MULTIPLE OBSERVER SITING ON TERRAIN WITH INTERVISIBILITY OR LO-RES DATA

W. Randolph Franklin & Christian Vogt
Rensselaer Polytechnic Institute, Troy, New York, USA 12180–3590
http://www.ecse.rpi.edu/Homepages/wrf/

KEY WORDS: Cartography, GIS, Mapping, Land, Vision, Analysis, Resolution, Terrestrial

ABSTRACT

We describe two current projects with our toolkit for siting multiple observers on terrain. (Both observers and targets are at some specified height above ground level. Observers can see targets, when not hidden by the terrain, out to a specified radius of interest.) Siting the observers so that they are intervisible, i.e., so that the visibility graph is a connected set, is the first project. The second project tests the effect, on the optimality of the multiple observer siting (w/o intervisiblity), of reducing the map cell's horizontal or vertical resolution. We lowered the resolution, sited observers optimally, then computed those observers' joint visibility index on the hi-res data. We observed that much less precise vertical resolution is ok, but that reducing the horizontal resolution by even a factor of two leads to an observer siting with significantly reduced joint visibility index, when evaluated on the hi-res data. Applications of multiple observer siting include siting radio towers and mitigating visual nuisances.

1 INTRODUCTION

The results reported here are part of our extended project that might be called *Geospatial Mathematics*, to understand and process terrain data, which means elevations in this context. Previous results have included • a *Triangulated Irregular Network (TIN)* program that can completely and quickly tin a 1201×1201 level-1 USGS DEM, (Franklin, 1973, 2001; Pedrini, 2000), • *Lossy and lossless compression* of gridded elevation databases, (Franklin and Said, 1996), • *Interpolation from contours to an elevation grid*, (Childs, 2003; Gousie and Franklin, 1998, 2003; Gousie, 1998), and • a siting toolkit for *Viewshed and visibility index* determination, (Franklin, 2002; Ray, 1994). Current components of this effort include researching new, compact, terrain representations, such as a "scooping" operator, and approximation from known points with an overdetermined Laplacian PDE. We are also studying operations on terrain, such as lossy compression while maintaining important properties, including gradients and visibility.

For visibility, this project has moved beyond viewshed and visibility indexes to study their applications, such as multiple observer siting, and limitations caused by finite horizontal or vertical resolution. This paper extends our earlier visibility work, including our siting toolkit, described in Franklin (2000, 2004a,b) and Franklin and Ray (1994), which also survey the terrain visibility literature. Notable related research includes the analysis of the effect of terrain errors on the computed viewshed in Nackaerts et al. (1999) and Fisher (1992), and the relation of visibility to topographic features studied in Lee (1992), and the pioneering work of Nagy (1994). Line-of-Sight Technical Working Group (2004) compared various LOS algorithms. Caldwell et al. (2003) computed a complete intervisibility database, the viewshed of every point in a 466×336 DEM. Champion and Lavery (2002) studied line-of-sight on natural terrain defined by an L_1-spline.

Consider a terrain elevation database (map cell), and an observer, $\mathcal{O}$. Define the *viewshed* as the specific terrain visible from $\mathcal{O}$ that lies within some radius of interest, $\mathcal{R}$, of $\mathcal{O}$. The observer might be situated at a certain height, $\mathcal{H}$, above ground level, and might also be looking for targets, $\mathcal{T}$, also at height $\mathcal{H}$ above the local ground. Note that if the observer and target heights are different then visibility is not symmetric.

Since the line of sight from $\mathcal{O}$ to $\mathcal{T}$ generally falls between adjacent elevation posts, some interpolation rule is necessary. Small changes in the interpolation rule can cause large changes in the computed viewshed. That subject still requires research since the correct choice depends on the assumed terrain model. Assuming terrain to be C^∞ (i.e., its derivatives of every order are continuous), is false. Indeed among the most important terrain features are cliffs, which are discontinuous. However, some assumption has to be made.

Define the *visibility index* of $\mathcal{O}$ as the fraction of the points within $\mathcal{R}$ of $\mathcal{O}$ that are visible from $\mathcal{O}$. The *single observer siting problem* is to site (i.e., find the location of) $\mathcal{O}$ so as to maximize its visibility index. The *multiple observer siting problem* is to site a set of observers so as to maximize their joint visibility index, i.e., the area of the union of their individual visibility indexes. We may find either the minimum number of observers to cover a specified area, or else the maximum area covered by a given number of observers. Covering all the terrain is probably impractical because of isolated single points that are lower than all their neighbors, and so are hidden from distant observers.

This multiple observer case is particularly interesting and complex, and has many applications. A cell phone provider wishes to install multiple towers so that at least one tower is visible (in a radio sense) from every place a customer's cellphone might be. Here, the identities of the observers of highest visibility index are of more interest than their exact visibility indices, or than the visibility indices of all observers. One novel future application of siting radio transmitters will occur when the moon is settled. The moon has no ionosphere to reflect signals, and no stable satellite orbits. The choices for long-range communication would seem to include either a lot of fiber optic cable or many relay towers. That solution is the multiple observer visibility problem. Mars also has no useful ionosphere, tho its satellites do have stable orbits. A related application is to site the observers to minimize their visibility, which is appropriate if they are visual nuisances that we wish to hide.

The terrain data structure used here is a matrix of elevations, often a 1201×1201 USGS level-1 Digital Elevation Model cell. The relative advantages and disadvantages of this data structure versus a triangulation are well known, and still debated; the competition improves both alternatives. This current paper utilizes the simplicity of the elevation matrix, which leads to greater speed and small size, which allows larger data sets to be processed.

For distances much smaller than the earth's radius, the terrain elevation array can be corrected for the earth's curvature, as follows. For each target at a distance D from the observer, subtract $\frac{D^2}{2E}$ from its elevation, where E is the earth's radius. (The relative error of this approximation is $\left(\frac{D}{2E}\right)^2$.) It is sufficient to process any cell once, with an observer in the center. The correction need not changed for different observers in the cell, unless a neighboring cell is being adjoined. Therefore, since it can be easily corrected for in a preprocessing step, our visibility determination programs ignores the earth's curvature.

The radius of interest, $\mathcal{R}$, out to which we calculate visibility, has no relation to the distance to the horizon, but is determined by the technology used by the observer. E.g., if the observer is a radio communications transmitter, doubling $\mathcal{R}$ causes the required transmitter power to quadruple. If the observer is a searchlight illuminating diffusely reflecting targets, then its required power is proportional to $\mathcal{R}^4$.

In order to simplify the problem under study enough to make some progress, this work also ignores factors such as vegetation that need to be handled in the real world. We assume that it's possible, and a better strategy, to incorporate them only later.

The ability to process large, hi-res terrain datasets is important. Demonstration programs may be useless for real applications if their time and space requirements limit them to toy datasets. Nevertheless, it is worthwhile to measure how good are the results computed on low-res datasets. For instance, the current state of battery technology limits the computation and communication speeds of portable devices, regardless of other hardware advances. For instance, subject to various caveats, doubling a microprocessor's speed doubles its power consumption.

2 OUR SITING TOOLKIT

This toolkit, whose purpose is to select a set of observers to cover a terrain cell, consists of four core C++ programs, supplemented with zsh shell scripts, Makefiles, and assorted auxiliary programs, all running in linux.

1. VIX calculates approximate visibility indices of every point in a cell. VIX takes several user parameters: $\mathcal{R}$, the radius of interest, $\mathcal{H}$, the observer and target height, and $\mathcal{T}$, a sample size. VIX reads an elevation cell. For each point in the cell in turn, VIX considers that point as an observer, picks $\mathcal{T}$ random targets uniformly and independently distributed within $\mathcal{R}$ of the point, and computes what fraction are visible. That is this point's estimated visibility index. $\mathcal{T} \approx 20$ to 30 appears sufficient.

2. FINDMAX selects a manageable subset of the most visible tentative observers from VIX's output, to be called the *top observers*. This is somewhat subtle since there may be a small region containing all points of very high visibility, such as the center of a lake surrounded by mountains. Since multiple close observers are redundant, we force the tentative observers to be spread out as follows.

   (a) Choose an appropriate value for $\mathcal{L}$, the desired number of top observers. Experimentally, $\mathcal{L} \approx 800$ suffices to cover 80% of the terrain, while a 95% coverage requires $\mathcal{L} \approx 3000$.

   (b) Partition the map cell into about $\mathcal{L}/\mathcal{K}$ equal-sized smaller blocks. Experimentally, $\mathcal{K} \approx 2$ is good.

   (c) In each block, find the $\mathcal{K}$ points of highest approximate visibility index (as determined by VIX). If there are more than $\mathcal{K}$ points with equally high visibility index, then select $\mathcal{K}$ at random (using a multiplicative hash function of the point's coordinates as a secondary sort key), to prevent a bias towards selecting points all on one side of the block. If a block has fewer than $\mathcal{K}$ points, then return all its points.

3. VIEWSHED finds the viewshed of a given observer at height $\mathcal{H}$ out to radius, $\mathcal{R}$. The procedure, based on Franklin and Ray (1994) and Ray (1994), goes as follows.

   (a) Define a square of side $2\mathcal{R}$ centered on the observer.

   (b) Consider each point around the perimeter of the square to be a target in turn.

   (c) Run a sight line out from the observer to each target calculating which points adjacent to the line, along its length, are visible, while remembering that both the observer and target are probably above ground level.

   (d) If the target is outside the cell, because $\mathcal{R}$ is large or the observer is close to the edge, then stop processing the sight line at the edge of the cell.

   One obvious "improvement", when the target is outside the cell, would be to move the target in to the edge of the cell before running the sight line. However, this would cause the computed viewshed to depend slightly on $\mathcal{R}$, which looks bad.

Various nastily subtle implementation details are omitted. The above procedure is an approximation, but so is representing the data as an elevation grid, and this method probably extracts most of the information inherent in the data. There are combinatorial concepts, such as Davenport-Schintzel sequences, i.a., which present asymptotic worst-case theoretical methods.

4. SITE takes a list of viewsheds and finds a quasi-minimal set that covers the terrain cell as thoroughly as possible. The method is a simple greedy algorithm. At each step, the new tentative observer whose viewshed will increase the joint (or cumulative) viewshed by the largest area is included, as follows.

   (a) Let $\mathcal{C}$ be the joint viewshed, or set of points visible by at least one selected observer. Initially, $\mathcal{C}$ is empty.

   (b) Calculate the viewshed, $\mathcal{V}_i$, of each tentative observer $\mathcal{O}_i$.

   (c) Repeat the following until it's not possible to increase $area(\mathcal{C})$, either because all the tentative observers have been included, or (more likely) because none of the unused tentative observers would increase $area(\mathcal{C})$.

      i. For each $\mathcal{O}_i$, calculate $area(\mathcal{C} \cup \mathcal{V}_i)$.

      ii. Select the tentative observer that increases the joint area the most, and update $\mathcal{C}$. Not all the tentative observers need be tested every time, since a tentative observer cannot add more area this time than it would have added last time, had it been selected. Indeed, suppose that the best new observer found so far in this step would add new area $\mathcal{A}$. However we haven't checked all the tentative new observers yet in this loop, so we continue. For each further tentative observer in this execution of the loop, if it would have added less than $\mathcal{A}$ last time, then do not even try it this time.

3 MAINTAINING INTERVISIBILITY

This project sites the observers close enough that each observer can "see" another observer, that is, contains some other observer in its viewshed. Since the observer and target heights are constrained to be equal here, this relation is symmetric. Further, we constrain this *intervisibility graph* to be connected. The algorithm modification is that the greedy algorithm selecting observers selects only observers that are in the joint viewshed of the already selected observers. Intervisibility forces the observers to be closer together; therefore the joint visibility index of any given number of observers is smaller, as plotted in Figure 1 on the following page. Figures 2 and 3 show the joint viewsheds for 60 observers. Circles indicate the radius around each observer. In Figure 3, a line joins each pair of intervisible observers.

The map cell is the USGS Lake Champlain W level 1 DEM, with an elevation range of 1576m. The radius of interest is 100 posts; the observer and target heights are 30m. Franklin (2004a) contains videos showing the joint viewshed growing as observers are added.

4 REDUCING VERTICAL RESOLUTION

Terrain data is available in different resolutions, both horizontal and vertical. Processing higher resolution data should give more accurate results, but at a computational cost. This section examines that tradeoff, on a rough and mountainous 1 arc second National Elevation Data Set (NED) downloaded from the USGS *Seamless Data Distribution System*, with bounds $(41.2822, 42.4899)$, $(-123.8700, -122.6882)$, on the California-Oregon border. The original size of the map was 4256×4349, of which we used the first 1201 rows and columns. The test scene, rendered with Povray, is shown in Figure 4. We rounded the floating point data to a 0.1 meter vertical resolution and converted it to integer. We reduced the resolution as needed by dividing each elevation by the reduction factor, rounding, and multiplying by the reduction factor. In contrast to section 3 above, in these experiments, we do not require that the observers be intervisible.

We tested the following combinations of $(\mathcal{R}, \mathcal{H})$: (80,10), (100,5), (100,10), (100,30), (100,50), (300,10), (500,50). We reduced the vertical resolution by the following factors: 1, 5, 10, 25, 50, 100, 250, and 500, resulting in resolutions ranging from 0.1m to 50m. We studied the joint visibility index resulting from siting 100 observers, and how much worse was the result from using lo-res data, repeating each experiment 5 times as follows.

1. Site observers on the hi-res data, computing their joint visibility index, v_h. Call this set of observers $\mathcal{O}_h$.

2. Site observers on the lo-res data, computing their joint visibility index, v_l. Call this set of observers $\mathcal{O}_l$.

3. Finally, transfer $\mathcal{O}_l$ back to the hi-res data and compute their joint visibility index, v_t, on the hi-res data, to see how much smaller v_t is than v_h. That difference is the effect of lowering the vertical resolution.

Figure 5 samples our observations, showing the joint visibility index of 100 observers sited for $\mathcal{R}$ =100, $\mathcal{H}$ =5. The *hi res* line plots v_h, which would be constant except for the Monte Carlo algorithm for estimating visibility indices. The *lo res* line plots v_l as the vertical resolution is lowered. The *transferred res* line graphs v_t, also as the vertical resolution is lowered. Two conclusions follow.

1. The difference between the lo and transferred lines shows how the *accuracy of the joint visibility index computation* is affected by using lo-res data. For a 10m resolution, there is no difference. Even for a very poor 50m resolution, the difference is only a few percent.

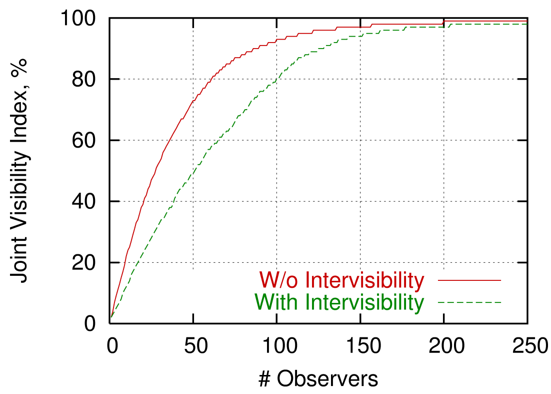

Figure 1: Joint Visibility Indexes, With and W/o Requiring Intervisibility, as Observers are Added

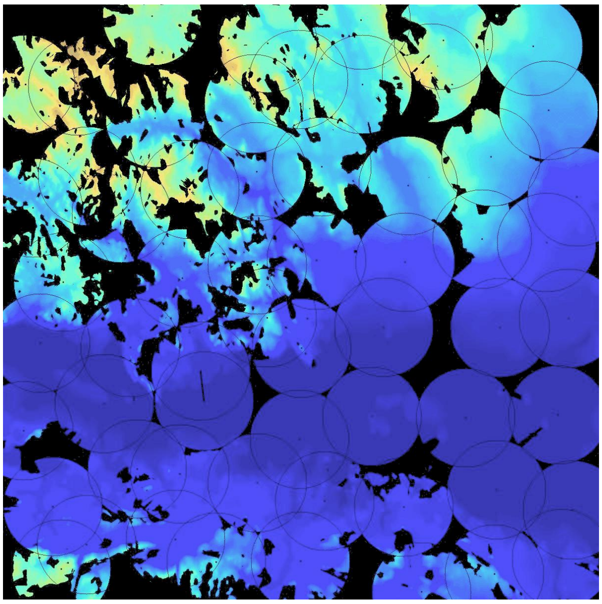

Figure 2: Joint Viewshed of 60 Observers on Lake Champlain W, Without Intervisibility

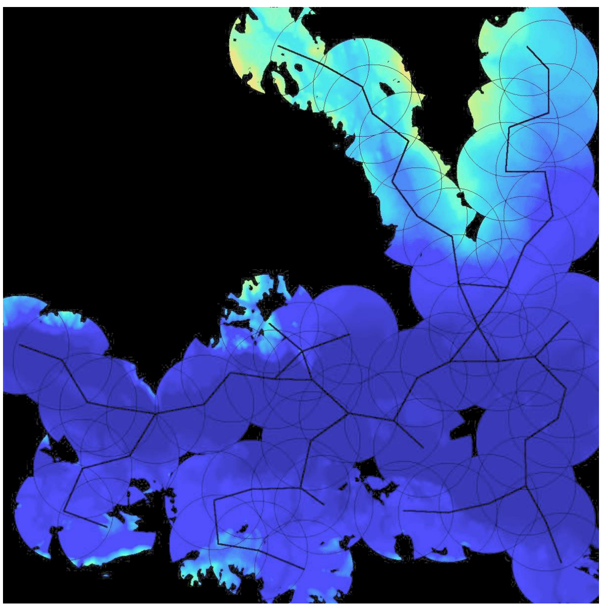

Figure 3: Joint Viewshed of 60 Observers on Lake Champlain W, With Intervisibility

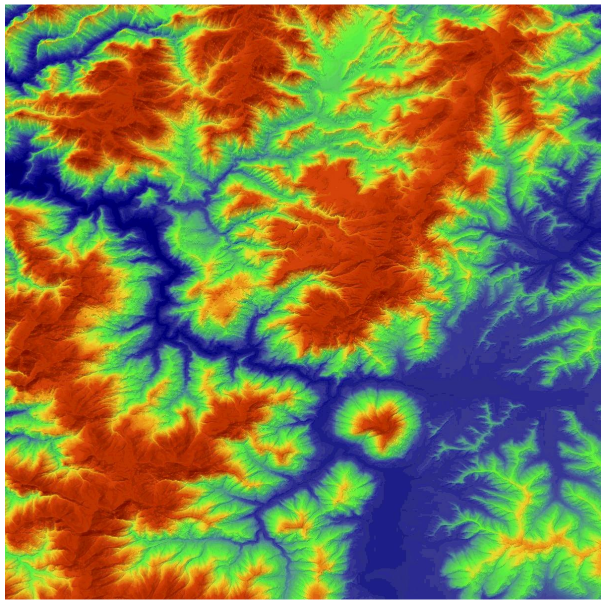

Figure 4: Cell Used to Test Varying Horizontal and Vertical Resolutions

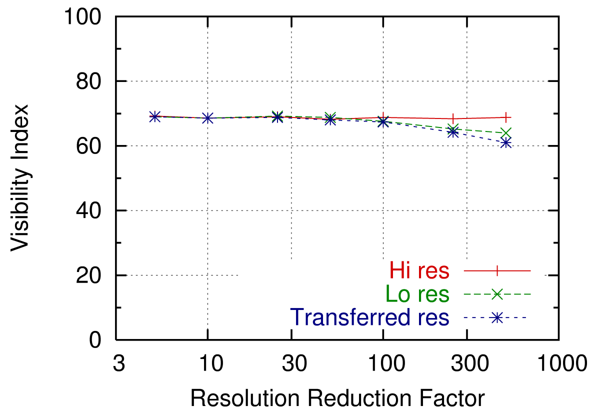

Figure 5: Effect of Reducing Vertical Resolution

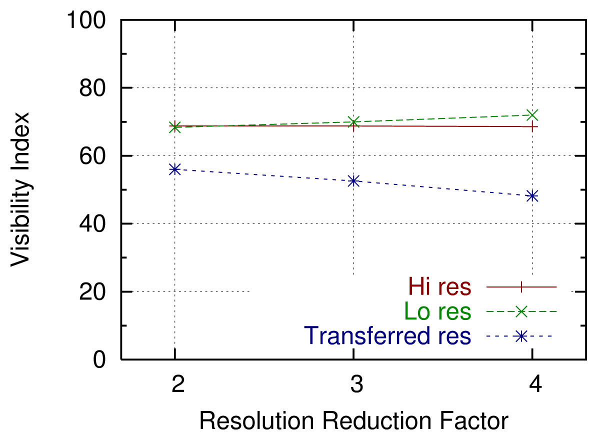

Figure 6: Effect of Reducing Horizontal Resolution

2. The difference between the hi and transferred lines shows how the *quality of the observer siting* is affected by using lo-res data. Even for a very poor 50m resolution, the difference is only ten percent; siting with lo-res elevations is poorer, as we would expect.

5 REDUCING HORIZONTAL RESOLUTION

Next, we tested the effect of reducing the map's horizontal resolution, using bilinear interpolation in Matlab, from 1201 rows and columns to 600, 400, and 300. We tested these combinations of ($\mathcal{R}$,$\mathcal{H}$): (80,10), (100,5), (100,10), (100,30), (100,50), (300,10), (500,50). When transferring each observer sited on the lo-res map back to the hi-res map, if possible, we placed it in the the lower right corner of the window of possible observers.

Figure 6 samples our observations, showing the joint visibility index of 100 observers sited for $\mathcal{R}$ =100, $\mathcal{H}$ =5. (In this and the reduced vertical resolution tests, we set FINDMAX to return 1000 top observers, and set the block size so each block had 2 top observers). For this experiment, the 100 observers jointly could see 70% of the map cell. As the cell was gradually scaled down from 1201×1201 to 300×300, the computed joint visibility index changed little, and often rose slightly. Perhaps the lo-res data has fewer small hidden dips. However, when the observers sited on the 300×300 cell were tested on the 1201×1201 data, a surprising phenomenon became apparent. The observers' joint visibility index fell from 70% to 50%. Even the 600×600 computation was 55% compared to 70%. That is, how much a set of observers can see depends strongly on the resolution at which that their siting is computed. Even a factor of 2 reduction in linear resolution is serious. This shows two different things:

1. Computing viewsheds with lo horizontal resolution data is inaccurate.
2. Effective observer siting requires hi-res data.

We are now considering whether slightly perturbing the lo-res observers' locations when they are transferred back to the hi-res map might increase their joint visibility index.

6 CONCLUSIONS

We have demonstrated multiple observer siting with intervisibility, and experimented on multiple observer siting w/o intervisibility while reducing the vertical and horizontal resolution of the map cell.

We observed that even considerably reducing vertical resolution (from 0.1m to 10m) does not worsen multiple observer siting. If this observation generalizes to other datasets, then expensive efforts to maximize vertical resolution are not justified, at least for observer siting.

However, reducing the horizontal resolution had the opposite effect. Siting observers on a cell with even a factor of 2 lower resolution produced a noticeably poorer joint visibility index, when measured on the hi-res data. That is, if our observations generalize, visibility index computations and observer sitings must be computed on map cells of the highest horizontal resolution possible.

Finally, the experiments that we report here, and many other unpublished tests, are possible only because of our very efficient (in both time and space) siting toolkit, which has easily handled map cells with up to 2000×2000 elevation posts.

7 THE FUTURE

These experiments, and other experiments with our toolkit, demonstrate the value of moving beyond mere viewshed computation to multiple observer siting. Indeed often large errors in viewshed computations do not significantly affect the siting. This illuminates a great opportunity: how far can this idea be pushed, to create faster, yet just as good, applications of visibility?

Assorted small extensions are also possible, such as computation of the joint viewshed for observers traveling along specified routes, multiobserver siting so that each target is covered by at least K observers, for $K > 1$, and trajectory planning of an observer to minimize or maximize the total viewshed.

Another area for investigation is the connectivity of either the viewshed, or its complement. Indeed, it may be sufficient for us to divide the cell into many separated small hidden regions, which could be identified using the fast connected component program described in Nagy et al. (2001).

There is also the perennial question of how much information content there is in the output, since the input dataset is imprecise, and is sampled only at certain points. A most useful, but quite difficult, problem is to determine what, if anything, we know with certainty about the viewsheds and observers for some cell. For example, given a set of observers, are there some regions in the cell that we know are definitely visible, or definitely hidden?

Finally, the proper theoretical approach to this problem would start with a formal model of random terrain, which is usually formed by running water. E.g., local maxima are common but local minima almost nonexistent. Then we could at least start to ask questions about the number of observers theoretically needed, as a function of the parameters. Until that happens, continued experiments will be needed.

REFERENCES

Caldwell, D. R., Mineter, M. J., Dowers, S. and Gittings, B. M., 2003. Analysis and visualization of visibility surfaces. In: Proceedings of the 7th International Conference on GeoComputation, http://www.geocomputation.org/2003/, University of Southampton, UK.

Champion, D. C. and Lavery, J. E., 2002. Line of sight in natural terrain determined by L_1-spline and conventional methods. In: 23rd Army Science Conference, http://www.asc2002.com/summaries/o/OP-09.pdf, Orlando, Florida.

Childs, J., 2003. Development of a two-level iterative computational method for solution of the Franklin approximation algorithm for the interpolation of large contour line data sets. Master's thesis, Rensselaer Polytechnic Institute. http://www.rh.edu/~chil7317/Final_Report.pdf. See also http://www.terrainmap.com.

Fisher, P. F., 1992. 1st experiments in viewshed uncertainty — simulating fuzzy viewsheds. Photogrammetric Engineering and Remote Sensing 58(3), pp. 345–352.

Franklin, W. R., 1973. Triangulated irregular network program. ftp://ftp.cs.rpi.edu/pub/franklin/tin73.tar.gz.

Franklin, W. R., 2000. Applications of analytical cartography. Cartography and Geographic Information Systems 27(3), pp. 225–237.

Franklin, W. R., 2001. Triangulated irregular network computation. http://www.ecse.rpi.edu/Homepages/wrf/sw.html#tin.

Franklin, W. R., 2002. Siting observers on terrain. In: D. Richardson and P. van Oosterom (eds), Advances in Spatial Data Handling: 10th International Symposium on Spatial Data Handling, Springer-Verlag, pp. 109–120.

Franklin, W. R., 2004a. Automated observer siting on terrain, showing intervisibility. http://www.ecse.rpi.edu/Homepages/wrf/research/siting-jan2004/.

Franklin, W. R., 2004b. Efficient observer siting on large terrain cells. (submitted).

Franklin, W. R. and Ray, C., 1994. Higher isn't necessarily better: Visibility algorithms and experiments. In: T. C. Waugh and R. G. Healey (eds), Advances in GIS Research: Sixth International Symposium on Spatial Data Handling, Taylor & Francis, Edinburgh, pp. 751–770.

Franklin, W. R. and Said, A., 1996. Lossy compression of elevation data. In: Seventh International Symposium on Spatial Data Handling, Delft.

Gousie, M. and Franklin, W. R., 1998. Converting elevation contours to a grid. In: Eighth International Symposium on Spatial Data Handling, Dept of Geography, Simon Fraser University, Burnaby, BC, Canada, Vancouver BC Canada, pp. 647–656.

Gousie, M. and Franklin, W. R., 2003. Constructing a DEM from grid-based data by computing intermediate contours. In: E. Hoel and P. Rigaux (eds), GIS 2003: Proceedings of the Eleventh ACM International Symposium on Advances in Geographic Information Systems, New Orleans, pp. 71–77.

Gousie, M. B., 1998. Contours to digital elevation models: grid-based surface reconstruction methods. PhD thesis, Electrical, Computer, and Systems Engineering Dept., Rensselaer Polytechnic Institute.

Lee, J., 1992. Visibility dominance and topographic features on digital elevation models. In: P. Bresnahan, E. Corwin and D. Cowen (eds), Proceedings 5th International Symposium on Spatial Data Handling, Vol. 2, International Geographical Union, Commission on GIS, Humanities and Social Sciences Computing Lab, U. South Carolina, Columbia, South Carolina, USA, pp. 622–631.

Line-of-Sight Technical Working Group, 2004. Line-of-sight (LOS) compendium. http://www.tec.army.mil/operations/programs/LOS/LOS%20Compendium.doc.

Nackaerts, K., Govers, G. and Orshoven, J. V., 1999. Accuracy assessment of probabilistic visibilities. International Journal of Geographical Information Science 13(7), pp. 709–721.

Nagy, G., 1994. Terrain visibility. Comput. & Graphics.

Nagy, G., Zhang, T., Franklin, W., Landis, E., Nagy, E. and Keane, D., 2001. Volume and surface area distributions of cracks in concrete. In: C. Arcelli, L. Cordella and G. S. di Baja (eds), Visual Form 2001: 4th International Workshop on Visual Form IWVF4, Lecture Notes in Computer Science, Vol. 2051/2001, Springer-Verlag Heidelberg, Capri, Italy.

Pedrini, H., 2000. An Adaptive Method for Terrain Approximation based on Triangular Meshes. PhD thesis, Rensselaer Polytechnic Institute, Electrical, Computer, and Systems Engineering Dept.

Ray, C. K., 1994. Representing Visibility for Siting Problems. PhD thesis, Rensselaer Polytechnic Institute.

ACKNOWLEDGEMENTS

This paper was supported by National Science Foundation grants CCF 03-06502 and CNS 03-23324.

May 11, 2004, 15:51