

Issues in Iterative TIN Generation to Support Large Scale Simulations

Michael F. Polis*
David M. McKeown, Jr.
Digital Mapping Laboratory
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

Large scale distributed simulation has been used to support training for battalion level operations within the DoD/Army's SIMNET facility. One of the key components for realistic combat training, through the use of hundreds of low-cost, high realism simulators, all linked using a variety of telecommunications technologies, is the fidelity of the underlying shared environment. For ground-based simulations the environment includes the terrain, road networks, buildings, and natural features such as drainage, forests, and surface vegetation. Given the limitations in graphics rendering capabilities for the low-cost computer image generation systems associated with SIMNET, one must trade fidelity in terrain representation for feasibility of simulation.

Our work has focused on reducing the number of polygons that must be rendered in a real time simulation while improving the overall fidelity of the terrain visualization. In this paper we present a new method for the generation of a terrain representation beginning with a digital elevation model (DEM) and producing a triangular irregular network (TIN) suitable for scene generation. We compare our technique to two commonly used methods, VIP and LATTICETIN and provide a performance analysis in terms of residual error in the TIN, and in the number of points required by each method to achieve the same level of fidelity. We also briefly outline the use of selective fidelity to improve the level of detail in selected sub-areas of the DEM.

While motivated by the constraints of real-time distributed simulation, this research has applications in many areas where an irregular network is a more computationally efficient representation than a dense digital elevation model.

INTRODUCTION

The Triangular Irregular Network (TIN) has distinct advantages over the more traditional Digital Elevation Model (DEM) grid for many applications such as visualization and simulation. The irregular point placement of a TIN permits the resolution of the model to adapt to terrain complexity. Many computations are simplified by the polygonal structure, such as drainage modeling and computer image rendering.

We first discuss previous approaches to TIN generation. Then we present a new method for generating a TIN from a DEM. The method is iterative, with the initial TIN containing the four corner points of the original DEM. At each step, an approximate DEM is interpolated from the TIN and an error surface is generated. We add correction points by detecting maximal errors in this surface and generate a new triangulation. The TIN is evaluated at the end of each step by a user defined stopping criterion, such as

*This research was primarily sponsored by the U.S. Army Topographic Engineering Center under Contracts DACA72-87-C-0001 and DACA76-91-C-0014 and partially supported by the Avionics Lab, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U. S. Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-90-C-1465, Arpa Order No. 7597. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Topographic Engineering Center, or the Advanced Research Projects Agency, or of the United States Government.

point count or maximum global/RMS error.

Since our algorithm generates a series of approximations to the DEM we can generate a hierarchy of triangulations simply by storing intermediate resolution TINs as they are generated. This hierarchy provides multiple levels of detail useful in visualization and modeling. A large area can be broken into tiles, with successively lower levels of detail for tiles successively more distant from the viewer. This allows the display of a large area in a limited polygon budget.

After describing our algorithm, we compare our method with two commonly used techniques, VIP [2] and LATTICETIN. We then discuss pragmatic issues in the use of our triangulation method for computer simulation and limitations imposed by computer image generation using a large real world terrain database.

PREVIOUS WORK

Interest in the generation and utilization of irregular terrain representations dates back nearly twenty years. The work of Peucker [6, 7] introduced the TIN terminology and outlined the basic TIN construction problem in terms of sampling constraints, as well as manual and automated techniques for triangulation. They also described the importance of topological structure in selecting points for inclusion in the TIN. Over the years a variety of techniques for manual and automatic point selection have been proposed and implemented.

Traditionally, points have been selected from stereo imagery by a skilled operator [1]. Recently, work has been done on automatic TIN generation from other terrain models already rendered in digital form. For example, Chen and Guevara [2] generate a TIN from raster data, and Christensen [3] generates a TIN from contours on a digitized topographic map.

One of the earliest manual point selection methods is the one used in the ADAPT system. This system, mentioned in Peucker [6], uses man-machine interaction to aid the user in producing a TIN which accurately represents the original terrain. The system goes through an iteration process in which it indicates inadequate parts of the TIN, which the human operator then improves.

Fowler and Little [5] select points in two stages. First, DEM points are automatically classified as peaks, pits, passes, etc. "Surface specific" points are selected based on this classification. A three-dimensional extension of a line simplification algorithm is applied to the points of ridge and valley lines. Additional support points are then selected adaptively until error is within a user defined tolerance. The complexity of implementing this algorithm has prevented its widespread use.

DeFloriani et al. [4] present a hierarchical triangulation method and data structure. With Hierarchical Triangulation (HT) the area is initially divided into triangles. There are many variants of this scheme, but typically, a rectangular area will only be divided into two triangles. Each triangle can then be subdivided by adding a point interior to the triangle and connecting this point to the triangle vertices. Typical criteria for this are: triangles are subdivided only when their maximum error exceeds a tolerance; and the new point is selected by testing all points, and using the one which gives the most improvement. This produces a hierarchy of approximations to the DEM, where later members of the hierarchy include all the points and triangle edges of earlier members. Sharing triangle edges is useful for operations like point location, but it is undesirable in simulation. First, error along shared edges cannot be improved, and these lines will be visible in the simulated terrain. Second, never dividing edges means that the number of sliver triangles is very high in later TINs. Finally, note that DEM boundaries are triangle edges, and can therefore never be improved. All these factors lead to artificial looking terrain.

Chen and Guevara [2] present an automatic point selection procedure called VIP for selecting points directly from the DEM. VIP is essentially a high-pass filter which gives an importance value based on the distance a DEM point is from the 4 lines connecting its diametrically opposed neighbors. This would seem to be a good solution to the point selection problem; all points in the DEM are ordered in terms of importance, which permits rapid selection of a point set of any given size. The procedure does have several problems which stem from the local nature of the importance criterion. First, the peak of

a small, sharp hill will be considered more significant than the peak of one which is large, yet slopes gently. Second, the VIP procedure chooses nearly all the points of valleys and ridge lines. This is desirable in the sense that these are important features to capture. However, this is clearly unnecessary in those places where the ridge or valley follows a straight line, since we can represent any straight line with two points. Both of these problems arise because only the 8 neighboring points of any DEM point are considered. An additional disadvantage is that the filter is undefined at the DEM boundaries. The definition can be extended, or another point selection method can be used for the boundaries, but some special handling is required.

Our algorithm, like HT and Fowler and Little, improves the TIN until it meets user defined criteria. Unlike Fowler and Little, the algorithm is simple, and point addition begins with a simple TIN containing only 4 corner points, although the algorithm could also begin with "surface specific" points. Like HT, we produce a hierarchy of triangulations, but in this hierarchy, only points are shared allowing error to be improved everywhere. Unlike both HT and VIP, the boundary is approximated in the same manner and to the same accuracy as the rest of the DEM.

ITERATIVE TIN GENERATION

We begin with a set of initial points. This always includes the corners of the DEM, and may include other points which the user wants to have included in the final TIN. A Delaunay triangulation [9] of these points provides the initial approximation. The TIN produced can be interpolated at every point of the DEM to produce an approximate DEM. Subtracting the original and approximate DEMs gives an error at every point. We find the maximum absolute value of the error and select points with this error value as candidate correction points. We then add the candidate points to the TIN if they are not within a specified vertical and horizontal tolerance of any point already in the TIN. The values used are 300m for the horizontal tolerance and 20m for the vertical tolerance. This step simply makes more efficient use of the point budget by keeping simultaneous corrections apart. The algorithm functions tolerably without it. Hence the precise values are not critical, but the values given work well for a wide variety of terrain. If all candidate points fail this criterion, then they are only tested against each other, so that at least the first will be added. The points are again triangulated. The process of selecting and adding points can be repeated until some stopping criteria such as RMS error, point count, or maximum error is met.

This TIN generation method is very flexible. The initial point set can include points from a variety of sources. These may be points selected manually by the user, boundary points required to match an adjacent terrain model, or the points selected by another TIN method. This means that our algorithm can be used to improve any existing TIN, such as one generated manually or by selecting peak and ridge points.

The point selection process can be modified to use an importance mask. This gives an importance value for each DEM location, which the error at that point is multiplied by before points are selected. High importance can be assigned to a region to concentrate points there. Individual points or entire regions can be excluded from point selection by setting an importance of zero.

We believe that this method makes very effective use of its points. Every point is added to correct a specific error, and the adjacency criterion applied to the candidate points usually ensures that only one point is used to correct the error.

In comparison to our previous algorithm based on the contouring of the error DEM at each iteration [8], the present one is greatly simplified. The iterative approach is retained, but fewer points are selected each iteration. The increase in the number of iterations requires a faster triangulation. Our previous method used a greedy triangulation which was based on the error measured along a candidate edge rather than its length. By considering the fit to the underlying terrain, superior triangulations were produced. The Delaunay triangulation tries to produce a triangulation with minimum total edge length and equilateral triangles, but does not consider the underlying terrain. It is, on the other hand, very fast, even when triangulating a large number of points, and it can add points incrementally. Its speed makes selecting only a few points per iteration practical, and the improvement in point quality compensates for the loss in triangulation quality, resulting in a better overall algorithm.

As the TIN generation process proceeds, the algorithm generates a hierarchy of triangulations. Any number of these can be saved to provide multiple levels of detail. The hierarchy will share points, but not triangle edges, unlike HT. This is desirable because not sharing edges means that error along edges in a triangulation can be improved. Hierarchies like that of HT also tend to include many long thin triangles which poorly represent the terrain and give it an artificial appearance during visualization.




Figure 1: VIP 2.891% TIN




Figure 2: 7 iteration LATTICETIN




Figure 3: 5000 point Iterative TIN



Figure 4: Relief of Mt. Tiefert DEM



Figure 5: VIP 2.891% TIN



Figure 6: 7 iteration LATTICETIN



Figure 7: 5000 point Iterative TIN



Figure 8: Mt. Tiefert DEM

PERFORMANCE ANALYSIS OF TIN GENERATION METHODS

In this section we compare the performance of our algorithm to two commercially available algorithms. The DEM used in Figure 4 is of Tiefert Mountain, which is located in the California desert in the National Training Center.* The DEM is composed of Defense Mapping Agency DTED Level 2 data having a 30 meter UTM grid spacing, and containing 172,960 points. The TIN methods will be compared by having each method select points from the DEM for two TINs, of about 2500 and 5000 points. The 5000 point TINs are compared visually and all the TINs are compared statistically.

We compared our algorithm to ESRI's tools for TIN generation in ARC/INFO. There are three TIN generation methods included: HIGHLOW, VIP, and LATTICETIN. HIGHLOW selects local minima and maxima. Triangulating these produces a TIN bounded by a convex polygon smaller than the DEM border. Since HIGHLOW does not generate a TIN which covers the whole DEM, its accuracy cannot be properly compared with the other methods, especially since the ability to represent the DEM boundary is important. The VIP implementation deals with the undefined boundary values by giving boundary points a higher importance than any other points. Thus they are included in the TIN regardless

*Tiefert Mountain is the tallest mountain in the 50km by 50km DEM described in Section 6.

of the percentage of points selected. The TIN shown in Figure 1 was generated with VIP using 2.891% of the DEM (5000 points). A lower resolution TIN was produced using 1.4455% of the TIN (2500 points). LATTICETIN starts with regularly spaced boundary and interior points and then iteratively adds additional points, apparently* by adding points inside triangles which do not meet an error tolerance. After running for four iterations, a TIN of 2794 points was produced. Continuing to the seventh iteration produced the 5188 point TIN shown in Figure 2.

Our algorithm produced the TIN shown in Figure 3 having 5000 points selected from the DEM and a 2500 point TIN for statistical comparison. Boundary points were selected in the exactly the same manner as interior points. The difference is apparent when Figures 1 through 3 are compared. By choosing all (or even a large fraction) of the edge points, both VIP and LATTICETIN give an unnatural representation near the boundary. LATTICETIN also has the flaw of beginning with a regular arrangement of points. This reduces its flexibility, and causes it to perform poorly with a small point budget. The VIP TIN has characteristically selected nearly all its points along sharp ridge and valley lines. This has two effects. First of all, the ridges are not represented efficiently, leaving little of the point budget to represent anything else. Second, since much of the low area is smooth, no points are selected there, and the TIN connects ridgelines.

Error Analysis

The accuracy of the TINs are shown in Table 1. Errors are given as elevation differences in meters. The column labeled *equivalent points* shows the number of points the iterative algorithm took to have a lower RMS error than the listed method. VIP is done some injustice here since it is forced to include all the boundary points, but there are only 1668 of these, and even excluding these points, the iterative algorithm outperforms VIP in both cases. The error figures for the 5000 point LATTICETIN are close to our method, but this is deceptive. Error improves slowly with this many points, so the difference in RMS error of less than a meter represents well over a thousand points. Both LATTICETINs trail by nearly the same number of points.** This is evidence of the efficiency of our method.

Performance Results for Mt. Tiefert DEM				
Method	TIN Points	RMS Error	Maximum Error	Equivalent Points
VIP	2500	48.5342	216	90
LATTICETIN	2794	6.25136	81	1332
Iterative	2500	4.08376	28	N/A
VIP	5000	36.2015	143	107
LATTICETIN	5188	2.94568	31	3719
Iterative	5000	2.26767	19	N/A

Table 1: Numerical Accuracy

Comparison of Static Visualization

A more subjective comparison can be made by using each of the three TIN methods to construct a perspective view of the terrain. While dynamic motion in a real-time simulation may tend to blur detail and lessen the effect of these static views, such a comparison is still worthwhile. In Figures 5 through 8 we show the DEM and each of the TINs from an near ground view looking south. VIP has not selected any points in the area between the near ridgeline and Tiefert Mountain, so the triangulator has directly

*The documentation I have seen does not explain the algorithm well.

** 1462 for the 2794 point TIN and 1469 for the 5188 point TIN

connected the two. Both LATTICETIN and our iterative algorithm have approximated the terrain well, although there are some details which are better defined in our TIN. Several minor ridges are visible only in the iterative TIN.

Figures 9 through 12 show a near vertical view looking from the south at Tiefert Mountain and the pass just to its east. Once again, VIP has selected few points in the low-lying areas. Both LATTICETIN and our algorithm capture the terrain well, but our algorithm captures slightly more detail. The ability to capture additional detail can be very valuable in ground vehicle simulation, where micro-terrain provides cover in combat.



Figure 9: VIP 2.891% TIN



Figure 10: 7 iteration LATTICETIN



Figure 11: 5000 point Iterative TIN



Figure 12: Mt. Tiefert DEM

GENERATION OF LARGE-SCALE TINs FOR SIMULATION

A digital elevation model shown in Figure 13 was constructed to support a SIMNET training exercise. The original DEM, provided to us by the U.S. Army Topographic Engineering Center (USATEC), covers an area 50 kilometers on a side (2500 square km) including the National Training Center (NTC), Fort Irwin, California. The range of elevations in the DEM is approximately 1500 meters. This area is primarily desert, with some highly eroded mountainous areas and intricate alluvial fans running to the desert floor. The sheer size of the area presents significant problems. The DEM consists of 1979x1979 points, nearly 4 million elevation posts. To maintain the desired polygon density for the SIMNET computer image generation systems, only 76,500 points (less than 2%) were to be selected for the TIN.




Figure 13: Shaded Relief of NTC 50x50km DEM

SIMNET Terrain Model

The usual SIMNET terrain model is a triangulated 125m DEM. The hardware is capable of using any polygonal terrain model, but a DEM has been used for ease of generation. The terrain and cultural features are divided into 500m square load modules. While the simulator is running, a block of load modules around the observer's position are loaded and displayed. This prevents the gaming area from being limited in size by hardware rendering capabilities. Polygons are not allowed to cross load module boundaries. If they do, they must be clipped at the boundary, which increases the total number of polygons. The 125m spacing of points in the original model was chosen to avoid this problem; load module boundaries fall along triangle edges, so no clipping is necessary. When a TIN is used for the terrain model, clipping is necessary, and given the small load module size, will increase the number of polygons markedly, although some of the additional points can be useful. At the load module corners, where the clipping lines intersect, we know points will be added. These points can be added to the initial TIN, and give a higher quality TIN in the same effective point budget. Another possible improvement to TIN accuracy would be to use elevations interpolated from the DEM for the points which result from clipping. This was not done. Instead, triangles which resulted from clipping were kept coplanar with the original triangle. By making the load module boundaries less apparent, this made the terrain less artificial in appearance.

The point budget is based on a number of factors. We want the same average polygon density as the 125m DEM, 32 per load module, 320,000 in the whole TIN. This bound should apply after clipping, which has a somewhat unpredictable effect. Also, certain

local limits must not be exceeded. Load modules cannot contain more than 100 polygons total; excess polygons must be put into a limited global pool. If there are too many polygons in the field of view of the simulator, it will not be able to display all of them. It was assumed that the latter constraint would be satisfied if no 4km square contained over 4096 triangles. All these factors suggested that the point budget should be chosen conservatively. Each load module in a 125m DEM has 9 interior points, 4 corner points and 12 other edge points. The point budget was 9 points per load module, 90,000 total, plus corner points. This violated the constraints, so the budget was further reduced 15% to 76,500 points. Points around the border of the DEM spaced 125m apart were also added so that the load modules would mesh with surrounding DEM data. The other border points were marked as unavailable by using a mask with their importance set to 0.

Selective Fidelity

An additional complication for the TIN generation process was the desire for reduced fidelity in the mountainous areas to permit increased detail in the areas of alluvial fans and on the desert floor. This was primarily driven by the fact that mountainous areas are not accessible to ground vehicles (simulated or otherwise); yet, due to their height and complexity, they tend to accumulate a large number of TIN points. This decreases the budget available for other areas of the terrain. An overlay indicating the mountainous areas was provided by USATEC, and this provided importance values in the mask. Importance for the low lying areas was set at 8 times that in the mountains. We smoothed the importance grid with a 7x7 Gaussian filter to avoid problems that might result from a discontinuity at the boundary of the mountainous area.

Computational Issues

The NTC test area is also large enough that generating the TIN for the whole area at once is impractical. Since there are local constraints on point density, this makes it doubly useful to divide the area into tiles which will each have a TIN generated separately. The area was divided into 36 tiles. This brings up the issue of consistency along the borders between tiles. A variety of approaches are available; points could be spaced regularly at 125m intervals, or they could be selected by a 1D line approximation algorithm. In both cases, an arbitrary decision as to the number of points along the border would have to be made. Instead, each TIN was generated by starting with the border points of adjacent TINs which already existed, allowing it to select border points naturally along the remaining borders. This permits the generation of separate tiles in parallel as long as adjacent TINs are not generated at the same time. Three DECStation 5000/200s using a common network file system worked in parallel to generate the TINs in 6 hours. The TINs were combined by re-triangulating, which made the tile boundaries invisible.

Final Results

Figure 14 shows a shaded relief terrain representation of the TIN produced by our method. The TIN was generated using selective fidelity in the mountainous areas. Using less than 2% of the original DEM points we were able to construct a TIN which shows an RMS elevation error of under 3 meters when compared to the original DEM. From a qualitative standpoint it appears that the major topographic features are generally preserved and that detail in the alluvial fans and desert floor areas are also quite good. This impression was confirmed using the SIMNET system at USATEC and driving an M1 tank (simulated) through the terrain.

Table 2 presents statistics about the TIN. The DEM actually covers an area slightly larger than 50x50km. TIN points were not selected in this area, and it is ignored in the statistical analysis. The ridges and mountains were represented well enough to allow their use for navigation, while the low lying areas were represented at a high level of detail. In many places, the drainage and alluvial fans are almost indistinguishable from the original.

The model is a significant improvement over the original 125m DEM in terms of function as well as appearance. Scout vehicles usually seek cover in drainage features. Since these features were small, they were smoothed over in the 125m DEM. The TIN was able to preserve these features, making the simulation more realistic.




Figure 14: Shaded relief of NTC TIN

Performance Results for Mt. Tiefert DEM				
Area	Mask Points	TIN Points	RMS Error	Maximum Error
Mounts	1,313,410	28,532	5.26791	52
Plains	2,587,215	64,939	0.86912	24
Overall	3,900,625	93,471	2.99929	52

Table 2: Numerical Accuracy

CONCLUSIONS

In this paper we have described a new iterative TIN generation method and compared its results to two popular methods. Our method has several advantages: it is simple and flexible, it generates a hierarchy of TINs, and it makes efficient use of its point budget. We also demonstrated the ability of our method to construct a large-scale TIN with improved terrain fidelity to support a SIMNET training exercise. Future work will focus on the integration of road and drainage networks compiled from high resolution digital maps into the coarse resolution TIN. We will continue to exploring the use of selective fidelity to preserve fine terrain detail.

ACKNOWLEDGEMENTS

We thank Mr. Doug Caldwell of the Autonomous Technologies Division, Research Institute, U.S. Army Topographic Engineering Center (USATEC), Fort Belvoir, VA for his help in describing the selective fidelity terrain generation problem and in generating the VIP and LATTICETIN examples used in this paper. Ms. Linda Graff of USATEC produced the mountainous area overlay importance grid used for the NTC example. This paper benefitted from careful readings by Jeff Shufelt and Steve Gifford. Our colleagues in the Digital Mapping Laboratory maintained a high level of humor as their workstations were commandeered for the NTC TIN construction experiment.

REFERENCES

- [1] J.R. Carter.
Digital Representations of Topographic Surfaces.
Photogrammetric Engineering and Remote Sensing 54(11):1577-1580,
November, 1988.
- [2] Zi-Tan Chen and J. Armando Guevara.
Systematic Selection of Very Important Points (VIP) from Digital Terrain Model
for Constructing Triangular Irregular Networks.
In N.R. Chrisman (editor), *Proceedings of the Eighth International Symposium on
Computer-Assisted Cartography*, pages 50-54. 1987.
- [3] A. Christensen.
Fitting a Triangulation to Contour Lines.
In N.R. Chrisman (editor), *Proceedings of the Eighth International Symposium on
Computer-Assisted Cartography*, pages 57-67. 1987.
- [4] L. DeFlorani, B. Falcidieno, and C. Pienovi.
Structured Graph Representation of a Hierarchical Triangulation.
Computer Vision, Graphics, and Image Processing 45:215-226, 1989.
- [5] R.J. Fowler and J.J. Little.
Automatic Extraction of Irregular Network Digital Terrain Models.
In *SIGGRAPH*, pages 199-207. 1979.
- [6] T.K. Peucker, R.J. Fowler, J.J. Little, and D.M. Mark.
*Digital Representation of Three Dimensional Surfaces By Triangulated Irregular
Networks (TIN)*.
Technical Report 10, Office of Naval Research, Geography Programs, 1976.
- [7] T.K. Peucker, R.J. Fowler, J.J. Little, and D.M. Mark.
The Triangulated Irregular Network.
In *Proceedings of the ASP Digital Terrain Model Symposium*. 1978.
- [8] M. Polis and D. McKeown.
Iterative TIN Generation from Digital Elevation Models.
In *Proceedings of the DARPA Image Understanding Workshop*, pages 885-897.
Morgan Kaufmann Publishers, Inc., San Mateo, CA, January, 1992.
- [9] F.P. Preparata and M.I. Shamos.
Computational Geometry.
Springer-Verlag, 1985.