

Computational Cartography

(Dagstuhl-Seminar 9645)

Organizers:

Christopher Gold (Université Laval, Québec, Canada)

Jack Snoeyink (University of British Columbia, Vancouver, Canada)

Frank Wagner (Freie Universität Berlin, Germany)

November 4–8, 1996

The topic of *Computational Cartography* is intended to be broad enough to cover Computer-assisted Cartography, Geographic Information Systems (GIS), and Computational Geometry applied to these and related subjects. It is clear that in the last few years these types of applications have attracted more and more interest in the computer science community, and also that GIS and related applications are being influenced by developments in computer science more than used to be the case. In addition, the whole concept of managing space within the computer is opening up new possibilities not covered by the traditional titles. It is obvious that there will be significant change in the next few years, and it is at get-togethers like this that the situation becomes clearer. On one hand the cartographers and GIS people made clear that the complexity of their applications asks for the contribution of theoreticians in form of algorithms which are simple, easy to implement and to integrate into existing systems. On the other hand the participating theoreticians from computational geometry are strongly interested in not only claiming the applicability of their GIS-oriented research in the abstracts of their papers but in really doing useful things. The workshop took place at the right moment and brought together people that were open to each other.

The participants (30) came from Germany (10), the rest of Europe (9), North America (10), and Asia (1); most of them from universities (researchers from the Computational Geometry, Geoinformatics and Cartography communities), a few from GIS companies or institutions using such systems. During the workshop 25 talks were given. Main topics of interest were *Simplification/Generalization*, *3D-Modeling of terrains*, *Map Labeling*, *Multiresolution Modeling*, *Map overlay* and *user-oriented aspects of GIS design*.

In addition a panel discussion on *Problems in Practice*, a debate entitled *Approximately Precise*, a software demonstration, and an open problem session took place.

Participants

Pankaj Agarwal, Duke University, Durham, USA
Jochen Albrecht, Universität Osnabrück, Germany
Helmut Alt, Freie Universität Berlin, Germany
Wolfgang Bitterlich, ESRI, Redlands, USA
Leila De Floriani, Università di Genova, Italy
William Evans, University of Arizona, USA
Andrew Frank, Technische Universität Wien, Austria
Wm. Randolph Franklin, Rensselaer Polytechnic Institut, Troy, USA
Ulrich Freitag, Freie Universität Berlin, Germany
Stefanie Gerke, University of Oxford, Great Britain
Christopher Gold, Université Laval, Québec, Canada
Martin Heller, Universität Zürich, Switzerland
Klaus Hinrichs, Universität Münster, Germany
Christopher Jones, University of Glamorgan, Great Britain
Rudi Krämer, Stadt München, Germany
Hans-Peter Kriegel, Universität München, Germany
Werner Kuhn, Universität Münster, Germany
Michael McAllister, University of British Columbia, Vancouver, Canada
Jürg Nievergelt, ETH Zürich, Switzerland
Enrico Puppo, Università di Genova, Italy
Thomas Roos, ETH Zürich, Switzerland
Jörg-Rüdiger Sack, Carleton University, Ottawa, Canada
Stefan Schirra, Max-Planck-Institut für Informatik, Saarbrücken, Germany
Jack Snoeyink, University of British Columbia, Vancouver, Canada
Tiow-Seng Tan, National university of Singapore
Marc van Kreveld, Universiteit Utrecht, The Netherlands
Jeffrey Vitter, Duke University, Durham, USA
Frank Wagner, Freie Universität Berlin, Germany
Alexander Wolff, Freie Universität Berlin, Germany
Weiping Yang, Université Laval, Québec, Canada

Abstracts

Surface Simplification

Pankaj Agarwal

Computer Science Department, Duke University, Durham, USA

Surface simplification, especially terrain simplification, is a central problem in GIS. This talk surveys the known techniques for polygonal simplification in the plane, and for terrain simplification in 3D. It also presents a few new algorithms for polygon and terrain simplification. Furthermore, we discuss some open problems.

Scalability of GIS Operations

Jochen Albrecht

Institute for Spatial Analysis and Planning in Areas of Intensive Agriculture (ISPA), University Vechta, Germany

full version: <http://www.ispa.uni-vechta.de/staff/jochen/papers/gislis96.html>

This work relates to the problem area of building pluggable geoprocessing applications. Such an application is virtual in the sense that it is only created upon request. The benefit of this approach is that a high degree of customization of a system is reached and thus flexibility is attained. This means that no longer is it necessary to build and acquire huge 'one size fits all' systems. A further advantage is that the user is relieved from the burden of learning a huge number of system functionalities. The flexibility is attained by providing services at such a level of granularity that the 'right-size' application can be built. To determine a reasonable degree of granularity required from service providers, we need to analyze geoprocessing functions and determine how they might be built up from lower-level functions.

This paper is a preliminary discussion on an operator algebra that must be the foundation of such a pluggable environment. The reason why we require an algebra is that we have to be able to describe the required operators using such a degree of exactness that the results of applying operators are predictable. This requires that the input and output parameters and side effects of an operator are fully determined. Such operator algebra must go beyond spatial data models and representations. The paper presents ideas to map basic spatial operators to the hierarchy of conceptual spaces that different spatial models rely on.

(The full paper at the URL given above is not yet completely clean, i.e., some of the math formulas need to be inserted as graphics. But the reader should be able to extract the gist of it.)

Fast Automated Map Labeling

Wolfgang Bitterlich

ESRI, Redlands, USA

full version: <http://maps.esri.com/>

Presented is a system for automatic point, line and area labeling. The emphasis is on speed, iterative methods are not considered. Labels which the algorithm cannot place are either ignored or (optionally) displayed in different color to be moved manually by the user. All produced labels are straight (not curved) along a line which is horizontal for point and area labels and can have any angle for line labels. The general approach is to sequentially identify for each label a set of strips in which the current label can be placed. Strips may be cut into shorter strip intervals by the position of already placed labels. Among the remaining strip intervals the "bestposition is chosen (or the label declared unplaceable if no sufficiently long interval remains). All label requests are entered and sorted before any placement begins. Lines with the same label name are grouped in order to produce only one label for each connected component. The labeling system is implemented in ESRI's ArcView, MapObjects, VISA ATM Locator and Internet Map Server.

A Formal Approach to Multiresolution Modeling

Leila De Floriani

Dipartimento di Informatica (DISI), Università di Genova, Italy

joint work with Paola Magillo, Università di Genova, and Enrico Puppo, CNR, Genova

The aim of this work is to provide a framework for multiresolution geometric modeling which is independent of both the dimension of the spatial objects under consideration and of the specific application. We present a formal model, called Multiresolution Cellular Complex (MCC), capable of capturing the characteristics of most models known in the literature. We provide an analysis of the relationships among the intrinsic structures of different multiresolution representations proposed, and a definition of a set of basic, application-independent operations on a multiresolution model. Finally, major data structures used to encode multiresolution models are reviewed, as well as algorithms which implement the basic operations on each data structure.

Right Triangular Irregular Networks

William Evans

Department of Computer Science, University of Arizona, USA

joint work with David Kirkpatrick, University of British Columbia

Geographic information systems strain the computational ability of computer hardware to its limit. Only recently have advances in computer hardware made possible the real-time three-dimensional visualization of large geographic data sets. Even with these advances, common geographic data sets are too detailed to permit real-time graphical

manipulation – some simplification and approximation of the data must be done. We describe a hierarchical data structure for representing a digital terrain (height field) which contains approximations of the terrain at different levels of detail. The approximations are based on triangulations of the underlying two-dimensional space using right-angled triangles. The methods we discuss allow the approximation to precisely represent the surface in certain areas while coarsely approximating the surface in others. Thus, for example, the area close to an observer may be represented with greater detail than areas which lie outside their field of view. We discuss the advantages of this method in terms of space usage and rendering time as compared to other methods based on more general triangulations. We also describe additional benefits such as rapid point location and neighbor calculation.

GIS in the Real World

Andrew Frank

Abteilung Geoinformation, Technische Universität Wien, Austria

Starting with the observation that commercial Geographic Information Systems (GIS) very seldom use algorithms from Computational Geometry (CG). To the practitioners the CG algorithms

- appear as very complex to implement,
- often assume real number arithmetic, for which computers provide only an approximation, and
- do not integrate with other parts of a system, e.g., database management system, memory management; if the efficiency of CG algorithms depends on a particular form of memory management, it will be very difficult to integrate it in a practical system of any complexity.

Additional difficulties are caused by the relative isolation of the two communities; the approaches used, the channels of communication and the languages used to describe the results are different and thus communication between GIS implementors and CG researchers is hindered. A most revealing example is the use of the term *generalization* by cartographers.

It is somewhat surprising to discover that a simple and fundamental problem such as the computation of the overlay of two polygonal meshes (i.e. partitions of space, so that all areas are jointly exhaustive and pairwise disjoint) to compute all resulting polygons, is still a practically open problem. The GIS industry has spent more than 100 person-years to program a general solution, but the best available programs still occasionally fail to produce a result for a new dataset and are continuously *fixed*. It is not clear if and with which restriction a solution is possible using a finite approximation for real numbers.

The industry will only use simple algorithms which are robust and produce results for any input (or fail in a graceful way). Industry will shy away from algorithm which cannot be integrated in a large program where data is stored in a database. The art of

a practically successful CG researcher is to abstract from reality, but stop before the abstraction does not contain the original problem any more (Occam's razor). Small improvements, especially in the worst case (behaviour) is not important to the industry; large performance improvements are possible by exploiting regularities in the data given by their meaning in the real world.

On Hierarchies

Andrew Frank

Abteilung Geoinformation, Technische Universität Wien, Austria

Using Hierarchical Methods for spatial reasoning is a popular research topic. Hierarchical spatial data structures, especially quadtrees, are used in many implementations of GIS and have proved their efficiency. Operations on hierarchical spatial data structures are effective to compute spatial relations, but do not automatically lead to hierarchical spatial reasoning. Hierarchical spatial reasoning is using coarser, less detailed representations to produce only an approximative answer if the quality of the approximation is sufficient for the task at hand. Therefore, hierarchical spatial reasoning is closely related to computing approximative results and estimation of their errors; it is related to *any time* algorithm, which produce a result immediately, but can produce better results if more time is available.

Humans use hierarchical methods often; in space they are often based on a container hierarchy. This is best observed in situations where hierarchical reasoning is misleading and humans make predictable errors (e.g., "Which city is more to the west? San Diego (California) or Reno (Nevada)? Humans seem to associate with most values a sense of their quality (precision and other data quality aspects) and violating these rules is ridiculed (as in jokes about accountants, reporting the national deficit as \$13,459,345,678.13, when nobody regards the 13 cents as realistic). This paper explores two spatial reasoning operations and deduces a general definition of *hierarchical spatial reasoning*. Although the examples are very simple - computation of area and intersection - and applied to a raster representation, the definition is general. Compared with other definitions it captures much of the essence of hierarchical spatial reasoning. This sets the framework in which general rules may be deduced, when hierarchical spatial reasoning can be employed.

For hierarchical spatial reasoning the following is required:

- a coarsening method to produce less detailed representation from the most detailed one,
- a function f of interest, which can be computed from these representations of varying level of detail,
- a corresponding function f to assess the error of computing f from a less detailed representation (in comparison with the computation on the most detailed representation).

Hierarchical spatial reasoning starts computing the value of interest from the least de-

tailed representation and progresses till a result is found for which the error assessment satisfies the requirements of the task at hand. Then no more details are necessary and the refinement can stop. Hierarchical data structures are useful for hierarchical reasoning, but they can be transformed to a more efficient *incremental hierarchical structure*, e.g. an incremental quadtree. On these, incremental hierarchical spatial reasoning algorithms are using previously computed values to compute the next approximation and are therefore as efficient as a direct calculation with the same error bound.

Elevation Data Operations

Wm. Randolph Franklin

ECSE Dept., Rensselaer Polytechnic Inst, Troy, USA

full version: <ftp://ftp.cs.rpi.edu/pub/franklin/dagstuhl.ps.gz>

We describe several past and proposed projects for converting elevation data to different formats, such as gridded, TIN, and spline, lossily compressing, applying to visibility and drainage to test the compression quality, and integrating everything into a test suite.

Collaboration between GIS and Computational Geometry

Christopher Gold

Centre de Recherche en Géomatique, Université Laval, Canada

This presentation concerns the conjunction of GIS and Computational Geometry (CG). From the GIS side, we need to become more familiar with data structures, algorithms and robustness issues. From the CG side there needs to be more realization that spatial data management includes much more than algorithms – even at the technical level it is important to know what the customer wants, and what application problems are important. A list of topics would be a help to the CG community – but even more, there should be a realization that the future requires people with an understanding of the applications problems as well as an understanding of the basics of the appropriate algorithms.

It is not a question of the mathematicians providing a simple answer and then stopping – most of the problems involve a complex interaction between the real world, its computer representation, the appropriate analysis tools, and the perceptions and needs of the users. We can no longer tell the users to conform to us: we must understand the basics of their world-view. My particular concern, for example, has been the idea of interaction – the user is continually updating the data or requesting different analyses in near real-time. This changes considerably the idea of the user interface, as well as the update algorithms. Yet this is only infrequently mentioned in CG. We need much closer collaboration in order to develop the next generation of spatial data management tools.

Map Overlay

Klaus Hinrichs

Fachbereich Informatik, Westfälische Wilhelms-Universität Münster, Germany

joint work with Ulrich Finke

Applications in cartography, computer graphics and computer aided engineering require the representation and manipulation of planar subdivisions. Topology-oriented approaches represent the relations existing between the vertices, edges and faces of a planar subdivision. We extend this standard approach by decomposing the faces into trapezoidal views. By defining a neighborhood relation among the views we obtain the view graph which forms the basis of the quad view data structure - a new data structure for representing planar subdivisions. We present some basic operations on this data structure, and we show how it can be used to traverse a planar subdivision. Furthermore we sketch an algorithm for overlaying planar subdivisions represented by the quad view data structure. This algorithm is optimal for simply connected planar subdivisions: its time and storage requirements are linear in the size of the output subdivision. Finally we apply this overlay algorithm to implement locational vector-based map overlay operations.

A Triangulated Data Model to Support Map Generalization

Christopher Jones

Department of Computer Studies, University of Glamorgan, Great Britain

Map Generalization is concerned with reducing the level of detail of a map object subject to constraints of scale and purpose. At a geomtric level it includes processes of shape simplification, object elimination, merging, reduction in dimensionality, typification, exaggeration and displacement of map features.

A triangulated spatial model is presented, along with a set of operations which facilitate implementation of generalization processes. The model is based on constrained Delaunay triangulation and is motivated by several requirements: manipulation of parts of map objects and parts of the space between objects; high precision; easily measured; maintenance of topological relations; efficient determination of proximity relations between neighbouring objects; and malleability to assist in determining the nature of interactions between objects following generalization operations.

Several implemented generalisation operators are described. These include the detection of conflicts and their resolution by displacement, through monitoring of triangle inversions due to overfolding of the triangulation; dimensionality reduction by generating a skeleton from the triangulation; building boundary simplification through flipping of right-angled external triangles; and merging of nearby objects by collapse or re-attribution of the free space triangles between the objects.

The Lettering of Maps and Related Problems

Rudi Krämer

City Authorities, Munich, Germany

In this talk, an example for the problem of map labeling was presented: in Munich we have about 19.000 drill holes on a square of about 20 km \times 20 km. These points (or a subset of them) are to be labelled on various maps in a such way that a (rectangular) label is attached to its site in one of its four corners without overlapping other labels or points.

A short analysis of this problem was given; it was mentioned, that checking all possibilities for lettering these points would be of order $\mathcal{O}(4^n)$, where n is the number of points to be labelled; a simple trial-and-error algorithm with a reasonable running time (and a performance guarantee of 25%) was presented. We would like to find a method with a better quality guarantee if such a method exists.

If all points must be labelled (i.e., we can blow up the map), then the problem can be solved approximately by a fast (non-trivial) algorithm with a running time of $\mathcal{O}(n \log n)$, suggested by Wagner and Formann. Unfortunately this method produces insufficient results in practice.

Another practically very relevant problem is the following. Given the height of the groundwater level as well as the kind and thickness of the different geological strata found in our drillholes, we would like to develop an algorithm which is able – to produce a 3D model of these layers,

- to point out areas of inconsistency given a set of rules, and
- to draw cross-sections of the model between any two points?

An example for such a rule would be “there cannot be layer A above layer B” or “if layer A contains water in some drillhole, then it must also contain water in a neighbouring drillhole if it is not at a higher ‘altitude’ there”.

Parallel Processing of Spatial Joins

Hans Peter Kriegel

Institut für Informatik, Universität München, Germany

full version: http://www.dbs.informatik.uni-muenchen.de/index_e.htm

joint work with Thomas Brinkhoff, Mannesmann AutoCom, Düsseldorf, and Bernhard Seeger, Fachgebiet Informatik, Universität Marburg

In this work, we show that spatial joins are very suitable to be processed on a parallel hardware platform. In particular, we use a shared virtual memory architecture which is well-suited for the design and implementation of parallel spatial join algorithms. We start with an algorithm that consists of three phases: task creation, task assignment and parallel task execution. In order to reduce CPU- and I/O-cost, the three phases are processed in a fashion that preserves spatial locality. Dynamic load balancing is achieved by splitting tasks into smaller ones and reassigning some of the smaller tasks to idle processors. In an experimental performance comparison, we identify the advantages and disadvantages of several variants of our algorithm. The most efficient

variant shows an almost optimal speed-up under the assumption that the number of disks is sufficiently large.

A Human-Computer Interaction View of Cartography

Werner Kuhn

Inst. for Geoinformatics, Univ. of Münster, Germany

full version: <http://gio.uni-muenster.de/>

Computational cartography involves concerns beyond geometry and databases. This talk takes a perspective on cartography from the field of Human-Computer Interaction (HCI). It presents some approaches and methodology that are used in interaction design and that appear useful for tackling some problems in modern, interactive cartographical settings. The major conclusion is that HCI's focus on usability, based on careful analyses of tasks and user communities, has the potential to make mapping both more straightforward and more useful in practice.

Outline:

1. A Recent History of Cartography

- Robinson
- Behavioral Studies
- Communication Studies

2. The Current Situation

- Cognitive Approaches
- Cartography faces new challenges...
- ... but is entrenched in 'AI' complete problems

3. Entailments of an HCI View

- Task orientation
- User centered design
- Dynamic interaction
- More modalities
- Map engineering
- Cooperation Support
- Help Systems

4. Conclusions.

Drainage Network and Terrain Generalization
Michael McAllister
University of British Columbia, Vancouver, Canada

We look at the problem of automatically extracting watersheds from a triangulated irregular networks (TINs) and hydrology information. When considered individually, each of these sources of data can define watersheds. For TINs, you must first derive the river system from the surface and can then trace ridges and river divides as bounding curves for the watersheds. For hydrology information, you can use the Voronoi cells of the river edges as an approximation to the edge's watershed. However, each of these methods admit errors: the drainage on the TIN may not match the known river network and the Voronoi cells of each river edge does not necessarily respect the surface.

The combination of the TIN and hydrology yields better watershed boundaries. First, we use an approximation to the medial axis to collapse all the lakes and river banks in the hydrology into tree structures that reflect the flow of water in the system. Next, we embed the reduced river network into the TIN as break lines to remove TIN artifacts that contradict the drainage system on the surface. Finally, we trace river divides from each river junction on the TIN, making sure that divides do not cross rivers, and combine the divides with the ridges on the TIN to form the watershed boundaries.

A Topological Model for the Generalization and Multiresolution Representation of Geographic Maps

Enrico Puppo

Institute for Applied Mathematics, IMA-CNR, Università di Genova, Italy

full version: <ftp://ftp.disi.unige.it/pub/person/PuppoE/PS/ssd95.ps.Z> and [sdh96.ps.Z](ftp://ftp.disi.unige.it/pub/person/PuppoE/PS/sdh96.ps.Z)

joint work with Giuliana Dettori, IMA-CNR, and Michela Bertolotto, DISI

A 2D geographic database made of points, lines, and regions is considered, a combinatorial and geometric model to represent it is proposed, and changes of resolution are studied within such framework.

Spatial entities are constrained to form a weakly disjoint covering of the plane, i.e.: no two lines can cross; no two regions can overlap; isolated points and dangling lines completely contained inside regions are admitted; regions may have holes. Topological relationships are highly simplified.

Concepts concerning changes of resolution are divided into two classes: detail (involving topology) and accuracy (involving metrics). A purely combinatorial structure, the abstract cell complex (ACC), is used to model all topological aspects of data, and changes of detail are studied as continuous functions between ACCs. Formal tools to control the consistency and graduality of generalizations between ACCs are proposed. Some model-oriented generalization operations are revisited in this framework. Metric aspects are treated through epsilon-homotopies of lines.

Multiresolution is obtained from collections of ACCs pairwise related through consistency functions. Either multilevel or tree models of multiresolution can be built in such a framework.

Fully Dynamic and Kinematic Voronoi Methods in GIS

Thomas Roos

ETH Zürich, Switzerland

full version: <http://www.inf.ethz.ch/personal/roos/>

joint work with Christopher M. Gold, Peter R. Remmele (ETH Zürich)

We give a survey of static, dynamic, and kinematic Voronoi diagrams as a basic tool for Geographic Information Systems (GIS). We show how the Voronoi diagram with its dual, the Delaunay multigraph, can be used to maintain the topology of the map objects of a GIS. The presented method allows the insertion, deletion, and translation of points and line segments in a Voronoi diagram of n generators. All elementary operations are available in $\mathcal{O}(\log n)$ expected time using expected linear storage.

The Voronoi approach also greatly simplifies some of the basic traditional GIS queries and allows even new types of higher level queries. The concept of a persistent, locally-modifiable spatial data structure that is always complete provides an important new approach to spatial data handling that is not available with existing systems.

Parallel Neighbourhood Modelling

Jörg-Rüdiger Sack

Carleton University, Ottawa, Canada

full version: <http://www.scs.carleton.ca/~gis>

joint work with PARADIGM group, Carleton University

It has been observed that in recent years Geographical Information Systems (GIS) and Spatial Information Systems have gone through substantial changes with respect to users, problems, problem domains, and data. The GIS community is rapidly expanding to include users from different sectors of the economy; along with this come different demands regarding the type, required speed, scope and scale of applications. Users in decision making positions require rapid, close to instantaneous responses even to complex queries. In addition, users today have access to an unprecedented amount of high resolution and high-quality data through scanners, satellites, range finders, medical equipment and other devices.

The effect of these changes is a rapid and huge increase in the computational demands placed on GIS. Processing a large raster often takes hours or even days; (processing a large raster, say of size 6000x6000 cells, at a speed of 1000 cells per second, would take 10 hours). To keep up with the computational demands without sacrifice (i.e., reduction in resolution or scope of model), parallel computing appears to be the only solution. Parallel hardware is readily available at a good price-to-performance ratio (in the small to medium range).

While parallelism in GIS appears to be necessary, the task of providing parallelism is highly challenging. Our primary research and development objective is to enable users, researchers and developers within GIS to use parallel computers without paying the high price of having to deal with the complex issues inherent to it (we provide *transparent parallelism*).

Here we report on our first major milestone towards achieving this objective: the design and prototype implementation of an environment for parallel raster-based NEighbourhood MOdelling (NEMO). NEMO is primarily intended for (but not necessarily restricted to) coarse-grained parallelism with a limited number of processors.

We demonstrate the validity of our approach and the efficiency of the system through the implementation of a variety of applications: several forest fire modelling applications (in collaboration with Forestry Canada and K. Clarke, UC Santa Barbara), earth quake modelling, ice-tracking for the Canadian shipping (Canada Remote Sensing), Image processing, Cartography (implementation of nearly all of Tomlin's Map Algebra functions). (NEMO has been delivered to our industrial partner ALMERCOCO Inc. for commercialization.)

A Computational Geometry's Algorithmic Library (CGAL)

Stefan Schirra

MPI für Informatik, Saarbrücken, Germany

full version: <http://www.cs.ruu.nl/CGAL/>

CGAL is an ESPRIT LTR project of Utrecht University, Tel Aviv University, Free University Berlin, ETH Zurich, INRIA Sophia-Antipolis, RISC Linz, and MPI Saarbrücken. A goal of CGAL is to make computational geometry available for (industrial) applications, e.g. in geographic information systems. To this end a library of geometric objects and algorithms is developed, in C++. The library is called CGAL, too.

The first part of the talk is a brief survey on CGAL, especially on its kernel.

The second part is concerned with precision and robustness problems in the implementation of geometric algorithms. Exact geometric computation is recommended as the most promising way to get reliable implementations of algorithms originally designed for arithmetic over the real numbers (this is my personal view and not necessarily the view of CGAL). Further we discuss the number type `real` in LEDA (<http://www.mpi-sb.mpg.de/LEDA/leda.html>) and techniques used there for exact computation, more precisely, floating-point filter, lazy evaluation, and separation bounds. The number type `real` provides exactness for all computations involving operations $+$, $-$, \cdot , $/$ and $\sqrt{}$, starting from integral or rational operands.

The third part is on (exact) geometric computation in CGAL. CGAL objects are parametrized by a number type, which gives you a lot of flexibility and lets you choose CGAL components according to your specific needs.

How to Send a TIN

Jack Snoeyink

University of British Columbia, Vancouver, Canada

joint work with Marc van Krefeld

Although it takes $\Omega(n \log n)$ time to build a Delauney triangulation of n points, once one is built, you can find, in linear time, a permutation of the points such that the Delauney triangulation can be reconstructed from the points in $\mathcal{O}(n)$ time.

This theorem gives a method to send a TIN model over a network that allows “progressive sending” for unstructured surface meshes.

Infrastructure Previewing

Tiow-Seng Tan

Dept. of Information Systems & Computer Science, National University of Singapore, Singapore

full version: <http://www.iscs.nus.sg/~tants>

We present an overview of an infra-structure previewing system currently under research and development in our Computer Graphics Research Lab. The aim of the prototype system is to preview interactively computer models of new infra-structure, and assess the impact of the new infra-structure on existing landscape. The various components of the system that are relevant to this workshop include, for example, landscape modeling tools, model simplification algorithms, collaborative design module, and augmented reality interface. Also, we describe some recent techniques in computing good triangulations that may be of interest to the GIS community.

A Note on Models for Mountains

Marc van Krefel

Universiteit Utrecht, The Netherlands

In physical geography there are several qualitative definitions of peaks and mountains. These definitions, however, don't really reveal how a definition should be formalized to a form that is unambiguous for a terrain surface. Such a definition is needed if mountains are to be identified automatically by a system. Possible applications are spatial data analysis, and terrain generalization.

In geomorphometry, a subfield of geomorphology, no such definitions can be found either. There are, however, a couple of quantitative descriptions of specific landforms that give pointers as to what parameters can be used, and what other shape descriptors are of interest. These include mean slope, ratio of diameter and depth, profile convexity, the hypsographic curve, and so and so forth.

If we suppose that a quantitative definition has been given, the computation of mountains and peaks can start. Computational geometers can develop efficient algorithms that can be used in an implementation. It may be possible that the best algorithm

that is found is not efficient enough for the particular application, for example if the definition requires complex optimization criteria or the algorithm has to be used in interactive situations.

The example of mountains partly served as an illustration of the process of automating a geographic (and geometric) task. It should start with geographers trying abstract certain geographic issues, if possible. In the development of a quantitative definition geographers and computational geometers can work together. Computational geometers have a good feel for the implications of quantitative definitions concerning geometry in general. If the definition has been decided upon, algorithms design can start, by computational geometers. Finally, the algorithms can be implemented to complete the automation.

Efficient Algorithms for Processing Line Segments in External Memory, with Applications to Databases and Geographic Information Systems

Jeffrey Vitter

Duke University, Durham, USA

full version: `file://cs.duke.edu/pub/jsv/Papers/ArV96.interval_management.ps.gz`
and `AVV95.SegmentGIS.ps.gz`

In the design of algorithms for applications that deal with massive data sets, it is important to consider the problem of minimizing the I/O communication. Geographic information systems (GIS) and constraint and object-oriented database systems often operate on very large spatial data sets. The geometric nature of the data arises either directly (as in geographic data) or because it is a Euclidean representation of relationships among objects or classes.

In this talk, we survey the known I/O results of two types: The first type deals with batch problems, in which each data set is not preprocessed and must be considered in full. We list I/O-optimal algorithms for batch problems such as trapezoidal decomposition, multiple planar point location, Delaunay triangulation, Voronoi diagram, conversion from gridded representation to multiple contours, red-blue line segment intersection, and general line segment intersection. The number of I/Os used in each case is

$$O\left(\left(\frac{N}{B} + \frac{K}{B}\right) \log_{M/B} \frac{N}{B} + \frac{T}{B}\right),$$

where N is the input data set size, M is the internal memory size, B is the block transfer size, K is the number of concurrent queries (where applicable), and T is the output size.

The second type of results pertains to dynamic problems that support insert, delete, and query operations. We present a new and practical external memory version of the interval tree that uses a novel weight-balanced B-tree structure, and we use it to give an I/O-optimal algorithm for stabbing line queries and interval intersection. The algorithm uses $O(\log_B N + T/B)$ I/Os per query, $O(\log_B N)$ I/Os per insert and delete, and $O(N/B)$ blocks of disk space. We also survey recent work in more generalized forms of dynamic range searching.

An Efficient and Effective Approximation Algorithm for the Map Labeling Problem

Alexander Wolff

Institut für Informatik, Freie Universität Berlin, Germany

full version: <http://www.inf.fu-berlin.de/map-labeling/papers.html> or esa.ps.gz
joint work with Frank Wagner

The *Map Labeling Problem* is a classical problem of cartography. Our formalization is as follows: Given n distinct points in the plane. Find the supremum σ_{opt} of all reals σ such that there is a set of n closed squares with side length σ , satisfying the following two properties.

- Every point is a corner of exactly one square.
- All squares are pairwise disjoint.

For this problem there is an approximation algorithm A which is theoretically optimal. It runs in $\Theta(n \log n)$ time and guarantees a label size of 50 percent of the maximum. Unfortunately A is useless in practice as it typically produces results that are intolerably far off the optimal size. On the other hand there are some heuristics with good practical behaviour which are used in real applications.

We presented a new approximation algorithm B which keeps A 's quality guarantee and runtime, and at the same time yields practical results closer to the optimum than the best heuristic known so far.

The sample data used in the experimental evaluation consists of three different classes of random problems and a selection of problems arising in the production of groundwater quality maps by the authorities of the City of Munich.

All heuristics, approximation algorithms, and exact solvers for the Map Labeling Problem we have implemented so far, can be tested at this URL:

<http://www.inf.fu-berlin.de/map-labeling/>

A System Approach to Computational Cartography

Weiping Yang

Centre de Recherche en Géomatique, Université Laval, Québec, Canada

Previous cartographic research and development for generalization are basically isolated, static ones, and impose a uniform resolution on one map. More intelligent uses of digital maps require that maps are designed and represented dynamically to fit the human mental models of spaces. This paper proposes a system approach to tackle problems involved in automated map generalization. The objectives are to combine database generalization and dynamic object generalization capabilities in the system, and to superimpose a map agent on top of a map object which constructs cognitive maps, performs tasks on behalf of, and communicates with users. The object classes are topologically and geometrically structured with the dynamic VMO-tree (Voronoi based Map Object Tree).