

Chapter 1.

Introduction

We live in a three-dimensional world. Inevitably, any application that analyzes or visualizes this world relies on three-dimensional data. Inherent characteristics of this data — its multi-dimensionality and typically large volumes — make this an interesting problem. Yet as spatial data play ever greater roles in real-time applications, demands on these data also become greater. The issue is then how to maintain the delicate balance between performance and accuracy, realism and flexibility, in spatial data representations.

This dissertation describes a unique three-dimensional surface model — and techniques for developing, manipulating, and merging this model with other representations— that improves on accuracy of the data without impacting performance. The key to success is twofold. First, accuracy is enhanced by using irregular triangular patches that adaptively conform to the surface. Although numerous triangulation algorithms do consider the significance of points on a surface, my approach is unique in that I also consider critical lines on the surface for the placement of triangle edges. Second, performance of both visualization and analysis operations is enhanced by using the model’s hierarchical structure for data generalization and filtering. I present several novel approaches to spatial

operations that take advantage of the triangulation hierarchy. Of particular significance is a polygonal line sweep algorithm for finding all triangles within a polygonal area.

1.1. Issues

The fact that many systems currently rely on spatial data structures is evidence that a great deal of work has already been done in this area. Yet despite the progress, persistent issues remain to be solved. Of six issues requiring further research, as described by Guenther and Buchmann [GB90], three apply directly to my work here:

- Support of multiple scales and precisions in the data without losing critical information,
- Filtering techniques to improve performance of spatial operations, and
- Merging techniques that allow different data representations to exist separately yet work together, so different representations may be used to their best advantage.

Spatial data representations, taken as a whole, encompass a very wide variety of technologies. This dissertation does not presume to cover all of spatial data structures. Instead, it focuses on three-dimensional surface models for use in visualization and analysis applications where performance is critical. These applications include visual simulators, scientific visualization, virtual reality, and military planning systems. For this reason, the issues presented here — and the remainder of this dissertation — are discussed in this narrower context.

1.1.1. 3D Surface Modeling

Spatial databases for real-time applications must support multiple scales and precisions without losing critical information. High-precision — and hence high-volume — data supports applications in which details are critical. Yet in some cases performance is more important than precision. For those cases, retrieving lower precision — and lower volume — data should be correspondingly faster.

Many of the stringent requirements detailed in a report on real-time display of digital cartographic data [TSDB88] apply equally to other real-time applications. These requirements are: 1) accuracy and adaptability, 2) economy of information, 3) multiple levels of detail, 4) structure facilitating fast search, retrieval, and display, and 5) rapid automated terrain model generation.

Data accuracy is essential to the extent that major features must retain their correct dimensions and characteristic shapes used to identify them. For example, mountain peaks and channel lines must not shift out of view or disappear altogether simply because the critical points don't conform to a regular grid. Distracting and misleading artifacts such as false features (e.g. ridge lines, added by Delaunay triangulation, that cross channel lines) gaps and holes (introduced by improper meshes), and triangles with extremely acute angles are unacceptable. Data structures that adaptively conform to the unique characteristics of a piece of terrain can provide greatest accuracy. Data structures that readily accept the addition of pertinent data in any resolution help to maintain or enhance accuracy.

Economy of information provides models of required accuracy using minimal quantities of data. It relieves data retrieval, display and processing bottlenecks caused by huge volumes of cartographic data. It enables simulation of larger gaming areas with greater accuracy without affecting display time. Adaptive cartographic models help to achieve this goal by eliminating unnecessary information such as hundreds of points representing an area of constant slope.

Different levels of detail or resolution support different applications as well as different requirements within a single application. Mixed resolution models allow greater detail to be used in areas of greater importance without increasing display and processing time. Although lower resolutions contain less data, they must represent major features as accurately as possible, retaining feature positions and overall shapes. A smooth transition between levels of detail eliminates discontinuities and prevents major features from moving, radically changing shape, or popping in and out of dynamic displays as level of detail changes. Links between the levels of detail facilitate rapid retrieval of — and smooth transitions between — multiple resolution models from the data base.

Fast spatial search, retrieval, and display are essential to real time applications. Area searches find data covering only the region specified. Location searches find information needed to animate a moving threat or target on the terrain surface. Fast neighborhood searches rapidly locate information for adjacent areas along some path. Structures facilitating fast display enable a real time application to display more polygons at one time. Triangulated geometric models, for which fast rendering algorithms and hardware are available, are excellent structures for rapid display.

Rapid, automated construction of the data model ensures that missions are not stalled for want of data. Automating the data base generation process, and designing efficient algorithms to do so, also reduces costs in terms of manpower, time, and actual price.

The problem here is to define a three-dimensional surface representation that meets these requirements, and specify techniques for fully automating its creation. Further discussion of this issue is found in chapter 2. Solutions to this problem are discussed in chapters 3 and 4.

1.1.2. Spatial data filtering

Spatial operations are typically compute-intensive due to the large volumes of multi-dimensional data that they must consider. Performance of these operations needs to be improved for real-time applications that rely on spatial data. Algorithms that use filtering techniques to prune the data, or data generalizations to approximate answers, could greatly improve performance of these operations.

The problem here is to describe how a hierarchical data structure can be used in the development of faster algorithms, showing how filtering may improve performance of the operations. This issue is addressed in chapter 5, where four spatial operations employing the hierarchical structure for data filtering are described.

1.1.3. Spatial relationships

Different spatial data representations are designed to serve different purposes. For example, vector or graph data is best for measuring distances and areas, whereas raster structures are better for fast display. Yet neither vector nor raster representations are ideal for three-dimensional modeling. Rather than force all data to conform to one data representation, merging data representations will allow applications to make use of all data representation capabilities without requiring the casual user to be cognizant of such differences. Such techniques will also be useful for heterogeneous databases, merging information from already existing databases [SSU90].

The problem here is to design algorithms for merging various representations. Chapter 6 describes how the hierarchical surface model may be merged with two popular two-dimensional data representations: vector or graph, and raster. Because most applications will need to accomplish this in real time, the goal is to find algorithms that operate quickly in reality as well as asymptotically. They must also use a minimum of additional storage space for structures facilitating the merge.

1.2. Historical Perspective

Spatial data representations have been a topic of research since the advent of the first graphics program. Diverse texts and papers on spatial data representations for specific applications — including computer graphics, computer aided

design, geographic information systems, image processing, simulation, and finite element methods — attest to the importance of this topic.

Amongst all this research there have been numerous attempts to define spatial data models constructed from critical points on the surface. Many of these are discussed in chapter 2. Yet none of these previous attempts pays adequate attention to the critical edges of the surface: ridge lines, channel lines, and break lines such as the edge of a cliff or a box.

The work described here remedies this oversight by presenting several algorithms that do pay attention to critical edges on the surface being modeled. Of particular importance are two triangulation algorithms. The first is an adaptive hierarchical triangulation that approximates critical edges as it adds detail to the model. It produces a tree structure where depth of a node corresponds to a level of detail that meets specified accuracy constraints. This tree structure supports adaptive pruning and is well suited to surface generalization and data filtering for improved performance of spatial operations. The second triangulation algorithm of importance is a curvature equalization method. This improves initial triangulations by moving vertices and edges so that small triangles cover areas of high curvature, and large triangles approximate relatively smooth patches.

This work also contributes to the literature by describing how the hierarchical triangulation may be used to improve performance of common spatial operations and data merging problems. Here the primary contribution of this work is the polygonal sweep line, based on topological sweep. Used in conjunction with the adaptive hierarchical triangulation, this search method can find all of the triangles within a polygonal region without having to examine each triangle individually.

1.3. Thesis overview

This dissertation is organized as follows. Chapter 2 provides background for this work. It discusses criteria for judging surface models, and the surface models themselves. It also addresses what surface features are considered critical, and where the data comes from.

Chapters 3 and 4 describe techniques for constructing surface triangulations. In chapter 3, the focus is generation of a hierarchical triangulation in which each level of detail retains the most critical surface information. Two strategies that consider surface topology in the construction are presented. Chapter 4 discusses a technique for improving existing triangulations by ensuring that the triangles approximate areas of nearly equal roughness.

Chapter 5 presents algorithms for spatial operations that use the hierarchical structure to filter the data for improved performance. Three representative data manipulations were selected for this task: zoom, multi-resolutions display, and line-of-sight calculation.

Chapter 6 describes techniques for merging the hierarchical triangulation with two-dimensional vector and raster data representations. These techniques also take advantage of the hierarchy for filtering, and build upon established methods in computational geometry and geographic information systems.

Chapter 7 concludes with an overview of the accomplishments described in the earlier chapters. It also outlines directions for future work in this area.