
HYBRID MODELLING AND ANALYSIS OF UNCERTAIN DATA

Michael GLEMSER, Ulrike KLEIN

University of Stuttgart, Germany

Institute for Photogrammetry

Michael.Glemsers@ifp.uni-stuttgart.de

Ulrike.Klein@ifp.uni-stuttgart.de

Working Group IV/III.1

KEY WORDS: Data Uncertainty, Object-Oriented Modelling, Hybrid Analysis Methods, Uncertainty Propagation.

ABSTRACT

An essential requirement for a comprehensive use of hybrid data is the consideration and processing of its uncertainty. Erroneous interpretations of analyses can be avoided if uncertainty is integrated as a mandatory component, stored and considered in all operations. In this contribution, a probabilistic approach is presented for modelling geometric and thematic uncertainty. By means of a flooding forecast as an example application, the enhancements which are necessary in data modelling and analysis functions are explained.

1 INTRODUCTION

Over the years, GIS have evolved to powerful tools which solve various tasks in a variety of different applications. The ongoing integration of new and improved analysis methods enables users to process even more complex application tasks with combinations of geometrical, topological and thematic aspects using hybrid data, i. e. raster data as well as vector data. A further step in development will be achieved by the opening of the architecture to interoperable systems (McKee and Kuhn, 1997, OGC, 1999, IOGIS, 1999). In such an open environment, the problems are distributed among different systems which act like single components. The final result is obtained in this teamwork-like approach by a combination where each component adds its part.

Data uncertainty plays a special role in such an environment. Although all data is uncertain to a particular degree, the necessary integration in GIS has not taken place so far. The consequence is that uncertainty is actually not taken into account during data processing. This kind of disregard can partly be justified only in the case of a closed system, where the user has full control over all steps from input to presentation. As the user is aware of the data quality, he is able to verify the result using his knowledge and experience. This might be seen as a reason why uncertainty has not been a relevant topic in GIS yet. But the control gets completely lost in an interoperable system. This starts already within the data input step. Today, a great number of existing databases offer a variety of data sets like topographic information, cadastral data, statistical data or digital orthophotos and satellite images. Data collection is changing widely from digitising own data to retrieving and transferring it from existing data bases. For such data sets, the user requires an uncertainty description which has to be added by the producer of the data set as a kind of meta description. This information allows the data to be checked in order to avoid the risk of combinations of data without any practical use. Furthermore, the uncertainty description serves as additional input parameter for the analysis process and thus enables the propagation of the uncertainty onto the result. For this purpose an appropriate uncertainty model has to be developed and integrated in GIS. In particular, the analysis methods require to be extended in order to propagate the uncertainty automatically as a kind of background process.

The emphasis of the research is on the description and propagation of the uncertainty through analysis in a GIS. In this article, the uncertainty integration is demonstrated by an example application in order to allow the reader an easy access to the developed approach. As an example application, the forecast of the spatial extent of a possible flooding is selected.

2 DATA UNCERTAINTY

Spatial objects have a geometric and a thematic component. The geometric component defines the position and the extension of the spatial phenomenon, and the thematic component includes all descriptive information (Bill and Fritsch,

1991). All data contain numerous types of uncertainty due to limitations of measurement techniques and the degree to which the data model approximates, or fails to approximate, the real world. In the following, the description of the uncertainty is given separately for the two components. The developed concepts are merged in section 3 by defining an uncertainty model for hybrid data.

2.1 Geometric Uncertainty

In GIS two basic approaches are used for describing the geometry of spatial phenomena: the vector model and the raster model. In the vector model basic geometrical elements are point, line and polygon. A point is unambiguously defined by its coordinates in the chosen spatial reference system. A line is formed by connecting two points. For objects, any desired complex structure can be built using these two structural elements. In GIS, implementations are often based on boundary representation. In this case geometry is determined by the boundary of the object. In the raster model the basic elements of the geometry description are regular cells. In practical applications, square raster cells are mostly used. One or several neighbouring raster elements form the geometry of an object. For the definition of the geometry of the raster, reference point, orientation and cell size must be given.

The handling of the geometric uncertainty differs for the two geometric approaches according to their characteristics. In case of the vector model, uncertainty in geometry refers to the variation in position of geometric primitives which represent a spatial object. The amount of various influences is responsible that the data cannot be collected completely error-free. Limited accuracy of measurement techniques as well as the fuzziness of the phenomena play an important role. The fuzziness becomes obvious if we compare the representation of an object with its origin in nature. The model requires a description in discrete form that separates an object unambiguously from its neighbouring objects. A sharp line observed in the field is generally of human origin, because nature's boundaries are always fuzzy. They correspond to extended transition zones that exist to varying degrees.

The integration of geometric uncertainty into data processing requires its description by the use of an appropriate model. Since the coordinates of points result from measurements, a stochastic approach is suitable to model the variation in the primitives. It treats each point as continuous random variable that varies according to a distribution function (e. g. Bill and Korduan, 1998). In general, the standard Gaussian distribution is applied for this purpose. A point is sufficiently characterised by the mean values μ_x, μ_y and the variances σ_x^2, σ_y^2 of its coordinates x, y . These parameters have to be determined for each point. The coordinates derived from measurements can be seen as appropriate estimates of the mean values. The variances can be determined by prior knowledge or estimated individually through comparison with reference data. This approach has to be extended since all points belonging to a boundary represent stochastic variables. In this sense measured points do not have a special status compared to other line points. All points are treated in the same way. For each point of the boundary the mean values of the coordinates and their variances should be observed to obtain a complete modelling of geometric uncertainty. Since we have to deal with an infinite number of points, an individual specification for each point is impossible. Therefore one individual value for each line or one value for the whole object should be specified.




Figure 1. Examples for probability matrices

In the stochastic approach, the modelling of the geometric uncertainty with mean value and variance can alternatively be defined by probabilities (Kraus and Haussteiner, 1993). They can be calculated for every position in space indicating the belonging to the object. The values depend on the local distribution function of the geometric primitive and the distance between the certain position and the object. The formulas for the calculation of the probabilities vary for the different types of primitives (Glemser and

Klein, 1999). For area objects e. g. the probability is calculated as

$$p(x, y) = p(d) = \int_{-\infty}^d f(t) dt = F(d)$$

with d as the distance of the position to the object, $f(t)$ as the density function and $F(d)$ as the distribution function. The formula defines a spatially continuous probability function which has to be approximated by a discrete raster. Then the probabilities are calculated for each cell. The raster forms a probability matrix that can also be used to produce a

graphical visualisation of the geometric uncertainty. For this purpose the probabilities are transformed into grey values. Figure 1 shows examples of probability matrices for different types of objects.

In case of the raster model, geometric uncertainty can be excluded. Spatial data describing the earth's surface or other parameters (e. g. heights or temperatures) are normally measured at irregular distributed points. To obtain a continuous representation of these data, values are calculated on a regular grid by interpolation algorithms. Form, size and position of the raster cells are fixed per definition in dependence on the application. Consequently, the geometry can be assumed as certain and therefore uncertainty of the data refers only to the assigned thematic attribute (e. g. in the heights or temperatures). Using remote sensing data a special case occurs, if the geometry of the raster results from rectification and geocoding of the data in a preprocessing step (Richards, 1993). Each raster cell shows differences in position compared to the real world. These differences have to be interpreted and modelled as geometric uncertainty of the raster cell. The thematic values primarily assigned to the raster cells consequently do not refer to the actual position of the cells, but to a random one. Thus the consequences of rectification and geocoding can alternatively be interpreted as thematic uncertainty. In that case the geometry of the raster is again considered as fixed while the thematic value in each raster cell is uncertain to a certain degree.

2.2 Thematic Uncertainty

The non-spatial thematic component of an object is described by attributes. The attribute values can be classified in discrete and continuous values depending on the scale they belong to. The modelling of the uncertainty of the attributes differs accordingly.

In case of discrete attribute values the degree of truth that the specific attribute possesses the assigned value has to be evaluated. Since the assignment of a value is not always unambiguously possible, specifications about alternatives must additionally be given. The assessment of the thematic uncertainty of an object comprises the estimation of the occurrence of all alternative attribute values. Probabilities are suitable as parameters. They can be managed in a vector for each object. The determination of these probabilities is shown by means of an application with raster and vector data.

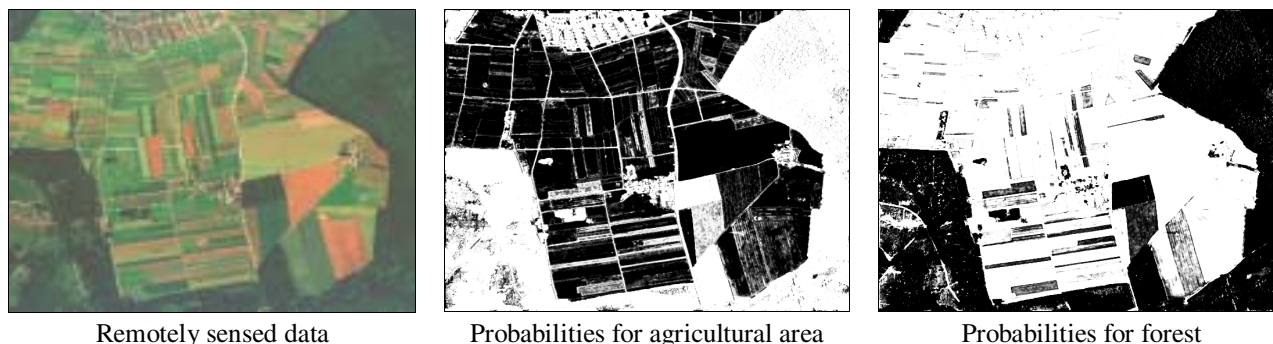


Figure 2. Test area and representation of the probabilities for different land cover types

Sources of raster data can be found in the field of remote sensing for example where classification algorithms produce land use values from satellite images. First, a set of land use types has to be defined. One problem is that a specified number of classes has to be chosen from a multitude of possible land use types. Subsequently, each pixel of the data set is classified using statistical methods, for example by a maximum likelihood classification (Richards, 1993). This results in the assignment of each pixel to one of the defined land use types. One particular type of classification problem is that of a so-called mixed pixel: a pixel which has several land uses within its boundaries causes an averaged radiation intensity to be received by the sensor. This pixel can then no longer be unambiguously assigned to a particular class. Errors in both steps lead to the uncertainty in the assigned thematic that can be evaluated by probabilities. Two different strategies are possible for the calculation of the probabilities. Comparing the classification result with reference data leads to the confusion matrix (Stehman, 1997). For each defined land use type a probability can be derived, but the most fundamental disadvantage of these global measures is that no information is available about the spatial distribution of the uncertainty within the classes. Alternatively, one might consider the class membership probabilities that are produced by statistical classification as indicators of the reliability of land use assignment in each pixel. The storage of the probabilities is managed by a raster matrix for each land use type. Figure 2 shows the probability matrices of agricultural area and forest for a test area (Klein et al., 1998).

In the case of vector data, data acquisition is performed by measuring the extension of the objects belonging to the predefined discrete attribute values. The attribute values that are assigned to the measured geometry are simply assumed

to contain no uncertainties. In reality this restriction is not always warrantable since the acquired objects can be inhomogeneous with regard to the thematic aspects or the assignment of the attributes can be uncertain. For a complete modelling of thematic uncertainty the percentage of the respective attributes, that can be interpreted as probabilities, should be specified.

In connection with thematic uncertainty, it is furthermore important to extend the uncertainty description to continuous attribute values so that e. g. heights or temperatures can be qualified. Also these data contain a certain degree of uncertainty since they are derived from measurements. The variation of the attribute value characterises the thematic uncertainty. Analogous to the geometric uncertainty a stochastic approach can be chosen to model the continuous attribute values as random variables. Then the variation of the attribute value is represented by mean value and variance. Thematic data existing in irregular distributed points are normally converted to a regular grid. For this purpose, it is important to determine the variances of the interpolated values by statistical methods, too. The storage of the interpolated values and their variances can again be managed by matrices i. e. as raster data (section 4.1).

3 HYBRID DATA MODEL INTEGRATING DATA UNCERTAINTY

A hybrid data model (Glemser and Fritsch, 1998) builds the basis for the integration of the uncertainty. The approach rests upon objects as exclusive representations of spatial phenomena. As a consequence, an object building procedure

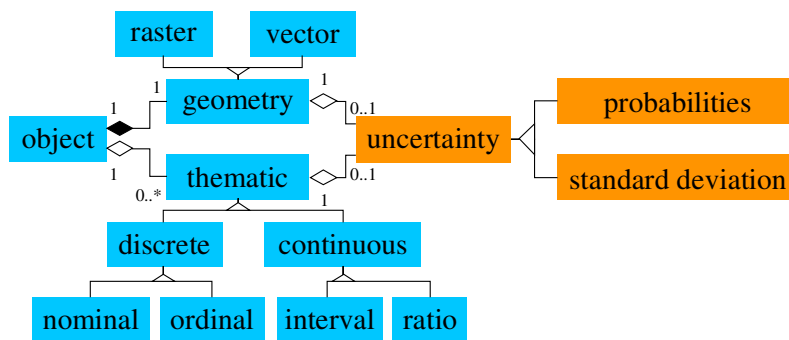


Figure 3. Hybrid data model integrating the uncertainty model

schematic graphical representation of the hybrid data model in UML-syntax (Booch et al., 1999). Since both the geometry and the thematic of an object can be uncertain, the hybrid data model has to be extended by a description of the uncertainty. As already explained, probabilities and standard deviations can be utilised. Since they are used similarly for the description of uncertainty in the geometric component, they can be seen as alternative possibilities. But in the thematic component they are used independent from each other. The uncertainty of discrete attribute values is described by probabilities, the uncertainty of continuous attribute values by standard deviations. If it is possible to assign alternative attributes to an object, each attribute requires its own uncertainty description. Thus the description of an object could comprise e. g. one probability raster for the geometric uncertainty and several measures for the thematic uncertainty of its attributes. Figure 3 shows the defined uncertainty model on the right side as an extension of the initial hybrid model.

has to be applied in some cases before the data can be integrated (Klein et al., 1998). The geometric component allows either a raster or a vector representation of the object. Such a hybrid model enables the user to process data without ever noticing whether the actual type is raster or vector. In the thematic component multiple attributes of any kind are attached to the object. Attribute values can be distinguished according to the scale they belong to. Figure 3 shows on the left side a

4 COMPLEX ANALYSIS

Since all data only have a limited accuracy it is important for the user to know how the result of the analysis can be assessed with respect to its accuracy. This problem can only be adequately solved if the uncertainty of the input data sets is inserted in the analysis process and propagated onto the result. As an example application, the forecast of the spatial extent of a possible flooding of a river is selected. The special task to be solved is to detect those parts of settlement areas which are flooded with respect to an assumed high water level. It defines a complex problem in the way that the whole process consists of two successive steps. In the first step a flooding model has to be developed. In the second step, the derived outcome is geometrically overlaid with topographic information, namely the settlement areas. The major goal is to determine the uncertainty of the result in addition to the traditional output without regard of data uncertainty.

4.1 Data Sets

Basically, the analysis requires different input data sets. The topographic surface strongly influences which areas are flooded. It can be described by a digital terrain model (DTM). As source of the flood, the location of the river is needed as well as the location of the settlement areas which might be affected by the flood. The necessary land use information can be taken from a topographic data set.

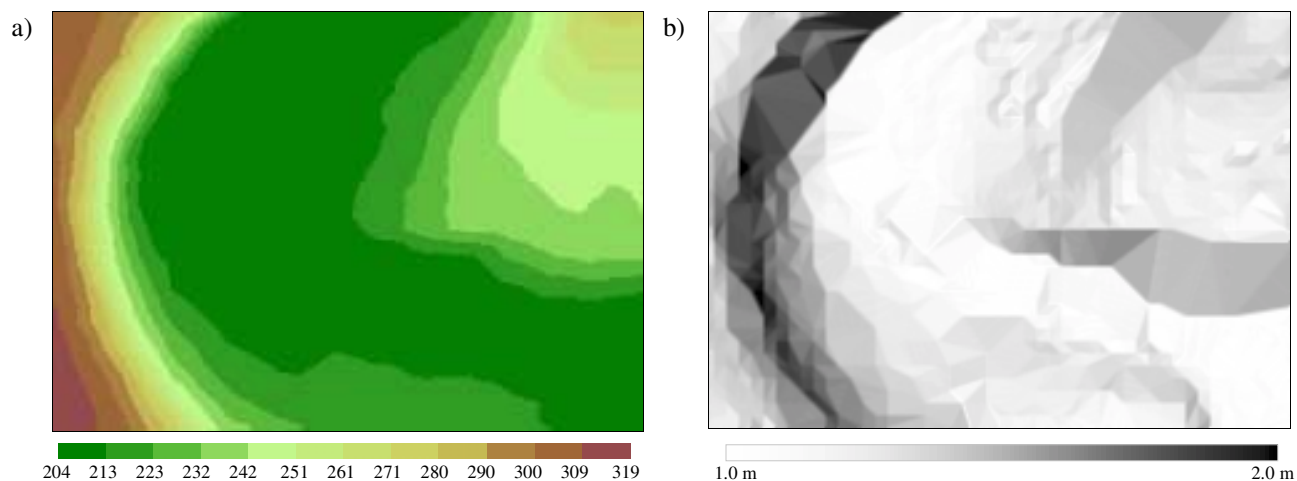


Figure 4. Digital terrain model (a) and its uncertainty given by variances of the heights (b)

The DTM for the study area is shown in figure 4a. The heights are structured in a regular grid and therefore they form a raster data set with a ground pixel size of 5 m. Each raster cell carries a height value as an attribute belonging to the continuous scale. For the representation of the uncertainty a variance for each raster cell is required. In principle the accuracy of the originally measured values has to be transformed to raster data. Since this information is only available sparsely, often qualified models are formulated. A simple approach would be to assume the same height variance for each cell. A more complex determination is based on the dependency of the variance on the slope. In this application a slope dependent approach is chosen (figure 4 b). The accuracy of the heights varies between 1 m and 2 m.

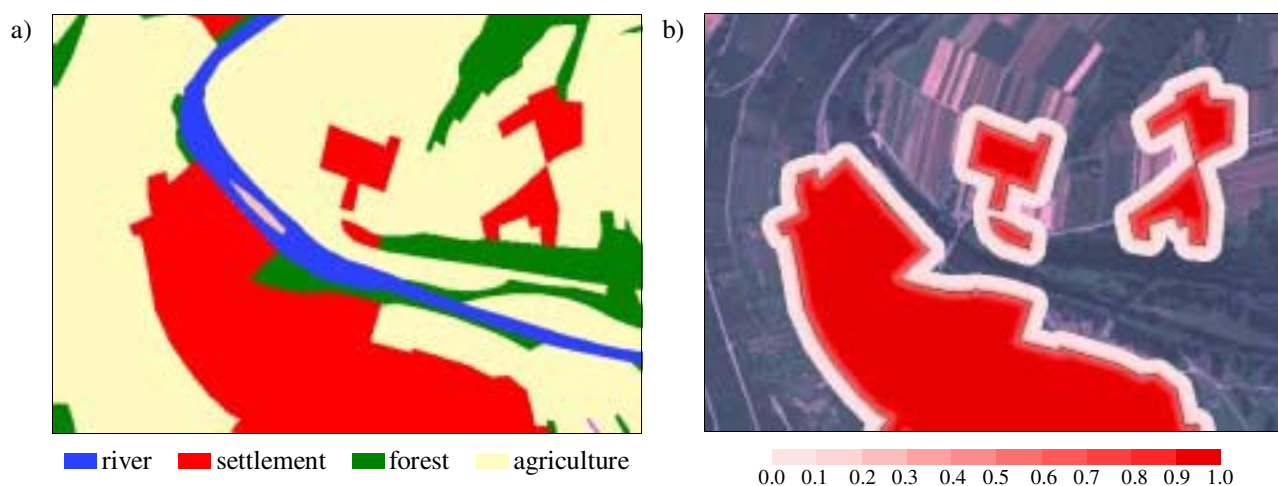


Figure 5. Land use data (a) and the graphical representation of the probabilities of the settlement data set (b)

The land use data set (figure 5a) is extracted from an existing digital topographic database, called ATKIS, which is the German topographic cartographic spatial database (ATKIS, 1997). The database defines a digital model of the physical landscape. The relevant contents of the landscape are modelled as objects which are successively aggregated to object classes (e. g. roads), object groups (e. g. road transportation) and finally to object domains (e. g. transportation). The scale is 1:25.000. All data is collected in vector format with a positional accuracy of 3 m for important objects (e. g. roads, railways, rivers) and a reduced accuracy of 10 m for all others. For the flooding analysis two object classes are relevant: stream and residential area. The study area contains one stream object (the river Enz) and five residential areas that build together the settlement area of interest. For visualisation of the uncertainty and implementation of the analysis, the description in terms of variances of the settlement areas is transformed to a representation in terms of probabilities (figure 5b). The probabilities indicate the degree of membership to the object.

4.2 Thematic Analysis

The determination of the flooded area requires the formulation of a flood model. For each raster cell of the river it is examined how the increasing water-level affects the neighbouring raster cells. Thus, an individual flooded area is determined for each river cell and the overlay of these partial areas results in the complete area. The result of the forecast of the flooding is presented in figure 6a. Since the DTM only has a limited accuracy it is important to know how the result of the flooding can be assessed with respect to accuracy. A major difficulty is that the functional context between the digital elevation model and the flooded region can only be modelled in a complex way. Thus, the variance propagation based on a functional context is not suitable to determine the uncertainty description of the flooded area. Simulation techniques are used alternatively to solve the problem.

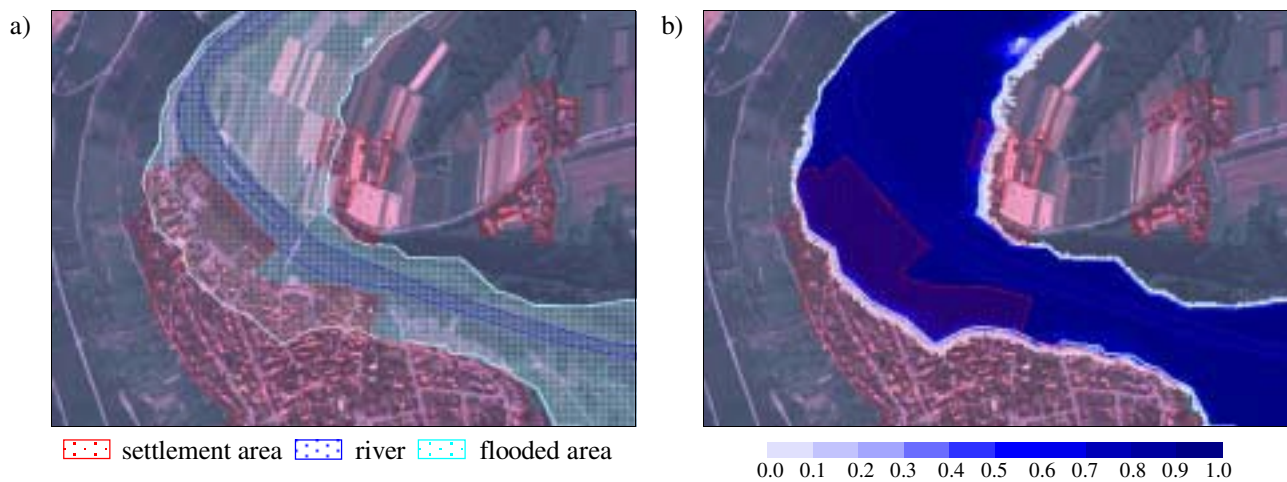


Figure 6. Outcome of the simulated flooding (a) and the representation of the probabilities of the flood region (b)

Simulation can be described as generation of realisations of a mathematical system by variation of the input variables (Piehler and Zschiesche, 1976). If stochastic data are used, simulation is called stochastic as well. In the present application this concerns the generation of the flooded region due to the stochastic behaviour of the elevation model. The problem can be solved by using Monte Carlo methods (Sobol, 1985). Monte Carlo methods are special procedures for the simulation of stochastic problems generating random samples by random numbers. The stochastic behaviour of the samples results from a predefined probability density function. For a given stochastic input variable, the outcomes of the simulation are the realisations of a random variable. The digital elevation model builds the basis of the simulation considering the heights as random variables. Assuming a Gaussian or Normal distribution, the heights are described sufficiently by height value and variance.

The determination of the uncertainty of the flooded region is achieved as follows: in a first step n realisations of the elevation model are generated using the probability density function for each height. Then, according to the formulated flood model, the spatial extent of a possible flooding is determined for each of the n realisations. The outcome of the simulation are n different and independent realisations of the flooded region. For a representative conclusion from the sample to the population the volume n of the sample has to be large enough (e. g. $n = 1000$). The quality of the derived statement about the uncertainty of the flooded area depends primarily on the volume of the sample. The more realisations are produced, the more reliable is the result - but computing effort increases at the same time.

For the analysis of the simulation, the n flooded regions are overlaid and for every raster cell the frequency of flooding is determined. Frequency can be interpreted as probability which indicates the membership of the cell to the object 'flooded region'. The complete probability matrix can be considered as outcome of the simulation. Grouping all cells with probabilities larger than 0.0 results in the maximum extent of the flooded area (figure 6b).

Concerning the integration of uncertainty, the simulation result covers a larger area in comparison with the traditional solution. Not only a modified extension is derived, but also the probability of a flooding is spatially quantified. The graphical representation of the probabilities shows clearly that the uncertainty of the heights causes the uncertainty of the boundary, namely the spatial extent of the flooded area.

4.3 Geometric Analysis

In the next step the flooded area is geometrically overlaid with the settlement areas. The result is a set of new objects consisting only of the intersection parts but possessing attributes of all the input objects. It is easier to implement the overlay operation with raster data than with vector data. In the raster domain it is a simple boolean And-operator on the

raster cells, while in the vector domain the computation of points of intersection is very time-consuming. To take advantage of this fact in hybrid systems, polygon overlay is realised using raster overlay techniques. Therefore, the settlement areas given in vector format have to be transformed to objects in raster format by a vector-to-raster conversion. The realisation of a hybrid GIS requires that the conversions run automatically as a kind of background operation without any user interaction, even without being noticed by the user (Fritsch et al., 1998).

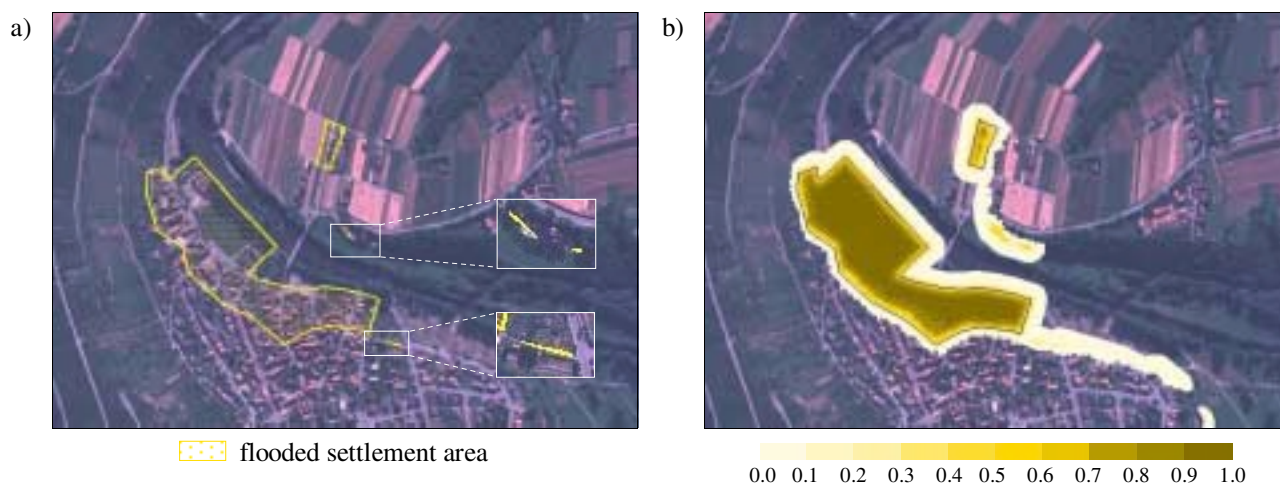


Figure 7. Flooded settlement area (a) and the representation of the probabilities of the flooded settlement area (b)

The outcome of the overlay is shown in figure 7a. The shapes of the five resulting objects can be divided into areas and lines. For a better identification, the regions with the three line objects are shown in an enlarged inset of the orthophoto. The line objects appear when settlement areas and flooded region touch or have only a very small overlay. With respect to the result of the conventional overlay analysis such sliver polygons are considered as irrelevant and removed from the database. However, it is doubtful if doing so is warrantable. Since all data is uncertain to a certain degree, the spatial extension of the objects respectively their boundary is to be considered to be uncertain, too. To obtain reliable results for the analysis the uncertainty of the objects needs to be modelled and integrated in overlay. Since the overlay operation can be interpreted as a simple boolean And-operator, the propagation of the probabilities is defined by

$$p(A_1 \cap A_2) = p(A_1) \cdot p(A_2) \quad .$$

$p(A_1)$ and $p(A_2)$ denote the probabilities of the two raster cells possessing the assigned attribute values A_1 and A_2 , in this case settlement area and flooded area.

The comparison of figure 7b with figure 7a shows that the integration of the uncertainty results in a larger endangered area. Therefore the sliver polygons can no longer be considered as irrelevant. They should be removed from the database only after analysing their probabilities. Disregarding the uncertainty leads to a misinterpretation of the results. This example application shows clearly that it is important to integrate the uncertainty in data processing as well as to include the uncertainty into the visualisation of the results.

5 CONCLUSIONS

The paper focuses on the extension of a hybrid data model by an uncertainty model. Since both the geometry and the thematic of an object can be uncertain, description of uncertainty has to be taken into account in both components. Probabilities and standard deviations are utilised as parameters for the uncertainty. The developed model establishes the basis for the integration of uncertainty in interoperable applications. An example shows in a demonstrative way which modifications become necessary and how results of analysis are to be interpreted with respect to their uncertainty. During analysis processes, either typical mathematical error propagation methods perform the estimation of the requested uncertainty or simulation techniques are utilised for complex applications. But it is not sufficient to consider data uncertainty only in individual analysis methods. In fact all available functions must provide appropriate extensions. A second source of uncertainty can be detected in the analysis model applied. A decisive characteristic of a model is that it defines only a kind of approximation of the reality. For example, the used flooding model does not correspond exactly to the real situation of a flooding because of the complexity of the process itself. This aspect has not been covered so far.

ACKNOWLEDGMENTS

This research work is funded by the German Research Foundation (DFG) within the joint project IOGIS (Interoperable Offene Geowissenschaftliche Informationssysteme).

REFERENCES

- ATKIS, 1997. Amtlich Topographisches-Kartographisches Informationssystem (ATKIS). Arbeitsgemeinschaft der Vermessungsverwaltungen der Länder der Bundesrepublik Deutschland (AdV). <http://www.atkis.de>.
- Bill, R., Fritsch, D., 1991. Grundlagen der Geo-Informationssysteme. Band 1: Hardware, Software und Daten. Wichmann, Karlsruhe.
- Bill, R., Korduan, P., 1998. Flächenverschneidung in GIS – Stochastische Modellierung und Effizienzbetrachtung. Zeitschrift für Vermessungswesen, Vol. 123, Part 1 in No. 8, pp.247-253, Part 2 in No. 10, pp.333-338.
- Booch, G., Rumbaugh, J., Jacobson, I., 1999. The Unified Modelling Language User Guide. Addison-Wesley.
- Fritsch, D., Glemser, M., Klein, U., Sester, M., Strunz, G. (1998): Zur Integration von Unsicherheit bei Vektor- und Rasterdaten. GIS, Vol. 11, No. 4, pp.26-35
- Glemser, M., Fritsch, D., 1998. Data Uncertainty in a Hybrid GIS. In: Fritsch, D., Englich, M., Sester, M.: GIS – Between Visions and Applications, IAPRS, Vol. 32, Part 4, Stuttgart, Germany, pp.180-187.
- Glemser, M., Klein, U., 1999. Hybride Modellierung und Analyse von unsicheren Daten. Schriftenreihe der Institute des Fachbereichs Vermessungswesen, Universität Stuttgart, Report No. 1999.1, pp.27-43.
- IOGIS, 1999. Interoperable Offene Geowissenschaftliche Informationssysteme (IOGIS). http://ifgi.uni-muenster.de/3_projekte/4dgis/texte/iogis/IOGIS.html
- Klein, U., Sester, M., Strunz, G., 1998. Segmentation of Remotely Sensed Images Based on the Uncertainty of Multi-spectral Classification. In: IAPRS, Vol. 32, Part 4, Stuttgart, Germany, pp.299-305.
- Kraus, K., Haussteiner, K., 1993. Visualisierung der Genauigkeit geometrischer Daten. GIS, Vol. 6, No. 3, pp.7-12.
- McKee, L., Kuhn, W., 1997. The OpenGIS Consortium's Purpose and Technical Approach. In: Fritsch D., Hobbie, D. (Eds.): Photogrammetric Week '97, Wichmann, Heidelberg, pp.237-242.
- OGC, 1999. Open GIS Consortium (OGC). <http://www.opengis.org>.
- Piehler, J., Zschiesche, H.-U., 1976. Simulationsmethoden. Mathematik für Ingenieure, Band 20. Teubner Verlag, Leipzig.
- Richards, J. A., 1993. Remote Sensing Digital Image Analysis. Springer-Verlag.
- Sobol, I. M., 1985. Die Monte-Carlo-Methode. Deutsche Taschenbücher, Band 41, Verlag Harri Deutsch, Thun und Frankfurt am Main.
- Stehman, S. V., 1997. Selecting and Interpreting Measures of Thematic Classification Accuracy. Remote Sensing of the Environment, Vol. 62, pp.77-89.

PRACTICAL EXPERIENCE WITH SPATIO-TEMPORAL GIS IN GEOPHYSICAL RESEARCH

Thomas Knudsen

thk@kms.dk

National Survey and Cadastre—Denmark, Geodetic Dept., 8 Rentemestervej, DK-2400 Copenhagen NV

KEY WORDS: Geophysics, GIS, Metadata, Software, Spatio-temporal.

ABSTRACT

Busstop, a spatio-temporal GIS prototype for geophysical studies, is presented. It includes a metadata representation system, a visualisation subsystem, and a spatio-temporal extension of a classical map algebra. It also takes an active approach to integration with existing non-GIS data handling and analysis programs. *Busstop* has been used for minor geophysical case studies with some success, but needs further refinement if it should move towards the role of a generic work horse for geophysical research. Potential ways to go include the transition to XML based metadata representation and integration with a full scale scripting language, such as Python.

1 INTRODUCTION

In current geophysical and geodetic research, data from Earth observation satellites play an important role: only by utilizing such data it is possible to get global observational coverage with a repeat rate enabling the study of ocean dynamics on local and regional scales. A detailed understanding of these phenomena is interesting in its own right for oceanographic purposes, but it is also interesting as ancillary information for geodesists trying to establish a reliable estimate of the long term mean sea level, i.e. the geoid.

Satellite based radar altimetry is a major data source for studies of ocean dynamics. Current altimeters are able to observe the sea level with a precision of 3-20 cm, but as they are nadir looking instruments they only provide information about the part of the ocean right below the satellite track. This results in a relatively coarse resolution, spatially and temporally: typical values (from the TOPEX/POSEIDON satellite) is 7 km along track and up to 150 km across track distance between observations, and approximately 10 days between repeated observations of the same point.

Hence, to get a more complete view of the sea surface, it is desirable to combine the altimetry with ancillary observations from satellite based scanning radiometers, which potentially give a much higher resolution, spatially and temporally (figure 1).

Unfortunately, the wildly differing observational geometries between the two types of sensors makes the interoperability a bit involved. Evidently some sort of GIS based handling system would be useful for relieving this by providing a streamlined framework for georeferencing. However, since we are dealing with dynamic phenomena it is not enough to focus on the spatial geometry: to be of real use, we must also take the differing temporal reference systems into account. This implies the need for a spatio-temporal GIS.

Except for a few notable studies (Beller et al., 1991, Mason et al., 1994, Wesseling et al., 1996, e.g.) the majority of work done in the field of spatio-temporal GIS has been focused on methods for updating and tracking man made objects changing their positions and/or properties as time goes (Langran, 1989, Langran, 1992, Peuquet and Duan, 1995, etc.). Conceptually, such “moving and transforming objects” are quite different from temporally evolving physical fields, which are the key “objects” of interest in geophysics.

Hence, to study and define the needs for spatio-temporal data handling functionality in the geophysical sciences, it makes sense to build a focused spatio-temporal GIS prototype from scratch. This was the rationale behind the development of the *Busstop* package (Knudsen, 1997, Knudsen, 1998).

2 DESIGN

Busstop (an acronym meaning Basic Utilities for Spatial and Spatio-Temporal OPERations) was developed using an iterative design approach, allowing reshaping and restructuring as experience was gained from actual use and implementation of the package. The main requirements of the final design can be summarized in 5 points:




Figure 1: 10 days worth of observations from the TOPEX/POSEIDON satellite altimetry mission (white cross-hatch pattern) overlaid on data from a single pass of the NOAA/AVHRR radiometer instrument. It is seen that AVHRR (which observes this area several times per day) has a much more complete areal coverage. On the other hand, it is limited to observe the ocean in areas where clouds are absent. The altimeter, which is radar based, looks right through the cloud cover and provides observations under all weather conditions. The two data types complement each other: each has its weak and its strong sides, and both are needed to get a reasonable idea of the state of the ocean.

1. Support a wide range of existing file formats.
2. Ensure a high degree of interoperability with existing data processing systems
3. Make the system highly platform independent
4. Avoid dependence on proprietary tools and libraries
5. Avoid software megalomania: keep the system modular

These requirements have resulted in a system based on 5 main ideas:

1. A metadata handling system
2. A parameterized data reader
3. A visualization system
4. A spatio-temporal map algebra and scripting language
5. A blueprint for modular extensions

Busstop was designed to integrate seamlessly with a large range of existing programs used in the geophysics and geodesy community. This ensures the availability of a sufficient amount of analytical functionality as Busstop works by providing *added value* to existing packages, rather than by attempting to provide a *complete set* of tools for every task anyone would want to carry out.

2.1 Integration

Busstop's integration with existing programs is based on its capability of reading a wide range of existing file formats. This is accomplished through the parameterized data reader, which is initialized with a description of the data file to be read, i.e. the metadata for the datafile.

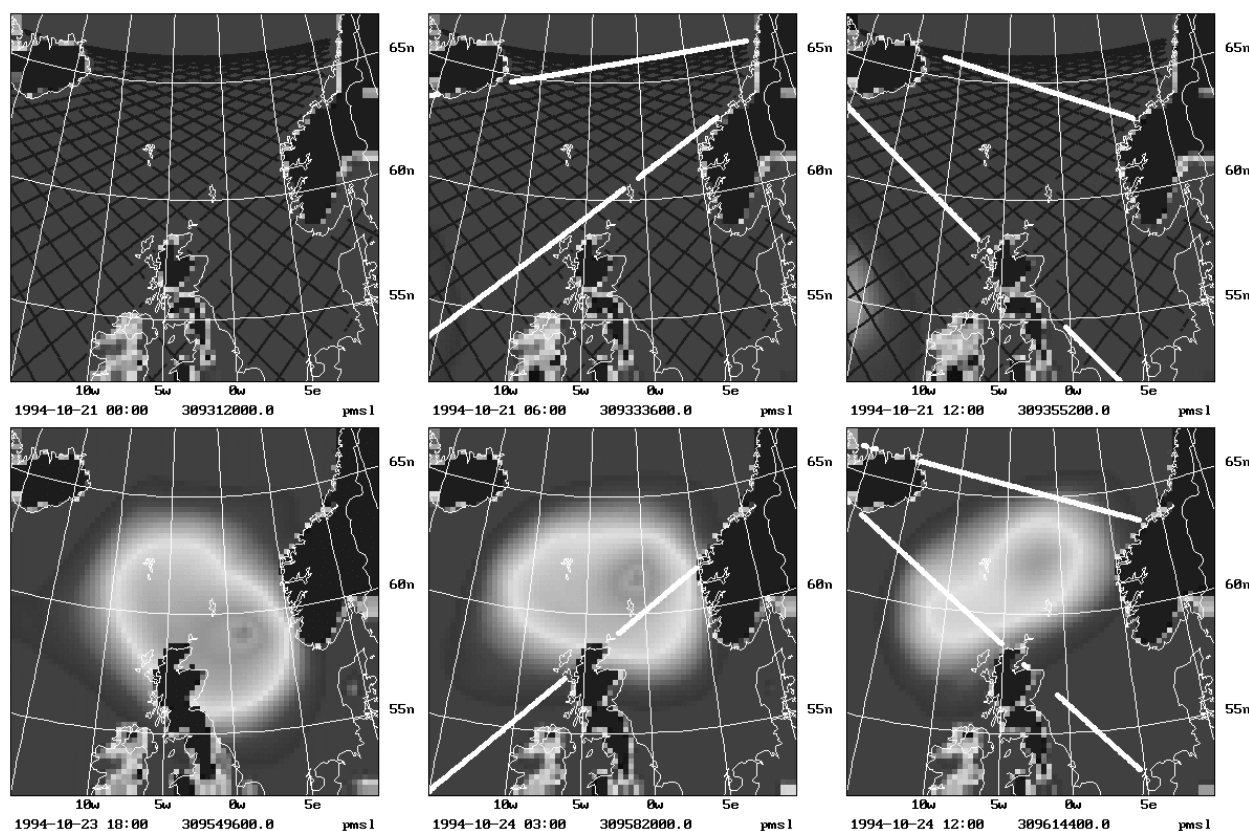


Figure 2: Selected time steps from an animation used to detect whether the centre of an atmospheric low pressure system was observed by a satellite altimeter. **Top left:** Start of the animation—the system has not entered the area yet. The dark hatch pattern shows the positions of the altimeter observations available from the entire study period. **Top centre:** Reveal automatically highlights the altimeter data observed within ± 1 time step (here: ± 3 hours) from the current time. **Top right:** The low pressure system is just about to enter the area at its western border. **Bottom left:** The system has now entered the area and matured fully. The dark altimetry overlay has been toggled off in order to get a better view of the meteorological data. **Bottom centre:** A satellite pass hit the system right in the eye. **Bottom right:** 9 hours later: the depression has weakened somewhat, but we get two more satellite passes close to the centre.

Evidently, the data reader is not able to read each and every file format imaginable. Rather, it has been designed in order to read a large number of file formats *likely* to be written by programs to be integrated with Busstop. In practice, this means that Busstop should be able to read any file format which is reasonably easy to write from a Fortran 77 program, as Fortran 77 is the computational workhorse *par excellence* in Busstop's intended user community. Here we take "easy to write from a Fortran 77 program" to mean *a file consisting of a variable sized header followed by any number of unspecified (but fixed) format records*. This definition is used for text and binary format files alike.

2.2 Data model and metadata

Busstop's data access is based on a synthesis of the traditional point and grid models into a more general data model. This model is based on the *generalized field concept* (Knudsen, 1997). Generalized fields can represent either coordinates, physical observables or constants. In cases where the actual geometry of the dataset is unimportant all types of generalized fields are treated transparently and identically.

The connection between the actual data file and its corresponding generalized view is formed by a *descriptor file*, containing a collection of metadata elements: For the *physical characteristics* of the data file, we need metadata elements describing the size of the file header and the number of generalized fields. Additionally, for each generalized field we need to know its kind, and its data type (e.g. byte, integer, real ...). For the *logical characteristics*, we need metadata elements describing the reference systems associated with the data set, e.g. its map projection, datum and temporal reference system (epoch). Busstop supports 119 map projections and 41 reference ellipsoids, through a transparent link to the USGS *proj* system (Evenden, 1990, Evenden, 1993).

The representation of metadata in Busstop descriptor files is based on a hierarchical system of plain text blocks, intended to be easily accessible for a human user. Inside the Busstop programs, a tiny embedded compiler compiles the text format representation into an internal binary representation, which is easily accessible for a computer program.




Figure 3: Eliminating error prone data from a gravimetry database: A case study based on free air anomaly data from the Skagerrak strait between Denmark and Norway. **Top left:** A suspiciously looking jump of +15 mgal (highlighted by the white box) was detected in this gridded dataset. This is an indication that a set of point gravity observations has slipped through the quality control and into the gravimetry database. **Top right:** The effect is easier to track down visually when looking at the gradient of the original grid. Now another parallel structure becomes clear to the west of the original one. **Bottom right:** Overlaying the positions of the original point data observations, and using a sufficient zoom level, it is possible to digitize two narrow polygons surrounding the error prone data. Subsequently the polygons are used to split the point data. **Bottom left:** Regriding after removal of the error prone observations shows that the +15 mgal jump has been eliminated.

2.3 Visualization

Reveal, the Busstop visualization sub-system can be used to create static visualizations of ordinary spatial datasets and to create animations of spatio-temporal datasets. Its list of features is relatively modest—the most important is perhaps the direct link between the visualization and the underlying data. This enables a continuous readout of the data value below the cursor, which is extremely useful when investigating small scale anomalies. *Reveal*'s functionality also includes:

- interactive zooming and panning
- interactive contrast enhancement

- on-screen digitizing of points, polylines and polygons
- on-the-fly resampling to any resolution
- automatically generated graticule overlay, for enhanced spatial overview
- overlay of point, vector and polygon data sets (e.g. shorelines)
- conversion of the visualizations to graphics file formats (GIF, Encapsulated PostScript)

Figure 2 illustrates a range of these features, using a dataset consisting of output from a meteorological model in combination with observations from a satellite radar altimeter.

Reveal, in its current version, runs on PC platforms using any comparatively recent graphics adapter. Preliminary studies (Knudsen, 1997, Knudsen and Brandt, 1998, Knudsen, 1998), have indicated that Reveal is rather convenient and sufficiently capable for a range of geophysical visualization tasks.

2.4 Bush

Bush—the BUSstop SHEll, is an environment for scripting and for algebraic operations applied to geodata. In traditional raster-GIS terms, *Bush* is a *map algebra*: a programming language optimized for geodata operations at the dataset level, rather than at the record or field level.

The concept of map algebrae was introduced by C. D. Tomlin during his work with the *Map Analysis Package*, MAP (Tomlin, 1983, Tomlin, 1990). Unfortunately the Tomlin map algebra is really a way of writing algebraic expressions in a non-algebraic language. To avoid duplicating this somewhat inconsistent approach, *Bush* has *not* been based on Tomlin's pioneering work, but rather on later work by Michael Shapiro: the *r.mapcalc* module for the *GRASS GIS*. This module implements a map algebra expressed in algebraic terms (Shapiro and Westervelt, 1992, Larson et al., 1991, USACERL, 1993). Hence, *Bush* (like *r.mapcalc*) supports the ordinary arithmetical, relational and logical operators, as well as a range of mathematical and trigonometrical functions. But contrary to the earlier map algebra implementations, which have been inherently 2 dimensional, *Bush* includes support for 0 dimensional point data and 3 dimensional spatio-temporal grids as well. This has made it much more useful for rapid implementation of solutions to problems involving dynamic phenomena (Knudsen and Brandt, 1998, Knudsen and Olesen, 1998).

The main point for *Bush* is its "understanding" of the data models it operates on. *Bush* knows how to pick ancillary information from the metadata. This also enables *Bush* to document the output data file properly, i.e. to generate a new set of metadata, including information about grid size, geographic and temporal coverage, map projection, and a description of how the file was produced.

3 EXPERIENCE

Busstop has been used for a handful of medium sized geophysical case studies. Below, two of those are reviewed. The first is based on purely spatial data, and simply demonstrates the seamless integration of Busstop and external programs. The second, however, is fully spatio-temporal and involves the combination of two datasets with entirely different georeferencing, chronoreferencing, and geometry.

3.1 Quality Control of Gravimetry

A paper by Knudsen and Olesen (1998) describes a Busstop based, semi-automated method for post-processing of gravimetric observations. The method (summarized in figure 3) is based on visual detection of suspicious areas in a grid, based on the original gravimetric point data set.

For the specific case in question, a set of shipborne gravimetry tracks showed up to be in error, and gave rise to a sudden jump in the data values in an otherwise smooth area. The flawed data were first hinted at when visualizing the gridded data, using Busstop Reveal. The extent and exact position of the area under influence was pinpointed by generating a new grid, showing the local gradient of the original one: during the visualization of this grid we used Reveal's overlay functionality to add a set of points indicating the positions of the original point data observations. By toggling this overlay on and off in rapid succession, the position of the error prone data became extremely clear. Under a sufficient zoom level, the Reveal on-screen digitizing feature made it a simple thing to digitize a slim envelope for these data. They were subsequently removed from the data base using the point-inclusion-in-polygon algorithm by Franklin (1972). Finally, the data set was re-gridded to confirm that the remedy had worked. As this was actually the case, the entire process was integrated as a *Bush* script for further automation.

The gravimetry used in the study was stored in the format used by the *Gravsoft* package for gravimetry and altimetry analysis (Tscherning et al., 1992), and the Gravsoft component *geogrid* was used in the gridding process. The integration of Gravsoft and Busstop was simply a matter of providing two metadata files: one describing the raw gravimetry, and one describing the grid format used by geogrid. Hence, this study is a neat example of the ease of integration supported by Busstop. Additionally, the study illustrates that Reveal's relatively limited visualization functionality is actually enough to ensure its practical value.

3.2 Atmospheric Corrections for Satellite Altimetry

The original inspiration for the study summarized below came when using Reveal to check the availability of satellite altimeter observations of the center of a severe atmospheric low pressure system. (the task used for illustrating Reveal's capabilities in figure 2)

To be of any practical use, satellite altimetry must be corrected for a large number of atmospheric effects, including refraction (the *tropospheric correction*) (Saastamoinen, 1972) and the *inverse barometric effect* (Wunsch and Stammer, 1997). These two effects depend solely on the atmospheric pressure in the area where the observation is made.

Corrections for these effects are usually based on output from relatively coarse resolution global meteorological models. Knudsen and Brandt (1998) used a high resolution regional atmospheric model to verify by comparison, that the commonly used approach really is justifiable, even in the presence of extremely severe weather phenomena.

The altimetry used in the study were collected by the TOPEX/POSEIDON satellite (Fu et al., 1994), and postprocessed by the NASA Oceans Pathfinder project (NASA, 1996). We used a subset of the Pathfinder data: All data from the region around the North Sea for a 10.75 days period in October 1994. This amounts to a total of 11554 individual point observations.

The Oceans Pathfinder data comes in two versions: One version with fully corrected height observations, and one version which is not corrected for inverse barometric effects.

We used data from both versions, and removed the tropospheric correction by utilizing the fact that it is proportional to the inverse barometer correction. This was done using a small Bush script.

The resulting data set, where both corrections were excluded, could now be re-corrected. The new corrections were computed using high resolution sea surface pressure data generated from *MM5V1*, the Penn State/NCAR Fifth Generation Mesoscale Model (Grell et al., 1995). In our setup, MM5V1 was configured to run in a 70×70 polar stereographic grid, with a grid spacing of 25 km. Data were dumped once per hour (model time), so for the 10.75 days study period, we ended up with a 3D spatio-temporal grid of $70 \times 70 \times 259$ points. The file format, and stereographic projection parameters of this file were documented in a metadata file, and we were ready to roll:

First task was to get interpolated pressure values for the positions and times where we actually had altimeter observations. This was carried out using the Busstop module *Tegrin* (TEMPoral GRid INTERpolation). With this data set available, the new set of corrections could be computed and applied easily, using a Bush script.

Final step was to compute the internal statistics of the original fully corrected data set, and the new re-corrected data set. This was done using the Busstop module *Basta* (BASic STATistics). Comparison of the two sets of statistics showed, as might be expected, that they were virtually identical, and hence that the common practice of using corrections based on coarse resolution global models seems quite justifiable.

The work flow described above is not totally identical to the one originally published: We remade the study several times during Busstop's development cycle, and were gradually able to move more and more parts of the process from ad-hoc programs to Busstop. The task was seen as a *typical* example of a task that is especially suitable for implementation using spatio-temporal GIS methods, since:

- It includes heterogeneous data sets: Grid vs. point
- It involves different coordinates for georeferencing: Geographical coordinates for the point data, vs. polar stereographic coordinates for the grid data
- It is based on two different chronoreferencing systems: Continuous time for the point data, vs. discrete time for the grid data

The gravimetry used in the study was stored in the format used by the *Gravsoft* package for gravimetry and altimetry analysis (Tscherning et al., 1992), and the Gravsoft component *geogrid* was used in the gridding process. The integration of Gravsoft and Busstop was simply a matter of providing two metadata files: one describing the raw gravimetry, and one describing the grid format used by geogrid. Hence, this study is a neat example of the ease of integration supported by Busstop. Additionally, the study illustrates that Reveal's relatively limited visualization functionality is actually enough to ensure its practical value.

3.2 Atmospheric Corrections for Satellite Altimetry

The original inspiration for the study summarized below came when using Reveal to check the availability of satellite altimeter observations of the center of a severe atmospheric low pressure system. (the task used for illustrating Reveal's capabilities in figure 2)

To be of any practical use, satellite altimetry must be corrected for a large number of atmospheric effects, including refraction (the *tropospheric correction*) (Saastamoinen, 1972) and the *inverse barometric effect* (Wunsch and Stammer, 1997). These two effects depend solely on the atmospheric pressure in the area where the observation is made.

Corrections for these effects are usually based on output from relatively coarse resolution global meteorological models. Knudsen and Brandt (1998) used a high resolution regional atmospheric model to verify by comparison, that the commonly used approach really is justifiable, even in the presence of extremely severe weather phenomena.

The altimetry used in the study were collected by the TOPEX/POSEIDON satellite (Fu et al., 1994), and postprocessed by the NASA Oceans Pathfinder project (NASA, 1996). We used a subset of the Pathfinder data: All data from the region around the North Sea for a 10.75 days period in October 1994. This amounts to a total of 11554 individual point observations.

The Oceans Pathfinder data comes in two versions: One version with fully corrected height observations, and one version which is not corrected for inverse barometric effects.

We used data from both versions, and removed the tropospheric correction by utilizing the fact that it is proportional to the inverse barometer correction. This was done using a small Bush script.

The resulting data set, where both corrections were excluded, could now be re-corrected. The new corrections were computed using high resolution sea surface pressure data generated from *MM5V1*, the Penn State/NCAR Fifth Generation Mesoscale Model (Grell et al., 1995). In our setup, MM5V1 was configured to run in a 70×70 polar stereographic grid, with a grid spacing of 25 km. Data were dumped once per hour (model time), so for the 10.75 days study period, we ended up with a 3D spatio-temporal grid of $70 \times 70 \times 259$ points. The file format, and stereographic projection parameters of this file were documented in a metadata file, and we were ready to roll:

First task was to get interpolated pressure values for the positions and times where we actually had altimeter observations. This was carried out using the Busstop module *Tegrin* (TEmporal GRid INterpolation). With this data set available, the new set of corrections could be computed and applied easily, using a Bush script.

Final step was to compute the internal statistics of the original fully corrected data set, and the new re-corrected data set. This was done using the Busstop module *Basta* (BASic STATistics). Comparison of the two sets of statistics showed, as might be expected, that they were virtually identical, and hence that the common practice of using corrections based on coarse resolution global models seems quite justifiable.

The work flow described above is not totally identical to the one originally published: We remade the study several times during Busstop's development cycle, and were gradually able to move more and more parts of the process from ad-hoc programs to Busstop. The task was seen as a *typical* example of a task that is especially suitable for implementation using spatio-temporal GIS methods, since:

- It includes heterogeneous data sets: Grid vs. point
- It involves different coordinates for georeferencing: Geographical coordinates for the point data, vs. polar stereographic coordinates for the grid data
- It is based on two different chronoreferencing systems: Continuous time for the point data, vs. discrete time for the grid data

REFERENCES

- Ascher, D., Hinsien, P. F. D. K., Hugunin, J. and Oliphant, T., 1999. Numerical python. Technical report, Lawrence Livermore National Laboratory, Livermore, California. <http://numpy.sourceforge.net/numPy.pdf>.
- Beller, A., Giblin, T., Le, K. V., Litz, S., Kittel, T. and Schimel, D., 1991. A temporal GIS prototype for global change research. In: ASPRS (ed.), GIS/LIS '91 Proceedings, vol.2, ASPRS, Bethesda, Maryland, pp. 752–765.
- Bray, T., Paoli, J. and Sperberg-McQueen, C. M. (eds), 1998. Extensible Markup Language (XML) 1.0. W3C Recommendation, World Wide Web Consortium. <http://www.w3.org/TR/REC-xml>.
- Clark, J., 1998. expat—XML parser toolkit. <http://www.jclark.com/xml/expat.html>.
- Evenden, G. I., 1990. Cartographic projection procedures for the unix environment—a user's manual. Technical report, U.S. Geological Survey, Woods Hole, MA, USA. USGS Open File Report 90-284.
- Evenden, G. I., 1993. Cartographic projection procedures, release 4, interim report. Technical report, U.S. Geological Survey, Woods Hole, MA, USA.
- Franklin, W. R., 1972. PNPOLY—point inclusion in polygon test. <http://www.ecse.rpi.edu/Homepages/wrf/geom/pnpoly.html>.
- Fu, L.-L., Christensen, E. J., Charles A. Yamarone, J., Lefebvre, M., Ménard, Y., Dorrer, M. and Escudier, P., 1994. TOPEX/POSEIDON mission overview. *Journal of Geophysical Research* 99(C12), pp. 24369–24381.
- Grell, G. A., Dudhia, J. and Stauffer, D. R., 1995. A description of the fifth generation Penn State/NCAR mesoscale model (MM5). Technical report, National Centre for Atmospheric Research, Boulder, Colorado, USA. NCAR/TN-398+STR.
- Gudmandsen, P. (ed.), 1998. Proceedings, EARSeL symposium 1997: Future Trends in Remote Sensing. A. A. Balkema, Rotterdam, The Netherlands.
- Knudsen, T., 1997. Busstop—a spatio-temporal information system. PhD thesis, University of Copenhagen.
- Knudsen, T., 1998. Adding GIS functionality to specialized analysis programs for remote sensing data: Gravsoft and the busstop project. In: (Gudmandsen, 1998), pp. 373–378.
- Knudsen, T. and Brandt, J., 1998. GIS and high resolution atmospheric modelling for corrections of sea surface observations from the TOPEX/POSEIDON satellite altimeter. In: (Gudmandsen, 1998), pp. 125–132.
- Knudsen, T. and Olesen, A. V., 1998. A GIS based strategy for quality control of gravimetry. *Bulletin of the International Geoid Service* 8, pp. 85–92.
- Langran, G., 1989. A review of temporal database research and its use in GIS applications. *International Journal of Geographical Information Systems* 3, pp. 215–232.
- Langran, G., 1992. States, events and evidence: The principle entities of a temporal GIS. In: ASPRS (ed.), GIS/LIS '92 Proceedings, vol.1, ASPRS, Bethesda, Maryland, pp. 416–425.
- Larson, M., Shapiro, M. and Tweedale, S., 1991. GRASS tutorial: r.mapcalc. Technical report, U.S. Army Corps of Engineers, Construction Engineering Research Laboratories, Champaign, Illinois, USA.
- Mason, D. C., O'Connell, M. A. and Bell, S. B. M., 1994. Handling four-dimensional geo-referenced data in environmental GIS. *International Journal of Geographical Information Systems* 8(2), pp. 191–215.
- NASA, 1996. NASA ocean altimeter pathfinder data set, vol 1, colinear data: TOPEX/POSEIDON version 2. CDROM.
- Peuquet, D. J. and Duan, N., 1995. An event-based spatiotemporal data model (estdm) for temporal analysis of geographical data. *International Journal of Geographical Information Systems* 8(2), pp. 191–215.
- Rossum, G. V., 1995. Python Tutorial. Open Source Library. <http://www.python.org/doc/current/tut/tut.html>.
- Rossum, G. V., 1999. Python Library Reference. Open Source Library.
- Saastamoinen, J., 1972. Atmospheric correction for the troposphere and stratosphere in radio ranging of satellites. *Geophysical Monograph*, Vol. 15, American Geophysical Union.
- Shapiro, M. and Westervelt, J., 1992. r.mapcalc: An algebra for GIS and image processing. Technical report, U.S. Army Corps of Engineers, Construction Engineering Research Laboratories, Champaign, Illinois, USA.
- The MathWorks Inc., 1996. Using MATLAB. The MathWorks Inc.
- Tomlin, C. D., 1983. Digital Cartographic Modeling Techniques in Environmental Planning. PhD thesis, Yale University.
- Tomlin, C. D., 1990. Geographic information systems and cartographic modeling. Prentice Hall, Englewood Cliffs, NJ, USA.
- Tscherning, C., Forsberg, R. and Knudsen, P., 1992. The gravsoft package for geoid determination. In: Proceedings of the First workshop on the European Geoid, Prague, Springer Verlag, New York, pp. 327–334.
- USACERL, 1993. Grass 4.1 Reference Manual. U.S. Army Corps of Engineers, Construction Engineering Research Laboratories, Champaign, Illinois, USA.
- Vogel, R. L., 1998. Why scientists have not been writing metadata. *Eos* 79(31), pp. 373+380.
- Wall, L., Christiansen, T. and Schwartz, R. L., 1996. Programming Perl, 2nd Edition. O'Reilly & Associates.
- Wesseling, C. G., van Deursen, W. P. A. and Burrough, P. A., 1996. A spatial modelling language that unifies dynamic environmental models and GIS. In: NCGIA (ed.), Proceedings, Third International Conference/Workshop on Integrating GIS and Environmental Modeling, Santa Fe, NM, January 21–26, 1996, National Center for Geographical Information and Analysis, Santa Barbara, CA, USA. http://www.ncgia.ucsb.edu/conf/SANTA_FE_CD-ROM/main.html.
- Wunsch, C. and Stammer, D., 1997. Atmospheric loading and the oceanic “inverted barometer” effect. *Reviews of Geophysics* 35(1), pp. 79–107.