: Only the shape of the object has changed.
    - Change Attribute and Geometry: The attribute values and the shape have changed.

After the update process of a lower level is completely finished all changes are propagated to the next level according to the stored links. The update propagation has to be understood as calling an appropriate generalisation method for the linked object at the next level. For the update process the system will provide a set of generalisation operators which are selected by a rule based system. That means the user has to maintain a set of rules which describes what operator should be used under certain conditions. These conditions are the type of the update event, the type of the updated object and the object which has to be changed, and thresholds for the attribute and geometry changes. Conflict situations which can not be handled automatically by the system are stored and after the processing graphically presented to an user in order that he can decide to solve a conflict situation interactively or by changing the update rule set. The update propagation will be repeated until we have reached the final representation level (in our case the DLM1000).

6 CONCLUSION AND OUTLOOK

In the presentation, two projects dealing with the establishment of an MRDB have been presented. The projects have different needs concerning the use of such a linked data structure: whereas in the GiMoDig project, a quick visualisation is of major concern, where the links are needed for "information drilling", as well as interpolation objects between different scales, the major issue in the WIPKA-project is the establishment of a consistent data structure for topographic data sets in Germany, and the possibility of propagating updates and thus greatly simplifying the update process. Concerning GiMoDig, open questions to be addressed in the near future are the development of additional generalisation algorithms for the transition of building/settlement objects to smaller scales, as well as the access to the MRDB using an extension of the Web Feature Server. In the WIPKA project, the next steps concentrate on the implementation of the update concept for selected objects.

7 REFERENCES

1 T. Devogele, J. Trevisan, and L. Raynal, Building a Multiscale Database with Scale-Transition Relationships: In Proc. Of the 7th Int. Symposium on Spatial Data Handling, Advances in GIS Research II, pp. 6.19-6.33, Delft, (1996)

2 R. Weibel and G. Dutton, Generalising spatial data and dealing with multiple representations. In: P.A. Longley, M.F. Goodchild, D.J., Maguire & D.W. Rhind, editors, Geographic Information Systems – Principles and Technical Issues, volume 1,. John Wiley & Sons, 2 edition, 1999, pp 125–155, (1999)

3 Y. Bédard and E. Bernier, Supporting multiple representation with spatial databases views management and the concept of "VUEL", ISPRS / ICA Joint Workshop on Multi-Scale Representations of Spatial Data, Ottawa, Canada, (2002)

4 T. Sarjakoski, L.T Sarjakoski, L. Lehto, M. Sester, A. Illert, F. Nissen, R. Rystedt, and R. Ruotsalainen, Geospatial Info-mobility Services - A Challenge for National Mapping Agencies. Proceedings of the Joint International Symposium on "GeoSpatial Theory, Processing and Applications" (ISPRS/Commission IV, SDH2002), Ottawa, Canada, July 8-12, 2002, 5 p, CD-ROM, (2002)

5 T. Badard, On the automatic retrieval of updates in geographic databases based on geographic data matching tools, in: `Proceedings of the 19th International Cartographic Conference of the ICA', Ottawa, Canada, (1999)

6 A. Cecconi, R. Weibel, and M. Barrault. Improving automated generalization for on-demand web mapping by multiscale databases, 10th International Symposium on Spatial Data Handling, Ottawa, Canada, (2002)

7 C. Vangenot, C. Parent, and S. Spaccapietra, (2002). Modelling and Manipulating Multiple Representations of spatial data, 10th International Symposium on Spatial Data Handling, Ottawa, Canada, (2002)

8 N. Kreiter, Multirepräsentationsdatenbank für die Kartographie. Vertiefungsblockarbeit. Eidgenössische Technische Hochschule, Zürich. www.karto.ethz.ch/teaching/VTB-Kreiter1.pdf, (2002)

9 N. Kreiter, Multirepräsentationsdatenbank als Basis von topografischen Landeskarten. Diploma thesis, unpublished, Institute of Cartography, Swiss Federal Institute of Technology Zurich, (2002)

10 A. Cecconi, Integration of Cartographic Generalization and Multi-Scale Databases for Enhanced Web Mapping, PhD Thesis, University Zurich, (2003)

11 Y. Gabay and M. Sester, Forming and utilizing communication between two spatial representations at different scales - a demonstration. Submitted to Geoinformatica, (2002)

12 T. Kilpeläinen, Multiple Representation and Generalization of Geo-Databases for Topographic Maps. Kirkkonummi, PhD Thesis, Helsinki University of Technology, (1997)

13 L. Harrie, L. and A.-K. Hellström, A Case Study of Propagating Updates between Cartographic Data Sets, in: `Proceedings of the 19th International Cartographic Conference of the ICA', Ottawa, Canada, (1999)

14 M. Sester, K.-H. Anders, and V. Walter, Linking Objects of Different Spatial Data Sets by Integration and Aggregation, GeoInformatica 2(4), 335-358, (1998)

15 D. Mantel, Konzeption eines Föderierungsdienstes für geographische Datenbanken, Diploma Thesis, University Hannover, (2002)

16 A.-M. Nivala and L. T. Sarjakoski, Need for Context-Aware Topographic Maps in Mobile Devices. To be published in the Proceedings of the 9th Scandinavian Research Conference on Geographic Information Science ScanGIS'2003, June 4-6, Espoo, Finland, (2003)
17 A. Sheth and J. Larson, Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases. ACM Computing Surveys 22:3, 183-236, (1990)
18 S. Conrad, Föderierte Datenbanksysteme : Konzepte der Datenintegration. Springer-Verlag, Berlin, (1997)
19 PostgreSQL, http://www.us.postgresql.org, (2002)
20 PostGIS, http://postgis.refractions.net, (2002)
21 OpenGIS® Simple Features Specification for SQL at http://www.opengis.org/techno/implementation.htm, (2002)
22 W. Staufenbiel, Zur Automation der Generalisierung topographischer Karten mit besonderer Berücksichtigung großmaßstäbiger Gebäudedarstellungen, PhD Thesis, University Hannover, (1973)
23 M. Sester, Generalization based on Least Squares Adjustment, in: 'International Archives of Photogrammetry and Remote Sensing', Vol. 33, ISPRS, Amsterdam, (2000)
24 M. Bader, Methoden zur Erkennung und Lösung von metrischen Konflikten in der Generalisierung von Polygonmosaiken, MSc Thesis, Department of Geography, University of Zurich, (1997)

ACKNOWLEDGEMENT

The funding of the European IST FP5 (Information Society Technologies Fifth Framework Programme) and the Federal Agency of Cartography and Geodesy, Germany, is gratefully acknowledged.