Visual Analysis of Urban Terrain Dynamics

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

Department of Computer Science, University of North Carolina at Charlotte

Abstract

This paper presents a framework for a general approach to multisource urban terrain data and models. The approach considers sources of uncertainty and error, how to carry them along, and methods to include them efficiently and effectively in terrain analyses. Multiple levels of detail with strict error bounds and confidence measures can be derived from this approach for both terrain models and terrain analyses. The approach can also be generalized to fully 3D dynamic terrain. We then focus on specific 3D urban models with tens to hundreds of thousands of structures and how to organize them. This leads to ideas about urban morphology and how to effectively organize the dynamics of urban environments.

Introduction

Modern urban areas are places of continuous change. Over periods of time of months, buildings may be torn down and new ones started; streets can be altered and new ones constructed; railways or other means of urban transport may change. The models of urban areas must be able to accommodate these changes. This is especially so since models are significantly higher resolution than previously and cover wider areas. Now models can typically have imagery and elevation data at resolutions of 1 foot to 1 meter with certain features (on key buildings, for

example) that may be at higher resolution. At these resolutions and for certain applications, even small changes can be noteworthy.

In addition, urban terrain models can come from many sources. These include varieties of sensors such as LIDAR, satellite imagery, airborne oblique photography, ground-based depth and appearance fields, SAR, and so on. Automated, semi-automated, or manual techniques are used to reconstruct the urban model. In the latter case, users may use 3D design software to create individual building or streetscape models from combinations of photographs, measurements, and building plans. A comprehensive, dynamic model should be able to handle contributions from any and all of these sources. In many cases, such as urban planning, civil engineering, or military applications, a lower resolution model of an urban area of interest will be augmented with higher resolution data, which may come from a sensing source other than the original data. These data must be embedded into the context of the existing model and often there is not time (nor for certain applications should it be necessary) to reconstruct the whole model based on the new data.

There has been a great deal of work developing interactive visualization and terrain analysis methods for large-scale, high resolution terrain. However, most of these methods treat the terrain model as a 2D surface and in many cases just as a height field. But for current urban terrain applications, there is a significant need to treat the terrain as a 3D model with multivalent heights and non-genus-0 topologies. Overhangs, subways, subterranean rooms and passages, bridges, etc. are all of interest in these models. In addition, the variety of data sources mean that a comprehensive modeling approach must deal with overlapping patches or volumes from

different sensors collected at different times. The modeler must deal with how to use this combination of resources (e.g., does one merge based on an analysis of overlapping patches, choose the most recent or "best" patches, etc.), deal with missing data and error, and keep track of the temporal history of the evolving terrain.

With respect to this last point, certain applications need to treat urban terrains as models with significantly faster dynamics. Military applications for urban combat zones, for example, must consider sudden damage to or complete destruction of buildings, roads, bridges, etc as well as cratering of the terrain surface and collapsing of subterranean structures.

In this paper, we present the initial steps in a comprehensive approach to organizing and using large-scale, high resolution, and dynamic 3D urban terrains. This approach can incorporate terrain data from all sources, including those described above, in the form of meshes, sampled point clouds, depth and appearance images, implicit surface models, volumetric models, and others. Errors, uncertainties, and confidence measures, both for the terrain models and for analyses based on the models, can be propagated in a multiresolution, hierarchical structure. The different types of dynamic effects described above can be handled in an efficient and effective way. A hallmark of this approach is fast update in localized regions where dynamic changes take place without impacting the rest of the representation.

In this paper we also focus on two types of applications: terrain analyses such as line-of-sight, trafficability, and penetrability; and urban planning involving both long term planning and large-scale urban projects. The former application is of prime interest to the military but is also of

interest to engineers and others. The latter application addresses issues such as dynamic zoning and how large-scale projects fit in the overall plan, which are issues of increasing importance to city planners.

Comprehensive Volumetric Approach

Our approach embodies not only the sample points, resulting mesh, etc. that describe the terrain, but also the inherent idiosyncrasies and shortcomings that are characteristic of the various methods used to collect terrain samplings. In addition, the geologic qualities of the terrain itself (such as surface composition, roughness, bogginess, etc.) are also taken into account during the calculation of the final terrain models derived from samplings. For generality, we develop a volumetric representation for the terrain, which embeds the uncertainty/error from both the sampling techniques and the terrain's physical qualities. The final volumetric representation, essentially formed by upper and lower bounds, can then be considered to encompass all of the possible physical surfaces that could have resulted in the original set of samples. It can also be extended to encapsulate different types of samplings from different sources (e.g., different sensors) into a single comprehensive representation. Finally, the representation can be extended to include fully volumetric terrain, as will be discussed further below.

The upper and lower bounds of our volumetric representation, derived from sampling error and geologic variations (for the given terrain type), can be determined for models that are connected or not, regularly or irregularly sampled, or that have multivalued heights or full 3D structure.

Our approach for transforming the sample data into a volumetric representation is based on the

voxelization techniques in Zelinka's Permission Grids (Zelinka and Garland 2002). Zelinka uses this approach to provide a precise upper bound for multiresolution mesh simplifications. We have extended the approach to families of terrain models with their own characteristic errors and geologic variations. In Zelinka's original algorithm, a volume is created entirely within a static specified distance (ε) of the original triangular mesh. We have replaced this static ε with a dynamic, non-uniform distance metric that adapts to each location on the terrain. For each sample point in our data, we apply our sampling error metric to create a volume around each sample point, bounded by the positional, elevation, and other errors for our desired confidence level. In order to compute confidence levels, ε is specified in terms of characteristics of the error distribution. (For example, ε is specified as the 1σ or 2σ distance for a Gaussian distribution.) In general, ε is a vector since the error will be asymmetric. (With LIDAR, for example, the error is significantly larger in the lateral direction than in the vertical direction.) Thus, for greater efficiency, our voxels can be non-cubic.

For all regions between sample points, we combine both the sampling error and the possible geologic error, discussed below, to determine values for ε at all points on the terrain. This allows us to fill in the volumetric model between the available sample points. By understanding the limitations and errors inherent in a sampling technology, we can achieve a volume that reflects a desired confidence level and thus can be traversed and evaluated very efficiently.

We implement the volume in a 3D data structure similar to an octree. Final minimum voxel size can be chosen at runtime and is limited based on available computational and storage resources. The ratios of the minimum voxel size and different values of ε determine the precision of the

volume. By decreasing the minimum size of the voxels we increase the precision of the voxels as approximate fits to the confidence bounds. During the volume creation phase, we recursively subdivide the volume until it reaches, for each local region, the necessary voxel size. After this process completes, the structure then proceeds to recursively remove redundant detail.

This hierarchical structure permits the creation of multiresolution models for different applications with fast access and minimum storage and memory requirements. The model can also be split apart into smaller models collected from different sources and/or processed in parallel.

Sampling Errors

Sampling errors, which vary from point to point, depend on the characteristics of the methods used to acquire data and their effects across different terrain regions. For example, LIDAR, which uses a pulsed laser that is scanned from an aircraft, returns a depth component that is quite accurate. However, the horizontal component, generated using a combination of GPS and inertial navigation updates (INU), can be considerably less accurate. In addition, scattering from corners of buildings and other effects can produce significant degradation in the depth reading. Deviation of the model from the actual terrain is also affected by the nature of the terrain itself. For example, as the slope of the surface increases, errors in the horizontal components significantly affect the accuracy of the vertical component. Figure 1 shows this effect in detail. Our approach was specifically designed to account for these types of errors.

Geologic Variations

In cases where surface models are tessellated, the areas between sample points are linearly interpolated. Such an approximation is more or less accurate depending on the nature of the underlying terrain. If the sampled terrain is a prairie/grassland or a smoothed city terrain (Figure 2(a)), we would say that the possible vertical difference (geologic variation) of the actual terrain from the linear interpolation between sample points would be rather small. However, if the sampled terrain is craggy and prone to unpredictable protrusions and pits, such as shown in Figure 2(b), the geologic variation between points sampled at the same density could potentially differ quite significantly from the interpolation between the sample points.

Terrain cover, such as tall grass or other vegetation, could also make a difference for certain applications. One can obtain the size and nature of these effects from GIS layers giving the terrain type and properties and also from statistical evaluations of terrain variability in the region of the sample points. Our approach takes account of these factors. We show in Figure 3 how total error, which in this case is sampling error plus geologic variation (assumed Gaussian here), is determined. One could have additional errors or more complicated representations.

Volumetric Terrain Analyses

These volumetric terrain models are quite suitable for a number of terrain analyses, including line-of-sight, penetrability, and trafficability.

For line-of-sight (LoS) applications, the integrated error bounding allows certainty measures to be tied to visibility calculations. This is of particular importance for time-critical military LoS

calculations, where terrain data acquisition is likely to have been rushed, perhaps haphazardly, or even partially incomplete. By placing a conservative bound on the possible errors, we can do more than simply report which regions are visible and invisible to units, but also calculate what regions are of questionable visibility and the degree of certainty as to their status.

Because the terrain model is volumetric, the results do not just indicate the visibility (as areas) on the ground (as most other methods do), but instead indicate visibility volumes. This is important when dealing with the hiding or discovery of units/objects of substantial size or those in flight. An example of a situation where this is useful is the calculation for an optimal flight path of an unmanned aerial vehicle over a combat zone that will allow it to produce useful reconnaissance of unobserved areas and possible vehicle movement, whilst avoiding visual detection from known enemy positions.

 Another benefit of a volumetric representation (over 2.5D methods) is that instead of being limited to calculating visibility from single eye-points (i.e. "point-to-point" and "point-to-area"), the user can do "volume-to-volume" calculations. This permits the calculation of the visibility of a unit's entire patrol area, a group of units, or a complex with multiple observation points. Figure 4 shows this capability and Figure 6 shows the underlying multi-resolution data structures.

The multi-resolution/hierarchical nature of our volumetric terrain models permits our applications to dynamically adjust the balance between accuracy and speed of calculations. When calculation time is not an issue, the system can use the highest resolution terrain data for maximal accuracy and confidence in the results. However, in a time-critical situation where the

user desires a result quickly, the applications can lower the resolution of the terrain and calculate orders of magnitude faster. Because each resolution level has an inherent confidence level associated with it, the applications can inform the user just how inaccurate these "rushed results" may be. Conversely, the user can specify a bound on the confidence and the system can adaptively switch to resolutions that provide the desired accuracy while calculation times are kept to a minimum. Because of the ability to maintain local as well as global errors in the terrain model simplifications, the user or automated manager can control where computational effort and thus accuracy is concentrated. A simple but effective method for concentrating computational power and accuracy is to define regions-of-interest, shown in Figure 4 as yellow boxes of highlighted terrain. Once a region-of-interest is defined by the user, the system automatically loads in the highest resolution terrain data it can locate for that particular region and produces a volumetric model of the highest resolution allowable under the current memory/time constraints. These methods are quite effective as shown in Table 1 for the scenario in Figure 5. Further discussion of this work can be found in (Butkiewicz et al 2007).

When dealing with a large-scale, high resolution terrain model, one can achieve a significantly higher data quality to storage ratio by identifying features of the terrain that are of importance for a specific application and storing these areas at a higher resolution than the surrounding terrain of less importance. A good example of this concept is the identification and preservation of ridgelines for terrain models that will be used in line-of-sight or other visibility applications. Because the volumes of visibility over and under horizons are almost always determined by the ridgelines of a terrain, it is imperative they be preserved at the highest resolution possible.

Penetrability and permeability of a terrain model are also important analytical concerns. Here penetrability refers to the actual physical entrance of the terrain (e.g., through digging or an explosion) whereas permeability refers to the ability to see through the terrain (or its foliage) visually or by sensor radiation. Ground cover, primarily vegetation, can be detected by and is somewhat permeable to scanning technologies such as LIDAR, which produces returns both on and in the vegetation/canopy and the ground itself. These returns can be classified as canopy, bare ground, buildings, etc. By treating vegetation cover as a volumetric layer above the ground, and assigning a density to these volumes, we can define the permeability of the vegetation from air to ground.

Various methods of measuring the earth's composition with different radiation bands exist. One such technology is synthetic aperture radar (SAR), currently used to measure terrain structure for purposes including study of geological structures such as volcanoes, active faults, landslides, oil fields, and glaciers. SAR that maps areas of the Earth's surface with resolutions of a few meters can provide information about the nature of the terrain and what is on its surface These data can give insight as to the penetrability of terrain and also for applications such as trafficability or flooding.

This research is of particular interest for military applications. Below ground structures such as fortified military bunkers, fuel delivery lines, and utility infrastructure are considered challenging targets, with "bunker-buster" weapons receiving much attention. By studying and understanding the complex nature of both the penetrability and permeability of the earth and that which covers it (buildings, vegetation, etc) tactical and strategic decisions can be made about

these traditionally difficult targets. We are now extending the 3D volumetric approach presented here to address both surface and 3D terrain models, traditional buildings, and subterranean structures.

Dynamic Terrain

It is most desirable to extend past a static single-sourced terrain model, utilizing multi-data-sources for the creation and maintenance of a comprehensive terrain representation. The terrain and structural elements of large population centers receive frequent sampling and scanning from LIDAR, satellite photography, and other techniques. Databases of building footprints and models are constantly updated for insurance and tax purposes. It is crucial to develop a terrain system that is not only capable of recognizing and integrating as many diverse, overlapping datasets as possible but that also possesses an understanding of both the age of and errors present in each source.

A system that allows fast integration of new data and removal of out-of-date data from the current amalgamated model is necessary for wide scale terrain models where single data sources are insufficient and for terrains that are modified or resampled often. Construction sites for new residential developments can replace wilderness and can then be replaced with the final buildings upon completion. Military commanders need to be able to easily remove destroyed buildings or add craters to their models so that changes can be immediately propagated to ground units. Thus we need to be able to handle terrain dynamics on scales of days to hours, especially for high resolution terrain.

Terrain models that contain temporal history are also of importance for a number of reasons. For example, they can enable city planners and historians to view county-wide areas and visualize changes and development over time. Developments can be evaluated for their environmental and/or aesthetic impact upon the surrounding land. The storage and retrieval process requires spatio-temporal access methods (STAM). The STAM research is found in many disciplines including databases (with sub-specialties in temporal (Zahniolo et al 1997), spatial, spatio-temporal (Bohlen et al 1999) databases); GIS and computerized cartography (Voisard et al 2002); and computer graphics and visualization (Shen et al 1999). Our approach is to develop an event or feature-based structure based on the concept of interval trees (Edelsbrunner 1980). Features could be ridgelines, subterranean structures, the urban morphology features described below, or any structure of importance. Changes in these structures can be followed over time with significant changes (events) noted. The events are then visualized along a timeline, with associated lower level features at any level of detail required. This event structure is typically much smaller than the original data and may in certain cases be orders of magnitude smaller. The interval tree then permits efficient retrieval of full data when needed.

Testbed Dataset

Large-scale, specific, and comprehensive terrain data are hard to obtain. There must be data from several sources, including multiple types of capture and modeling. There must also be 3D data, data subsets that can be added or deleted, and regions where different sources cover the same area (sometimes in contradiction with one another). To insure having data freely available for our research, evaluation, and application development, we have our own multisource, dynamic dataset. We plan to continue adding to this dataset to make a comprehensive source. This dataset

will be used for evaluation, to try out ideas on data organization, terrain analyses, etc. and will be shared freely with collaborators and other interested parties.

Our initial testbed dataset is a collection of data sources covering the entire county of Mecklenburg, NC (where Charlotte is located). (See the excerpt in Figure 7.) Terrain data are from 3 sources: USGS topological surveys in DEM format at 30 meter resolution, a more recent (2003) DEM at 20 foot resolution, and a 20000x20000 (~4 meter resolution) LIDAR return collection. The LIDAR returns, which cover the entire county, are also classified (building, vegetation, bare earth, etc.) While buildings are easily distinguishable in the LIDAR returns, we also possess footprint data for all 370,000+ buildings and structures in the county. By comparing these footprints to the LIDAR returns, we can associate heights with them and recreate simple models for each. Landmark buildings, campus buildings and other structures have been manually modeled by architecture students. To provide 3D test data for subterranean structures, a notional subway system was placed under downtown Charlotte, as well as bunkers and other underground construction, such as utility infrastructure. We have organized the dataset so that different sources can be added or subtracted and, in particular, so that data acquired from real locations can be separated from notional data. This organization will also permit us to study dynamic effects

Representing Complex Urban Models using Urban Legibility

Complex urban models can be composed of hundreds of thousands of buildings laid out on high resolution terrain elevation maps covered with ortho-rectified imagery. The building models can be generated automatically from a combination of footprint and height data (the latter from

LIDAR, for example) with generic or more specific textures. More detailed specific buildings would then be generated with semi-automated methods and embedded in the database. Ultimately these models should also be represented in the volumetric approach described in the last section. However, as we will see, if the high-level applications for these models are different, the multiresolution representations must be different, too.

Although automatically-generated building models are often very simple in geometry due to the fact that they are 2.5D protrusions of footprints, interactive viewing and manipulation of a large number of these buildings can still easily exceed the available memory on most computers and the capabilities of modern graphics cards. For interactive viewing and manipulation of large number of building models, a simplification scheme is essential. Unfortunately, traditional decimation techniques (Garland and Heckbert 1997) (Luebke 2001) do not work well for simplification of urban models (see Figure 8) as the decimation process often creates models that no longer resemble the originals.

To simplify urban models in a meaningful way, we apply the principles of urban legibility. Urban legibility is a concept that has been used for many years in the area of city-planning as well as computer graphics (Dalton 2002) (Ingram and Benford 1995), and is defined as the ease with which parts of the city can be recognized as a coherent pattern and thus contribute to an accurate and usable mental map of the city. In his book The Image of the City, Kevin Lynch (Lynch 1960) further categorizes the elements of urban legibility into 5 groups:

- Paths: streets, walkways, railroads, canals, etc.;
- Edges: boundary elements such as shorelines, walls, edges of developments;

•   Districts: medium to large sections of the city, such as a specific residential district, that have their own existence to the observer;

•   Nodes: strategic spots of intense activity (e.g., Times Square);

•   Landmarks: recognizable structures that are distinct to the observer.

To make large-scale urban models interactive and navigable, we aggregate these large models that may contain hundreds of thousands of buildings using the above principles (see Figure 9). Hierarchical textures are applied to the aggregated models in a manner similar to imposters (Sillion et al 1997)(Maciel and Shirley 1995), and landmarks are treated with special care so that they are maintained throughout the simplification process in order to retain the skyline (Figure 10). The result is a general approach that produces multiresolution models of cities that are legible, understandable, and navigable at all levels of detail (Chang et al 2006). (See Figure 11).

Figure 12 shows an excerpt from one model at different levels of detail, and Figure 13 shows the number of polygons and frame rates using different levels of detail in a flythrough scene. From these two figures, we see that by using urban legibility as the basis of urban model simplifications, we can achieve drastic increase in rendering speed by decreasing large numbers of polygons without sacrificing too much visual realism. Within the context of dynamic terrain, we can now consider the factors below.

Discovering elements of legibility

Using our methods, we can create clusterings and groupings of urban models such that the groupings obey the general rules of urban legibility. However, this doesn't actually identify the exact legibility elements themselves that define a particular urban model. Finding out which of

the five urban legibility elements are creating the logical districts and groupings is important for both theoretical and practical purposes. From a theoretical stand point, finding the elements of legibility can help urban planners quantify the major features of a city. From a practical stand point, knowing which elements are important in a city can assist visitors or inhabitants of a city to more easily generate a mental image of the city (Darken and Sibert 1993).

To find the legibility elements of a city, we overlay the result of the clusters and groupings over an existing urban dataset that contains information on paths, edges, districts, nodes, and landmarks. By using decreasing levels of detail in the urban model, we are able to correctly identify and rank the elements hierarchically from the most visually important to the least. The ranking allows us to identify the main features that define an urban model globally to minor features that shape the model locally.

Given the ranking of the elements of a city, it is now possible for us to define a city's urban form in a quantitative manner. Figure 14 shows the result of applying decreasing levels of detail of the same urban model overlaid on a satellite image of the city. As can be seen, when using a coarse level of detail of the urban model, only the most important features of the city are retained; whereas with the finer levels of detail, more localized legibility elements become visible.

Urban morphology – Comparing two different cities

Given the ranking of the legibility elements in a city, we are now able to represent and describe a city in a quantitative manner. More importantly, this quantitative description of a city allows us to compare and identify the differences between different cities in an analytical fashion.

For example, New York City has a grid-like structure, and is generally defined by its boundaries to the river and the ocean. Washington DC has a more radial structure with "rays" (or roads) emanating from the White House and the Congress. In contrast, a growing city such as Charlotte has a more "sprawling" sense to it as new developments are made without the rigorous global planning such as New York or Washington DC (Figure 15).

Although most urban planners agree that New York City, Washington DC, and Charlotte are very different, they can only describe how these cities are different in a qualitative and subjective way. It has not been possible for the urban planners to identify the exact elements that make these cities different or to communicate these differences analytically. With our capabilities, we are starting to categorize cities based on their legibility elements. It may then be possible to start finding quantifiable elements that distinguish American cities from European cities, or older cities from newer cities.

Urban morphology – How a city changes over time

Cities change over time and the legibility elements that define the urban form of the city change over time as well. It is not difficult to imagine a set of legibility elements that distinguishes an urban model in the past, but eventually become obsolete as the city grows.

In many small European towns, the towns start with a crossroad, a church, and a town hall, and the intersection of the two roads is the defining "node" of the city. However, as the town grows lager and becomes a city, the church, town hall, and the crossroad no longer remain the defining

legibility elements for the city. The city of Charlotte is no different. In 1919, when the city's

population was less than 60,000, the entire city only occupied what is now the downtown area.

At that time, the layout of the city was a relatively structured grid with numbered streets similar

to New York City (Figure 16). However, as the city grew, more roads were built, and what now

defines the urban form of Charlotte is no longer the grid-like structure, but a relatively

unstructured network of roads (Figure 15).

Using the same framework discussed above for quantifying and ranking elements of urban

legibility, we can identify the urban elements from different phases of an urban model over time.

We can imagine finding a legibility element that distinguishes an urban model in the past, but

eventually becomes obsolete as the city grows. Our urban legibility structure will permit us to

effectively and compactly organize these elements over time.

Visualizing Dynamic Urban Environment

As newer buildings are created in a growing city, older buildings are often destroyed to make

space. From a visual and cognitive sense, a single creation or destruction of a building often does

not affect the overall legibility of the urban model. This minor visual change, however, needs to

be reflected within the underlying hierarchy of the legibility model. In other words, buildings in

a city should not be re-clustered simply because one single building is created or destroyed.

Instead, the newly created or destroyed building should try to obey the existing urban form from

a visualization and representation perspective. However, at some point the legibility model does

change enough that re-clustering is necessary. These may also be significant points in time when

the conceptual view of the city changes as well. It could add significant new nodes, for example,

or new landmarks and paths. We have devised methods to identify and account for these changes.

Acknowledgments

References

Bohlen, Michael H., Jensen, Christian S. and Scholl, Michel O., editors. *Spatio-Temporal Database Management: International Workshop Stdbm'99*. Springer-Verlag New York, Inc., 1999.

Butkiewicz, T., Chang, R., Wartell, Z., and Ribarsky, W. "Analyzing Sampled Terrain Volumetrically with Regard to Error and Geologic Variation". Proc. of SPIE Volume 6495, *Visualization and Data Analysis*, 2007

Chang, R., Butkiewicz, T., Ziemkiewicz, C., Wartell, Z., Pollard, N., and Ribarsky, W. 2006. "Hierarchical Simplification of City Models to Maintain Urban Legibility," Charlotte Visualization Center Tech. Rep. CVC UNCC 06-01.

Dalton, R. C., "Is Spatial Intelligibility Critical to the Design of Large-Scale Virtual Environments?" *Journal of Design Computing 4. Special Issue on Designing Virtual Worlds*, 2002.

Darken, R., Sibert, J. "A Toolset for Navigation in Virtual Environments". *Symposium on User Interface Software and Technology*, p157-165, 1993

Edelsbrunner, H. 1980. Dynamic Data Structures for Orthogonal Intersection Queries. Tech. Rep. F59, Inst. Informationsverarb., T.U. Graz, Graz, Austria.

Garland, M., Heckbert, P. "Surface Simplification using Quadric Error Metrics". *Proceedings of SIGGRAPH*, p209-216, 1997.

Ingram, R., Benford, S. "Legibility Enhancement for Information Visualization". *IEEE Conference on Visualization*, 1995.

Lynch, K. 1960. *The Image of the City*. The MIT Press.

Luebke, D. "A Developer's Survey of Polygonal Simplification Algorithms". *IEEE Computer Graphics and Applications*, p24-35, May/June, 2001.

Maciel, P. and Shirley, P. "Visual Navigation of Large Environments using Extended Clusters". *Symposium on Interactive 3D Graphics*, p95-102, 1995.

Shen, H., Chiang, L., and Ma, K. "A fast volume rendering algorithm for timevarying fields using a time-space partitioning (tsp) tree". *IEEE Visualization'99*, pages 371–377, 1999.

Sillion, F., Drettakis, G., Bodelet, B. "Efficient Impostor Manipulation for Real-Time Visualization of Urban Scenery". *Computer Graphics Forum*, p 207-218, Volume 16, 1997.

Voisard, A., Rigaux, P., Scholl, M. *Spatial Databases with Applications to GIS*. Morgan Kaufmann, Inc., San Francisco, California, 2002.

Zahniolo, C., Ceri, S., Faloutsos, C., Snodgrasss, R. T., Subbrahmanian, V.S., and Zicart, R. *Part II Temporal Databases in Advanced Database Systems*. Morgan Kaufmann, Inc., 1997.

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

Figure 1. An example of non-uniform sampling error taken from LIDAR. As the slope (θ) of the terrain increases, the error in horizontal position has increasing effect on the accuracy of vertical measurements.

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

Figure 3. The total error is a combination of the error resulting from the sampling processes and the error due to geologic variations between sample points. The shaded areas underneath the curves depict the bounded error volume that encapsulates the surface.

Figure 4. Example results of both point-to-point and point-to-volume visibility calculations in our LoS application. Volumetric results are depicted as black boxes in a yellow region-of-interest. (Note that the voxel sizes here have been enlarged for illustrative purposes.) Point-to-point visibility from the red objects are shown here as connecting lines-of-sight, but can also be represented as icons above visible or invisible units.

Figure 5. Line-of-sight scenario consisting of two teams of 53 units each across a 20km x 14km terrain. Statistics for this scenario at different levels of accuracy are given in Table 1.

Figure 6. The user's view presents a simplified terrain mesh, while underneath the calculations are performed on the volumetric models. The application can adaptively switch between multiple

resolutions (each with their own confidence measures) maintaining the desired balance between computational speed and accuracy or confidence of the results.

Figure 7. Excerpt from testbed dataset with collection of building models, underlying DEM, and classified LIDAR point cloud.

Figure 8. Comparison of our method (right) to original data (top left) and to a well known simplification method (Qslim (Garland and Heckbert 1997), bottom left). Both methods have the same number of polygons, but Qslim loses all sense of the original model. After texturing, our greatly simplified model (bottom right) retains the features of the original model.

Figure 9. Outline of an urban model after aggregation. The left image shows the original outline of the aggregated models, the right shows the simplified outline. The simplification process preserves the legibility elements in the city model, thereby creating a simplified model that remains legible and understandable.

Figure 10. Maintaining the landmarks and the skyline. The image on the left shows the original model (243,381 polygons), the image in the middle shows a simplified model with landmark preservation (15,826 polygons), and the image on the right shows an unpreserved skyline (13,712 polygons). Note that the quantitative difference between turning on and off landmark preservation is merely 2,114 polygons in model complexity, but the visual difference between the two is very significant.

Figure 11. Clustering buildings in a city. The left image shows clustering results that follow the urban legibility element Paths. The right image shows the result of a more traditional distance based clustering.

Figure 12. Original textured 3D model (left) of Xinxiang, China; simplified model (right) with only 18% of original number of polygons and aggregated textures. View-dependent rendering is applied to the hierarchical multiresolution structure on the right.

Figure 13. Frame rate and polygon counts. Using simplified models with different levels of detail in a flythrough scene (a high ε value denotes large amount of simplification), we see that our simplification method can provide drastic speedup by decreasing the number of polygons rendered to the screen.

Figure 14. Finding elements of urban legibility. To correctly identify the elements of urban legibility and hierarchically categorize them from the most to the least important, we overlay different levels of detail of the urban model over GIS data that contain information on roads, districts, etc. The images from left to right show the urban model in decreasing LoD. The coarsest level of the model retains the most significant elements in the city such as the waterway that cuts through the city diagonally, the park in the center of the city, and the main road that runs east-west through the city; whereas the finest level of detail contain much finer elements that are only significant in a localized way.

Figure 15. Views of three cities: New York City (left), Washington DC (middle), and Charlotte

(right). These three cities have distinctively different layouts; New York City resembles a grid-like structure, Washington DC is radial with roads emanating from the Congress and the White House, and Charlotte is mostly unstructured with strong "sprawling" sense to it.

Figure 16. Map of Charlotte in 1919, when Charlotte was a small city with a clear grid-like set of roads.

Accuracy Level	Max Error Allowed	Voxel Size	Time	Voxels Traversed
Low	150m	85.7m	0.12s	164,519
Medium	100m	50.0m	0.19s	257,649
High	50m	16.6m	0.56s	716,466
Best	30m	10.0m	0.92s	1,172,485

Table 1. Statistics for arbitrary accuracy levels determining balances between model accuracy and computation time. Calculations were done on a terrain of size 20km x 14km generated from a 30 M resolution data source. Establishing visibility information for the two teams (each of 53 units), required 2809 inter-unit visibility calculations.