PROCEEDINGS OF SPIE

Analyzing sampled terrain volumetrically with regard to error and geologic variation

Thomas Butkiewicz, Remco Chang, Zachary Wartell, William Ribarsky

SPIE.

Analyzing sampled terrain volumetrically with regard to error and geologic variation

Thomas Butkiewicz, Remco Chang, Zachary Wartell, and William Ribarsky

University of North Carolina at Charlotte

ABSTRACT

Most terrain models are created based on a sampling of real-world terrain, and are represented using linearly-interpolated surfaces such as triangulated irregular networks or digital elevation models. The existing methods for the creation of such models and representations of real-world terrain lack a crucial analytical consideration of factors such as the errors introduced during sampling and geological variations between sample points. We present a volumetric representation of real-world terrain in which the volume encapsulates both sampling errors and geological variations and dynamically changes size based on such errors and variations. We define this volume using an octree, and demonstrate that when used within applications such as line-of-sight, the calculations are guaranteed to be within a user-defined confidence level of the real-world terrain.

Keywords: Terrain, Visualization, Analytics, Geology, Sampling, LIDAR

1. INTRODUCTION

Numerous applications utilize models generated from sampling of real-world terrain. Triangulated irregular networks (TIN) created from digital elevation models (DEM) and other traditional techniques for creating these terrain models, often for the purposes of visualization, lack an analytical understanding of the first steps in the process, namely, producing a model from actual terrain. Yet, this understanding is of vital importance for line-of-sight, trafficability, and other terrain analysis applications. In line-of-sight, for example, the confidence measure for the spatial analysis of where, from a given vantage point, objects would be hidden or visible is much more important than the perceptual goodness of the terrain model for visualization. This confidence measure must be derived from an analytical understanding of the relation between the sampled model and the actual terrain. In this paper we combine very efficient techniques from visualization and computer graphics with analysis methods to achieve this understanding. We then apply this visual analytics approach to terrain analysis applications, detailing a line-of-sight application and discussing other applications.

Our approach considers not only the sample points, but the sampling methods used to collect them. The geologic qualities of the land are also taken into account when calculating the variations possible between sampling points.

We calculate a volumetric representation for our terrain model, based on both the sampling techniques and geologic properties of the region. This volume encapsulates the family of all possible surfaces which could have produced the sample data.

We directly utilize this volume for calculations in our applications. In terrain analysis applications, the particular analysis can be applied directly to the bounding volume of the terrain. For example, in line-of-sight, the visibility calculations can be applied directly to the volume, producing a result with a conservative bound on errors and uncertainties introduced during sampling, accounting for possible variations between sample points either locally or globally.

Terrain analysis has the following three general challenges: being able to quickly and accurately analyze large areas of densely sampled terrain; being able to sustain measures of error and uncertainty appropriate to the application throughout the analysis process; and providing an analytical framework for understanding the behavior of the particular terrain models used. This paper addresses the first two challenges and provides a foundation for considering the third.

Computer Science Department, UNC Charlotte, Charlotte, NC
{tjbutkie, rchang, zwartell, ribarksy}@uncc.edu

In this work we have produced the following novel capabilities and significant results:

- A new, general approach for carrying forward bounded error and uncertainty measures throughout an analysis process, regardless of the terrain model.

- A non-uniform, distance-based voxelization with a compact and efficient hierarchical structure

- A fast line-of-sight algorithm that operates within the hierarchical voxelization and without regard to a particular mesh.

2. RELATED WORK

Several areas of related work are relevant to this paper: terrain and mesh rendering in interactive 3D computer graphics, modeling uncertainty in terrain GIS, and methods for implicit surfaces found in computer graphics, computational physics, and related fields.

Multi-resolution terrain and mesh rendering in interactive 3D computer graphics is a mature field [12] with seminal works such as [6][8][11] among many others. Advances in efficient rendering of terrain systems continue to be made [9][21][22][23]. This paper extends work by Garland [6] and Zelinka [5] by adding probabilistic notions to the mesh vertex locations. This concept is not found in other rendering literature

The GIS community continues to investigate approaches to modeling uncertainty in terrain and its effect on common GIS computations such as line-of-sight computations, hydrology simulations, etc. Santos et al [14] present a method for incorporating uncertainty in terrain modeling by expressing elevations as fuzzy numbers and they construct surfaces that incorporate the uncertainty. They generalize some classic interpolators (linear versus splines, etc.) and compare them qualitatively. Given a set of fuzzy sample elevation points $(x_i, \tilde{z}_i)$ where xi is a point on the 2D plane and $\tilde{z}_i$ is a height value represented as a fuzzy number. Santos et al use Matlab to numerically solve for a triple of surfaces that represent an interpolation of the fuzzy points. The upper surface is an upper fuzzy boundary, the lower surface is the lower fuzzy boundary, and the middle surface is the modal surface. This work is closely related to our work and is similar in spirit to simplification envelopes [20]. However, as pointed out by Zelinka permission grids appear to be faster and simpler to compute with than simplification envelopes. Further, Santos et al is limited to height fields whereas permission grids are not. Anile et al [1] propose the use of fuzzy terrain models to incorporate uncertainty and unpredictable variability in the landscape due to factors such as vegetation, presence of unmapped human artifacts, etc. They incorporate this fuzzy terrain model into line of sight computation over DEMs using discrete lines-of-sight. Their work is basically an extension of Cohen-Or and Shaked [3]. Again, our work is more general since it is not limited to height fields.

Finally, our work has some relation to methods of implicit surfaces. This work is found in computer graphics and computational physics. Blinn [2] is perhaps the first to introduce implicit surface methods from computational physics to the computer graphics community. Implicit surface techniques can be used for interpolating and approximating surfaces from polygon soup [15], while specifying an error tolerance within which the implicit surface should lie relative to the original data. The latter general mesh approach along with level set methods [17] can certainly be applied to the domain of terrain visualization and analysis. However, permission-grid-based approaches appear computationally simpler and faster, especially if one desires to carry the errors in the original sample points into a representation for the errors in the interpolated mesh.

3. SOURCES OF ERROR

Terrain data collected through sampling inherently contains errors and uncertainty [4][13]. Traditional methods address this uncertainty by linearly interpolating between the sample points (such as triangle faces in a TIN or regions between topographic contour lines). In this section, we address sampling errors and geologic variations through an analytical understanding of the errors, and incorporate these errors into our model based on their characteristics.

3.1. Sampling errors

Many different methods for collecting terrestrial sample points exist, each with their own characteristic output as well as idiosyncrasies which must be accounted for when generating a terrain model from the samples.

An illustrative example is the use of LIDAR. These systems make distance measurements from aircraft to the ground using a pulsed laser. The vertical component of each sample point, being a measurement of the speed of the beam's return, is highly accurate. However, the position/horizontal component of each sample point is generated from a GPS system which is limited in its accuracy and update rate. To compensate for the update rate, an inertial navigation unit (INU) is used to estimate the aircraft's location between GPS updates, which introduces additional error into the reported positions of samples [4] [13].

The errors introduced from these LIDAR systems become more complex when applying the system to actual terrain. As the slope of the sampled surface increases, the positional error begins to substantially affect the vertical components. Figure 1 illustrates this behavior where positional error affects vertical accuracy based on the terrain's slope. Notice that because the positional data is acquired through the less accurate GPS and INU, while the vertical or "height" information is acquired using an accurate laser, the combined position from these two instruments can result in a change from point A (where the laser actually struck the surface) to point B (due to the error in the GPS and INU).
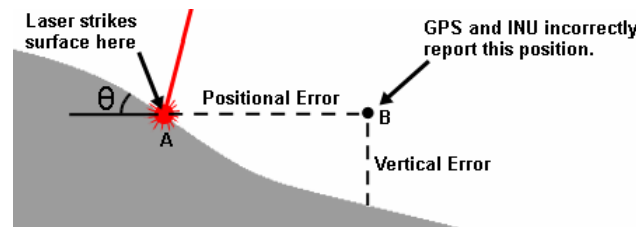
Figure 1. As the slope (θ) of the terrain increases, the error in reported positions from the GPS and INU systems has an increasing effect on the accuracy of vertical measurements. The laser strikes the surface at point A, but because the positioning system introduces a positional error, the sample is recorded as point B.

This effect is minimized when the positional error falls along contour lines (where slope=0), and maximized when across contour lines (where slope is at its maximum). However the displacement of the positional error is assumed to be random and thus we must account for the maximum amount possible. Koppé's formula [13] defines the effect of positional error on vertical error as:

$$Error_{Total\ Vertical} = Error_{Vertical} + (\tan(slope) * Error_{Positional}) \tag{3.1}$$

We can then determine a conservative bound on the total error introduced by the sampling process:

$$Error_{Sampling} = \sqrt{(Error_{Total\ Vertical})^2 + Error_{Positional}^2} \tag{3.2}$$

It is crucial to consider these non-uniform sampling errors for applications that depend on knowing where the actual terrain could reside, such as line-of-sight or trafficability, and where one desires an accurate estimation of the certainty of calculations.

3.2 Geologic variations

Traditionally the areas between sample points are linearly interpolated from the surrounding points (such as the faces in a TIN or the regions between topographic contour lines). This method can incorrectly treat the terrain as a uniform surface constrained only by the sample points. Real terrain behaves erratically, and the data points are only a sampling of a varied terrain. One must instead examine the geologic characteristics of the region and judge to what extent the terrain may diverge from its predicted location based on its distance from a known measurement.

Consider a terrain sampled at 5 meter intervals. If the sampled terrain is a prairie/grassland such as that depicted in Figure 2(a), we would say that the possible vertical difference between two sample points (geologic variation) would be very small, say 0.5 meters. However if the sampled terrain is craggy and prone to unpredictable protrusions, as is shown in Figure 2(b), the geologic variation between two sample points could easily be as much as 4 meters.

(a)

(b)
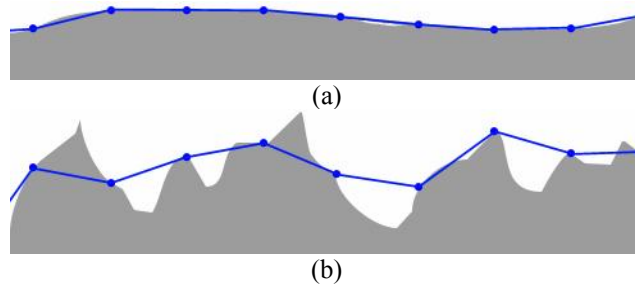
Figure 2. (a) The linearly interpolated surface between the sampled points on the top grassland terrain varies only slightly from the actual terrain, while in (b) the rocky outcroppings between sample points protrude far past the interpolated surface.

These land characteristics might be manually selected from a list of terrain types, entered parametrically for a specific area, or generated automatically from inspecting the variation of the region's height measurements. Analyzing the terrain to automatically generate these geologic properties is beyond the scope of this paper, but we refer to Santos et al [14], which utilizes fuzzy sets to approximate these properties. We choose to implement a distance-based formula representing geologic variation as a Gaussian distribution about each sample point. Doing so allows us to easily associate a confidence level with distances from sample points based on a user-defined value for the terrain's geologic variation (Δ). We define the maximum geologic error as:

$$Error_{Geologic} = (1 - P(d)) * \Delta \tag{3.3}$$

Given a point on the terrain that is at distance (d) from a sample point, $P(d)$ is the probability that this point actually lies on the interpolated surface between the sample points. (See Figure 3). $P(d)$ is further defined as:

$$P(d) = Normalized \quad (\frac{1}{\sqrt{2\pi}} * e^{(f(\Delta)*\frac{-d^2}{2})}) \tag{3.4}$$

Where $f(\Delta)$ is a function relating the maximum rise value for the geologic variation to sharpness of the normal distribution around the sample points. $f(\Delta) = \Delta^2$ provides a simple but stable relationship as outlined in Figure 4. By normalizing the function we ensure that for a distance d, $P(d)$ approaches 1 as d approaches 0, and $P(d)$ approaches 0 as d approaches infinity.

As can be seen in Figure 4, if the terrain is craggy (if Δ is large), the confidence level in the interpolated surface drops exponentially. However, if the terrain is mostly flat, the confidence level decreases more slowly. In all cases, the maximum geologic error approaches Δ as the distance (denoted as d) to the nearest sample point increases.
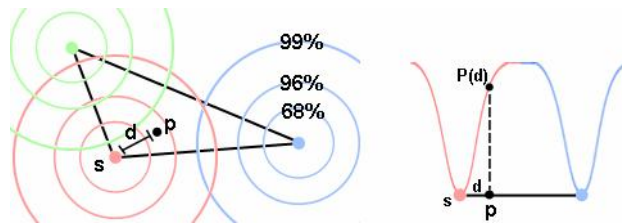

Figure 3. The distance (d) between point (p) and the nearest sample point (s) is used with the normal distribution curve for that sample point to determine the confidence level of the linearly interpolated surface at point (p).
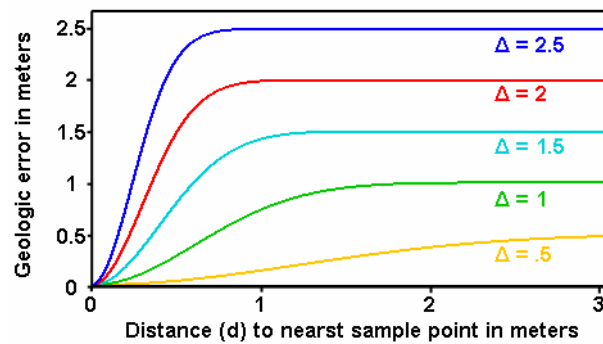
Figure 4. Results of using f(Δ) = Δ² for the Gaussian distribution in Equation 3.4. Larger values for geologic variation (Δ) provide a sharper rise in error as distance from sample points increase as well as a higher maximum error bound when approaching zero confidence in the linearly interpolated surface.

As illustrated in Figure 5, The sampling error from Equation 3.2 and the geologic error from Equation 3.3 can be summed together to find the total error:

$$Error_{Total} = Error_{Sampling} + Error_{Geologic}$$ (3.5)
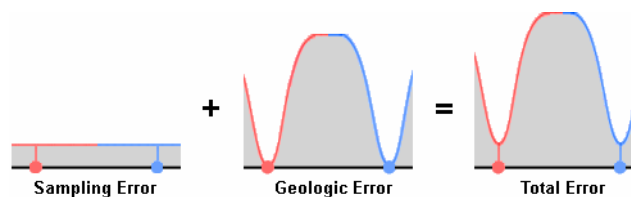
Figure 5. The total error is a combination of the error resulting from the sampling processes and error due to geologic variations. The areas underneath the curves depict the bounded error volume around the surface.

4. IMPLEMENTATION

In order to accommodate both sampling error and geologic variations, we adopt a volumetric representation for the terrain model. The volume bounds the sampling error and geologic variation around the model and represents all possible surfaces of the terrain. The volume is created as a hierarchically discretized voxel grid, similar to the Permission Grids proposed by Zelinka et al [5].

4.1 Sample data

Our input is formatted as a triangulated irregular network (TIN). We use a TIN because of the relative ease of creating it from digital elevation models (DEM) and other sources of terrain data. TINs preserve the areas of sparse sampling as such, while DEMs usually have these areas filled in with interpolated values which, in the final DEM image are not distinguishable from actual sample points. The provided connectivity information is also useful for dealing with irregularly sampled data such as that resulting from LIDAR. We consider the vertices to be the set of sample points and the faces to be a linearly interpolated guess as to the surface in between.

4.2 Grid creation

We model and store our surface conceptually as a simple 3D voxel-based volume slightly larger than the bounding box of the terrain. This approach is based upon the concept of Permission Grids [5], which Zelinka uses to define a volume around a surface to bound error during mesh simplification.

The volume's resolution and size must first be chosen. The two specified parameters that guide its creation are a confidence value (ε) value and a precision value (α). The value ε is the positional error amount for the desired

confidence level. (The ratio between positional error and elevation error is provided during voxelization.) For example if the sampling system gives positions accurate within 1 meter 95% of the time, then ε = 1m would result in a 95% confidence level. The precision value α determines how many voxels are used for that particular ε. (Voxel size = ε / α) Zelinka shows that the minimum effective value for α is $\sqrt{3}$ [5].

The volume, while conceptually a simple 3D array of voxels, is implemented in a hierarchical data structure similar to an octree. During the voxelization phase, the volume recursively subdivides itself until it reaches the necessary voxel size as determined by ε and α. This hierarchical structure allows the creation and use of Permission Grids with minimum storage and memory requirements.

A major benefit of using this discretized volume for our models is the ease of splitting the model into smaller models. This allows calculations on large portions of terrain to be done in parallel, and also facilitates easier storage and level of detail. Additionally, once a volume has been calculated it is easily saved for future reuse.

4.3 Voxelization

In their original form, Permission Grids create a discretized volume guaranteed to be entirely within a constant distance of the mesh. However, by modifying the existing voxelization algorithms for faces and edges we create volumes around the surface using adaptive distances based on the sampling density and geologic qualities. The volume is then distorted in a manner consistent with the sampling errors (see Figure 5).

We voxelize our TIN one face at a time. For each face, we first determine a bounding box of voxels around the triangle. Based on the size of the triangle, we expand the bounding box to be large enough so as to accommodate any possible sampling and geologic errors (see Equation 3.5).

Each voxel in this box is then tested individually to determine if it should be occupied or empty in the final volume. The process of determining each voxel's status begins by first evaluating if the voxel is closest to an edge or the face. If a voxel is within an acceptable distance (as defined by Equation 3.5) from its closest component, it is considered to be occupied. Note that the acceptable distance is adaptively calculated based on the behavior of sampling errors and local geologic properties.
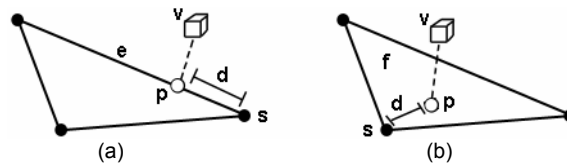

Figure 6. Given a voxel (v), figures (a) and (b) demonstrate the cases where the voxel is closest to an edge or a face respectively. In both cases, the voxel's position is projected onto its closest component (the projected position is denoted as p) and the distance (d) to the nearest sample point (s) is determined.

Possible cases are shown in Figure 6. For the case where a voxel is closest to an edge (see Figure 6(a)), we determine the projected point (p) of the voxel onto that edge. We then find the distance (d) from point (p) to the nearest end (sample) point of the edge. The value (d) is then used in Equation 3.5 to obtain the total error from the interpolated surface at the point (p). If the total error given is less than the length of the vector from the voxel to the edge, then the voxel is considered to be occupied (see Figure 7).

Similarly, for the case where a voxel is determined to be closest to the face itself (see Figure 6(b)), we find the nearest point on the face to that voxel. Once again we find the distance from that point to the nearest sample point and use that value in the same fashion as the edge case.

As illustrated in Figure 7, if the resulting total error is greater than the distance from the voxel to its nearest component, then the voxel is occupied in our model as it must be considered a possible location of the surface.
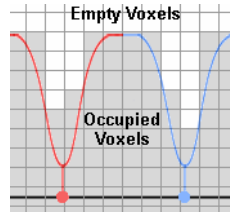
Figure 7. The volume surrounding the surface is discretized and any voxels on or under these error bounds (see Figure 5) around the surface are occupied in our model and considered a portion of the possible terrain.

Figure 8 shows the effect of our volumetric approach given an arbitrary equilateral triangle for a range of values for Δ. The blue lines represent the voxel grids, which in turn represents the bounding volume. Notice the increasing semi-circles emanating from the sample points; this is consistent with our probabilistic model shown in Figure 3. As Δ increases, the thickness of the volume perpendicular to the face increases accordingly. The bounds surrounding the edges can be seen behaving similarly to the lines in Figure 4.

(a)

| Δ | $\Delta+0.5$ | $\Delta+1$ |

(b) $\Delta+1.5$ $\Delta+2$ $\Delta+2.5$
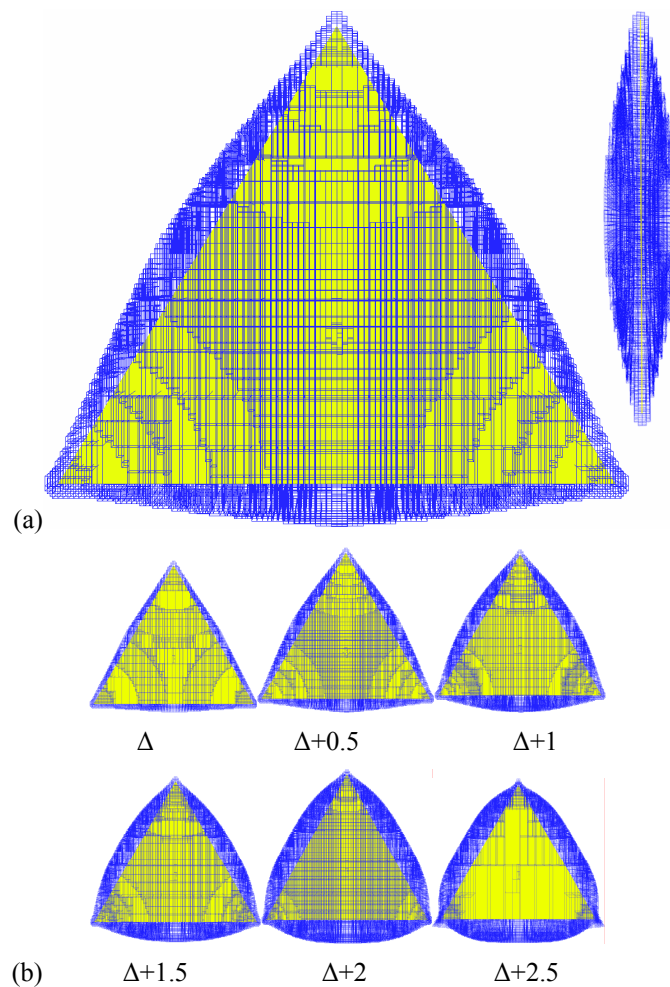
Figure 8. (a) Top down and side view of the volumetric result (shown in blue) of the voxelization process conducted on an arbitrary equilateral triangle. (b) Results for different geologic variation (Δ) values. Notice in all instances the thickness increases farther from sample points.

5. APPLICATIONS

There are many applications in terrain analysis that do not depend on visual goodness, but instead require a more precise, analytical representation of the model. Our volumetric representation of the terrain guarantees that all surfaces within the volume are mathematically bounded to a given confidence level based on the sampling error and geologic variations. Through our algorithm, our line-of-sight application can provide precise bounded probabilities to the status (visible or invisible) of an object from multiple vantage points.

## 5.1 Line-of-sight / visibility

Unlike most traditional line-of-sight or visibility algorithms (see [18] for a survey of some of these algorithms), the goal of generating a volumetric encapsulation model is to place a conservative bound on the possible real-world surface that could have produced the sampling data. This is important for all terrain analyses and, in particular, for applications concerned with visibility calculations. We demonstrate that our algorithm integrates easily and efficiently within a line-of-sight application for determining areas of invisibility to military sentry units. Unlike most existing line-of-sight applications where visibility is considered to be either visible or invisible, our application gives confidence levels to the visibility. Due to the errors and geologic variations in the sampled data, such confidence levels are a more realistic representation.

We conduct our visibility calculations entirely within the hierarchical data structure, without regard to any triangle mesh. We store our volumetric visibility output in a separate, identical structure.

By volumetrically calculating line-of-sight we gain the ability to specify volumetric patrol areas (volumes) instead of exact eye points similar to [19]. The volumetric output allows us to consider the hiding of larger extended units such as heavy artillery. By varying our error tolerance during voxelization, we can produce volumes that provide any desired certainty measure. For large and complex terrain models, the line-of-sight calculations can be done in parallel as a speed-up due to the independent structure of our voxel grid.

First, voxels that represent the view volumes are determined. A standard 3D DDA algorithm [7] is used to traverse the voxels along rays cast from that eye point outward in all directions to the edges of the bounding box. All occupied voxels are considered as occluders, and voxels with an occluder between them and the view voxel are considered to be occluded.

Figures 9 and 10 show the resulting volumes, voxelized in all invisible regions, displayed over a simplified mesh of the terrain. These visibility volumes can then be interactively viewed. Additional eye points (or view volumes) can be added and their visibilities calculated and subtracted from the original invisibility volume. In this way the user can interactively place units to minimize unobserved areas on the map.
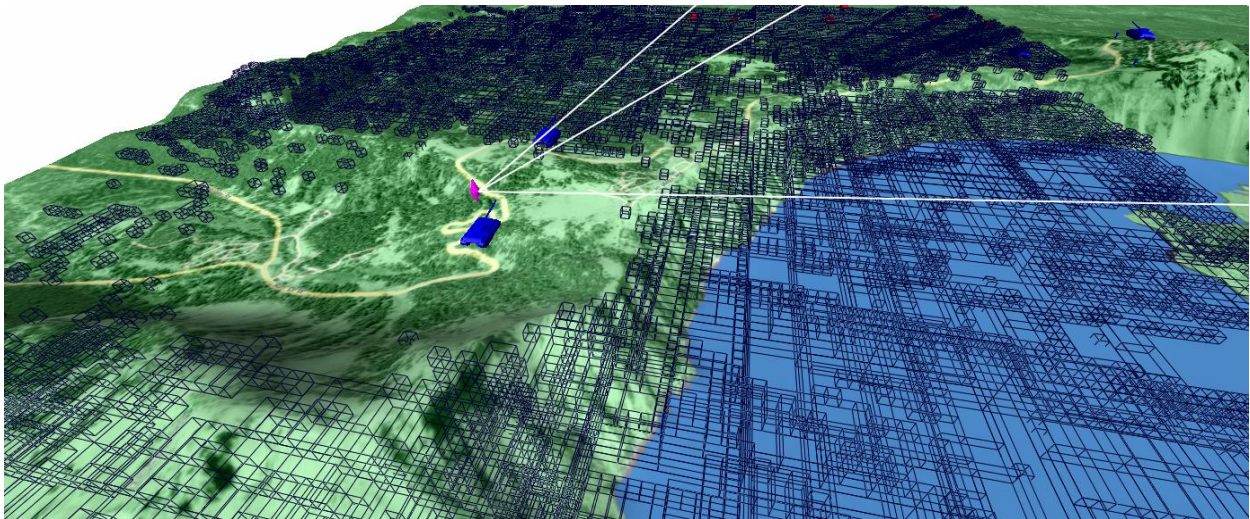
Figure 9. Example visibility output for a region-of-interest showing the voxels found to be invisible from the selected eye-point (the blue tank with the pink marker)
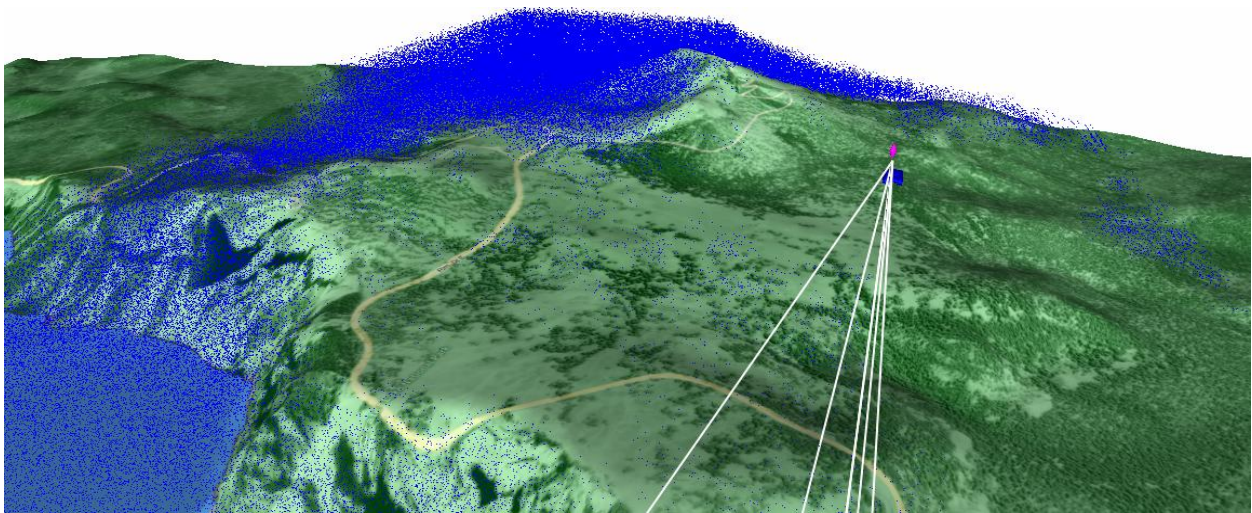
Figure 10. Visibility calculations displayed as a "fog" over the terrain, indicating the volumes unseen by the blue tank with the pink marker. Displaying invisible regions as volumes instead of areas permits assessment of visibility for units that are airborne or of considerable height.

6. FUTURE WORK AND DISCUSSION

Although representing geologic errors as Gaussian distributions provides an intuitive formulation for terrain, we believe that better models can be created to accommodate both terrain and urban models. Urban models differ drastically from terrain in that variations in urban models could be abrupt whereas variations in terrain are more natural and gradual. For example, in sampling urban models, one sample point could fall on the ground, while the next is on the top of a sky scraper. We theorize that formulation for such drastic variations could better be modeled using step-wise functions instead of Gaussian distributions.

For the line-of-sight application, we would like to make use of non-boolean occupancy of the voxels. Such voxels would no longer be outright occluders, but occlude adaptively based on desired confidence levels. Strengthening the line-of-sight algorithm to detect when rays pass through, then below the surface, and emerge once again, could provide output with core volumes guaranteed to be invisible with 100% certainty. Adaptively subdividing the voxels during the visibility tests could also lead to a finer level of detail in the output.

7. CONCLUSION

We combine techniques from visualization and computer graphics with analysis methods to achieve an understanding of the errors involved with sampling a terrain and generating a model from said samples. We apply our approach to terrain analysis in the form of line-of-sight determination.

We demonstrate that by using a volumetric representation for terrain models, we can encapsulate all possible surfaces for the sample points given sampling errors and geologic variations. We further demonstrate the effectiveness of this volume representation when used along with existing techniques as well as new analytical methods such that all analysis is performed entirely within the voxel-based data structure. Using our technique we can determine point-to-point, point-to-area, and area-to-area volumetric visibility on the terrain with an associated confidence level given the sampling errors and geologic variations.

ACKNOWLEDGMENTS

REFERENCES

1. Marcello A. Anile and Primo Furnoa and Giovanni Galloa and Alessandro Massolob. *A fuzzy approach to visibility maps creation over digital terrains.* Fuzzy Sets and Systems. Volume 135, No 1. Pages 63-80. 2003.

2. James F. Blinn. A Generalization of Algebraic Surface Drawing. ACM Trans. Graph. Vol 1: No 3. Pages 235-256. 1982.

3. Daniel Cohen-Or and Amit Shaked. *Visibility and Dead-Zones in Digital Terrain Maps.* Computer Graphics Forum. Vol 14, No 3. Pages 171—180. 1995.

4. M. E. Hodgson, P. Bresnahan, "Accuracy of Airborne Lidar-Derived Elevation: Empirical Assessment and Error Budget", Photogrammetric Engineering and Remote Sensing, p 331-340, Volume 70 Part 3, 2004.

5. Steve Zelinka, Michael Garland, "Permission Grids: Practical, Error-Bounded Simplification", ACM Transactions on Graphics, p 207-229, Volume 21 Part 2, 2002.

6. Michael Garland, Paul S. Heckbert, *Surface Simplification Using Quadric Error Metrics*, Proceedings of SIGGRAPH 1997, p209-216, August 1997.

7. Akira Fujimoto, Takayuki Tanaka, Kansei Iwata, "Arts: Accelerated Ray-Tracing System", IEEE Computer Graphics and Applications, p 16-26, Volume 6 Number 4, April 1986.

8. Hoppe, H. *View-dependent refinement of progressive meshes.* In Computer Graphics (SIGGRAPH '97 Proceedings) (1997), vol. 31, pp. 189-198.

9. F. Losasso, H. Hoppe. *Geometry clipmaps: Terrain rendering using nested regular grids.* ACM SIGGRAPH 2004, 769-776.

10. Koppé, C. *Über die zweeckentsprechende Genauigkeit det Höhendarstellung in topographischer Plänen und Karten für allgemeine technische* Vorarbeiten. Z. VermessWes., p34, 1902.

11. P. Lindstrom, D. Koller, W. Ribarsky, L. F. Hodges, N. Faust, and G. A. Turner. *Real-Time, Continuous Level of Detail Rendering of Height Fields.* Proc. ACM SIGGRAPH 96, pp. 109-118. 1996.

12. David Luebke, Martin Reddy, Jonathan D. Cohen, Amitabh Varshney, Benjamin Watson, Robert Huebner. *Level of Detail for 3D Graphics* (The Morgan Kaufmann Series in Computer Graphics) Morgan Kaufmann; 1st edition (July 22, 2002)

13. D. H. Maling, *Measurements from Maps*, Pergamon Press, New York, N.Y., p154-155, 1989.

14. Jorge Santos and Weldon A. Lodwick and Arnold Neumaier. *A New Approach to Incorporate Uncertainty in Terrain Modeling.* Geographic Information Science: Second International Conference, 2002. Vol 2478. Pages 291-299.

15. Chen Shen and James F. O'Brien and Jonathan R. Shewchuk. *Interpolating and Approximating Implicit Surfaces from Polygon Soup.* Proceedings of ACM SIGGRAPH 2004. Pages 896-904.

16. Suzanne Wechsle. *Effect of DEM Uncertainty on Topographic Parameters, DEM Scale and Terrain Evaluation.* State University of New York. College of Environmental Science and Forestry, Syracuse, New York. 2000.

17. H.K. Zhao, S. Osher, B. Merriman, M. Kang. *Implicit and Non-parametric Shape Reconstruction from Unorganized Points Using Variational Level Set Method*, Computer Vision and Image Understanding. Vol. 80, 2000, pp 295-319.

18. D. Cohen-Or, Y. Chrysanthou, C. Silva, *A survey of visibility for walkthrough applications*, Proc. of EUROGRAPHICS'00, course notes, 2000.

19. Gernot Schaufler, Julie Dorsey, Xavier Decoret, François X. Sillion, *Conservative Volumetric Visibility with Occluder Fusion*, ACM SIGGRAPH, p229—238, 2000

20. J. Cohen, A. Varshney, D. Manocha, G. Turk, H Weber, P. Agarwal, F. Brooks, W. Wright, *Simplification envelopes,* SIGGRAPH Proceedings, p119-128, 1996.

21. Cignoni, P.; Ganovelli, F.; Gobbetti, E.; Marton, F.; Ponchio, F. & Scopigno, R. (2003), *Planet-Sized Batched Dynamic Adaptive Meshes (P-BDAM),* in 'VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)', IEEE Computer Society, Washington, DC, USA, pp. 20.

22. Duchaineau, M.; Wolinsky, M.; Sigeti, D.E.; Miller, M.C.; Aldrich, C. & Mineev-Weinstein, M.B. (1997), *ROAMing terrain: real-time optimally adapting meshes*, in 'VIS '97: Proceedings of the 8th conference on Visualization '97', IEEE Computer Society Press, Los Alamitos, CA, USA, pp. 81--88.

23. Hwa, L.M.; Duchaineau, M.A. & Joy, K.I. (2004), *Adaptive 4-8 Texture Hierarchies*, in 'VIS '04: Proceedings of the conference on Visualization '04', IEEE Computer Society, Washington, DC, USA, pp. 219--226.