On the Education of GIS Algorithm Design

Marc van Kreveld

institute of information and computing sciences, utrecht university

technical report UU-CS-2003-029

www.cs.uu.nl

On the Education of GIS Algorithm Design*

Marc van Kreveld

Institute of Information and Computing Sciences
Utrecht University
E-mail: `marc@cs.uu.nl`

Abstract

To teach the design of algorithms in an algorithmic GIS course, it is suggested to complement lectures with a particular type of project. How such a project can be set up is discussed, as well as a number of interesting results from the projects themselves. The projects address relative position of two regions, island group labeling, non-contiguous area cartograms, urban street network generalization, and diagram placement in administrative regions. The suggested approach of algorithm design is generally useful in research, not only for educational projects.

1 Introduction

Since a number of years I have taught a GIS course as part of an algorithmic specialization in computer science. The course used to be a fourth year doctoral study course, but with the introduction of the bachelor-master system in the Netherlands, it became a master course of the master program *Geometry, Imaging and Virtual Environments* [6].

The students that follow the course have a thorough background in algorithms, including the standard sorting and searching topics (bubble sort, merge sort, quicksort, balanced binary search trees, hashing, complexity analysis), and also some more advanced algorithms like network flow and shortest paths on graphs. Since they haven't had any courses in the spatial sciences, the GIS course I teach includes a number of the basic ideas and concepts from GIS.

- Geographic data, scales of measurement, aggregation
- Geographic data representation
- Data sources and acquisition, types of error
- Geographical analysis, buffers, overlay, models
- Automated cartography: label placement, map generalization
- Data structures: quad trees, R-trees
- Digital elevation models and algorithms

---

[*] Presented at the Dagstuhl seminar on Computational Cartography and Spatial Modeling, September 28 to October 3, 2003. `http://www.dagstuhl.de/`

The lectures of the course are algorithmically inclined, and algorithms are given for various computational problems that are encountered: raster-vector conversion, buffer computation, map overlay, line simplification, label placement, generalization, algorithms for building and searching in quad trees and R-trees, and algorithms for elevation models. Most algorithms are vector-based. Although the lectures provide a good overview of GIS and algorithms, it lacks active participation of the students in the design of algorithms for GIS.

In my opinion, the major difference between algorithms for GIS, and algorithms research in computer science is that the problems to be solved are ill-formulated. They are not formalized to clean, easy-to-state computational tasks. For example, label placement for a river label requires many different requirements to be satisfied at once, like the label being close to the river, not intersecting with the river, following the bends of the river, not intersecting other map features or labels, being as much horizontal as possible, not having too large curvature, and so on and so forth. Any river label placement method that uses only two of these requirements can generate placements that are unacceptable on maps. Even a list of requirements as just given doesn't lead to a well-defined computational task easily. For example, the closeness of the label and the river can be measured in many different ways: smallest distance between the label and the river, smallest distance of center-of-gravity of the label and the river, Hausdorff distance (directed or undirected), and so on. To provide the students with the knowledge of different possible formalizations and distance measures, one additional lecture of the course treats these issues: summary statistics for spatial objects, size measures, distance measures, and similarity measures. This lecture was given as early as possible in the course, so that the project could start with only little delay after the start of the lectures.

The observations above lead to the idea that the study of formalizations for a GIS problem is an essential aspect of the algorithm design. This is the basic principle of the student projects that complement the lectures of the GIS course. This paper describes the set-up of these projects and the results that were obtained during the GIS class of the Spring of 2003. The next section describes the two main phases of the projects in more detail. These phases are: (1) requirements study and problem formulation, and (2) algorithm design and efficiency analysis. Section 3 describes several projects and the results that were obtained. Several observations and methods are interesting in their own right. During the course, which was 160 hours worth of credits to be done in eight weeks, there was no time to include implementation and testing on real-world data.

As noted, the suggested approach of algorithm design for GIS is generally applicable in GIS research. It seems to be useful primarily for cartographic computational problems, although applications to other GIS tasks exist as well.

2 Set-up of projects

All projects were done by a pair of students. With two students, it is possible to share ideas and do some brainstorming, whereas with three students, the chances are that one of the students doesn't play an active role. Every pair was handed out two pages with a description of the project and the activities and assignments that are part of the project, together with a time line. The time line included three appointments with the lecturer to help with the generation of useful ideas and to adjust possible undesirable deviations from the direction I had in mind.

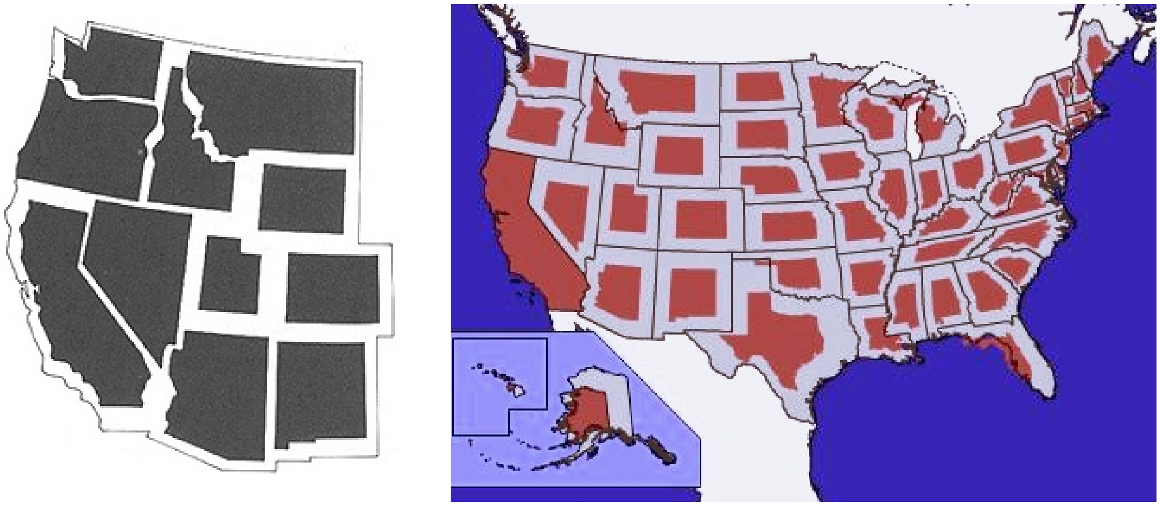

Figure 1: Left, non-contiguous area cartogram without inner boundaries. Right, with inner boundaries (from [5]).

Below one of the project descriptions is given, namely for non-contiguous area cartograms (minimally adapted for this paper). An example of such a map is shown in Figure 1, see also [13]. The projects were handed out on May 15. The first meeting with each pair of students separately was set on May 22, to ascertain that the project description is clear and what is expected for the phase one assignments.

Non-Contiguous Area Cartograms

Cartograms are a type of map in which the size of administrative regions (countries, states, provinces) is changed. This is done with purpose of not representing the actual surface area, but some other geographical variable like population. For certain applications this is more relevant. For example, for maps that show political preference the population of regions is more relevant than the size of the regions. Non-contiguous area cartograms are a type of cartogram in which the shape of the regions is preserved, but not the size or scaling factor. This implies that the connectivity (adjacency) of the regions cannot be kept, because otherwise overlap would occur.

Figure 11.3 from the book of Borden Dent, *Cartography: Thematic Map Design,* shows the difference between a normal cartogram and a non-contiguous area cartogram. Figure 11.11 also shows a non-contiguous area cartogram, although the region shapes are somewhat simplified.

In this project the construction problem of non-contiguous area cartograms will be studied. We assume that an administrative subdivision is given, which contains for every region the value of a geographic variable to be mapped, like the population of the region.

Assignment (phase 1): Read Chapter 11 from the above-mentioned book. Study Figures 11.3 and 11.11 carefully, and think about how such maps could be constructed automatically. To this end, first analyze which aspects play a role in this type of visualization.

Then start brainstorming, imagine which (geometric) criteria are of influence on the quality of such maps. The shapes of the given regions are fixed, but the positioning and the global scaling factor are not fixed (the *relative scaling factors* obviously are fixed!). Try to come up with as many criteria as possible that influence whether the resulting map is a good non-contiguous area cartogram. A minimum is four criteria. The criteria may be somewhat intuitive (like 'density on the map', although this criterion will not be so relevant in this project). At the same time, try to avoid vagueness. Make a list of the different criteria. It will form the first part of the first assignment to be handed in.

Next, figure out which criteria seem to be conflicting and which seem to cooperate (seem to imply or enhance each other). Some pairs may be largely independent. Do this for every pair of criteria. This study of pairs of criteria forms the second part of the first assignment.

As the next step, design measures with which you can determine the quality of the cartogram. Do this separately for every criterion. The possibly intuitive criteria given before will become quantified. Use measures from the lectures or of your own creation, whichever seems more appropriate. Often it is possible to quantify a criterion in several different ways. Think about this thoroughly, give the possibilities, and make a well-founded choice. Describe the measures and choices made, forming the third part of your first assignment.

The thinking and notes for the three parts is what should be finished before June 2, and will be discussed with the lecturer on June 2. Bring your notes!

During this second project meeting on June 2 it will also become clear what the geometric algorithmic problem statements can be. These will be developed further into two problem specifications, which form the fourth and last part of the first assignment.

On June 12 the report of first assignment must be handed in. It will consist of the four parts described above. You don't have to give an introduction or text about the usefulness of cartograms; restrict yourself to the four parts that lead to the geometric algorithmic problem statements. This report should have a length of 1500–2500 words (3–5 pages).

Assignment (phase 2): The report of the first phase will be discussed with the lecturer during the third project meeting on June 16. At the same time we will look ahead on the second phase. In the second phase the two geometric algorithmic problem statements are taken in order to design algorithms. In case of two conflicting criteria, you must choose how to respect both to a certain degree. This may involve setting a minimum requirement on the one criterion and optimizing the other one, or vice versa.

Algorithms that try to optimize one or more criteria can be iterative (like evolutionary algorithms) or combinatorial (like genetic algorithms, or incremental greedy algorithms). Which approach is more suitable for the two problem statements should be considered. You should make notes of your ideas concerning this issue.

One of the two problem statements will be chosen, and an algorithm for it designed and analyzed. The efficiency analysis should not only be a worst-case analysis; with iterative algorithms this isn't possible to begin with. However, you should then analyze the efficiency of any iterative step of the algorithm. For genetic algorithms you should analyze the efficiency of computing the fitness. Also, you should try to judge whether a worst-case efficiency bound is expected to be typical in practice. Make practical assumptions about the input data and use these to obtain better, typical case running time estimates.

In the time between the third project meeting (June 16) and the handing in of the second assignment no fourth project meeting is planned. However, it is recommended to visit the lecturer in case you have questions or need help with the algorithm design and analysis; this part need not be done independently.

The notes on conflicting criteria, the ideas about iterative or combinatorial algorithms, the algorithm itself, and the efficiency analysis form the basis of the second report to be handed in. Furthermore, your report should contain a short review (10–20 lines) in which you discuss to what extent your algorithm will result in a qualitatively good map. In other words, to what extent you expect to fulfill the listed criteria. The second report should also be 1500–2500 words in length. The date for handing in is July 3.

The example project description above is divided into two phases that can be seen as familiarization and understanding, and design and analysis. The division is not so clear; the beginning of the second phase is concerned with fixing the problem statement for which an algorithm should be designed, and could be seen as part of the first phase too.

The subtask of studying pairs of criteria is mainly intended to give a better understanding of what a criterion really implies. It is easy enough to list a criterion without realizing its implications. The subtask of finding several formalizations of one criterion is important to see that there is not just one possibility, and it is important to think carefully which choice is probably best. It can be that one formalization results in high-quality maps, whereas another formalization for the same criterion sometimes gives mediocre maps.

It is necessary with relatively open assignments like these to have several meetings with the students. Several pairs of students were not quite sure what was expected of them. They, for instance, needed an example of one formalization before they could make formalizations themselves for the other criteria. The meetings were also used to give counterexamples to ideas for criteria or formalizations, so that the students could adjust or refine their ideas. Examples of this are given later in this paper. One week before handing in the second report, the students presented their projects to the class, with the criteria chosen and the algorithm developed, but without the analysis of the algorithm.

The other projects handed out were the construction of time-space maps, the placement of labels for groups of islands, urban street plan generalization, the placement of statistical diagrams in administrative regions, and the determination of the relative position of two countries. The project description above is typical and the set-up applies to the other projects as well. First there is a study of maps or situations to become acquainted with the choices of cartographers, and reverse engineering this into constraints and criteria. Second, there is a study of pairs of criteria to discover shared or conflicting 'interests' of the criteria. Thirdly, quantification of the criteria must be done. Fourthly, geometric problem statements had to be given. The second phase is also largely the same for all projects. Only the project on relative position was a bit different, which will become clear in the next section.

3 Results from various projects

In this section we describe for five of the projects the observations, ideas, and results from the GIS class of Spring 2003. Several of the ideas are interesting in their own right, and the projects further illustrate the use of this type of project in a GIS course. The time-space map project is omitted to avoid repetition. In previous years, projects on flow maps and various cartographic generalization operators were also used.

3.1 Non-contiguous area cartograms

The description of the non-contiguous area cartogram project was given in full in the previous section. We give a selection of the ideas and results from that project.

The pair of students who did this project identified six criteria: (i) the displacement of the scaled regions with respect to their original position should be minimized; (ii) the outer shape of the map (joint region) should be preserved; (iii) regions that are close on the original map should be close in the cartogram; (iv) the relative positioning of two regions with respect to each other should be preserved; (v) the scaled region should overlap as much as possible with the original area to keep recognizability; (vi) the overlap between different scaled areas should be minimal.

What seems to be missing is a criterion that controls the global scaling factor. With the six criteria above one could reduce the size of each region with some factor by scaling, which makes all criteria easier to respect. However, it is undesirable to have many small regions on the cartogram. Hence, the missing criterion should attempt to keep the regions large enough. Depending on the quantification of criterion (v), it can be used to keep the scaled regions large.

The criteria above do not mention internal boundaries. From Figure 1 one can see that there are non-contiguous area cartograms with and without internal boundaries. They also exist without internal and external boundaries, which explains the presence of criterion (ii).

It seems clear that criteria (i) and (iii) enhance each other, that is, attempting to respect (i) will automatically lead to respecting (iii), and vice versa. Similarly, (i) enhances (ii) and (iv). The missing criterion mostly conflicts with criterion (vi).

The students chose as the quantification of criterion (i) the distance of the centers of gravity of the region on the original map and the scaled version in the cartogram. To obtain a measure for the whole map, the distances can be summed, or the maximum distance over all pairs of corresponding regions can be chosen to represent the overall quality for the displacement criterion. For criterion (v) the students chose the percentage of the scaled region that overlaps with the original. This quantification gives the same weight to all regions. Choosing the actual area of overlap would tend to only place the large scaled regions on the actual, original position. However, using a percentage will cause all scaled regions to be small, and the missing criterion is necessary. The students seemed to assume that the size of each region for the cartogram was given. Ideally such a choice should only be made later and more explicitly in the geometric problem statement. The quantification of the other criteria is omitted here.

One of the geometric problem statements derived is a the computation of cartograms that minimize the sums of displacements of the centers of gravity (criterion (i)) under the restriction of minimal overlap of two different regions in the cartogram. This problem statement has two separate minimizations that will not give the same result. Therefore, the problem is ill-stated. It would have been better to enforce no overlap at all, or maybe allow each region to have overlap of at most 1% with other regions, and under this restriction, minimize the sum of displacements. A reversed problem statement could be to put a threshold on the maximum allowed displacement for each region in the cartogram, and minimize the overlap under this restriction.

In the second phase the students chose an evolutionary algorithm approach for a solution. The fitness is the total area of overlap of all the regions. The analysis of the algorithm to compute the fitness, based on map overlay [3], is coarse and only $O(n^4 \log n)$ time in the worst case and $O(n^3 \log n)$ time on the typical case is claimed. In reality both bounds are two orders of magnitude lower. Other aspects of the second phase are omitted.

3.2 Placement of labels with groups of islands

Two pairs of students (independently) studied the problem of label placement for groups of islands. The observations and ideas of the two pairs of students who did this project are similar and we discuss them at once.

In the project description, the students were asked to study the classical paper on label placement by Imhof [8], and Chapter 14 of Dent's book [4]. These sources already contain many criteria, albeit for point, line, and area labeling only. Island group labeling is similar to area labeling, but not quite the same. The students were also asked to study maps. The island group labeling problem was limited in the project to the placement of one label with a set of polygons that represent islands. No other map features or labels needed to be taken into account. A minimum of six criteria had to be listed.

One pair of students separated the criteria into two lists, one of hard constraints and one of soft constraints (in the style of Wolff et al. for river labeling [17]). The difference is that hard constraints must be respected, whereas soft constraints are only attempted to be taken into account. The hard constraints included a maximum spacing between the letters of the label, avoidance of twistings of the text, and an obvious association of the islands to the label. The soft constraints included easy readability, text following the shape of the group of islands, representation of the size and importance of the islands by the size of the label, positioning of the center of the label in the center of the group of islands, avoidance of letters on the islands, and preference for horizontally placed text. They also observed that there are three types of label placement with islands: (i) internal to the group and horizontal, (ii) internal and curved, and (iii) along the convex hull of the group and curved.

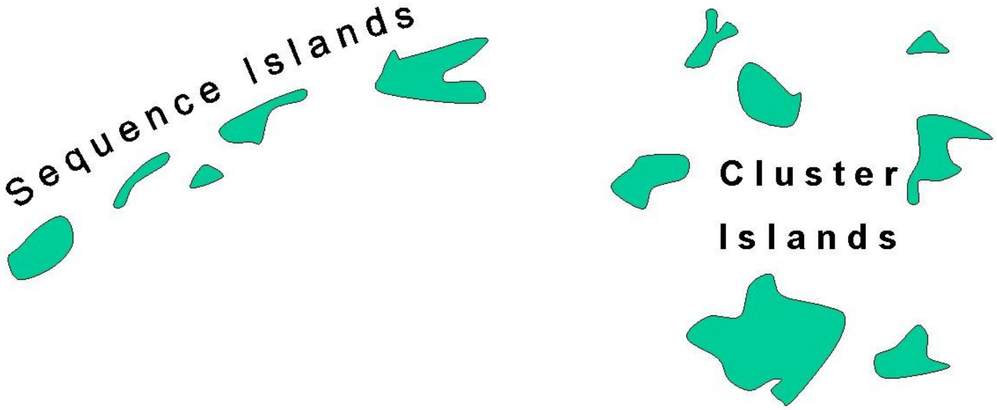

Figure 2: Labeling of groups of islands.

The other two students had somewhat different constraints on the shape of the text, spreading of the label, and association of the text and the islands. They also observed that either the label is between the islands, or, for a compact and small group, the label is to the right and slightly above the islands, similar to a point feature label. This leads to different criteria and different problem statements.

In the quantification, the minimization of the maximum distance from the label to the islands is important. In other words, the distance from the label to the furthest island should be as small as possible. Here the distance is measured between the centers of gravity of the label and the islands. Other choices are possible too.

After analyzing the criteria that enhance or are in conflict with others, the following problem statements were derived. For a given size of the label, assume horizontal placement, and determine the position that minimizes the maximum distance to an island under the condition that no overlap between label and islands occurs. Another derived problem is determining the longest stretch along the convex hull of the islands where the bending is limited to a prespecified value. A label along this stretch, outside the convex hull, and in the middle of this stretch, usually is a good placement according to various criteria listed.

Both pairs of students developed algorithms for the central placement version and used sampling and hill-climbing, or simulated annealing, to obtain a good placement in practical cases.

3.3 Placement of diagrams in regions on a map

Diagrams can be placed on regions of a map to show statistics of that region. For example, one pie chart per province can show political preference of that province, or one bar diagram can show labor in different sectors in that provence. For convenience the regions in which a diagram needs to be placed are referred to as provinces here. The shape of the diagram to be placed can vary, but can be abstracted to be a circle or polygon for a given instance. Reverse engineering shows that cartographers let the diagrams overlap with the province boundaries if needed. Furthermore, the sea and foreign countries are used as "overflow space" for diagrams that do not fit in a province, provided that the province is adjacent to the sea or foreign countries.

The students listed as the main criterion that two diagrams may not overlap. A second criterion is a central placement of the diagram in the province. Thirdly, partial placement of a diagram in a neighboring province is worse than partial placement in the sea or other non-province region (of course depending on the amount). Fourthly, placing diagrams on top of province boundaries must be avoided, especially for the junctions of three provinces. The students assumed that not only the shape, but also the size of the diagrams is given and fixed.

In the quantification of placement on top of province boundaries, the students chose to use relative covering over absolute covering. For a province, the percentage of its boundary length that is covered is used, not the length itself. This is natural, because if one can assure that for every province at least 70% of the boundary is still visible, one can say that every province is still recognizable.

The problem statement chosen by the students included the criteria of no overlap between diagrams, central placement in the province, and preference of overlap with the sea over overlap with other provinces. Their algorithm was based on hill climbing with random starting positions. They restarted the process several times to get several solutions of which the best one could be chosen.

3.4 Urban street plan generalization

The generalization of street plans in urban areas is a difficult problem [2, 9]. This is especially true if the streets do not have a classification by importance, which is common in non-urban regions. Two pairs of students considered versions of urban street plan generalization without initial classification in their projects. Both began by assigning an importance to each street, which is determined primarily by its length. One pair extended this by making streets more important if they connect two important streets (a bridge connecting two major streets is rather important). The only geometric criterion they included is that two streets may not be too close to each other, unless they have an intersection.

The other two students took a more cartographic approach by trying to maintain relative crowdedness or clustering of the subregions of the map. They also included the concept of incremental selection/elimination, as in [14, 15]. The idea is to guarantee that when zooming out (or in), only a subselection (extension) of the previous selection is made. Note that the early settlement selection algorithms (like in [10]) did not have this property [16]. Another geometric criterion listed is a minimum size for built up blocks, and coalescence for streets.

The quantification of crowdedness is the most tricky one of the quantifications. A practical solution, chosen by the students, is to partition the map space into a square grid of suitable cell size, and for every cell, considering the percentage of the area covered by streets. On a smaller map scale, the square cells are smaller, but streets keep their width, so a selection must be made to assure that the area of the cell covered by streets is roughly the same percentage as before. This idea is similar to requiring that the total street length inside the cell should decrease with some constant factor between 0 and 1 that should be the same for every cell. The global problem statement obtained is to preserve relative crowdedness while attempting to select important streets and not allowing coalescence (quantified as non-intersecting or adjoining streets that are too close). This statement is not completely formalized yet.

The algorithm works in three steps. First an importance is assigned to every street by length. Second, from low importance to high importance, a street is checked for visual collisions with other streets, in which case it is deleted. By the order, a street is deleted only if it visually collides with a more important street (even if that other street would also be deleted later). This results in the incremental selection/elimination needed for on-line zooming. Third, the relative crowdedness is considered. The algorithm considers all streets resulting from step two again, from low importance to high importance. For each street they determine whether the relative crowdedness in the cells through which it passes improves or deteriorates if we would eliminate it. If the improvement in some cells is more than the deterioration in other cells, the street is eliminated. The students remarked correctly that the third step does not preserve the incremental selection/elimination principle.

3.5 Relative position of two countries

The problem of determining the relative position between two different regions [11, 12] (which we call countries here) is a bit different from the others because the goal is not cartographic. Instead, the GIS user's intuitive notion of relative position should be captured. The result of the method to be designed is an angle that describes the relative position of one country with respect to the other. This angle can be classified to one of the main compass directions. In some cases no relative positioning is appropriate and in that case the method should detect this. Figure 3 shows three possible situations.

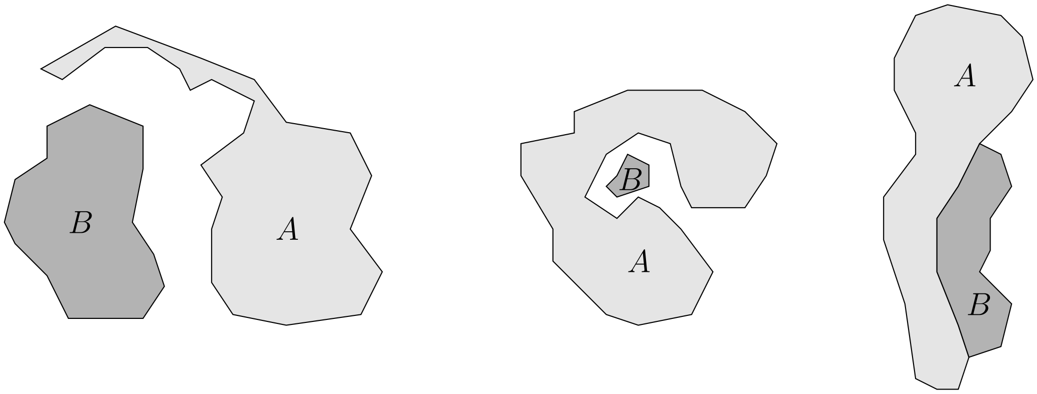

Figure 3: Three situations of two countries and their relative position. Left, B is West of A; middle, A and B do not have a clear relative position; right, B is East of A.

This project was done by two pairs of students. The project description encouraged the students to try many different definitions of relative position, and draw many examples and counterexamples (where the definition does not correspond with intuition). The most simple idea is to take the angle of the vector from the center of gravity of country A to the center of gravity of country B. This method does not indicate when a relative position should not be assigned, and it may also fail intuition in some cases (right in the figure; B would be considered South of A). Another idea is to use the average angle of the two outer common tangents, provided they exist (they don't in the middle figure). It is also possible to combine these ideas and consider two lines through the center of gravity of A and tangent to polygon B.

It appears that when the two countries are not close together, nearly every idea works. Some ideas, like the one with common tangents, are sensitive to long but small features of the shape of a country, whereas others are not (left in the figure; B would be considered Southwest of A).

An interesting approach of one pair of students is based on sampling. Choose a random point $a \in A$ and a random point $b \in B$, and consider the angle of the vector $\vec{ab}$. By sampling many pairs and averaging the angle, we get a reasonable choice of the relative position. The standard deviation of the angles gives a measure of how reliable the answer is. If the standard deviation is too high, then no relative position should be assigned.

The other pair of students extended their initial ideas on center of gravity and tangents by a preprocessing step that removes thin parts of the polygons, to circumvent the incorrectness in the left figure when tangents are used. It is natural to use the erosion, or Minkowski difference, with a disc of some well-chosen radius [7], and then compute the relative position of the smaller polygons A' and B'. In this case the common tangents approach has fewer problems. The students proposed a mixed approach, depending on the situation. Which situation occurs depends on a geometric criterion, which was included in the algorithm.

4 Concluding remarks

Projects are used often as part of a GIS course, but the type of project decribed here is new, to my knowledge. All projects turned out to be interesting research topics that allowed an initial geometric study of the problem, followed by a problem definition and algorithmic solution. In some cases the listing, analysis and quantification of the criteria was the more interesting aspect, in other cases the algorithm and analysis was more interesting.

In the set-up of having to finish the project within eight weeks, it is difficult to finish the whole project in all aspects. It would have been better to have an extra week, and an additional meeting to give suggestions on the algorithm and its analysis. For example, I could have suggested to the students who did island label placement that a combinatorial algorithm can also be used, based on furthest point Voronoi diagrams [1]. In several other cases too, this would have improved the final solution presented, or its analysis.

The assignments handed out, with explicit tasks to be done, written down, and handed in, proved to be important. For projects like these with an outcome that is not fixed beforehand, the students should have clear guidelines of what is expected of them, like the explicit request for listing of criteria, pairing and analysis of pairs, and drawing of sketches which had to be brought to the meetings. Nevertheless, it happened several times that students did not provide a sufficiently complete quantification of the criteria, or a formal statement of the computational problem in the reports handed in. The individual meetings were also essential to lead the projects into an interesting direction, help with the generation of useful ideas, and avoid that students are stuck.

One of the things I felt was missing in the assignment that was handed out is asking for an explicit listing of the "degrees of freedom" of the problem. For example, for island labeling this can be the position of the label, the shape of the label, the size of the label by scaling, and the length of the label by letter spacing. The criteria to be listed can be satisfied by using one or more of the degrees of freedom. After realizing what these degrees of freedom are, and how they affect the desired criteria, the students should be asked to select a subset of the most important degrees of freedom and criteria (to limit complexity of the problem), and assume that the other degrees of freedom are given and fixed.

As noted, the 8-week teaching term does not allow for the whole project trajectory with proper feedback. In Spring 2004 the course will be given again during 11 weeks, and slightly extended in hours (credits) for the students. This means that more attention can be given to the last, algorithm design and analysis part. It would also be interesting to include implementation, but this shifts the emphasis partially from algorithm design to testing, which is already present in another part of the GIVE curriculum.

Acknowledgemens. The author wishes to thank the students of the GIS class of 2003 who did the projects: Leon Bouquiet, Nils Breunisse, Maurizio Camporelli, Siu-Siu Ha, Tjerk de Greef, Bart de Groot, Roderik Heij, David Hörchner, Martijn Lafeber, Reinier van Leuken, Marlies Mooijekind, Jaap Jan Nagel, Michiel Rook, Dirk Smit, Paul Rijneveld, Tim Schlechter, Vincent Tieleman, and Christian Vossers.

References

[1] F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23(3):345–405, September 1991.

[2] J. Bobrich. Cartographic map generalization in urban districts. In *Proc. of GIS Research UK, 9th Annu. Conf.*, pages 513–515, 2001.

[3] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry – Algorithms and Aplications.* Springer-Verlag, Berlin, 1997, 2nd edition 2000.

[4] B.D. Dent. *Cartography - thematic map design.* McGraw-Hill, 5th edition, 1999.

[5] S. Fabrikant. Cartograpohic variations on the presidential election 2000 theme, 2000. `http://www.geog.ucsb.edu/~sara/html/mapping/election/map.html`.

[6] Center for Geometry, Imaging and Virtual Environments (GIVE). Institute of Information and Computing Sciences, Utrecht University. `http://www.give.nl/`.

[7] H.J.A.M. Heijmans. Mathematical morphology as a tool for shape description. In *Shape in Picture - Mathematical Description of Shape in Grey-level Images*, NATO ASI Series F, vol. 126, pages 147–176, Berlin, 1994. Springer-Verlag.

[8] E. Imhof. Positioning names on maps. *Amer. Cartogr.*, 2:128–144, 1975.

[9] B. Jiang and C. Claramunt. A structural approach to model generalization of an urban street network, 2003. 5th AGILE Conference on Geographic Information Science.

[10] G.E. Langran and T.K. Poiker. Integration of name selection and name placement. In *Proc. 2nd Int. Symp. on Spatial Data Handling*, pages 50–64, 1986.

[11] D.J. Peuquet and Z. Ci-Xiang. An algorithm to determine the directional relationship between arbitrarily-shaped polygons in the plane. *Pattern Recognition*, 20(1):65–74, 1987.

[12] S. Shekhar and X. Liu. Direction as a spatial object: A summary of results. In R. Laurini, K. Makki, and N. Pissinou, editors, *ACM-GIS '98, Proc. 6th Int. Sympos. on Advances in Geographic Information Systems*, pages 69–75. ACM, 1998.

[13] NCGIA / USGS. Cartogram Central, 2002. `http://www.ncgia.ucsb.edu/projects/Cartogram_Central/index.html`.

[14] M. van Kreveld, R. van Oostrum, and J. Snoeyink. Efficient settlement selection for interactive display. In *Proc. Auto-Carto 13: ACSM/ASPRS Annual Convention Technical Papers*, pages 287–296, 1997.

[15] P. van Oosterom. The GAP-tree, an approach to 'on-the-fly' map generalization of an area partitioning. In J.-C. Müller, J.-P. Lagrange, and R. Weibel, editors, *GIS and Generalization – Methodology and Practice*, GISDATA. Taylor & Francis, London, 1995.

[16] R. van Oostrum. *Geometric Algorithms for Geographic Information Systems*. PhD thesis, Institute of Information and Computing Sciences, Utrecht University, 1999.

[17] A. Wolff, L. Knipping, M. van Kreveld, T. Strijk, and P.K. Agarwal. A simple and efficient algorithm for high-quality line labeling. In Peter M. Atkinson and David J. Martin, editors, *Innovations in GIS VII: GeoComputation*, chapter 11, pages 147–159. Taylor & Francis, 2000.